

OXFORD

context
and
coherence

the logic and grammar of prominence

U N A S T O J N I Ć

Context and Coherence

Context and Coherence

The Logic and Grammar of Prominence

UNA STOJNIC

OXFORD
UNIVERSITY PRESS

OXFORD

UNIVERSITY PRESS

Great Clarendon Street, Oxford, OX2 6DP,
United Kingdom

Oxford University Press is a department of the University of Oxford.
It furthers the University's objective of excellence in research, scholarship,
and education by publishing worldwide. Oxford is a registered trade mark of
Oxford University Press in the UK and in certain other countries

© Una Stojnić 2021

The moral rights of the author have been asserted

First Edition published in 2021

Impression: 1

All rights reserved. No part of this publication may be reproduced, stored in
a retrieval system, or transmitted, in any form or by any means, without the
prior permission in writing of Oxford University Press, or as expressly permitted
by law, by licence or under terms agreed with the appropriate reprographics
rights organization. Enquiries concerning reproduction outside the scope of the
above should be sent to the Rights Department, Oxford University Press, at the
address above

You must not circulate this work in any other form
and you must impose this same condition on any acquirer

Published in the United States of America by Oxford University Press
198 Madison Avenue, New York, NY 10016, United States of America

British Library Cataloguing in Publication Data
Data available

Library of Congress Control Number: 2020945722

ISBN 978-0-19-886546-9

DOI: 10.1093/oso/9780198865469.001.0001

Printed and bound in the UK by
TJ Books Limited

Links to third party websites are provided by Oxford in good faith and
for information only. Oxford disclaims any responsibility for the materials
contained in any third party website referenced in this work.

Contents

<i>Preface</i>	vii
1. Introduction	1
PART I. SO, WHAT IS A CONTEXT AND HOW DOES IT WORK?	
2. The Model of a True Demonstrative: Extra-linguistic Effects on Situated Meaning	23
3. An Alleged Ambiguity and the Dynamics of Context-Change	33
4. Pointing Things Out: Prominence and the Attentional State of a Discourse	40
5. Context and Discourse Conventions	58
6. Interlude: Context and Common Ground	75
PART II. CONTEXTUALIZING CONTENT	
7. Content in Context	85
8. The Challenge: Non-propositionalism	95
9. Dynamic Propositionalism	113
PART III. THE LOGIC OF PROMINENCE	
10. Content, Context, and Logic	143
11. Prominence Semantics for Modality	153
PART IV. CONCLUSIONS AND DIRECTIONS	
12. The Grammar of Prominence	171

APPENDICES

Appendix A. Formal Definitions for the Attention-Coherence Approach to Pronoun Resolution	189
Appendix B. A Formal Language for Modality with Coherence	193
<i>Bibliography</i>	203
<i>Index of Names</i>	213
<i>Index of Terms</i>	216

Preface

Natural languages are vastly context-sensitive: what appears to be one and the same string of words can convey quite different meanings on different occasions of use. Yet, we are able to understand each other effortlessly, in real time. What exactly determines the meaning on an occasion of use, and how are we so good at recovering it? This book offers a perhaps surprising response: contrary to the dominant tradition, which maintains that the meaning of context-sensitive language is at least partially determined by non-linguistic features of utterance situation, I argue that meaning is determined entirely by grammar—by rules of language that have largely been missed, and the effects of which have been mistaken for extra-linguistic effects of an utterance situation on meaning. The interpretation is effortless because it is underwritten by our knowledge of these linguistic rules.

While the linguistic notion of context I develop sheds new light on the nature of linguistic content, and the interaction between content and context, its theoretical interest extends beyond the scope of theorizing about language. Context has become a standard tool in a philosopher's toolkit. It has been appealed to in numerous debates in virtually every sub-field of philosophy: in epistemology, ethics, value theory, metaphysics, metaethics, and logic, among others. But the notion of context that these debates assume is typically the traditional, non-linguistic one. If the arguments of this book are on the right track, however, this presupposition should be abandoned and the arguments re-examined in light of the linguistic constraints on context-sensitivity resolution. It is my hope that this book will provide new tools for constraining and evaluating these philosophical debates, and for re-thinking the arguments in the new light.

This book has grown out of a project that I have been working on for a number of years. The main ideas originated in the work that led to my doctoral thesis (Stojnić, 2016). The project has developed over the years, leading to a series of publications which have dealt with specific applications of the general theoretical framework I was developing. This book is an attempt to present and defend this general theoretical framework in detail, and to articulate its key theoretical underpinnings and consequences. The book has also benefited from the feedback this earlier work has generated. It presents an attempt to incorporate and address much of this feedback I have received over the years.

Some parts of the book include, and build on, some of this previous work, although in each case, they have been modified and extended to incorporate

the feedback and address criticisms that the earlier work has provoked. Parts of Part I, specifically Chapters 2–5, include parts of the work I have done in collaboration with Matthew Stone and Ernie Lepore, and that has been published as Stojnić, Stone, and Lepore (2017) in *Linguistics and Philosophy*. I owe tremendous gratitude to my co-authors. I also thank them for allowing me to incorporate some of our earlier work in the present monograph. Any faults introduced in extending, and building on, this work are, of course, entirely my own. Parts of Part II are based on Stojnić (2019) which was published in *Noûs*, and Part III builds on Stojnić (2017b) which was published in *Philosophy and Phenomenological Research*. Parts of Part IV overlap with Stojnić (2018) in the Oxford University Press volume *Beyond Semantics and Pragmatics*, edited by Gerhard Preyer, but most of it is novel. The two appendices, A and B contain the formal systems developed in Stojnić, Stone, and Lepore (2017) and Stojnić (2017b), slightly updated to accommodate some novel examples.

I have had the opportunity to present the material included in this book—sometimes in its very early form—at various conferences, workshops, and venues over a number of years, including the philosophy departments at The Australian National University, University of Barcelona, University of Belgrade, Boston University, Columbia University, CUNY, Dartmouth College, MIT, NYU, UMass Amherst, UCSD, Oxford University, Rutgers University, University of Pittsburgh, University of Toronto, the linguistics departments at NYU, University of Maryland, and UMass Amherst, the ILLC at the University of Amsterdam, the Institute for Language, Cognition and Computation at the University of Edinburgh, the University of Connecticut Logic Group, the 6th CSLI Workshop on Logic, Rationality, and Intelligent Interaction at Stanford University, PhLiP 2016, CUNY Cognitive Science Speaker Series, and 2015 Philosophy of Language and Linguistics Conference at IUC in Dubrovnik. I have taught parts of the book in a graduate seminar in philosophy of language at Columbia University in the Fall of 2017 and the Spring of 2019, a graduate seminar on context-sensitivity at Princeton University in the fall of 2018, in a series of four guest lectures at 2017 LSA Summer Institute, and at 7th NASSLLI. I am thankful to the audiences and participants at these events for their feedback.

I have also benefited from the helpful comments on earlier written versions of parts of this manuscript, sometimes in their very early form, from Daniel Altshuler, Johan van Benthem, Elisabeth Camp, Simon Charlow, Cian Dorr, Andy Egan, Thony Gillies, Michael Glanzberg, Gil Harman, John Hawthorne, Robert May, Stephen Neale, Cory Nichols, Jeff King, Alex Lascarides, Ernie Lepore, Jim Pryor, Stephen Schiffer, Matthew Stone, Gala Stojnić, and two anonymous readers for the OUP whose comments and suggestions significantly improved the final version of this manuscript. I likewise owe gratitude to a number of colleagues for extensive discussions of the material included in

the manuscript. I would particularly like to thank: Josh Armstrong, Nicholas Asher, Chris Barker, Bob Beddor, David Beaver, Maria Bittner, Justin Bledin, Paul Boghossian, David Braun, Fabrizio Cariani, Nate Charlow, Sam Carter, Jonathan Cohen, Sam Cumming, Kosta Došen, Janice Dowell, Peter van Elswyk, Branden Fitelson, Melissa Fusco, Graeme Forbes, Kathrin Glüer-Pagin, Jeremy Goodman, Simon Goldstein, Gabe Greenberg, Alan Hájek, Dan Harris, Ben Holguín, Dan Hoek, Luca Incurvati, Andy Kehler, Angelika Kratzer, Magdalena Kaufmann, Stephen Kaufmann, Karen Lewis, Barry Loewer, Brian McLaughlin, Eliot Michaelson, Matt Moss, Sarah Murray, Ethan Nowak, Barbara Partee, Paul Pietroski, Peter Pagin, Lewis Powell, Brian Rabern, Hannah Rohde, Jessica Rett, Daniel Rothschild, Craige Roberts, David Rosenthal, Paolo Santorio, Alex Sandgren, Jonathan Schaffer, Philippe Schlenker, Ted Sider, Adam Sennet, Brett Sherman, Mandy Simons, Giuseppe Spolaore, Will Starr, Kristen Syrett, Zoltán Szabó, Rich Thomason, Frank Veltman, Brian Weatherson, Timothy Williamson, Mark Wilson, and Seth Yalcin.

Aneta Stojnić offered invaluable assistance in indexing the book. I am also thankful to two felines, Betty and Yoca, for inspiring some of the examples.

I am thankful to the editor for the OUP, Peter Momtchiloff, for his support, patience, and guidance.

Finally, I owe immense gratitude to my family: my parents, Jasna and Danko Stojnić, and my sisters Aneta and Gala Stojnić. Without their unconditional love and support, and their tireless encouragement, I would have never been able to complete this manuscript.

Una Stojnić,
Department of Philosophy,
Princeton University,
May 2020.

1

Introduction

What I communicate with ‘Give me that,’ pointing at a book, differs from what I communicate by the very same string of words when 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 expressions like ‘that’ on an occasion of use and which resources permit us to recover it so easily?

Everyone agrees that what we can communicate is to some degree constrained by grammar. However, the received view, one that might seem *prima facie* obviously correct, is that grammar underdetermines content: most theorists believe that the role of grammar is limited, since the interpretation largely relies on general reasoning about the specific communicative situations and intentions of the speakers in those situations. What we say when we speak depends not solely on the words and expressions we use, but also on the circumstances in which these words and expressions are used, and the intentions behind such uses. For instance, while it is clear that what an utterance of ‘Give me that’ means is partly determined by grammar—by the meaning of words ‘Give,’ ‘me,’ ‘that,’ and the way they are put together in accordance with the rules of English—its meaning also seems to depend on the accompanying pointing gesture toward the desired object. Indeed, it is commonly assumed that seemingly non-linguistic features of utterances, such as demonstrative gestures, play a crucial role in determining the meaning of most context-sensitive expressions. To see what’s at stake, consider the difference between an utterance of (1) and (2):

(1) I am happy.

(2) She is happy.

Both (1) and (2) express different contents on different occasions of use: if I were to utter (1), I would have uttered something different than what you would have, had you uttered it, and similarly, if I were to utter (2) pointing at Mary, I would have said something different than if I did so pointing at Sue. And while ‘that,’ ‘she,’ and ‘I’ all have a specific linguistic meaning (their *character*, to borrow the terminology from Kaplan (1989a,b)), a lexically encoded rule that specifies a referent given a context, ‘she’ and ‘that’ require something beyond their linguistic

meaning in a context in order to pick out a particular referent, in a way that 'I' doesn't. That is, the meaning of 'I' fully determines its referent as a function of context: it doesn't matter whom the speaker intends to refer to, or whom she is pointing at, or who is salient in the utterance situation, her utterance of 'I' always denotes her. No matter how much I intend to use 'I' to pick out someone else, I cannot do so. By contrast, the linguistic meaning of 'she', or 'that', does not by itself completely determine the referent given a context. *Prima facie* at least, it does not fully constrain whom the speaker can refer to by using 'she' (or 'that'), even given a specific situation. After all, she can choose to point at any given female (or thing) in the contextual situation, or exploit some other clue to make her intention to refer to a specific individual manifest. And this affects the strategies that the audience exploits in interpretation: while all it takes to figure who 'I' picks out is to see who's speaking, even if you know that the linguistic meaning of 'she' requires that, roughly, the referent be a salient female, you still need to reason about the contextual cues and speaker's intentions to determine which referent is selected on a given occasion of use.¹

The expressions that share this feature with 'I' are known as *pure indexicals*, and those that behave like 'that' and 'she' are known as *true demonstratives*.² Pure indexicals are thought to be few (perhaps only 'I'). Most context-sensitive expressions are understood to be like true demonstratives in requiring some extra-linguistic, contextual supplementation. If this is correct, then the interpretation of (most) context-sensitive expressions is underdetermined by grammar; only once their linguistic meaning is supplemented by contextual,

¹ Isn't this distinction superficial? Isn't, after all, the fact that someone is speaking a non-linguistic feature of the context? And couldn't one think of the meaning of 'she' as automatically delivering the prominent female individual, in the way that 'I' automatically delivers whoever is speaking? This superficial similarity is misleading. The point is that while it is sufficient for the speaker to utter 'I' for it to successfully pick out the speaker, it isn't sufficient to merely utter 'she' to successfully determine a specific referent; some further work is needed to determine some female as the most prominent, in the relevant sense. And this is why, the standard account maintains, one has to reason from a complex set of diverse epistemic cues toward identifying who the speaker had in mind in uttering 'she,' but not in uttering 'I'. And while it is true that it is a non-linguistic fact that I am speaking when I utter 'I', this doesn't mean that non-linguistic features of utterance situation fix the meaning of 'I' on an occasion of use: as with 'I', one has to look into the world to determine who Mary is when the speaker utters 'Mary is smart'; that Mary is who she is is a non-linguistic fact, but not the fact that determines that 'Mary' refers to Mary.

² This terminology, and the corresponding distinction between these two classes of context-sensitive expressions is due to Kaplan (1989b). There is a significant controversy over whether any expressions are genuine pure indexicals. Even the pure indexical status of 'I' has been contested (cf., e.g., Nunberg, 1993). However, even if 'I' can allow for certain kinds of non-rigid uses, such uses are constrained in ways unlike anything we witness with demonstratives. Specifically, even in such cases, the interpretation of the first-person indexical is significantly constrained by its linguistic meaning, not requiring, or even allowing, the open-ended abductive interpretive process involved in recovering the intended referent of 'she'; speaker intentions cannot determine an arbitrary referent as the referent of 'I' by exploiting the flexibility of the open-ended set of contextual cues that would reveal such an intention.

extra-linguistic mechanisms—for example, speaker intentions, or other contextual cues—do they have a specific interpretation.³ So the question becomes: what mechanisms play the role of supplementation that fixes the meaning of true demonstratives?

The most common answer is that, ultimately, it is the speaker's referential or communicative intentions that play this role. The thought is that what an utterance of (2) means depends on whom the speaker intended to talk about, whom she intended to refer to or had in mind. On this understanding, various epistemic cues—for instance, pointing gestures, or salience of an individual in the utterance situation—serve as guides toward identifying this meaning-determining intention, but it is the intention that determines the meaning.⁴ The idea is, roughly, that if the speaker is intending to refer to Mary, she should utter (2) to convey that Mary is happy, just in case in this context, there are enough resources for her audience to recognize her intention, that is, just in case her intention is made appropriately manifest. Correspondingly, when tasked with interpretation, her audience can draw on any epistemic cue the utterance situation makes available in the course of identifying the speaker's intention. But it is because the speaker intended to talk about Mary that 'she' picks out Mary.

The idea that interpretation relies on mind-reading—on figuring out speakers' communicative intentions—goes back, at least, to Locke (1689). Locke, however, thought of mind-reading as mediated solely by expressions uttered, maintaining that communication is entirely a matter of encoding and decoding content through linguistic forms. This is compatible with the meaning fully being determined by the grammar. The current theorizing, by contrast, follows in the footsteps of Grice (1975) and Kaplan (1989a), in understanding the mind-reading to

³ Which expressions are context-sensitive is a matter of controversy; some believe only pure indexicals (e.g., 'I') and demonstrative pronouns (e.g., 'she,' 'that') are context-sensitive (Cappelen and Lepore, 2005), others believe that virtually all expressions are context-sensitive (Travis, 1989). Most would agree, however, that more than just a few, but fewer than (almost) all expressions are context-sensitive. Yet, there is still great disagreement over which particular expressions are (and in what way). (See, for instance, the debates over whether modals (Egan, Hawthorne, and Weatherson 2005; Yalcin 2007; von Fintel and Gillies 2008; Kolodny and MacFarlane 2010; Dowell, 2011), or predicates of personal taste (Cappelen and Hawthorne 2009; Egan 2010; MacFarlane 2014), or knowledge ascriptions (DeRose, 1995, 2009; Cohen, 1998; Hawthorne, 2004; Stanley, 2005; Schaffer and Szabó, 2013), are context-sensitive.) However these debates are settled, most authors maintain that insofar as there are some context-sensitive expressions, most of those behave like true demonstratives in the relevant sense—that is, they require extra-linguistic supplementation.

⁴ Of course, while most authors would maintain that intending to refer to an object *o* is necessary for fixing *o* as the referent of an utterance of a demonstrative, few would think that, unless qualified in some way, this would also be sufficient. Few would bite the bullet and say that if I point to the statue in front of me saying 'That's my favorite object,' it is sufficient that I intended the Moon to be the referent of 'that,' for my utterance to mean that the Moon is my favorite object. My intention has to be in principle recognizable. Similarly, some would argue the linguistic meaning of an expression constrains what one can reasonably intend with its utterance. For instance, I cannot reasonably intend to pick out Bill with 'she,' if he is a man. For further discussion on how to qualify the requisite intentions, see, for instance, Schiffer (1972, 1981, 2005); Neale (2004); King (2014b).

involve not mere decoding, but reasoning on the basis of non-linguistic features of utterance situation, because either such features, or the intentions they reveal, partly determine the meaning of expressions. On a Lockean understanding, the speaker who intends to convey a message should choose the linguistic form that expresses that message (in a given context), and the audience, in perceiving that a particular form was uttered, is to recover the message by consulting the linguistic rules that assign the message as the meaning of the form uttered. In other words, they infer that the speaker intended the message because they uttered the form that expresses it. On the Gricean/Kaplanian picture, instead, the audience has to determine the speaker intentions first in order to even determine what the message expressed by the form uttered in the context is. It is this Gricean/Kaplanian model that has dominated the discussions about context-sensitivity resolution, and it is this model that I will challenge in this book.

The view which associates speaker intentions with a meaning determining role is, indeed, by far the most dominant one.⁵ But it is important to note that even the views that dissent from it, still insist on the crucial role of extra-linguistic factors in the resolution of context-sensitivity. For instance, the most notable alternative maintains that the non-linguistic features of context—objective, speaker-independent, aspects of the real-world situation in which an utterance takes place—are what ultimately fixes the meaning of context-sensitive expressions (Wettstein, 1984; Glanzberg, 2007). On this kind of view the facts about the real-world situation—which gestures accompany the utterance, who or what is salient, or relevant, or other kinds of real-world knowledge the agents can draw on—don't merely serve as epistemic cues that help identify referential intentions, they themselves fix the meaning of context-sensitive expressions. So, the presence of a referential intention is not necessary for determining reference: if Mary is objectively most salient (where salience is spelled out independently of speaker's mental states), then the utterance of 'she' refers to her, regardless of whom the speaker had in mind.⁶ Glossing over the nuances differentiating various versions of the dominant approach, I shall call the type of view that maintains that the speaker's referential or communicative intentions are necessary for determining the meaning of a context-sensitive expression, the *Intentionalist* view. I shall call the view that maintains that the objective non-linguistic features of context can fix the meaning of a context-sensitive expression, the *Objectivist* account. As both types of account emphasize the

⁵ See, for example, Schiffer (1972, 1981, 2005); Kaplan (1989a); Neale (1990, 2004); King (2014a,b); *inter alia*.

⁶ See, for example, Wettstein (1984), Reimer (1992), Glanzberg (2007), Lewis (2019) for discussion of this type of view. One can also hold a hybrid view: sometimes the speaker's mental states, and sometimes the objective features of the real-world situation fix the semantic content, or perhaps, different mechanisms are relevant for different types of context-sensitive expressions.

underdetermination of meaning of a context-sensitive expression by grammar, and the need for extra-linguistic supplementation, I shall dub such accounts *Extra-linguistic*.

This book argues for a novel alternative to the extra-linguistic approach. I shall argue that context-sensitivity resolution is a matter of linguistic, rather than extra-linguistic, mechanisms; it is governed by linguistic rules, which determine particular values of contextual parameters on which context-sensitive items depend for their meaning at any given point in discourse, independently of speaker intentions, *and* non-linguistic features of utterance situation. Against the tradition, I shall argue that the reason we can interpret context-sensitivity so effortlessly is because language—grammar—itself is far more expressive and pervasive than has been assumed; the resolution of context-sensitivity is entirely a matter of linguistic convention. It is the linguistic rules that determine what ‘that’ or ‘she’ picks out on an occasion of use. And, moreover, in recovering this content it is this narrow set of linguistic cues that we exploit.

The view I shall develop and defend posits linguistic conventions that have gone unnoticed, because their principal domains are entire discourses, not just their constituent words and sentences. While it is not controversial that the way sentences are constructed depends on conventions of syntax and semantics which specify the rules by which individual expressions combine, 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; I argue they also determine the resolution of context-sensitivity. The important set of mechanisms I will draw on are the so called mechanisms of *discourse coherence*, which provide an implicit organization to the discourse that signals how individual utterances are organized into a coherent whole (Hobbs, 1979; Kehler, 2002; Asher and Lascarides, 2003). To illustrate what kinds of mechanisms I have in mind, consider (3):

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

Out of the blue, speakers tend to understand an utterance of (3) in one of two ways: either Liz’s action is a result of Phil’s (perhaps conveying that she’s acting in disapproval), or her action is described as similar to Phil’s. In one case, the discourse is organized around a cause-effect (i.e., event-result) relation, describing one eventuality as resulting from the other, and in the other around a resemblance one, comparing two eventualities. This, crucially, has an effect on the interpretation of the demonstrative pronoun ‘him’: in the former case ‘him’ is understood to refer to Phil, and in the latter to Stanley. This kind of correlation between establishing coherence in discourse and the resolution

demonstrative pronouns has been well confirmed.⁷ However most extant theories treat it as a pragmatic default. Standard coherence theoretic accounts interpret this correlation as evidence of an inferential relationship between a speaker's intention in organizing the discourse and her referential intentions. I argue this is a mistake: there is a tighter connection between discourse coherence and pronoun resolution, one underscored by linguistic convention. Moreover, I will argue that this tight connection extends to other types of context-sensitive expressions, beyond the resolution of demonstrative pronouns. In short, I argue that discourse relations that connect and organize utterances are a part of the grammar of a language, and that they govern the resolution of context-sensitivity as a matter of grammar, too.

While at first glance it may seem radical, I shall argue my account is both empirically superior to the extra-linguistic alternatives and has welcome theoretical consequences. Conversely, while the extra-linguistic approach *prima facie* appears obviously on the right track, and though its endorsement is almost universal, I shall argue that it is both empirically inadequate and leads to radical and troubling consequences for the understanding of the nature of semantic content, its interplay with context, and its relation to logic.

In this regard, I will examine the recent and growing trend in the literature which advocates abandoning the standard truth-conditional, propositional accounts of meaning.⁸ The traditional paradigm, according to which the meaning of a sentence is identified with a proposition that is, or determines, a truth-condition which represents the world as being a certain way has been hugely influential. The identification of meaning with representational, truth-conditional content goes back to at least Frege (1892, 1918), and is reflected in some of the most dominant research programs in philosophical and linguistic theorizing about meaning to this day. The motivation for the departure from the truth-conditional paradigm stems from the puzzling behavior of certain context-sensitive expressions, which appear to evade the standard, extra-linguistic, contextual treatment: context, as standardly understood, fails to determine

⁷ Consult, for example, Kehler (2002); Wolf, Gibson, and Desmet (2004); Kehler et al. (2008); Kaiser (2009); and references therein.

⁸ See, for instance, Gillies (2004, 2010); Swanson (2006); von Fintel and Gillies (2007); Yalcin (2007, 2015); Willer (2014); Moss (2015); Charlow (2015, 2020); Starr (2016); *inter alia*. Most of these accounts draw closely on the influential work of Veltman (1985, 1996), which offers a formal framework, as well as a linguistic argument, for abandoning propositionalism. (I note that Moss (2015) is somewhat of an outlier on this list, as her account, while non-propositionalist, maintains that the relevant bits of discourse express informational content, modeled probabilistically. I shall discuss this nuance in more detail below.) These neo-expressivists have their precursors in the work of early non-cognitivists, such as Ayer (1936); Stevenson (1937); Hare (1952), as well as more recently Gibbard (1981), but also in the work on conditionals stemming from Adams (1975). Unlike many of these predecessors, however, they attempt to give a linguistic argument against the truth-conditional accounts.

truth-conditional content for some utterances involving such expressions. Thus, the puzzling behavior is argued to motivate non-propositional accounts which deny these utterances propositional content. Instead, these accounts tend to take the characteristic dynamic effect an utterance has on the context to be the key aspect of its meaning. The dynamic effect of the problematic utterances—*their context-change potential*—is irreducible to propositional content, and is incompatible with such utterances expressing (non-trivial) propositional meaning.⁹ At the same time, the problematic expressions *prima facie* give rise to apparent failures of certain classically valid patterns of inference, suggesting a need for a departure from classical logic. This is taken as further support for the non-propositional treatment as the various implementations of non-propositionalism depart from classical logic by invalidating at least some of these patterns.

I shall argue that both of these types of puzzling behavior that motivate non-propositionalist, non-classical accounts can be traced back to the underlying assumption of the extra-linguistic model of context-sensitivity resolution. Philosophers, operating under the traditional assumptions about the resolution of context-sensitivity, were led to abandon the propositionalist paradigm and to conclude that there is a deep tension between classical logic and the semantics of a natural language discourse, once faced with a class of expressions that do not fit such contextualist treatment. But the problem is not, I shall argue, in the idea that these expressions express truth-conditional content; the problem is in the underlying assumption of how a context operates to determine these truth-conditions. Once we have a more nuanced understanding of the linguistic mechanisms driving context-sensitivity resolution, we will see that we can explain away the puzzling phenomena while still maintaining that utterances containing the problematic expressions express non-trivial representational, truth-conditional meaning, and, even more surprisingly, we can do so within a framework that preserves classical logic.

However, while on my account utterances express truth-conditional content, that's only one aspect of their meaning. The account shares with non-propositionalists the idea that an important aspect of meaning is the dynamic effect an expression has on a context in which it is uttered. But the non-propositionalists are mistaken both about the nature and the effect

⁹ It is worth bearing in mind that non-propositionalist accounts come in different flavours. Some non-propositionalists will want to argue, for instance, that the relevant bits of discourse express a specific type of non-propositional content, rather than merely inducing a particular dynamic effect on the context (see, e.g., Schroeder (2013) or Moss (2015, 2018)). In giving this preliminary, umbrella description of the non-propositionalist approach to meaning, I am thus skirting around some important distinctions between different proponents of non-propositionalism. I will discuss those in more detail in what follows. The preliminary, tentative characterization suffices at this stage.

of this dynamic aspect of meaning. The dynamic meaning encodes the effect of discourse conventions, specifying how they manipulate and change parameters of context, in a way that determines the propositional content expressed by the subsequent discourse.¹⁰ Thus, far from being incompatible with dynamic meaning, my account has it that the dynamic meaning plays a crucial role in determining propositional content. The account thus reconciles the dynamic approach to meaning—dynamic semantics—with the traditional representational, truth-conditional approach to content.

My account has further philosophical ramifications. Philosophers frequently identify (purported) context-sensitivity in philosophically interesting expressions like ‘know,’ or ‘good,’ and appeal to it to shed light on problems involving the concepts these expressions denote, often leading to radical conclusions about the underlying phenomena. But they also assume that these expressions get their values fixed in context partly through speaker intentions, or non-linguistic features of utterance context. And often the details of such arguments rely on these underlying assumptions about what context is, and how it interacts with meaning. If my account is on the right track, then a host of debates over contextualist treatments of various philosophically important expressions rests on a flawed extra-linguistic, flexible, conception context, and context-dependence in philosophical theorizing will have to be invoked quite differently than has become customary.

I begin by briefly outlining the key theoretical concepts and considerations, emphasizing the main departures from the traditional picture and the main theoretical consequences of the view I shall develop. I will delegate detailed development and arguments for my model to the later chapters.

1.1 Context

Traditionally, context has been understood to comprise features of utterance situation upon which the content of context-sensitive expressions depends, or those which can serve as cues toward recognizing speaker intentions that play meaning-fixing role. Agents take these cues—together with their general

¹⁰ Note that when I say that a discourse expresses propositional, truth-conditional, content, I do not intend to deny that certain bits of discourse, for example, imperatives or interrogatives, do not simply denote propositions. A wide range of research in linguistics and philosophy of language points to ways of integrating the idea that semantics of mood might require distinct types of content, with the idea that ordinary assertions of declarative utterances express propositional content (irrespective of the disagreement over whether all declarative utterances express propositional content, which is what non-propositionalists deny). For discussion, see Ginzburg (1994); Portner (2007, 2012); Condoravdi and Lauer (2012); Kaufmann (2012); Roberts (2012); Charlow (2014); Murray and Starr (2018); Ciardelli, Groenendijk, and Roelofsen (2019); Starr (2020); *inter alia*.

world knowledge—as input to abductive reasoning process aiming to recover the overall most plausible interpretation. Some parameters of context provide the value of a context-sensitive expression straightforwardly and automatically—if the speaker says ‘I am happy,’ the value of ‘I’ will be the speaker, automatically. But in most cases, both which aspects of the utterance situation are relevant—intentions, salience, various objective facts about the utterance situation, or some combination of those—and how they should be combined together and weighed against one another—will require abductive reasoning, the goal of which is to recover the overall most plausible interpretation. For instance, suppose I utter (4):

(4) She is fast.

To understand (4), you will have to figure out who ‘she’ denotes. Even if you know that the meaning of ‘she’ depends on which female I intended to refer to, or which one is the most salient, you will still have to engage in reasoning as to which content this determines. To do so, you will have to take various things into account. Have I pointed at someone? Is some female referent already the unique, or most prominent, topic of discussion? Is some female referent (uniquely) salient in our surrounding? And so on. Similarly, you will have to figure out what was meant by ‘fast.’ Are we discussing a sport—swimming, or track and field? If so, are we talking about amateurs or professional athletes? Maybe we are discussing horses, or cheetahs? Or are we, perhaps, engaged in a task that would designate possession of some degree of speed significant for that task (e.g., reaching a drowning person within some limit of time)? Does the situation we are in otherwise make any one of these parameters salient, or does it make it plausible that the speaker intended a particular interpretation? These factors are complex, open-ended, and need not all point in the same direction; consequently, reasoning toward the most plausible interpretation is defeasible and open-ended.

Further, features of an utterance situation are not static—as we speak, different things can become salient, others recede in salience. The following example, from Lewis (1979) illustrates the point:

(5) [A] cat, Bruce, [...] has been making himself very salient by dashing madly about. He is the only cat in the room, or in sight, or in earshot. I start to speak to you: 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.

At the beginning of the discourse, Bruce is the most salient cat; however as the discourse unfolds and switches to the description of the New Zealand cat, Bruce recedes in salience, and the New Zealand cat becomes the most salient one. Consequently, the description 'the cat,' which denotes the most salient entity that is a cat, no longer picks out Bruce, but the New Zealand cat.¹¹

On the common understanding, tracking such features is likewise a matter of rational, abductive reasoning: the agents understand that the value of a particular parameter changes, by taking into account all sorts of cues, background knowledge, and considerations. (E.g., it only makes sense that the cat that's most recently mentioned and that we are currently discussing is the most salient when 'the cat' is interpreted, just as it makes sense that—if we have not mentioned any cats—the most salient one, or the one intended, is the one that's perceptually present in the utterance situation.) And such reasoning is by its nature holistic and defeasible.

So, on the traditional picture, the context carries both linguistic and non-linguistic information that one can equally draw on in interpreting a context-sensitive item, and that can interact in complex ways in determining the content of a context-sensitive item. The content of an utterance, thus, depends on variable and changing features of the utterance situation, encoded in such information. Of course, the standing linguistic meaning of an expression plays a role in fixing the content: that 'she' as a matter of meaning picks out the third-person, singular, female individual salient or intended, constrains who the expression can pick out in a context. But according to the traditional picture the non-linguistic information the context carries plays a key role in fixing the meaning, and so—on the interpretive side—the audience will have to engage in holistic, abductive reasoning about such features in order to recover the content.

While the traditional picture might be *prima facie* natural, it is, I shall argue, mistaken. According to the account of context and context-sensitivity I shall develop, the features of context that fix the meaning of context-sensitive expressions are *linguistically* determined, through the effects of *discourse conventions*, which trigger semantic effects of the linguistic items an utterance—or more generally the discourse in which it is embedded—comprises. The context, I shall argue, is a running record of *linguistic* information that is contributed by discourse-internal, linguistic cues; the record, in turn, fully determines the interpretation by fixing the content expressed by the discourse.¹² My conception of context is thus *linguistic*, rather than extra-linguistic.

¹¹ The complexities concerning the semantics of definites, though important in their own right, can be set aside for the purposes of this illustration.

¹² The idea of a context as a conversational record, or scoreboard, goes back to Lewis (1979). But while Lewis maintains that the context is an evolving record of contextually relevant parameters, he is not committed to the idea that the record is maintained and updated entirely through linguistic mechanisms.

Given this idea of context, a context-sensitive item receives its interpretation fully and automatically from the linguistic context, selecting the value of the relevant contextual parameter that the context records at a particular point in discourse at which the item occurs. To illustrate the general idea, consider the following discourse:

(6) Betty came in, and she sat down.

The content of the pronoun 'she' in (6), is not determined independently of what came before. The first conjunct sets the stage: it mentions a female antecedent, Betty, that the pronoun can then pick up on. On the traditional picture, we could explain this by saying that mentioning someone makes them salient, and so that 'she' naturally picks out Betty.¹³ But suppose that while (6) is uttered, a distinct female individual, Mary, is the most salient: she is jumping up and down, yelling loudly, and clearly making herself the center of our attention. Even so, 'she' is naturally understood as selecting Betty.¹⁴ How do we explain this? I shall argue it is a part of the linguistic contribution of (6) that it forces this particular interpretation: it is a part of the standing meaning of a proper name, 'Betty,' that it affects the context so that its referent is made prominent; and it is because (6) exhibits a narrative structure that Betty remains thus contextually prominent for the interpretation of the subsequent pronoun 'she.'

A discourse can, of course, be ambiguous with respect to its structure. Recall (3):

(3) Phil tickled Stanley. Liz poked him.

The discourse in (3) is ambiguous: it could be interpreted so that Liz poked Phil, or so that she poked Stanley.¹⁵ But the ambiguity, I shall argue, does not have a source in (potentially open-ended number of) possible contextual resolutions of the pronoun; rather different interpretations of the pronoun have their source in different coherence relations that connect the two sentences into a coherent

¹³ Admittedly, this is just a basic, bare-bones version of the traditional account. We shall see, however, that even the more sophisticated ways of spelling things out within the traditional framework mischaracterize the mechanisms of context-sensitivity resolution: it is linguistic, rather than extra-linguistic mechanisms that play the meaning-fixing role.

¹⁴ I am assuming that in (6), the pronoun is not stressed, or accompanied by a deictic gesture. One could get a reading according to which 'she' in (6) selects Mary, if the pronoun was appropriately stressed. But prosodic stress (and, I shall argue, deictic gestures), are precisely linguistic, language-specific, elements, rather than extra-linguistic bits of information that serve as inputs to abduction.

¹⁵ Again, I am assuming no pointing gestures are present in (3).

unit.¹⁶ If the discourse harbors Parallel, a relation signaling comparison between the events described by the two sentences, then the pronoun is resolved to Stanley, the object of the previous sentence, whereas if it is organized by Result relation, signaling that event described by the second sentence resulted from the first one, the pronoun is resolved to the subject of the previous sentence, Phil. These relations, I shall argue, are a part of the discourse in (3)—much like the proper name, ‘Phil,’ or the past tense, are—and, further, they dictate a particular resolution of the pronoun, by affecting the context in such a way that one or the other candidate referent is deemed the most prominent referent in the linguistic context in which ‘him’ is interpreted. In this way, (3) (and similarly, (6)) lacks dependence on non-linguistic context: the interpretation of the pronouns they feature is fully determined by the pronouns’ local linguistic context, which in turn is determined by the linguistic items internal to these discourses.

More generally, then, while on the standard picture an expression has a standing linguistic meaning which together with (potentially extra-linguistic) contextual information helps fix its content on an occasion of use, on my account expressions, as a part of their standing linguistic meaning, create a linguistic context that determines the content expressed by subsequent context-sensitive expressions. Somewhat more precisely, while on the standard picture expressions are assigned a character—which in a context delivers a content—on my account expressions are interpreted as instructions to update the context, potentially changing the values of contextual parameters from the input context to the output one, reflecting the changes brought about by the meaning of each expression. This sets the values of contextual parameters which help fix the semantic, truth-conditional content of an utterance just as they would on the standard picture, by consulting the linguistically assigned character of each expression (e.g., given that, per its character, ‘she’ picks out the third-person, singular, female individual, an utterance of ‘she’ will select the most prominent element satisfying this character as recorded in the current state of the conversational record). However, since the contextual parameters are linguistically maintained, the semantic content expressed is linguistically determined, through and through. In this way, any contextual effects are those of this linguistically maintained context, traced back to the effect of some linguistic item within the discourse.

Notice that, on my account, discourses are rich: they comprise not just individual sentences (along with all their overt and covert elements), but also discourse relations—coherence relations between individual sentences. While this might appear strangely unparsimonious, I shall argue that there are good

¹⁶ See Hobbs (1979); Kehler (2002); Asher and Lascarides (2003); *inter alia*). Much more will be said about discourse coherence in subsequent chapters.

empirical reasons to embrace this richness.¹⁷ Each of these linguistic elements, in turn, is associated with an instruction on how to update the conversational record, thus gradually building the linguistic context which determines the interpretation. It is such instructions that I call *discourse conventions*. Discourse conventions determine the *prominence* of potential interpretations of context-sensitive items: they underscore the grammar of prominence. Notice, the notion of prominence I argue is relevant for the interpretation of context-sensitive expressions is linguistic; as we have seen in our discussion of (6), it is distinct from, and cross-cuts, the worldly salience emphasized by the traditional approach.

1.2 Content

Ultimately, we want to explain how we communicate so effortlessly with context-sensitive language. We want to explain how it is that we are able to convey our thoughts through language, and how the information conveyed through communication can come to guide our action. For instance, suppose I want to convey to you my thought that it is currently raining in NYC. I utter (7):

(7) It is raining (in NYC).

Upon hearing my utterance, you come to understand it, and if you take me to be reliable and sincere, you might come to believe what I said, and this might impact your action: you might decide to bring an umbrella. A similar exchange might happen if I utter, for instance, the following:

(8) That is your umbrella.

while pointing at the umbrella by the door. How do such exchanges succeed? A simple idea is that my thought, and my utterance, express a content, which represents the world as being a certain way, for example, as such that in it, it is raining in NYC (at a particular time), or that the object I pointed at is the addressee's umbrella. Both my thought and my utterance are true just in case the world is indeed this way. It is because you understand that my utterance expresses this content—and understand what it would take for my utterance to be true—that you can understand my utterance and, if you deem me reliable and sincere, you can come to believe what I asserted. This might further impact your

¹⁷ In doing so, I shall build on the literature on discourse coherence. See Hobbs (1979, 1990); Kehler (2002); Asher and Lascarides (2003); Webber et al. (2003); *inter alia*.

action (you might decide to grab your umbrella). Propositions are taken to be (or at least determine) this truth-conditional, representational content.

But if my utterance of (7) or (8) expresses a proposition, then the question is how is this proposition connected to the semantic content of the sentence in (7), that is, (8). The semantic theory, of course, should play some role in determining the content expressed. Presumably, that I manage to convey my thought with (7), or (8), has something to do with the semantic content that the correct semantic theory for English assigns to (7), that is, (8) (relative to a context).¹⁸

Now, the traditional picture, we have seen, maintains that the proposition expressed by an utterance such as (8) isn't fully determined by grammar; the input of extra-linguistic features of the utterance situation is needed to determine (at least) what 'that' denotes.¹⁹ But even so, one could subscribe to a simple and direct way of connecting the compositional semantic content with the content expressed embodied in what one might call the *identification thesis*. The identification thesis maintains that the compositional semantic content of a sentence (of a language) in a context just is the content asserted by it in that context; that is, that the compositional semantic content of a sentence (of a language, on an occasion of use) just is its so-called *assertoric content* (on the same occasion of use). Indeed, traditionally, identification has been assumed, where both the compositional semantic content (of a sentence) and the assertoric content have been identified with the notion of *propositional* or *truth-conditional* content. Propositions are what we believe and assert; and they are the compositional semantic content sentences express relative to contexts.

Assuming that grammar underdetermines the resolution of context-sensitive items and that the input of the extra-linguistic context is needed to determine the proposition expressed, how can one maintain identification? The idea is that one can relativize assignments of semantic content to expressions to contexts. A semantic theory for a language $\mathcal{L}$ assigns a compositional semantic content to each expression of $\mathcal{L}$ and specifies the recursive rules whereby these contents are compositionally combined to form contents of larger expressions of $\mathcal{L}$. But to handle context-sensitivity, one maintains that compositional semantic content is assigned only relative to a context. In this way, the input of the non-linguistic context plays a role during semantic composition, and thus context-sensitivity can be resolved during the course of semantic-composition, so that in a context, a sentence expresses a compositionally determined proposition.²⁰

¹⁸ See Stojnić (2017a) for a more detailed discussion of this point.

¹⁹ Indeed, the same can be said of (7), if it, too, harbors context-sensitive elements: for instance, if we adopt a contextualist treatment of tense.

²⁰ Not everyone would agree that context-sensitivity should be accommodated by relativizing *semantic* content to contexts. Some authors would maintain that semantic content is rather minimal and (largely) context-insensitive, and that the apparent contextual-sensitivity is due to non-semantic

Identification would thus provide a straightforward account of the role of semantic content in the theory of communication. Unfortunately, it has been forcefully criticized, as propositional content turned out to be inadequate for the role of compositional semantic content, due to its failure to account for all the embedding behavior of compositional semantic content.²¹ But even among the theorists that reject identification, most would still expect that semantic content somehow determines the informational, propositional content, the content of assertion. Instead of identification, they subscribe to *determination*, the thesis that semantic content (at least partially) determines the content of assertion and belief—that is, that the latter can be recovered from the former.²² One then has to spell out how the propositional content is recovered. The broad range of research on context and context-sensitivity has focused on spelling this out. In spelling it out, context has been presumed to play the crucial role, where

effects of context (Bach, 1994). Others argue that there's relatively little interest in positing the notion of semantic (as opposed to pragmatic) content altogether (Recanati, 2004). Such authors, *ipso facto* wouldn't endorse *identification*. For discussion, see King and Stanley (2005).

²¹ For criticisms along these lines, see Evans (1979); Lewis (1980); Kaplan (1989b); Dummett (1993); Stanley (1997a,b); Ninan (2010); Rabern (2012); *inter alia*. See King (2003); Stojnić (2017a) for further discussion. A simple way to see the point is as follows (cf. Lewis (1980); Kaplan (1989b)). Consider the sentence 'John has five coins in his pocket'. According to the standard account, such a sentence is true in a context *c*, just in case it expresses a true proposition in *c*. Since propositions—the kinds of things we believe and assert—are true/false relative to worlds (they represent the world as being a certain way), this sentence is true in *c*, just in case the proposition it expresses in *c*, namely *that John has five coins in his pocket*, is true (at a world in which it is uttered). Now, consider 'It is possible that John has five coins in his pocket'; the truth of this sentence (in a context *c*) depends not on whether John has five coins in his pocket in the world of utterance, but on whether there is *some* (accessible) possible world in which he has five coins in his pocket (after all, that it is possible that he has five coins in his pocket, doesn't mean that he does). One can account for this by treating 'It is possible that' as an intensional operator, which takes the compositional semantic content of the embedded sentence—that *John has five coins in his pocket*—and delivers a proposition that is true in all the worlds *w* such that in some world *w'* accessible from *w*, John has five coins in his pocket. In other words 'it is possible that' shifts the world of evaluation for the proposition expressed by the embedded sentence. But this only works if the content of the embedded sentence is sensitive to worlds of evaluation; otherwise, the contribution of the operator would be vacuous. Now, suppose there were operators that shift not the world of evaluation, but some other parameter on which the truth of the content of the embedded sentence depends. For instance, suppose that 'it's always the case that' in 'It's always the case that John has five coins in his pocket' is an operator that takes the content of the embedded sentence and checks whether it is true relative to *all* times, not just the time of utterance. For this to work, the content expressed by the embedded sentence has to be true/false relative to times, not just worlds: if the content were only sensitive to worlds, then the contribution of the operator 'always' would be vacuous. But if this is right, then the content of the sentence cannot just be a proposition traditionally construed, since propositions traditionally construed do not vary in truth-values relative to times. I do not mean to endorse that 'always' is a temporal operator; whether it is or isn't is a matter of controversy. But the recent literature suggests English contains at least some intensional operators beyond the world-shifting ones, be it modals, antecedents of conditionals, quantifiers, temporal operators (Yalcin, 2007, 2015; Ninan, 2010; Rabern, 2012). These arguments seem to suggest that propositions—traditionally construed—cannot play the role of compositional semantic content.

²² See Lewis (1980); Kaplan (1989b).

it is again the extra-linguistic model of context, the one that I challenge in this book, that has been largely assumed.²³

At the same time, a different, potent challenge to the traditional view of content and context-content interaction has emerged in recent literature. The challenge comes in the form of *non-propositionalism*, which maintains that not only are propositions not suited to play the role of compositional semantic content, but they are not even suited to play the role of the assertoric content. Non-propositionalists challenge both identification and determination as typically construed. They do so insofar as they maintain that, for at least certain kinds of (declarative) utterances—notably, modal claims—their compositional semantic content does not determine propositional informational or assertoric content, because such utterances have no assertoric, informational content to begin with.²⁴

My account provides an intervention in these debates from two perspectives. First, I argue that the contextualist accounts have missed discourse conventions that guide the resolution of context-sensitivity. As described, on my account,

²³ Assume again, for the sake of the argument, that ‘It is always the case that’ is a temporal operator. Then, the compositional semantic content of sentences has to be true/false relative to worlds *and* times, not just worlds. But one can recover a proposition from this content by supplying the time of the context as the time of evaluation. So, in a context *c*, ‘he [pointing at John] has five coins in his pocket,’ expresses a proposition that is true at a world just in case the individual pointed at, that is, John, has five coins in his pocket at that world, and *at the time of c*. Notice that here, the context helps determine the proposition expressed in two ways. First, it fixes the meaning of context-sensitive elements the utterance harbors—for example, determining that ‘he’ denotes John—and, second, it provides values for the parameters in the point of evaluation that are needed to recover propositional content—here, by fixing the time parameter in the point of evaluation to be the time of context. Finally, some would add a third way in which a context can help determine the assertoric content: through free pragmatic enrichment of semantically determined content, unconstrained by the presence of context-sensitive elements in the logical form of an utterance (e.g., Sperber and Wilson (1986); Carston (1988, 2002); Bach (1994); Stainton (1994); Recanati (2004), *inter alia*). Interestingly, the extra-linguistic model has been endorsed both by the proponents of free enrichment, as well as those who trace *every* effect of context down to some (context-sensitive) element in the logical form of the utterance (e.g., Stanley and Szabó (2000); Taylor (2001); Stanley (2004); King and Stanley (2005); *inter alia*). On either view, *non-linguistic features* of context are operative in fixing the asserted content.

²⁴ Cf. Yalcin (2011b, 2015, 2012b); Willer (2013); Bledin (2015); Charlow (2015); Starr (2016); *inter alia*. Note that, as mentioned earlier, not all non-propositionalists deny identification: some maintain it, but only by rejecting the idea that assertoric content is ordinary *propositional*, representational content. See, for instance Kolodny and MacFarlane (2010); MacFarlane (2014); or Moss (2015, 2018). Similarly, Schroeder (2013) argues that while we have good reasons to think that a single notion of content plays a role of compositional content and of the objects of assertion and belief, and is a bearer of truth and falsity, we should reject the idea that this notion of content also plays the role of representational content, that is, that it plays a role in “carving up the world at its joints,” and is “associated with metaphysical commitment, and [is] the appropriate [object] of excluded middle” (Schroeder, 2013, p. 410). But Schroeder too maintains that the content of attitudes isn’t propositional in the ordinary sense (albeit he calls such objects ‘propositions’). Thus, all these views work with non-standard notion of assertoric content. As it will become apparent they likewise face the challenge of accounting for the linguistic data I discuss in Parts II and III. I shall discuss these issues in more detail in what follows.

expressions are interpreted as instructions to update the conversational record. This layer of semantic content encodes discourse conventions, which determine the state of the record at any given point. Discourse conventions specify an effect that each expression has on the conversational record during the course of semantic composition, and, since they determine the values of the contextual parameters that fix the meaning of context-sensitive items, they thereby build the truth-conditional, propositional content expressed by the discourse. In this way, on my account, semantic content, though not reducible to propositional, assertoric content, determines this content during the course of compositional derivation. Notice, contrary to the traditional picture, where the compositional semantics underdetermines assertoric content, the latter only fully determined with the aid of extra-linguistic supplementation, on my account, semantic content fully determines the propositional assertoric content, through a process of semantic composition. Thus, while my account does not subscribe to identification, it vindicates determination thesis in its strongest form, much stronger than traditionally assumed.²⁵

Second, we shall see, my account challenges the non-propositionalist attack on determination: I argue that once the effects of discourse conventions have been taken into account, and adequately captured, the seeming non-propositionality of modal discourse turns out to be an illusion. The account delivers ordinary propositional content even for discourses that *prima facie* seem to evade propositionalist treatment. More importantly, we shall see that a broader range of data suggests that such propositional content is needed to properly account for the range of interpretations such discourses allow. In this way, my account offers a novel and, I argue, empirically more adequate account

²⁵ As an anonymous reader points out, Yalcin (2015) suggests that the questions concerning the state of knowing the compositional semantic content of an expression, and the capacity that state underwrites, is quite different from the question of how this knowledge, and the relevant capacity, are employed in communication (if at all) to transfer informational content. That is, the knowledge of a sentence's semantic content and its assertoric content might be two very different kinds of knowledge. Note that I do not need to dispute this. On my own account, expressions have very rich compositional semantic contents—the dynamic meanings—which serves as an instruction for building propositional content. While I defend determination, I do not defend identification. Hence, the knowledge of the relevant linguistic conventions that serve to build propositional content of a sentence, or a clause, on an occasion of use can be distinct from the knowledge of the proposition actually built on a given occasion. I do however make a concrete proposal as to what role this rich compositional semantic content plays in determining informational content, and, so, a concrete proposal as to its role in a theory of communication. Now, of course, non-propositionalists like Yalcin would deny this vindication of determination is really desirable, since they deny the usefulness of the notion of assertoric content (or at least the assertoric content understood in terms of propositional content), to begin with. Yalcin (2015) in particular, argues that no useful notion of assertoric content need transpire in an adequate account of communication, and that, indeed, such a notion breaks down for at least some bits of ordinary (declarative) discourse. However, I precisely argue that such notion does not break down in the relevant cases, and, moreover, that an adequate account of communicative import of utterances for which it allegedly breaks down crucially appeals to propositional content. Hence, I submit, vindicating determination is desirable.

of the interaction between semantic content and assertoric, communicative content, one that presumes neither identification, nor underdetermination.

With these preliminary theoretical considerations in mind, we can now turn to the development of my view. The book is organized in four parts, followed by Appendices. In Part I, I outline dominant theories of context and context-sensitivity resolution in more detail, and situate my account in opposition to such theories, by focusing on the basic case of context-sensitivity—that of a demonstrative pronoun. In Part II, I develop in detail my theory of content and context interaction, focusing in particular on a specific challenge to the truth-conditional, representational accounts of content stemming from the puzzling behavior of modal discourse which seems to evade truth-conditional, propositional treatment. I show that once a proper account of dynamic effects of discourse conventions on context is captured, these effects will in turn determine a truth-conditional content expressed by constituent utterances. Specifically, we will see that dynamic semantic meaning, properly characterized, serves to compositionally determine propositional, representational meaning.

In Part III, I turn to the issue of the relation between natural language and logic. Recent literature posits a seeming tension between natural language discourses and classical logic. It has been argued that various classically valid inference patterns fail for certain fragments of natural language. At the same time, the apparent violations of certain classical patterns of inference have been propagated as another piece of evidence for alternative, non-representational, non-propositional semantics, which tend to invalidate some of these patterns.²⁶ I trace the seeming failures of classical inference patterns to the underlying assumption of the extra-linguistic model of context-sensitivity, arguing that once the effects of discourse conventions on the interpretation of context-sensitive language are taken into account, we will be able to explain, and, indeed, predict the appearance of putative counterexamples to classical inference patterns, while maintaining a semantics that preserves classical logic. At the same time, the account will offer a precise way of individuating logical forms of arguments presented in natural language discourse, one that doesn't require an artificial ban on context-shifting, but that also allows us to identify shifts in context that lead to equivocations.

Finally, in Part IV, I draw general conclusions for theorizing about the interaction between context and content in communication, and discuss directions for further extensions and applications of the framework.

²⁶ See McGee (1985); Veltman (1985); Gillies (2004, 2010); Yalcin (2007, 2012b); Kolodny and MacFarlane (2010); Willer (2013); Bledin (2015); Charlow (2015); Moss (2015); *inter alia*.

In developing my account in Parts I–III, I will primarily focus on two classes of expressions: on demonstrative pronouns and modal discourse. This choice is not accidental. One of the key arguments of the book can be seen as follows: the traditional understanding of context-sensitivity has it that most context-sensitive expressions behave like true demonstratives insofar as their meaning is partly fixed by extra-linguistic features of utterance situation. I argue this is a wrong view even for the paradigm case of true demonstratives—that of demonstrative pronouns. But not only is this view wrong even in the most basic case, its assumption leads to radical conclusions in philosophy more generally. The case of modality is particularly relevant and important here, as it has been at the center of the discussion about context-sensitivity where such radical conclusions have been drawn. In particular, considerations involving the interaction between modals (and conditionals) and context have been the cornerstone of the recent rise of non-propositionalism, as well as at the center of the debate concerning the relation between semantics for natural languages and logic.²⁷ I argue that the arguments within this debate have implicitly presupposed the extra-linguistic account of context, and that, once this presupposition is rejected, the arguments for non-truth-conditional accounts as well as the counterexamples to classical patterns of inference dissolve.

On my account, there is a deep connection between demonstrative pronouns and modality. Modals are like pronouns in that they are prominence-sensitive—sensitive to the prominence of a contextually supplied interpretation. This prominence is governed by linguistic rules—discourse conventions—and its underlying semantics can be shown to support a well-behaved, classical logic. In other words, it is underwritten by, and induces, a systematic grammar and logic. This is why it is particularly instructive to focus on modals: not only are they,

²⁷ Note, I do not argue that failures of classical forms of inference arise only for non-propositionalist semantics. A propositionalist semantics *can* invalidate the relevant classical patterns of inference, for example, *modus ponens* (see McGee (1985); see also Khoo (2013) for some discussion); and in the light of the data, we shall see it is in fact tricky to provide one that does not. Similarly, I'm not arguing that *in principle* one could not have a non-propositionalist account of meaning with suitable semantic definitions that would yield classical logic (although, this might depend on how permissive we are disposed to be regarding what propositions are). But a certain battery of data (specifically, a battery of data that allegedly shows that modal utterances, as a matter of empirical fact, fail to express truth-conditional content, or that they, as a matter of empirical fact, give rise to failures of classical patterns of inference) has been used in support of a kind of a package deal—a particular kind of non-propositionalist semantics along with a non-classical logic (a semantics broadly inspired by the tradition in dynamic semantics stemming from the work of Veltman (1985), as well as a certain understanding of context/content interaction going back to Stalnaker (2002, 1978)). I argue that this package deal is not supported by the data, and, moreover, that looking at a broader range of data, which uncovers the systematic effects of discourse conventions of the kind I postulate, we have reasons to favor a very different kind of package deal—a propositionalist account of meaning of the sort I develop and a classical semantics.

like pronouns, prominence-sensitive—but a failure to appreciate the import of contextual prominence has led to radical conclusions regarding the grammar, meaning, and logic of modal discourse.

It is, perhaps, worth pointing out that a different way of developing my project would proceed by a series of case studies, or applications, to various context-sensitive items, or by looking at a series of applications within philosophy (apart from the cases I presently discuss) where the traditional assumptions about context play a role in arguments for philosophically surprising or radical conclusions (in for example, epistemology, ethics, metaphysics). But such an approach would, I suspect, make the book prohibitively long, and at the same time, it would blur the key theoretical results concerning the import of discourse conventions both on meaning and on underlying logic of natural language discourse, which are highlighted in the debate about modality. I do, however, outline how my account is to be applied to various other cases, and discuss theoretical tasks and challenges for such application in Part IV.

PART I

SO, WHAT IS A CONTEXT AND
HOW DOES IT WORK?

2

The Model of a True Demonstrative: Extra-linguistic Effects on Situated Meaning

Since, intuitively, the meaning of an utterance seems to vary with utterance situation—one can express very different meanings with what appears to be one and the same string of words on different occasions of use—it is in a sense natural to understand context as comprising circumstances of an utterance situation that might affect its meaning and interpretation: the fact that it was uttered in a particular situation, to a particular audience, against a particular body of common background information, as a part of a particular discourse. Traditionally, these circumstances were understood as utterance-external, extra-linguistic conditions, and consequently, different theoretical notions of context were designed that purport to model different ways in which such conditions, understood as utterance-external, affect the interpretation.

A notion of context familiar from Kaplan (1989a,b) was designed to capture this type of dependence, whereby an expression's interpretation varies with some parameter of a real-world situation in which the utterance was made. A Kaplanian context is a sequence of parameters—representing features of the real-world situation in which an utterance is taking place—on which the interpretation of the utterance can depend: a world in which the conversation is taking place, the speaker of the utterance, the time and location of the utterance.¹ For instance, when one utters (1), repeated below, the indexical 'I' is interpreted as selecting the value of the agent- or speaker-parameter of the context, thus fixing the truth-conditions of the utterance. (1) is true in a context *c*, and a world of evaluation *w*, just in case whoever is the speaker of *c*, say *c*_s, is happy in *w*.²

¹ Of course, other parameters might be required. With this in mind, Lewis (1980, p. 79) writes that a Kaplanian context is "a location—time, place, and possible world—where a sentence is said. It has countless features, determined by the character of the location." Similarly, in "Afterthoughts," Kaplan maintains that "context is a package of whatever parameters are needed to determine the referent, and thus the content, of the directly referential expressions of the language" (Kaplan, 1989a, p. 591).

² The idea is that the context determines the proposition expressed by a sentence on an occasion of use. Then we want to say that a sentence is true in a context *c* just in case it expresses a true proposition in *c*. And assuming that propositions are true and false relative to worlds, we say that a sentence expresses a true proposition in a context, just in case the proposition it expresses in that

(1) I am happy.

The standing linguistic meaning of ‘I’—its character—can be modeled as a function that takes a context, and delivers the value of the speaker-parameter of the context: a function from context to content.

The Kaplanian model of context makes some substantive idealizations. In particular, it interprets an utterance against a single, static, unchanging context. Yet, the utterance situation evolves as we speak, the values of its parameters constantly changing. This change may affect the interpretation of what we subsequently say. For instance, if I follow your utterance of (1) with (9), I won’t thereby express a proposition that *you* are not happy, contradicting your previous statement.³ The context has changed between our utterances—once I spoke up, *I* became the new value of the speaker parameter:

(9) I am not happy.

One might think that at least each sentence occurs within its own unique context. But things can happen in the real-world situation as we speak, in a way that

context is true at the world of that context, that is, the world in which the utterance takes place. Note, however, that the truth of a sentence is officially relativized not just to a context, but also to another parameter, a world, which is sometimes called an *index* or a *point of evaluation*. Recall, as noted in footnote 21 of Chapter 1, this relativization to an index is needed because the truth of some sentences depends not simply on the world of the context, but on the truth of some *other* sentence in that context, relative to a world *other* than the world of the context. The idea is that propositional content isn’t merely playing the role of what is expressed by a sentence in a context, but is also the compositional semantic content that a sentence contributes compositionally to larger expressions it can be a part of. For the original motivation for double indexing, that is, relativizing truth to not only a single parameter, a context, but also an index, see Kamp (1971); Kaplan (1989b). Whether an index contains only worlds, or perhaps has to include some other parameters, such as a time, or a location, is a matter of controversy having to do with the properties of compositional semantic content. For now, we can bracket these issues, as they are orthogonal to our present concerns. What matters for our purposes is that given this specification of truth of (1) relative to a context and an index, we can then determine the truth of (1) in a context *simpliciter*, by saying that (1) is true in a context *c*, just in case whoever is the speaker of *c*, say *c*_s, is happy in the world of the context, that is, the world in which the utterance actually takes place, *w*_{*c*}.

³ Note that Kaplan distinguishes between the context as an abstract formal object, and the context as a representation of a real-world spatio-temporally extended situation, the latter notion belonging to what he calls the “speech act theory.” He defines truth and validity for sentences-in-context, rather than utterances, which are the real-world events that take time, as he maintains that validity requires a stable context across a sentence or an argument, while utterances, taking time, might involve a context-change. The worry is that one wouldn’t want to treat ‘Now is identical to now’ or ‘Today is identical to today’ as a counterexample to the law of self-identity just because utterances take time, so that the value of ‘today’ or ‘now’ might shift by the time the speaker tokens it the second time. So, while Kaplan is clearly aware that, in a sense, context changes as we speak, he believes that the *logic* of demonstratives requires a method of keeping the context constant (Kaplan, 1989b, p. 91). We will see, however, that in order to adequately analyze argument patterns in a natural language, we will precisely have to allow for a (systematic) change in context between sentences and even within a single sentence. This issue will be discussed in more detail in Part III.

affects the interpretation even during the course of an utterance of a single sentence. We can see this if we look at an utterance of a sentence with multiple occurrences of a true demonstrative, as in (10):

(10) He [pointing at Bill] is sad, because he [pointing at Tim] is leaving.

(10) features two occurrences of a context-sensitive expression, 'he,' and the two occurrences receive different resolutions—the first one picks out Bill, while the second one Tim. Intuitively, this has to do with the fact that the two occurrences are accompanied with two distinct pointing gestures, pointing at different individuals. So, if a single, non-changing context is operative in determining the meaning of (10), and if 'he' has a uniform linguistic meaning, how is it that the two occurrences of 'he,' receive different interpretations relative to the same context? In other words, how are true demonstratives to be incorporated into the Kaplanian model of context?

The basic form of Kaplan's answer is that the extra-linguistic supplementation (the demonstration) somehow combines with the linguistic meaning of a demonstrative to determine its content. One way to capture this is to model the linguistic meaning of a demonstrative, 'he,' as follows:

Definition 2.1 Where α is a singular expression, $he[\alpha]$ relative to a context c picks out the individual o iff o is denoted by α in c (and o is a male individual).⁴

$he[]$ is a rigidifying expression which combines with some singular term in a context, to fix its content. The singular term it combines with, in turn, is contributed by extra-linguistic mechanisms. Kaplan's idea is that demonstrations—for example, the pointing gestures—provide a descriptive content, a mode of presentation for the demonstrated object. For instance, in (10), the first pointing gesture provides some descriptive content, say, 'the individual the speaker is pointing at,' so 'he' as a result, in this context, picks out Bill. The second occurrence of 'he' in (10), accompanied by a distinct demonstrative gesture, receives a different interpretation. On this view, a demonstration—the act of pointing—plays the meaning-fixing role. Of course, not every use of a demonstrative is accompanied by a demonstration in any natural sense of the term, which is why Kaplan ultimately settled with the Intentionalist account: it is ultimately the speaker's intention which supplements the meaning of a demonstrative expression by providing the relevant descriptive content; the demonstrative gesture on this account, merely serves as an external manifestation of this

⁴ This account specifies the meaning of 'he' on the model of the meaning of *dthat*, Kaplan's artificially coined rigidifying expression. Kaplan defends this proposal in Kaplan (1989a).

intention, but one that plays no meaning-determining role (Kaplan, 1989a).⁵ The fact that the speaker demonstrated something only serves as an epistemic cue that can help with intention recognition.⁶

How does one understand the character—the linguistic meaning—of a demonstrative on this model? And, in particular, how does such a meaning interact with the context, so that multiple occurrences of a demonstrative within a single utterance can receive distinct interpretations? For a singular term α , the character of $he[\alpha]$ is just a function from contexts to the denotations of α in those contexts. So, for instance, the character of $he[the\ speaker]$ is just the same as the character of 'I'. On one natural understanding of this proposal, this predicts an ambiguity in the linguistic meaning of a true demonstrative: 'he' is represented with different expressions on different occasions of use, and is associated thus with different characters. And this ambiguity transpires in examples like (10).⁷ Specifically, if we do not want to maintain that the context changes between every two occurrences of a demonstrative, so that, for example, in (10), the two occurrences of 'he' are interpreted in different contexts, we have to allow that distinct occurrences of a demonstrative within a single context can be paired with different demonstrations. Since the demonstration partly determines the character of a demonstrative, the two occurrences of 'he' in (10) have different characters.⁸

An alternative way of incorporating demonstratives in Kaplan's account makes the ambiguity even clearer. On this approach, we incorporate the idea that a single context can contain multiple different demonstrations by expanding the context to include in it a (potentially infinite) sequence corresponding to distinct demonstrata, associated with distinct demonstrations, or distinct meaning-determining intentions. Correspondingly, instead of a single demonstrative, we now introduce a sequence of demonstratives, ' he_1 ', ..., ' he_n ', and associate

⁵ This view is the orthodox Intentionalist view of the sort we discussed in Chapter 1. Although Kaplan's view in "Afterthoughts" was constrained to perceptual demonstratives—he maintains that the directing intention that supplements the meaning of a demonstrative is the intention to demonstrate or refer to the *perceived* object—it is natural to extend his view, relaxing this requirement to allow for referential intentions more generally to play the meaning determining role. For this move, see, for example, Bach (1992a,b); King (2014a,b). For further discussion, see Reimer (1992).

⁶ As discussed in Chapter 1, one might endorse a hybrid view where both demonstrations and intentions can play a meaning determining role. (See, e.g., Reimer (1992).)

⁷ Although Kaplan himself maintains that demonstratives are ambiguous at least between "referential" and "bound" uses (more on this below), the type of ambiguity of relevance here is more radical—different occurrence of a demonstrative, regardless of whether the use is bound or referential, come with a new representation and character. Kaplan calls this the "exotic ambiguity," and deems it "harmless," but he also notes that this treatment raises difficulties for understanding the notion of validity. See Kaplan (1989a, p. 590). We will discuss concerns regarding validity below and, in more detail, in Chapter 10.

⁸ See Braun (1996) for discussion of this point.

each with a character of its own, so that 'he_{*n*}' selects the *n*th demonstratum of the context. This account, again, treats demonstratives as ambiguous—in fact, we have indefinitely many demonstratives with a distinct character. On this account, (10) would be represented as (11), featuring two distinct expressions 'he_{*m*}' and 'he_{*n*}', for some $m \neq n$.

(11) He₁ is sad because he₂ is leaving.

This ambiguity seems problematic: English has only one word 'he'; yet, the view would have it that we have indefinitely many indistinguishable expressions 'he'.

One could try to avoid ambiguity by saying that a demonstrative doesn't have a complete linguistic meaning and, further, receives no complete representation in the logical form until supplemented by extra-linguistic reasoning. The grammar underspecifies meaning and representation; consequently, different indices in (11) are contributed post-semantically, as a result of a pragmatic process that fixes the content. That is, the grammar underspecifies the full representation of (10), but its ultimate form is exactly the one in (11) as a result of the import of the two distinct pointing gestures, or speaker's intentions.⁹ However, while this view wouldn't maintain that a demonstrative is lexically ambiguous, notice that it would still maintain that different occurrences of a demonstrative receive different representations in the logical form; consequently, a sentence containing a demonstrative could be associated with distinct logical forms, where the difference between these forms would boil down to a difference in the representation of the demonstrative itself.

⁹ This approach would distinguish the character of a demonstrative, understood as a function from context to content, from its standing linguistic meaning. On this proposal, the linguistic meaning of a true demonstrative isn't its full-blown character; rather, it is incomplete and determines a character only once coupled with some (extra-linguistic) supplementation. A version of this view is defended by Braun (1996). The account would then maintain that demonstrative pronouns are linguistically underspecified, and that full meaning-representation requires a contextual supplementation. On this account, too, (10) would be represented as (i):

(i) He₁ is sad, because he₂ is leaving.

However, here the linguistic meaning of 'he' would be a function that takes the extra-linguistic supplementation, for example, a demonstration, represented by the subscripts '1,' and '2,' respectively, and delivers a character: relative to the first demonstration, it delivers the character that maps the context to the 1st demonstratum in the context, and relative to the second, the one mapping the context to the 2nd demonstratum in the context. On this view, while there is a single linguistic meaning associated with a demonstrative, there is still an underlying ambiguity in the representation of demonstratives resulting from the underspecification of their meaning by grammar.

Note that, while I sketch this strategy as a modification of the second alternative for integrating true demonstratives in a Kaplanian framework, one could appeal to a similar strategy with respect to the first alternative, maintaining that the character of a demonstrative expression is underspecified by its linguistic meaning, which is the meaning of the operator *he*[] specified in the Definition 2.1. The character in turn is only determined once the extra-linguistic contextual supplementation is supplied to the operator. Similar considerations would apply.

For this reason, all these ways of handling the representation of an occurrence of a demonstrative come at a cost. To see what's at stake, consider the following:

- (12) Ann came in, she sat down and she turned on the news. Therefore, Ann turned on the news.

Is (12) valid? On the face of it, it seems like an instance of conjunction elimination. Of course, if we imagine that the final occurrence of the demonstrative pronoun is uttered with a pointing gesture toward someone else, perhaps Betty, then (12) clearly will not count as an instance of conjunction elimination. Now, on the present proposal, the difference between (12) and its counterparts involving pointing gestures and deictic uses of pronouns is going to come down to either sheer lexical ambiguity of 'she' or it will amount to a post-semantic intrusion of general pragmatic reasoning into the logical form. In the latter case, this intrusion would only be reflected in the representation of the pronoun itself—in the associated index—leaving out the mechanisms that determined a particular choice of indexing. This, in turn, complicates the question of the relation between validity and grammar. Given the surface form in (12), there is no direct and systematic way of determining whether (12) is associated with a valid form or not. Indeed, since on the family of accounts presently under consideration even relative to a single, static, unchanging context, different occurrences of a pronoun can receive different representation, it seems that we would have to say that (12) is invalid, since there are contexts—namely, ones in which the second occurrence of the pronoun picks out someone other than Ann—in which the premise can be true while the conclusion false. Now, of course, you might say the *form* of an argument that allows for readings in some contexts according to which the truth of the premise doesn't guarantee the truth of the conclusion is a *distinct form* from the one according to which all occurrences of pronouns are co-indexed, and are co-indexed with 'Ann.' And the latter form is valid—in all contexts, the truth of its premise guarantees the truth of the conclusion.

That, of course, is true. But it is beside the point. The point is that either an argument is associated with indefinitely many logical forms, all differing only in the representation of an indefinitely lexically ambiguous expression, 'she,' where the extra-linguistic mechanisms simply help in disambiguating the form, *or* the context actually plays a double role in first fixing the form of an argument and *then* helping interpret this form. So, for instance, first the general reasoning plays the role of fixing the index on the occurrence of the pronoun, and then the thusly indexed pronoun gets its interpretation by looking at the parameter of the context corresponding to the index. But this seems misguided, getting things precisely backwards: it is because the general mechanisms provide

different interpretations to the different occurrences of a pronoun that we end up assigning them distinct logical forms, not the other way around.

Here is one way to dramatize the point. Recall that we introduced indexing in such a way so that 'he_{*n*}' is the *n*th demonstratum of a context. Of course, this involved a bit of an idealization. What makes something the *n*th demonstratum of a context? A natural response one might give is that it is the *n*th thing the speaker demonstrated (if there is one), or the *n*th thing she intended in that context (if there is one). But even assuming that we can neatly order demonstrations in the context (so that no two demonstrations are, say, simultaneous), this has an unwelcome consequence: namely, each new occurrence of a demonstrative would come with a fresh new index. So we couldn't, even if we wanted to, maintain that the two consecutive occurrences of a demonstrative are co-indexed. This would mean that we can never, even in principle, represent the following as a valid argument:¹⁰

(13) She is happy, therefore, she is happy.

More plausibly, we would want to allow for a distribution of indices that would allow for different occurrences of a pronoun to be co-indexed. But this would require that we assign the same index just in case the meaning-fixing mechanisms—the demonstrations or the meaning-fixing intentions—determine the same referent for the relevant occurrences of the pronoun. But it is difficult to see what determines in general that this is the case, other than the fact that the speaker is intending to demonstrate, or is demonstrating, the same, or different things. And it is difficult to see how we can recognize that this condition obtains, unless we first recognize the speaker intends to demonstrate, or is demonstrating, one thing rather than another. But if so, the indexing plays no role in determining interpretation; rather the indexing is determined by (the sameness or difference in) the interpretation. The process gets things backwards.¹¹

Is there a way to avoid the ambiguity or underspecification in the meaning of a demonstrative? Either an ambiguity or underspecification seems inevitable if we want to maintain that

- (i) utterances like (10) are interpreted against a single, unchanging context;
and

¹⁰ Kaplan is aware of this problem. He describes it as an "exotic ambiguity" in true demonstratives.

¹¹ Of course, there are grammatical constraints on co-indexing (see, e.g., Chomsky, 1981; Fiengo and May, 1994, 2006). For instance, in 'Ann likes her,' 'her' cannot be co-indexed with 'Ann' on pain of violating Principle B of binding theory (Chomsky, 1981). However, few would argue that the grammatical constraints go so far as to fully determine the distribution of indices for arbitrary deictic uses of a pronoun within a sentence.

- (ii) the two occurrences of the demonstrative in (10) have different semantic interpretations.

For if both (i) and (ii) hold, then either demonstratives are linguistically ambiguous or the grammar underspecifies their representation: only once we couple them with extra-linguistic supplementation do we get a full-blown logical form.

(ii) is uncontroversial, but why hold (i)? Indeed, we could maintain that a context shifts between any two occurrences of a demonstrative, so that a demonstrative picks the unique demonstratum *of that context*. The contexts would thus be thinly sliced to contain at most one demonstration, but it would be rapidly changing, even mid-utterance. This view would blur the difference between pure indexicals and true demonstratives. Each occurrence of a demonstrative could be uniformly represented, and have a complete linguistic meaning, selecting as its content the unique demonstratum of the context. While this would allow for a uniform representation of a demonstrative in the logical form, such a demonstrative wouldn't be different from 'I': its character would fully determine its semantic interpretation given a context (albeit the contexts relevant for its interpretation would be more ephemeral).¹²

More importantly, the view that would maintain that the context changes with every occurrence of a demonstrative would still owe us an explanation of how exactly the context changes and, even more importantly, how it does so in a way that determines the correct interpretation. For instance, we would need an account of how it is that in (10) the context changes in such a way that the semantic content of the second occurrence of 'he' in its shifted context is no one other than Tim.¹³ Presumably, in this example, the specific shift in context is associated with the pointing gesture that accompanies the utterance of the pronoun. However, since pointing gestures are not essential for demonstrative reference, one would need to allow that the context-shifts can be guided by other

¹² Braun (1996) makes this point in favor of the underspecification view, and against the context-shifting view.

¹³ Braun (1996) sketches a version of a context-shifty view maintaining that the context shifts automatically after every utterance of a demonstrative, but doesn't specify how it is that it shifts *in a particular way*. For Braun, the initial context is just equipped with a series of demonstrations one of which is designated as *focal*. After the first demonstrative is uttered, the context shifts, and the next demonstratum is made focal. But, of course, there is a question of what orders the demonstrations in any given input context, to begin with; another way of asking the same question is: what makes a specific demonstratum the new focal one?

mechanisms as well, whatever those may be. It is precisely the story of how such context-shifts are induced that we need.

One natural thought is to take the supplementive mechanisms that, on the traditional account, complete the meaning of a demonstrative to be instead the mechanisms that drive the context-change. But the whole point of drawing the distinction between pure indexicals and true demonstratives was to recognize that, for the latter class, the linguistic meaning fails to specify a content barring some supplementive mechanisms, the recognition of which requires open-ended, abductive reasoning, of the sort not involved in the interpretation of 'I'. Mistaking such supplementive mechanisms for mechanisms of context-*change* would be missing the point: it is precisely the underspecification of meaning by grammar within a *single* context that was the key feature of true demonstratives.

Relatedly, allowing for context-shifting makes our question about logical inference, if anything, even harder. For, while now we have all occurrences of a demonstrative uniformly represented in the logical form, we have to say something about how the potential change in context between any two occurrences of a demonstrative affects the validity of an argument containing multiple occurrences of a demonstrative. One reaction—a Kaplan inspired one—is to simply say that in evaluating an argument the context is stipulated to be fixed. One might worry, however, that such a stipulation involves too great of an idealization to be plausibly applicable in evaluating actual argument patterns in natural language discourse: unless we have a systematic theory of which mechanisms guide context-change, we cannot simply stipulate that those mechanisms, whatever they are, are not operative in a given case.

However, blurring the distinction between pure indexicals and true demonstratives is a difficulty for the context-shifty view *only if* we maintain that the linguistic meaning of a demonstrative is indeed incomplete in this way, and, in particular, if the contextual, supplementive mechanisms that determine the resolution can only be recognized through open-ended, abductive reasoning, even within a specified context. Whether this is so, of course, is an empirical question. If, by contrast, it turned out that contextual parameters evolve and change in systematic ways—if we can identify systematic mechanisms whereby the context changes in a way that determines specific interpretations of a pronoun, and if, given such an antecedently set context, the resolution of a demonstrative pronoun is automatic, rather than mediated by further extra-linguistic supplementation—then the context-shifty view would be vindicated. What would transpire, were this to be the case, is that there are moves in discourse that antecedently, but systematically, set up a context in a way that

allows a demonstrative pronoun to automatically select a referent, much as 'I' does; true demonstratives would indeed turn out to be like pure indexicals. On such a view, a demonstrative pronoun would receive both a non-ambiguous linguistic meaning and a uniform representation.¹⁴

In what follows, I develop and defend such an alternative, context-shifty account. A demonstrative pronoun, like a pure indexical, has a simple, unambiguous meaning: it selects the most prominent entity—one at the center of attention—in the current context (in the sense to be made precise shortly). The context, in turn, is determined, at any given point in the discourse, by *linguistic* mechanisms of context-change. Once we have a full account of these mechanisms, we shall see that we can give a systematic account of the meaning of a demonstrative pronoun, fully on analogy with pure indexicals. Furthermore, we will also see that we can give a systematic account of validity of argument patterns containing demonstratives, one that fits with our intuitive judgments, and does not involve idealized stipulations of a monolithic, unchanging context.

¹⁴ Of course, if the context, indeed, is shifty, then the question arises again of how we should characterize the notion of truth, validity, and entailment. I shall return to this question below, and in more detail in Part III.

An Alleged Ambiguity and the Dynamics of Context-Change

So far, we have entertained the view that a true demonstrative refers to a salient entity in the context that satisfies its linguistic meaning, where determining the salience of a particular entity requires some supplementation independent of that linguistic meaning. But there is an immediate problem with this account of the meaning of a demonstrative. To see the problem, consider the following:

- (14) She [pointing at Mary] is hungry.
- (15) Every woman drives the car she owns.
- (16) Mary came in. She sat down.
- (17) A woman came in. She sat down.

While in (14) and (16) ‘she’ refers to Mary, in (15) and (17), ‘she’ does not refer at all. Rather, its semantic interpretation varies with the indefinite ‘a woman,’ or the quantifier, ‘every woman.’ This means that what kind of semantic interpretation a pronoun receives varies not merely across different speech act situations, but across different linguistic environments in which it occurs. But this also means that the meaning of a pronoun cannot simply be such that it picks out the salient or demonstrated object, or the object the speaker intended to refer to in the speech act situation, for there is no such object in the case of (15) or (17). So, the linguistic meaning of a demonstrative cannot specify that its semantic content has to be a salient object.

Regarding the question of what the semantic content of an occurrence of a demonstrative pronoun is, then, the traditional reaction was to posit ambiguity: demonstrative pronouns are ambiguous between their referential and bound uses.¹ How does this explain that the pronoun ‘she’ in (15) co-varies with ‘every woman,’ while in (14), (16), and in (18), it refers to Mary?

- (18) Every woman drives the car *she* [pointing at Mary] owns.

¹ See Heim and Kratzer (1998) for a discussion.

One way to understand this proposal is to assimilate the difference between the occurrence of the pronoun 'she' in (15) and (18) to that between a bound and free variable. The idea is that a pronoun is just like a variable: in (15) it is represented by a variable that is bound by the quantifier 'every woman,' while in (18), it remains free, and hence is interpreted relative to a contextually specified assignment function.² Notice that on this account, as Kaplan (1989b) points out, the distinction between (operator-)bound and referential demonstrative pronouns gives rise to a structural ambiguity, reflected in the different choice of a variable.³

To make this proposal a bit more concrete, let us start with a simple account of the difference between bound and free variables, familiar from the standard *Tarskian* semantic framework. In this framework, a sentence can be represented with a formula which is interpreted relative to an assignment function of values to its variables. One can think of an assignment as a sequence of individuals, where the variable takes on as its value as a particular member of the sequence, specified by the index associated with the variable.⁴ For instance, a variable, x_i , on this model, is interpreted as selecting an i th element of the assignment. When a variable is bound within the scope of a quantifier, it receives a bound interpretation because the quantifier varies the assignment function within its scope. For example, a universal quantifier, $\forall_i$, varies the assignment function in its scope relative to the i th element, so affecting the interpretation of the variable x_i in its scope. Specifically, it requires that the sentence in its scope is true for all possible assignments of values to the variable x_i . Less abstractly, this allows us to capture that an utterance like (15), where a pronoun occurs bound by a universal quantifier, is true, just in case 'she is happy' is true under *all* assignments of values to the variable corresponding to 'she'; thus (15) is either true under all assignments, or none.

Roughly, then, (15) and (18) would be represented as (19) and (20), respectively:

(19) Every woman₁ drives the car she₁ owns.

(20) Every woman₁ drives the car she₂ owns.

² For this understanding of the ambiguity between bound and referential pronouns, see, for example, Kaplan (1989b,a). See, in particular, the discussion in Kaplan (1989b, p. 571–3).

³ As we shall see shortly, discourse-bound, E-type, interpretations of pronouns, as in (17), for example, complicate this issue somewhat.

⁴ I am simplifying a bit. More precisely, an assignment is a function, mapping a set of natural numbers, including 0, to a set of individuals. A variable is associated with an index, specifying the element in the assignment function it is sensitive to.

(19) would deliver the reading according to which every woman drives the car *that* woman owns, whereas (20) is interpreted as every woman driving the car that some specific, contextually specified, woman owns. The difference between the two interpretations comes down to a difference in the choice of a variable.

The Tarskian framework is limited insofar as the quantified expressions can only bind variables within their syntactic scope. This will not suffice to explain how it is that the pronoun in (17) can depend on, and co-vary with, the indefinite 'a woman.' What examples like (17) call for is a way of allowing for an expression to bind a variable across a discourse, beyond its scope. A familiar way of extending the Tarskian framework to allow for a discipline of binding across a discourse is offered by dynamic semantics.⁵ Dynamic semantics modifies the standard Tarskian framework by specifying truth-conditions of a formula not relative to a single assignment, but rather a pair of assignment functions, an input assignment function modeling the dependencies that are available for interpreting the formula representing the utterance, and an output one, modeling the dependencies that are potentially created by the formula representing the utterance. In other words, a formula is interpreted as an instruction to update the available interpretive dependencies. It creates an output assignment by augmenting interpretive dependencies from prior discourse reflected in the input assignment, with any new interpretive dependencies the formula creates for the interpretation of subsequent discourse. Roughly, this allows a variable within a formula to receive the value from the input assignment function for that formula, which reflects the interpretive dependencies created by the preceding discourse, and it allows the formula to pass down a (potentially updated) output assignment, encoding new dependencies for subsequent variables. 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. However, now we can define the existential quantifier in a way that allows for alternative *output* assignments, so that the interpretation of subsequent variables can receive a bound interpretation throughout the subsequent discourse. That is, while in the standard framework, a quantifier $\exists_i$ binds a variable x_i in its scope, now $\exists_i$ varies *output* assignment functions, thus binding a variable x_i in the subsequent discourse. A formula is true on a given input assignment just in case an output assignment makes it true.⁶

⁵ The dynamic approach to semantics stems from Kamp (1971) and Heim (1982). (See Dekker, 2011, for an accessible overview.) Employing dynamic semantics as a modeling choice allows for a particularly straightforward and elegant way of capturing interpretive dependencies across a discourse. Of course, dynamic semantics is not the only semantic framework available for interpreting discourse-bound pronouns. For a prominent alternative, see, for example, Elbourne (2005).

⁶ So, instead of treating an existentially quantified statement as true just in case there is an assignment function of values to variables that makes the embedded clause true, as in a standard

Here is then a rough sketch of how we can capture, within this framework, the interpretive dependence of the pronoun on the indefinite ‘a woman’ in (17): the first sentence is translated as an update formula, interpreted as a relation between assignment functions, updating potential dependencies for subsequent discourse. Since it contains an indefinite—which is interpreted as having an existential quantifier force—(17) is true on a given input assignment function, just in case there is an output assignment function according to which some, unspecified, woman (a witness for the quantifier) came in.⁷ The second sentence is translated by a formula interpreted, again, as a relation between assignments. It takes this output assignment as a new input one; consequently, ‘she’ in the second sentence can be interpreted as a bound variable that takes as its value whatever the witness for ‘a woman’ is in this new input assignment, and this is how (17) receives its anaphoric reading. The whole discourse is true on the initial input assignment just in case there is a non-empty output assignment it leads to, that is, just in case a woman came in and sat down.

The framework thus *allows* for both anaphoric and deictic readings, and for both operator and discourse-bound interpretations. But there is still a question: how does the pronoun ‘she’ on the anaphoric interpretation get bound by the existential quantifier ‘a woman’ in (17)? And how is it that, on the deictic reading, the pronoun fails to get so bound? In Tarskian and dynamic semantics alike, variables behave like temporary names. The name is temporary because a quantifier like $\exists x_i$ re-defines what x_i names within its scope (varying the output assignment it leads to); we can then interpret occurrences of x_i as local names for potential witnesses of the quantifier. So, if ‘a woman’ is represented as a quantifier $\exists x_i$, then ‘she’ gets translated with x_i , and interpreted as looking for the i th element of the input assignment; this element, given constraints imposed by the prior discourse, will be required to be a woman who came in.

Tarskian framework, we interpret it as true, given an input assignment function g , just in case there is some (non-empty) assignment function that results from an update with it. More precisely, in this framework, we interpret formulas as relations between assignment functions. A formula (usually called ‘an update’) is true relative to an input assignment, g , just in case it relates it to some (non-empty) output assignment, h . The simplest update, $[C]$, is interpreted as a partial identity relation between assignments, and contributes information fixed by the condition C . If g is an input assignment, and h an output one, $[[C]](g, h)$ iff $g = h$ and C holds on the interpretation of variables given by g . (So, for instance, $[[x_n \text{ is a woman}]](g, h)$ iff $g = h$ and whatever entity g assigns to x_n satisfies the predicate ‘is a woman’ in a given model; note that this ensures that in the output assignment h , x_n is constrained to be a woman.) In standard fashion, a sequence of updates H and K is represented by a single update $H; K$ that performs the update H followed by K . In addition to the simplest update, which contributes information specified by a condition, we will need more complex updates which update the values of the variables; they will be described presently.

⁷ Any output assignment according to which there is woman who walked in suffices. This is what is meant by saying that a witness for the existential is an “unspecified” woman.

By contrast, translating the pronoun with a different choice of a variable— x_n , for $n \neq i$ —gives an unbound, deictic reading: the variable selects whatever is stored at n th position. In this way we capture the two readings of (17). So, this is the sense in which the difference between unbound and anaphoric pronouns boils down to the difference in the choice of a variable.

But where does this leave us concerning our starting point? How does any of this relate to our earlier discussion of context-sensitivity resolution? We are told that where a pronoun is bound by some antecedent expression in discourse, it is interpreted as a bound variable co-indexed with the relevant antecedent. In turn, when an occurrence of a pronoun is unbound—when it is used deictically, say—it is represented as a free variable. Two questions still remain: what determines that a pronoun is interpreted as bound by a particular antecedent, as opposed to possible others, or as free? And, further, when a demonstrative pronoun is represented as a free variable, how is it that it receives a particular interpretation, as opposed to possible others? For instance, what is it by virtue of which an utterance of ‘She is happy,’ comes to mean something different if I’m pointing at Sue, and if I’m pointing at Mary?

In our earlier discussion, we presented this problem as a problem of how the context of utterance resolves the meaning of the pronoun ‘she.’ We can still see this problem in the same way. Let’s consider deictic uses first. On the current picture, a deictic pronoun is represented as a free variable. The sense in which its interpretation depends on a context of utterance can be captured by saying that the free variable receives its meaning by having the extra-linguistic context determine the value of the relevant index of the assignment.⁸ In this way, then, we can think of the interpretive dependencies encoded in the assignment as reflecting the current state of the context. The new addition is that we can capture the idea that linguistic items—say, an indefinite—can change the context by creating a new dependent interpretation for a pronoun. Further, this allows us to also capture the idea that the context can change—even mid-utterance—in a way that affects the interpretation of a subsequent pronoun. And insofar as the real-world speech situation can affect the current state of an assignment indicating how a particular free variable is to be interpreted, say, by allowing for a pointing gesture to set the value for the free variable, we can capture deictic uses of pronouns, and specifically, a difference in interpretation of distinct deictic uses of a pronoun, as in (10), repeated below:

(10) He [pointing at Bill] is sad, because he [pointing at Tim] is leaving.

⁸ In essence, the value of a free variable is fixed by a contextually specified assignment. See Kaplan (1989b). The idea is that the context fixes an assignment of values to free variables, just as it fixes the world in which an utterance takes place.

We just maintain that the two occurrences of the pronoun, 'he,' are represented with two differently indexed (free) variables, as in (21), thus allowing them to receive different interpretation:⁹

(21) x_1 is sad because x_2 is leaving.

Of course, simply saying that the context specifies an assignment of values to free variables does not give us an answer to the question of how a particular reading does get determined. While representing (18) with (20) captures the observation that the pronoun is unbound, it is the fact that the speaker pointed at Mary that plays the crucial role in determining the interpretation. The same goes for (10) and (21). As before, then, contextual supplementation is required to fix the meaning of the pronoun. And all the same questions about the mechanisms whereby a particular interpretation is determined remain open. And, however we answer those questions, such mechanisms are left implicit in the logical form of a sentence like the one in (18) or (10).

Similarly, turning now to bound cases, while the account allows for a changing context that constrains the interpretation of a bound variable, whether a particular interpretation is retrieved, even in the bound case, depends on the selection of indices. Specifically, it depends on a specific antecedent a pronoun is interpreted as co-indexed with. And while there are grammatical constraints on binding, they are typically taken as merely constraining, rather than determining the interpretation of a pronoun. For instance, consider the following:

(22) A woman met a girl. She greeted her.

Even if we allow that the two indefinites update the assignments, in a way that allows for the two subsequent pronouns to be bound, there is still a question of which pronoun gets co-indexed with which antecedent. While grammar is taken to constrain co-indexing—ruling certain interpretations out—extra-linguistic supplementation is still a crucial component of interpretation, and the linguistic meaning of a demonstrative is still incomplete: relative to a context, it underdetermines an interpretation. And, of course, whether a particular occurrence of a pronoun is bound or deictic is also often underspecified by grammar. (Again recall the difference between (15) and (18).)

Further, while we have allowed for the linguistic material to change the contextual dependencies by creating new dependent interpretations for subsequent variables, we have done so in a way that still gives raise to ambiguity worries.

⁹ Allowing but, of course, not requiring. Two distinctly indexed variables *can* have the same value, but there is no requirement that they do.

For instance, the representation in (21) represents different occurrences of one and the same word, 'he,' with distinct variables—the variation in interpretation boils down to this difference in the choice of a variable. This means that we have to either treat the word 'he' as ambiguous, each differently indexed variable being associated with a specific character, or that its meaning is incomplete: the grammar underdetermines the choice of an index on a variable and the extra-linguistic mechanisms enrich the form in (21), fixing the indices, and determining the interpretation of the two variables.

In sum, this story advances beyond the static Kaplanian picture insofar as it allows for the context to dynamically change with the evolution of a discourse. In this way it makes it possible for both antecedent linguistic, and non-linguistic material to affect the resolution of a subsequent pronoun. It thus allows for an account of both deictic and bound interpretations of pronouns. But the story, I shall argue, is not right. The meaning of a demonstrative pronoun isn't incomplete given a context, and the grammar does not underdetermine its interpretation. The logical forms that locate the ambiguity and variability of interpretation in the pronoun miss a slew of robust, linguistic mechanisms that affect its resolution, and incorrectly characterize the linguistic meaning of a pronoun as incomplete or ambiguous. Once these mechanisms are recognized and their import separated from that of the linguistic meaning of a pronoun, we will be able to provide an account that claims a number of empirical advantages: it allows for a uniform semantics for demonstrative, anaphoric, and bound pronouns, and provides straightforward resources for describing the linguistic constraints and preferences on the interpretation of pronouns. Furthermore, I shall argue that the account explains interpretive constraints that are not explained or captured by either the Kaplanian or the dynamic view we have seen: it explains why interpreters (as we shall see) often resolve pronouns in ways which yield implausible or inconsistent readings despite the apparent availability of pragmatically plausible interpretations, and it correctly predicts cross-linguistic differences in the distribution and resolution of pronouns, that are puzzling from the standpoint of the earlier accounts.

Pointing Things Out: Prominence and the Attentional State of a Discourse

Here is a story I shall defend. The meaning of a pronoun is simple, uniform, and unambiguous; as a first pass, a pronoun selects the most prominent candidate interpretation—what is “at the center of attention” at the point in discourse in which it occurs. At any given point in a discourse, the context provides a ranking by prominence of candidate interpretations for a pronoun, tracking what’s most prominent—that is, at the center of attention. The prominence ranking changes and updates, as the discourse unfolds, as a function of the meaning of linguistic items the discourse harbors, word by word. In this way, the resolution of a pronoun requires no extra-linguistic supplementation. It is linguistically determined, through and through: by its standing linguistic meaning, and the *linguistically* set up context. Let me start by presenting the account, and then offering a defense.

The basic idea is that we can represent candidate resolutions of a pronoun (on their bound and referential uses alike) as ranked according to their relative prominence encoded in the context: those higher in the ranking are preferred as interpretations of pronouns over those lower. Let us call this ranking the *attentional state* of an ongoing discourse.¹ To model the dynamics of the attentional state of a discourse, we will build on the dynamic framework introduced before.² We will still exploit the idea of an assignment, but will take assignments

¹ The idea that we can rank candidate referents according to prominence is familiar from *Centering theory* (Sidner, 1983; Grosz, Joshi, and Weinstein, 1995; Bittner, 2014). However, within Centering theory, the preference for higher-ranked entities is understood as a pragmatic preference, as one of many (potentially non-linguistic) cues we use to recover the intended referent of a demonstrative pronoun. By contrast, I shall model both the prominence ranking, and the preference for higher ranked elements, as semantic: the ranking is maintained and updated by linguistic rules, and the corresponding context-sensitive expressions select the most prominent interpretation (in the sense to be made precise shortly) as a matter of their meaning; that is, the linguistic meaning of a pronoun, its character, given this linguistically maintained ranking, determines its semantic value.

² The key idea of the account I shall defend, namely, that pronominal interpretation is fixed through linguistic rules, is independent of a particular formal implementation, and could in principle be captured in different semantic frameworks. However, not only does the dynamic semantics approach fit neatly with the goal of capturing the effects of discourse conventions on the interpretation of context-sensitive language but, as we shall see below, and in more detail in Parts II and III, the understanding of the import of discourse conventions as that of semantically specified rules for manipulating contextual parameters is difficult to preserve in a framework that

to be stacks, specifying a sequence of individuals—candidate interpretations for the pronoun—ordered by prominence.³ The most prominent individual is in the initial position, 0—the top of the stack—and ones in subsequent positions are deeper in the stack, thereby, receding in prominence. Since we are interested in prominence, we will treat variables not as temporary names, but rather as markers of prominence; similarly, we will organize the interpretive dependencies in discourse not by names, but by prominence. Correspondingly, an index on a variable tracks prominence, so that, for example, x_0 selects the most prominent interpretation, stored at the 0th position of the current assignment. While this change might seem inessential at first, it allows us to capture the dynamics of prominence-change across a discourse which, we shall see, affects the interpretation a pronoun receives at a particular point in discourse.

I shall argue that the attentional state is maintained through discourse conventions, the set of linguistic rules affecting prominence. The effects of these discourse conventions reflect the context-affecting aspect of meaning of various linguistic expressions within an utterance. The changes in prominence are, thus, induced as a matter of the meaning of linguistic expressions.

How can an expression have such context-affecting meaning? We have already seen how within the dynamic framework presented earlier: an utterance of an expression can bring about new interpretive dependencies for the interpretation of subsequent expressions, as in (17). As in the dynamic system presented above, then, on our current account, utterances will still be interpreted as instructions to update the available interpretive dependencies, but now we can capture how they do so by affecting *prominence* of potential interpretations of a pronoun.⁴ Specifically, unlike in the dynamic system seen earlier, they update potential

doesn't assign a level of semantic interpretation which governs and tracks context-change. Dynamic semantics models, and is motivated by, the idea that the context—the values of interpretively relevant parameters—changes with evolving discourse, and that this change is driven by semantic composition. This will be an important feature of the view I will defend as well. However, as will become clear in what follows, the specifics of my account, both formally and theoretically, will differ from the standard dynamic approach of the form we have seen in Chapter 3. The departure from the standard account will allow us to separate the dynamic layer of meaning from a representational, truth-conditional content, which will lead to a better understanding of context-content interaction.

³ Formally, a stack is a function mapping a set of natural numbers, including 0, to a set of individuals. Note that, correspondingly, variables are now markers of prominence, not temporary names. By interpreting utterances as relations between such stacks, we will still be able to capture unbound and anaphoric readings of pronouns alike, in a way that shall be made formally precise shortly. For other systems exploiting stacks for modeling discourse dynamics, see, for example, Zeevat (1992); van Eijck and de Vries (1992); Vermeulen (1993); Dekker (1994, 2011); Kaufmann (2000); van Eijck and Unger (2010); Bittner (2014); *inter alia*. See also Haug (2014) for a related approach to modeling discourse dynamics implemented in compositional Discourse Representation Theory (DRT) with partial assignments.

⁴ Officially, just as in the dynamic system presented above, utterances are still interpreted as relations between assignments (where assignments are now modeled as stacks).

interpretive dependencies not by introducing new temporary names, but rather by introducing new prominent interpretations, demoting others. The idea is that certain expressions, as a part of their meaning, bring certain interpretations into focus, essentially inducing a shift in attention. They do so by updating what's currently at the center of attention—at the position 0 on the stack. In this way, they create new dependent interpretations for subsequent pronouns: a pronoun is interpreted as selecting what's currently at the center of attention.

Just as before, in the current framework, an existential quantifier can update the available dependencies by varying the output assignment function, that is, the output stack. However, now it does so by pushing values onto the stack, inserting them at a specified position, and thereby decreasing the prominence of other candidates in the discourse by one position. In this way, it varies output assignment functions. So, for instance, as a first pass, an indefinite, such as 'a woman,' pushes a new value to the top of the stack (which is further constrained to be a woman), demoting others, thereby creating a new dependent interpretation for a subsequent occurrence of the pronoun, 'she,' reflected in the output stack.⁵

As a first pass, we can still interpret a pronoun as a variable, but we now treat a variable as a marker of prominence; we can represent the pronoun as 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. A pronoun so defined selects the top-ranked value, the top of the current assignment, understood as what is "at the *center* of attention" at a given point in discourse. The intuitive idea is that pronouns are prominence-sensitive: they always pick out the top-ranked candidate in the current attentional state.⁶

Now given this account, why is it that, say, in (17), the pronoun 'she' is naturally interpreted as bound by and co-varying with 'A woman'? And why is it that when the pronoun in (17) is uttered in tandem with a pointing gesture, say, as in (23), it receives a specific unbound, deictic interpretation?

(23) A woman came in. She [pointing at a female cat, Betty] sat down.

Finally, can we explain this difference in interpretation while representing the pronoun in a uniform unambiguous way? Let us take these questions in turn.

⁵ We will make this sketch formally precise shortly. A more informal sketch will serve us well enough for now. An indefinite, as before, will store an *unspecified* value at the center of attention: it will relate an input assignment, *g*, to an output assignment just in case some woman is stored at the center of attention in the output assignment, followed by all other elements of *g* in order (and no others). Again, any witness for the existential will do.

⁶ Of course, we will have to qualify this to account for the fact that different pronouns have different linguistic constraints—'he' means something different than 'she.' We will make such qualifications shortly. For now, let us just characterize prominence-sensitivity generally, disregarding such constraints.

On the present framework, (17) will be explained as follows. The discourse begins with a description involving an indefinite woman, putting her at the center of attention, and proceeds by developing a narrative about this woman. The pronoun in turn picks out what's at the center of attention, that is, the aforementioned woman. Hence the pronoun receives the anaphoric interpretation—'she' selects the aforementioned (unspecified) woman—the witness for the indefinite. Somewhat more precisely, we can describe this as follows. The indefinite makes its witness prominent by pushing a new value to the top of the stack, requiring that it be a woman. The rest of the utterance constrains this value, by requiring that the aforementioned woman came in. Finally, the pronoun selects the top-ranked element of the current input assignment, the one output by the first sentence, thus getting bound by the indefinite, 'A woman,' the item that most recently pushed a new value to the top.

Making this idea formally precise will help us showcase how we can deliver a uniform meaning and representation of a pronoun. Let @ be a variable, used as a marker of prominence, that is interpreted relative to an assignment g , as selecting its top-ranked element. So, given an assignment g , @ picks out whatever is stored at the 0th position of the assignment that is, what is at the center of attention. The meaning of @ is a first approximation to the meaning of an English pronoun.⁷ Then we can define an update [*woman*(@)]—interpreted as a relation between assignments—requiring that whatever is at the center of attention in the input stack is a woman. An assignment function passes this condition just in case its top-ranked element is a woman.⁸ Such basic update simply provides more information about a particular element in the sequence, in this particular case, the top-ranked one, but doesn't affect prominence ranking. For that purpose, we introduce updates that can affect ranking, updating what's at the center of attention. For instance, we want to capture that an indefinite can make its witness prominent for the interpretation of a subsequent pronoun (demoting other values in turn). So, as a first approximation to the meaning of the English indefinite article, we translate it as a dynamic existential quantifier, $\langle \alpha \rangle$, which changes the input assignment by introducing a new unspecified top-ranked individual at the center of attention, demoting all other elements one position in the ordering.⁹ This update does affect the prominence ranking, changing what

⁷ Obviously, this is merely a rough approximation. The present gloss does not yet differentiate meanings of different pronouns. As promised, further qualifications that rectify this shortcoming will be introduced in due course.

⁸ Recall, this simplest type of an update is understood as partial identity relation. Formally, the update represented as [*woman*(@)] relates two assignments, g and g' just in case $g = g'$ and g_0 is a woman. In other words, the update tests whether an assignment satisfies the condition specified.

⁹ This is a non-deterministic update. It relates an input assignment function g to an output one h just in case h potentially differs with g in the top position, and for every subsequent position i , $h_i = g_{i-1}$.

is at the center of attention. Putting all this together, (17), on its ‘anaphoric’ interpretation, then, is represented as (24):¹⁰

(24) $\langle \alpha \rangle; [woman(@)]; [came.in(@)]; [sit.down(@)]$

Here is a gloss on how the updates in (24) change the input stack: the dynamic existential quantifier, $\langle \alpha \rangle$, corresponding to the indefinite article, introduces its witness as the new top-ranked candidate resolution for subsequent pronouns; the subsequent updates constrain this witness to be a woman, to have come in, and to have sat down. Thus, (24) is true just in case some woman came in and sat down.¹¹ Note that the pronoun, ‘she,’ is thus bound by the indefinite, ‘a woman,’ as is intuitively correct. And we achieve this without representing the pronoun as carrying potentially distinct indices and so, without introducing ambiguity or underspecification.¹²

Now, what about the non-anaphoric use of ‘she,’ as in (23), repeated below, where the pronoun is used while pointing at a female cat?

(23) A woman came in. She [pointing at a cat, Betty] sat down.

The indefinite, just as in (17), makes an unspecified woman prominent. How is it, then, that the pronoun, if it just selects whatever is at the center of attention, isn’t bound in (23)? It is obvious that in this case the pointing gesture affects pronoun resolution. One intuitive way to describe this effect is to say that the pointing gesture induces another shift in attention; it makes a new entity prominent, demoting what was previously at the center of attention. As a result, the pronoun in the second sentence in (23) selects the cat demonstrated by the accompanying gesture, namely Betty. But how can we capture this effect of pointing on our current account, while still maintaining that the pronoun is uniformly represented as @?

Suppose that we associate the pointing gesture with its own attention-shifting update. It is easy to do so. For instance, to represent the effect of pointing, we can introduce a family of updates, $\langle \pi e \rangle$, where π corresponds to the act of pointing

¹⁰ As in the standard dynamic system, described in Chapter 3, the semicolon represents update sequencing: $\llbracket H; K \rrbracket(i, j)$ if and only if there is some sequence h such that $\llbracket H \rrbracket(i, h)$ and $\llbracket K \rrbracket(h, j)$.

¹¹ As in the standard dynamic system, described in Chapter 3, the formula is true for an input assignment g iff it is true for some output assignment.

¹² Strictly speaking, we have represented both the contribution of the pronoun, and the argument structure of the predicates in the first sentence in the same way, as @. This is inessential: we could separate the representation of the pronoun, from that of the arguments of a predicate, by reserving @ for the former, and using indexed variables (e.g., x_0) for the latter. Even so, each occurrence of the pronoun would be uniformly and unambiguously represented as @. The current choice of notation is mere convenience, and it will be refined in what follows.

and e names the individual being pointed at.¹³ This update stores the entity denoted by e as the top-ranked entity—the new center of attention—and pushes all others down a position in the ordering. So, (23) is represented formally as (25):

(25) $\langle \alpha \rangle; [woman(@)]; [came.in(@)]; \langle \pi b \rangle; [sit.down(@)]$

The first three updates in (25) proceed exactly as in (24), introducing a new top-ranked entity, and constraining it to be a woman who walked in. We still formalize ‘she sat down’ with $[sit.down(@)]$, which denotes an update constraining the top-ranked entity in the current attentional state to have sat down. But the update contributed by the pointing gesture, $\langle \pi b \rangle$, changes the attentional state which is the input to the final update, making Betty the top-ranked individual. Consequently, Betty is at the top of the assignment when $[sit.down(@)]$ is interpreted. In other words, a pointing gesture updates the prominence ranking, making Betty the center of attention. Thus we capture that the pronoun selects Betty, and we do so without exploiting any indices—without any ambiguity or underspecification in the pronoun.

It is, of course, easy to represent the pointing gesture in the logical form directly. But what is the justification for so doing? In doing so aren’t we unjustifiably introducing an extra-linguistic component into the logical form of an utterance? Pointing should be “off-the-record,” to use Kaplan’s wording (Kaplan, 1989a, p. 582), an extra-linguistic supplementation, not an independent semantically interpreted expression in the logical form.

The received wisdom is, indeed, that the effects of pointing gestures accompanying the use of demonstrative pronouns are precisely the kind of extra-linguistic, pragmatic effects that supplement the linguistic meaning of a demonstrative pronoun, but are not themselves reflected in the logical form of an utterance. They are off-the-record. They provide the extra-linguistic supplementation that was needed to supplement the meaning of true demonstratives. Or not even that, perhaps they are merely one among many (defeasible) epistemic cues interlocutors can exploit to figure out the speaker’s meaning-fixing referential intention.¹⁴ After all, one might think,

¹³ More generally, $\langle \alpha \rangle$ corresponds to the contribution of indefinites, and $\langle \pi e \rangle$ to definites, where e is a denoting expression. That said, it should be understood that different constraints are needed for different types of definites. I set this issue aside, as it is orthogonal to our present concerns.

¹⁴ Kaplan (1989a) endorses the Intentionalist version of the view, arguing pointings do not fix the meaning, but can help recognize speaker intention that does. Kaplan (1989b) holds an Objectivist account, according to which the pointing gesture is an objective, yet extra-linguistic, feature of context that *determines* the meaning of a demonstrative pronoun, regardless of the speaker’s intentions. Note that, on the Objectivist view, too, the role of a demonstrative gesture can be seen as defeasible, as just one among many features of context that can contribute to the resolution of a demonstrative pronoun. That is, while the Objectivist view can maintain that a pointing gesture

why else would the speaker choose to accompany her use of a pronoun with a pointing gesture if not because she intends to refer to an object in the indicated direction? The pointing gesture serves as an epistemic cue that helps manifest the speaker's intention *because* it makes sense to point at the intended referent, rather than to randomly point at something else. But the pointing gesture does *not* render the demonstrated object prominent as a matter of a convention or a linguistic rule: such rule is not needed, because rationality alone guides our understanding of what the intended effect of pointing is. This means that grammar underdetermines interpretation: a pronoun is a true demonstrative, only receiving an interpretation after an extra-linguistic supplementation. If so, then, representing the pointing gesture as contributing an attention-shifting update in the same way as, say, we represent the contribution of an indefinite is misleading. Such logical form would be no less specified by grammar than one that requires post-semantic, pragmatic processes to determine the index on a variable.

Contrary to the received view, I submit, we should represent the contribution of a demonstrative gesture, a pointing, explicitly, and interpret it independently of the interpretation of a pronoun, because there are good reasons to understand its contribution as underwritten by conventions that specify a particular semantic contribution. It is thus the received view that mischaracterizes the role of demonstrative gestures in interpretation. A demonstrative gesture is not a mere extra-linguistic supplementation on a par with other extra-linguistic resources. Indeed, it is not a part of extra-linguistic context, to begin with, but rather a part of utterance, an expression along others, with its own conventionally specified contribution, that of an attention-shifting update. This means that uttering a sentence featuring a demonstrative pronoun while pointing at different things is not like uttering the same sentence in various different circumstances; it is like uttering different sentences altogether.¹⁵

The key argument for representing gestures as linguistic items, with their contribution separately represented in the logical form of a discourse is that their effects exhibit hallmarks of linguistic conventionalization. The association

plays a meaning-fixing role, it needn't hold that it's the sole factor that plays this role. What is important for us, in any case, is that neither Intentionalists nor Objectivists see the contribution of a pointing gesture as a contribution of a linguistic item, that should be represented separately in the logical form.

¹⁵ While the idea of having pointing gestures represented as *bona fide* elements of an utterance might seem exotic at first, integrating the contribution of gestures within the process of semantic composition is neither entirely new nor outlandish. For a grammar integrating gesture and speech, see, for instance, Alahverdzhieva and Lascarides (2011). For more on interpretation of gestures within a semantic framework, see Lascarides and Stone (2009, 2006); Schlenker (2018). It's worth pointing out that these authors consider a wide range of speech-accompanying gestures, not a class of demonstrative gestures specifically.

between a form and shape of a gesture and its semantic effect is arbitrary, learned, and it varies across different linguistic communities. Kendon (2004) argues that, while English speakers can use a range of hand shapes to indicate an object, their particular choice of shape and form affects the semantic contribution they make with the gesture.¹⁶ So, a gesture with the index finger and thumb extended, the other fingers curled closed, uses the direction of the index finger to single out an object as an individual, distinguishing it from its alternatives, from other objects surrounding it (e.g., ‘That [pointing at a cat with an index finger extended] is Betty’). By contrast, a gesture with a palm up, 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 (e.g., ‘That [pointing at a cat with a palm up flat hand open toward the audience, fingers toward Betty, the cat] is a British Shorthair cat’). The distinction in form underwrites distinction in meaning.¹⁷ Moreover, it is important to remember that there are many possible gestures that are *not* typically used as deictic gestures by English speakers though they could well have been. For instance, 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.¹⁸ There is nothing inherent in the gesture that makes it unsuitable for demonstrating objects; we just don’t use it that way. Similarly, there’s nothing inherent in the gesture that makes it perform the specific function it does; we just use it that way.

¹⁶ Based on a corpus of examples collected from a number of linguistic communities (including Campania (Italy), Northamptonshire (UK), and the United States) Kendon (2004) identifies seven different specialized hand pointing gestures, distinguished by the combination of the hand shape and forearm orientation, each with a distinct semantic contribution. Further, he notes that there was some variation in the semantic import of particular pointing gestures across the linguistic communities studied. For instance, the positioning of the palm/forearm in the extended index finger pointing (down vs vertical) was semantically significant among the speakers in Campania, but not among the speakers of Northamptonshire in England.

¹⁷ See Kendon (2004, chapter 11) for further discussion. Kendon describes several different open-hand deictic gestures, distinguished by the combination of the hand shape and forearm orientation they display, each with its specific and distinct semantic contribution. However, there is a general difference in semantic contribution between the open-hand pointing gestures and extended index finger pointing gestures. While the latter single out a particular object under discussion, the former all serve to present an object not as a singled out particular which is the topic for discussion, but as relating to the main topic by exemplifying it in a certain way: e.g., as being an instance of a class, or as displaying features that are under discussion or that are to be attended to.

¹⁸ As Kendon (2004) describes, one can point with their thumb, but only when the object pointed to is to the side or to the rear of the speaker, and further, only when the identity and location of the object has either already been established (typically by a previous index finger pointing gesture), or is inessential to the main point of the utterance. The thumb pointing thus often appears as anaphoric on index finger pointing, and differs from it both in its semantic contribution, and in the conditions for its felicitous use. As Kendon describes, speakers do correct themselves when the incorrect choice is made in choosing between index-finger and thumb pointing.

The arbitrary fact that some spatially directed actions are taken to indicate objects, and to indicate them in specific ways—while analogous ones are not—reveals convention at work in English speakers' demonstrations. This is further supported by actual cross-linguistic variation: other languages, unlike English, allow speakers to indicate objects in other ways, for example, by deictic gestures of their lips, and not an extended index finger. Cuna, a Chibchan language spoken in southeastern Panama and northwestern Colombia, is one of many examples of a language of this sort.¹⁹ Similarly, Wilkins (2003) describes that among the Arrernte, a central desert Australian Aboriginal group, a pointing gesture with an extended index finger is used to indicate a single object, whereas a pointing gesture with all the fingers spread and the palm facing downwards is used when indicating a multiplicity of objects. Meanwhile, Kendon (1988; 2004) reports that among the Warlpiri and Warumungu of the central desert of Australia, the angle of the arm varies with the distance from the demonstrated object, so that the arm is angled almost vertically when pointing at very far away objects, and angled towards the ground if the object is very close. Such cross-linguistic variability reveals linguistic conventionality at work: the arbitrariness and variation show that the mapping of the form to meaning is not underscored by either general rationality or some universal principle of human psychology. In other words, you have to learn what the gesture means, just as you have to learn what the words mean. And had the meaning-fixing conventions of a specific language worked out differently, a gesture, like a particular word, could have come to mean something else entirely.²⁰

Furthermore, as Kendon (2004) points out, non-verbal means of indicating an entity are governed by rules sensitive to form, meaning, and the relationship with ongoing speech. So, English speakers count deixis as well-formed only when the pointing action is synchronized appropriately with the prosody of the accompanying utterance. They often repair utterances that fail to align speech and gesture in time. This requirement for synchronicity is similar to how the prosodic focus is grammatically constrained to appear at a particular position with a particular contour for it to have a proper semantic contribution. The misalignment is marked; for instance, it would be strange to utter 'I would like a piece of that cake,' where the pointing gesture accompanies the utterance of 'I,' but not 'that cake.'²¹ Finally, speakers likewise often correct their choice of gesture when the deictic gesture actually performed is not one with the desired semantic contribution (Kendon, 2004).

¹⁹ See Wilkins (2003) and Kendon (2004) for details and further examples and references.

²⁰ Convention, for present purposes, is understood *à la* Lewis (1969), though the point can be made without subscribing to all the details of the Lewisean account. What matters is that the form to meaning mapping is arbitrary.

²¹ See also Schlenker (2018).

In addition to these signs of conventionality, there are further good reasons not to treat deictic gestures as just one among many pieces of evidence that defeasibly point to a particular interpretation. In (23), 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 (and other things being equal), grammar commits the speaker to an anaphoric reading of the pronoun.²² For instance, even if Mary is jumping up and down in front of the interlocutors, attracting their psychological attention, and being clearly the most salient entity in the speech situation, unless an overt signal establishes her as the referent of ‘she,’ the anaphoric reading will still be the one recovered.²³

Indeed, some means of explicit signaling of a shift in 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. As is commonly noted, beginning a conversation with a pronoun is normally infelicitous, or odd, unless the referent is somehow specified, typically by a pointing gesture:

(26) She is happy.

Unless I point at a particular entity in uttering (26), or unless other resources already place a referent at the center of attention, (26) will be odd; the utterance will appear incomplete.²⁴ Note that it is irrelevant for my point that one can use (26), sometimes, perhaps for a surprising effect, even when such resources aren’t in place, as for instance, would be the case if (26) were the first sentence of a novel, or the first thing I uttered upon having met you for the first time. Though

²² The “other things being equal” clause is required since, as we will learn shortly, there are other conventional means that can signal a shift in attention. The thought is that (17), uttered without a focal stress on the pronoun, and out of the blue, receives naturally an anaphoric interpretation.

²³ Again, I’m assuming that (23) is not embedded within a larger discourse, and that the pronoun is not stressed. Intonation, like pointing, is an attention-shifting mechanism, and can consequently change interpretation. Similarly, embedding an utterance within a larger discourse can give rise to interpretive dependencies that re-structure prominence ranking, and thus affect the interpretation of a pronoun. I return to this point in Chapter 5.

²⁴ Aren’t some of those other possible resources extra-linguistic? At this point, I have not argued yet that they aren’t. But as will become clear in what follows, grammar specifies a diverse set of resources that affect prominence. Deictic gestures can raise entities to prominence, but so can specific intonation patterns, or, as we have seen earlier, indefinite noun phrases. We will see a range of further resources in Chapter 5, and Chapter 12. Recall, further, that deictic gestures are diverse, and need not exploit hand pointings; they can involve head gestures, eye gaze, etc. (see Kendon, 2004, ch. 11). Such mechanisms can contribute prominence-affecting updates even when an utterance occurs discourse initially, and the formalism presented here is expressive enough to make their contributions explicit in its logical form. The prediction of the account is therefore *not* that deixis without a pointing gesture is impossible discourse initially, but that it succeeds only for utterances that employ linguistic mechanisms that independently require construing a particular referent as the center of attention in the current state of the ongoing discourse. A more detailed account of deixis without overt accompanying demonstrative gestures is developed in Stojnić, Stone, and Lepore (2018). I will return to this issue in Chapter 12.

admittedly, even in such cases it might sometimes be possible for the audience to eventually “figure out” what the speaker had in mind, the process by which that occurs is markedly different from the seamless interpretation of pronouns on the fly that we see in normal circumstances. The potentially open-ended reasoning about what the speaker wanted to convey kicks in precisely after one is faced with the infelicity of the utterance. The utterance cannot be properly interpreted on its own, so some kind of repair is needed in order to help guide understanding. My account would simply maintain that a part of the linguistic material in the utterance is missing, and the reasoning is about which material one would have to posit to arrive at an utterance that has a plausible interpretation in this case. In essence, my suggestion is that, insofar as we arrive at a plausible interpretation in such cases, this is a matter of post-semantic repair. If by contrast we maintained that in this deviant case and in the typical case alike, general abductive reasoning were in play, we shouldn’t expect much difference between the two cases: the interpretation should be just as seamless.

Further, the presence of a deictic gesture is hard to override, which suggests that it does not merely serve as one piece of evidence on a par with other contextual, extra-linguistic cues. Suppose I want to say that Mary is my best friend, but due to some accident, perhaps a muscle spasm, or confusion, I point at Sue while saying ‘She is my best friend.’ While you might realize that a mistake of sorts happened, because, say, you might have good reasons to think that Mary is the one I in fact wanted to talk about, it is essential to the case that a *mistake* happened: I accidentally said something I didn’t mean, just as I would have said something I didn’t mean to say had I uttered a wrong word. To make the case more dramatic, suppose that I wanted to insult Ann, but accidentally pointed at Sue, while uttering ‘You are a jerk.’ It seems that in that case I would have to apologize to Sue. I couldn’t simply say that I didn’t say she was a jerk, or that I said that Ann was. Of course, I could say that I *meant* that, but the fact that I’m apologizing and making this excuse is precisely explained by the fact that I didn’t *say* what I meant.²⁵ Notice that the predicament is analogous to the one I would be in if I were to accidentally utter a wrong name. Suppose I say ‘Sue is a jerk,’ accidentally uttering ‘Sue’ instead of ‘Ann.’ You might, in this case, if you have enough evidence, conclude that I meant to say something about Ann, rather than Sue. (Perhaps you know I dislike Ann but not Sue.) However, this doesn’t make ‘Sue’ mean Ann on this occasion. Similarly, realizing that I made a pointing error doesn’t make my pointing any less an instance of pointing at the person I actually pointed at, Sue, and doesn’t make Sue any less the referent of the accompanying occurrence of ‘she.’

²⁵ The case is similar to the famous Carnap/Agnew case from Kaplan (1970). For some discussion, see Reimer (1992); Stojnić, Stone, and Lepore (2018); King (2014a).

Note that had a pointing gesture merely served as a kind of a (defeasible) cue indicating an underlying intention to single out a particular individual (or, alternatively, had the pointing gesture itself had a context-sensitive interpretation), one would expect that if the audience had enough evidence to figure out who the speaker had in mind, the fact that the speaker obviously didn't point at whomever she actually had in mind, wouldn't result in infelicity. As always the audience would just work out the overall most plausible interpretation given their evidence; from this standpoint, the speaker wouldn't have made an error. But this would be a wrong prediction: the speaker clearly would have made an error in such a case.

Finally, note that the linguistic contribution of a pointing gesture affects the inference patterns (17) licenses. In particular, while (27) is valid, (28) isn't.

(27) A woman came in. She sat down. Therefore, some woman sat down.

(28) A woman came in. She [pointing at Betty, the cat] sat down. Therefore some woman sat down.

On the standard account of pronoun resolution, the difference between the logical form of (17) and (23) is not underwritten by grammar of the language: the rules of language underspecify the logical form and it is only after the process of pragmatic interpretation—guided by the epistemic cues available—that a particular logical form is fixed. As a result the difference in form between (17) and (23) boils down to a difference in the representation of a pronoun, that is, to a difference in the choice of a variable that represents the pronoun: on the anaphoric reading, the pronoun is co-indexed with the preceding indefinite and, so, bound, while on the deictic reading it is represented with a free variable. This suggests that the ambiguity is in the pronoun itself. Further, it also suggests that, if we want to capture the intuitive difference between (27) and (28) as a difference in validity of the logical form, we would have to understand validity partly as a pragmatic notion, as it is the pragmatic reasoning that guides the choice of the representation of the pronoun.²⁶

²⁶ Since in neither (27) nor (28) is the pronoun syntactically bound, it is crucial for delivering the two readings to calculate in how both the indefinite, and the pointing gesture, can affect the interpretation of the pronoun in the second sentence. This means that one cannot simply capture validity by relativizing an inference pattern to a fixed context: it is crucial for the interpretation of 'she sat down' in (27) to calculate in how the first sentence affects the context, and the subsequent interpretation of the pronoun in the second sentence. Note that, for the same reason, it would be misleading to focus on context-sentence pairs, *à la* Kaplan: due to the interpretive dependencies that motivate the dynamic semantic framework, the information carried by individual sentences

Both of these suggestions seem misguided. First, as discussed earlier, English has only one unambiguous word ‘she’. Different occurrences of ‘she’ do, of course, receive different interpretations, as in (27) and (28), but this is because ‘she’ is a dependent expression, whose interpretation, given its unambiguous standing character, varies with the changes in the surrounding linguistic environment it is embedded in, for example, one involving an indefinite or a pointing gesture. Since pointing gestures and the effects of indefinites that create such dependent interpretations seem to be conventionalized, leaving them out of the logical form of a sentence, while assigning varying representations to different occurrences of the pronoun gets things exactly backwards. It represents the difference between (27) and (28) as a difference in the representation of a single ambiguous expression, rather than as a difference in the linguistic material the two utterances contain, which in turn leads to a difference in the interpretation of an unambiguous, but dependent linguistic element.

Further, because representing the effect of a gesture within the logical form allows us to capture the difference in the linguistic material the two discourses contain, we are able to capture the validity of (27) as a matter of logical form alone. Given our understanding of truth, we say that an argument is valid, just in case for any input assignment (on any model) it describes a true update. Then (24) entails that some man came in and sat down, as a matter of logical form; but (25), on the other hand, does not.²⁷ This is a welcome result: the validity of (27) is preserved under the substitution of non-functional expressions, which means it is underwritten by form alone.²⁸

Given that the effects of a demonstrative gesture are conventionalized and, moreover, since they affect licensing of inference patterns, they should be represented independently in the logical form. Representing them in this way allows us to preserve the tight connection between validity and logical form. At the same time, it allows us to vindicate the idea that pronouns are unambiguous: a pronoun, on my account, receives an unambiguous, uniform representation in both (27) and (28), as should be the case given the fact that there is only one unambiguous word ‘she’ in English. As the corresponding representations in

often depends on the dynamics of context affected by a prior discourse. I shall return to the problem of characterizing validity in a dynamic setting in more detail in Part III.

²⁷ More precisely, within our dynamic system, an 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 Appendix A. Note, before, we raised concerns regarding our understanding of validity and logical form within a framework that allows for a shifty context. We see now a way to capture validity as a matter of form, even given a shifty context. I will discuss the dynamic account of validity and how it relates to a Kaplanian ban on context-shifting at length in Part III.

²⁸ Of course, the pronoun and its antecedent have to agree in their ϕ -features—in person, number and gender. We will make this explicit in our logical form in Chapter 5.

(24) and (25) make clear, it is the discourse within which a pronoun occurs that is ambiguous, rather than the pronoun itself.²⁹

What I have argued, then, is that a pronoun accompanied by a deictic gesture has a dependent interpretation fixed by the conventional attention-shifting effect the gesture has on the attentional state of the discourse. This means that the meaning of a pronoun occurring in tandem with a pointing gesture is determined through this convention, and the interpretive process involved in the resolution of the pronoun lacks the open-ended abductive character typically assumed. But at this point one might wonder, isn't the kind of pragmatic reasoning traditionally understood as guiding the interpretation of a pronoun necessary in determining the *demonstratum* of a demonstrative gesture itself? And didn't we just then replace an intentionalist account of pronoun resolution with an intentionalist account of demonstrative gestures? That is, in interpreting a gesture, isn't it ultimately the speaker's intention that fixes its *demonstratum*, and reasoning about the speaker's intentions that guides the interpretation?

Even if we had to subscribe to an intentionalist, or more generally, extra-linguistic, account of demonstrations, given our previous considerations, the contribution of demonstrations should still be represented separately in the logical form. Even if the grammar were to underspecify the *demonstratum* of a pointing gesture, this underspecification shouldn't be represented as an

²⁹ An anonymous reader suggested that there may be meaning related conventions that are not a part of the logical form: for instance, they suggest that quotation would be such an example, as it is clearly conventional, and it clearly bears on meaning. Or, for a hypothetical example, suppose our linguistic practice included a conventional regularity that every time that speakers lie, they scratch their heads. It is worth pointing out that some of the most prominent accounts of quotation do indeed maintain that the effects of quotation are factored in the logical form (see in particular Cappelen and Lepore (2007), for discussion of these issues). The motivation for this, indeed, lies in part in the aim of capturing the logical properties of quotation. And it is reasonable to assume similar considerations would apply to the effects of scratchings should there be such a convention. But it is worth clarifying that the reason for representing an effect of a linguistic convention such as the one I argue is associated with pointing isn't simply that pointings are conventionalized. Rather, first, they are conventionalized in the same way other arbitrary form-to-meaning mappings are, in that other mappings could have been—and, indeed, in different linguistic communities, are—employed instead. Second, their conventionalized effects affect the interpretation of truth-conditions and underwrite logical inference. And, finally, they are subject to well-formedness constraints. A pointing gesture has to be synchronized with the speech in a similar way prosody has to be so synchronized in order for the utterance to count as well-formed. That there should be such constraints on well-formedness would be really surprising if pointing effects were not integrated in the logical form (indeed, it seems somewhat arbitrary to exclude the effects of pointings while including effects of prosody). Finally, a dialectical point: in a certain sense everyone has to represent the effect of a pointing gesture in the logical form if they are to capture the correct interpretation and inference licensing potential of examples like (28). On the standard account this is represented by associating a pronoun with a specific index. But this seems misguided: once we agree that pointings are conventionalized, and that their effects have to be somehow reflected in the logical form, it seems misguided to reflect them in the representation of an otherwise linguistically unambiguous pronoun. The pronoun and the pointing gesture make distinct, separate contributions to the logical form. I thank an anonymous reader for pressing this point, and for suggesting the analogy with quotation and the hypothetical example of a lying convention.

ambiguity or underspecification in the pronoun. If the fully specified logical form resolves the underspecification of the pointing gesture rather than the pronoun, then this resolution shouldn't be reflected as a choice of an index on the pronoun, but it should be represented as the contribution of the gesture itself. So, even if this were the right account of demonstrations, the pronoun would still receive a fully specified meaning given the antecedently set up context. That said, the extra-linguistic account of demonstrations should be rejected.

Intentionalism about pointings faces the same problem as intentionalism about pronoun resolution. Intentions are neither necessary, nor sufficient for securing the demonstratum. If you ask me who is a jerk, and in response I (accidentally) point at Mary, while intending to point at Sue, you will have understood me as saying that Mary is a jerk, regardless of the fact that I intended to point at Sue, and did not intend to point at Mary.

Notice that the deictic gesture in (23) is represented as an act of pointing at a particular *demonstratum*, Betty, not an act of pointing at x for some contextually determined x . That is, pointing is not semantically interpreted as having a context-sensitive meaning, which given a context (and together with potentially extra-linguistic resources that context makes available), determines a referent. Rather, pointing gestures are ambiguous between multiple possible forms, for example, *pointing at Betty*, *pointing at her tail*, *pointing at one of her whiskers*, etc. In this way, they are akin to names. There are numerous individuals bearing the name 'Betty.' When I say 'Betty is a gray cat,' you need to do some interpretive reasoning to disambiguate whose name—or, more accurately, *which* name—was actually featured in my utterance.³⁰ Of course, disambiguation can involve reasoning drawing on various resources, including reasoning about speaker intentions. But this means that such resources are exploited in recognizing the form uttered, not in contributing meaning to an underspecified form.³¹

³⁰ The question of individuation of proper names is complex, and going into it at length here would be a distraction. Let me briefly note that, while there is a sense in which individuals called 'Betty' have the same name, it is still important to note that names, on their standard uses, do behave ambiguously: 'I greeted Betty, and brushed her fur' has no reading according to which I greeted and brushed distinct individuals named 'Betty'. This suggests that either 'Betty' is lexically ambiguous (so there are different names 'Betty' for different individuals), or it is context-sensitive, and underspecified, so that each occurrence of 'Betty' has to be distinguished with a (contextually disambiguated) index. This is so, regardless of how we deal with the semantics of proper names. Here I have in mind the commonly assumed former view (for defense, see Kripke (1980, 7–9), Kaplan (1989b; 1990) and Soames (2002); c.f. Schoubye, (2017) for a dissenting view). But ultimately, however things stand with respect to names, what matters for my case is just that the ambiguity view is the one we should endorse for demonstrations. The point of analogy is that demonstrations are like names, insofar as names are ambiguous, directly referential, expressions.

³¹ For some further discussion, see Stojnić et al. (2020); see also Nowak and Michaelson (2020).

The same goes for pointings. The interpretive work in (23) comes in settling the *form* of the gesture: whether it is a pointing at Betty or a pointing at something else. But *disambiguation* is pre-semantic, in Kaplan's sense: it involves the interpretive work needed to settle the *linguistic form* of an utterance, not to assign content to the form (Kaplan, 1989a). Disambiguation is distinct from semantic interpretation: it is only once the form is disambiguated that it can be semantically interpreted. Semantics determines what an expression means, but not which expression was uttered. And, while general pragmatic reasoning about the speaker's intentions and available epistemic cues plays no role in semantic interpretation, it *can* play a role in guiding the audience to recognize *which form* of a pointing gesture was performed. But this is the role they can play in the disambiguation of *any* ambiguity, for example, in the disambiguation of a use of the word 'bank,' or a name, 'Betty.' And, as with other ambiguities, conventions governing demonstrative actions constrain possible disambiguations. A flat hand shape with the palm up, fingers toward the audience, allows for a certain range of interpretations, but not others; similarly, for an extended index finger, or the word 'bank,' or the name 'Betty.'

Taking stock, I have argued that the effects of an attention-shift triggered by an indefinite noun phrase or an act of demonstration are conventionalized and should be represented independently in logical form. These items come with linguistically specified contribution to the ranking of the prominence of candidate referents in a context, manipulating what is at the center of attention. Representing these contributions explicitly allows us to dispel the illusion of the ambiguity in the pronoun in the examples we have considered: we can assign one unambiguous meaning to a pronoun, uniformly represented in the logical form, that fully determines its resolution in a given context. To a first approximation, pronouns pick out the most prominent candidate interpretation, the one at the center of attention, at a given point in discourse. Consequently, 'she' behaves like a pure indexical just as 'I' does; the sole difference is that while 'I' is sensitive to the speaker-parameter of the context, 'she' is sensitive to the attentional state, that is, to the dynamically evolving, yet linguistically specified, prominence of candidate resolutions.

As promised, at this point, a few refinement are in order. We said that the meaning of @ is our first approximation to the meaning of a pronoun in English. Obviously, this cannot be the whole story. We want to be able to capture the difference in meaning between different pronouns. And while the meaning of @ captures what is shared by the character of different pronouns, for example, 'he' and 'she,' namely, that they are prominence-sensitive, it fails to capture the differences in meaning between different pronouns. Luckily, it's relatively straightforward to introduce necessary refinements to capture these differences.

First, pronouns come with person, gender and number requirements (the ϕ -features) that must be satisfied in pronoun resolution. ‘She’ cannot be interpreted as co-varying with ‘a boy,’ because ‘she’ must be resolved to a *female*, and it cannot be interpreted as co-varying with the witness for ‘some girls,’ because ‘she’ must have a *singular* interpretation. Similarly, ‘we’ and ‘they’ cannot be interpreted as co-varying since ‘they’ requires a third-person interpretation *disjoint* from speaker and addressee, while ‘we’ requires a first-person pronoun interpretation that includes the speaker. The person, number and gender requirements are *linguistic*, built into the meaning of pronouns: it matters whether a speaker utters ‘he,’ ‘she’ or ‘they,’ precisely because they differ in meaning.³²

Second, we need to model the syntactic constraints on interpretation. For instance, we want to capture that ‘her’ in (29) cannot semantically co-vary with Betty, as that would violate Principle B of Chomsky’s Binding Theory, which requires roughly, that a non-reflexive pronoun must not be bound by an expression in the same clause.³³

(29) Betty met her.

We can build these constraints into the character of a demonstrative pronoun: we say a pronoun denotes the value stored at the highest ranked position on the stack that respects grammatical features and applicable syntactic principles. The pronoun thus performs a search on the attentional state, considering each candidate in order until it finds one that fits its character. That is, as a matter of its character, the pronoun automatically selects the top-ranked entity in the given context which satisfies the constraints the character encodes. Definition 4.1 represents the character of ‘she,’ integrating the constraints of semantics and syntax. (We can define the characters of other demonstrative pronouns *mutatis mutandis*.)

Definition 4.1 (The standing linguistic meaning of ‘she’) Interpreted relative to an assignment g ‘she’ denotes g ’s highest ranked entity that is singular, feminine, 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.

³² A speaker can manage to ‘speaker-refer’ in the sense of Kripke, to a woman with the pronoun ‘he’ (or a definite ‘that man’) but this should not be confused with semantic reference.

³³ More precisely, the antecedent of a non-reflexive anaphoric pronoun must be free in its governing category; it must *not* be local, or c-command the pronoun (Chomsky, 1981).

We modify our formalization accordingly: while @ denotes the value stored at top-ranked position, where *she* is a predicate representing the constraints associated with the English third-person singular, female pronoun—requiring that it receives an interpretation which is singular, feminine, 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—@*she* is an individual expression denoting the highest ranked entity from the current assignment that satisfies the property denoted by *she*. Then, (24)–(25) become (30)–(31).

(30) $\langle \alpha \rangle; [woman(@)]; [came.in(@)]; [sit.down(@she)]$

(31) $\langle \alpha \rangle; [woman(@)]; [came.in(@)]; \langle \pi b \rangle; [sit.down(@she)]$

We now have a full characterization of the standing linguistic meanings, that is, characters, of demonstrative pronouns.

Of course, so far, we have only focused on a limited range of cases, specifically on the distinction between (27) and (28), and the way in which the pronoun is sensitive to prominence-affecting effects of indefinites and deictic gestures. But the effects of indefinites and deictic gestures are but an example of linguistic mechanisms governing prominence and affecting the resolution of pronouns. Language encodes a wide range of prominence-affecting mechanisms—discourse conventions—which in turn maintain the attentional state of the discourse. They determine what is at the center of attention at any given point in discourse, antecedently setting up the context for the interpretation of subsequent pronouns, so that a pronoun can just select what is at the center of attention at a particular point in the discourse. I now turn to uncovering the effects of these discourse conventions.

Context and Discourse Conventions

Consider (32) and (33):

- (32) A woman met Sue. She greeted her.
- (33) John was disappointed with Tim.
 a. He fired him.
 b. He disobeyed him.

In (32), there are two antecedents and two pronouns. We know that a syntactic constraint on binding explains why English speakers do not treat the pronouns in (32) as co-varying (violation of Principle B). But what explains why we recover a particular interpretation for ‘she’ and ‘her’? Specifically, English speakers would normally interpret ‘she’ in (32) as co-varying with the indefinite, ‘a woman,’ introduced in the first sentence, and ‘her’ to Sue.¹

Clearly, we have to allow that indefinites and names alike can update the stack, as we want to account for the examples like the following:

- (34) Sue came in. She sat down.

And obviously the same point is easily made with indefinites only, as witnessed by our earlier example, from Chapter 3, repeated below:

- (22) A woman met a girl. She greeted her.

If each indefinite pushes a new value to the top, then we would predict exactly the opposite pattern from what we actually witness.

Turning to (33), here we have one and the same sentence followed by two distinct continuations. Yet, in one case, ‘he’ in the subject position picks out

¹ I’m assuming (32) is uttered without accompanying demonstrative gestures. That the noun phrase in subject position is preferred as an antecedent for the subsequent anaphor over the one in object position is well-documented by corpus and psycholinguistic studies. (See Walker, Masayo, and Cote, 1994; Kameyama, 1996; Bittner, 2014; and references therein.) More precisely, these studies show that prominence is linked to the grammatical role in which the noun phrase is realized, where the subject takes precedence over the direct object, which takes precedence over the indirect object.

John and in the other Tim. But both Tim and John cannot be simultaneously the highest ranked individual in the attentional state of the discourse that results after updating with the first sentence. What can we make of this within the current framework?

Taking these challenges in turn, we start with (32). In English, attention ranking follows the grammatical role in which noun phrases are realized; the noun phrase in subject position takes precedence over the one in the direct object position, which takes precedence over the one in the indirect object position.² This is why, in (32), 'she' is resolved to the candidate referent introduced by the previous subject, 'a woman,' rather than one introduced by the previous object 'Sue.' From our standpoint, what is crucial is that this preference for referents introduced by noun phrases in subject position is a grammatical feature of English, a reflection of the fact that English is a so-called subject-prominent language. This is not a feature that is universally shared across languages. Other languages, topic-prominent ones, grammaticize prominence with specialized morphemes, like topic markers; languages with flexible syntax utilize word order.³ This variability is, again, a sign of a linguistic convention in play. It cannot be some universal feature of human psychology that underscores the preference.

Accordingly, 'her' in (32) likewise selects the most prominent entity satisfying its character. But while the ranking remains constant, the constraints imposed by the character are different. Due to Principle B, 'her' selects the highest ranked entity distinct from the one selected by 'she,' which is the one introduced by the noun phrase in the object position, that is, Sue. The same explanation applies to (22).

We can represent this formally by linking grammatical roles to specific designated positions on the stack of candidate interpretations: the subject corresponds to position 0, the direct object to 1, the indirect object to 2, representing receding prominence. Then, instead of letting noun phrases contribute updates that indiscriminately push values always to the top of the stack, let them update the stack in accordance with the grammatical role in which they occur: a noun phrase realized in the subject position updates position 0, reflecting comparative prominence of the referent introduced by the subject noun phrase, the one in the direct object position updates 1, the second-ranked one, and that in the indirect object position updates position 2, the third-ranked one. We can represent this explicitly in our formal language by expressions of forms $\langle \alpha n \rangle$ and $\langle \pi n c \rangle$ interpreted as updates that push referents to a specific position n in the stack (and

² This observation is well-documented. See, for example, Kameyama (1996); Kehler (2002); Bittner (2014).

³ See Walker, Masayo, and Cote (1994); Kameyama (1996); Bittner (2014).

assuming the value of n follows compositionally from the grammatical status of the expression being represented, e.g., 0 for the subject, etc.), pushing all other one position down, as before. Distinguishing the representation of pronouns from argument structure, we can use variables to represent the argument structure explicitly, so that, for example, $[woman(x_0)]$ is interpreted as requiring that the subject be a woman, and $[likes(x_0, x_1)]$ as requiring that the subject likes the (direct) object.⁴ Incorporating these observations, the first sentence of (17) and the first sentence of (32) are formalized as (35) and (36), respectively.

(35) $\langle \alpha 0 \rangle; [woman(x_0)]; [came.in(x_0)]$

(36) $\langle \alpha 0 \rangle; [woman(x_0)]; \langle \pi 1 s \rangle; [met(x_0, x_1)]$

Given this setup, we should also represent the fact that a pronoun is assigned a grammatical role in its own sentence (e.g., it can be in the subject, or direct or indirect object position). Thus, instead of simply interpreting a pronoun as a variable linked to a particular position, we represent it, as we do with all other noun phrases, as a further update to the attentional state of the discourse. So, for instance, a pronoun in the subject position selects the highest ranked value that satisfies its character at the current attentional state of the discourse, and promotes this value to the position 0, just as any other noun phrase in the subject position does. Since in (17), the pronoun ‘she’ is in the subject position, it should thus be represented with the update expression— $\langle \pi 0 @ she \rangle$ —interpreted as an update that searches for the top-ranked candidate referent that satisfies the property denoted by *she*, and promotes that entity as the new value of the top-ranked subject position, 0. So, (17) then receives the representation in (37).⁵

(37) $\langle \alpha 0 \rangle; [woman(x_0)]; [came.in(x_0)]; \langle \pi 0 @ she \rangle; [sit.down(x_0)]$

The same applies to the representation of (32). The pronoun in the subject position, ‘she,’ selects the top-ranked referent that satisfies the character of the pronoun—a woman introduced by the indefinite in the subject position of the previous sentence—and re-introduces it to the top. What about the object pronoun ‘her’? To represent the constraint imposed by Principle B in resolving the pronoun, let $@she^{x_0}$ be an expression that denotes the most prominent

⁴ Notice that this way of exploiting indexing on variables still uses variables as markers of prominence rather than as names.

⁵ Notice that the pronoun re-stores (i.e., pushes a copy of) an old candidate referent—the one selected by $@she$ —to the designated position on the list of prominent candidates. (If the entity is already at the top-ranked position, as in (37), the update associated with the pronoun will not change possible interpretations of subsequent ones.)

referent whose value satisfies the pronoun other than x_0 —the top-ranked female other than the subject. So, (32) is represented as (38):

- (38) $\langle \alpha 0 \rangle$; [$woman(x_0)$]; $\langle \pi 1 s \rangle$; [$met(x_0, x_1)$];
 $\langle \pi 0 @ she \rangle$; $\langle \pi 1 @ she^{x_0} \rangle$; [$greeted(x_0, x_1)$]

This allows us to incorporate the effect of subject-object preference on prominence. To see how the pieces fit together, let us work through the details of (38). First the indefinite introduces an (unspecified) witness to the subject position ($\langle \alpha 0 \rangle$), and the witness is further constrained to be a woman ($[woman(x_0)]$). Then the object noun phrase, ‘Sue,’ introduces the referent of the name, Sue, to the object position, 1 ($\langle \pi 1 s \rangle$). Further, we require that the subject met the object—some woman met Sue ($[met(x_0, x_1)]$). Now we reach the second sentence. The pronoun ‘she’ searches for the top-ranked entity satisfying the character in the current attentional state of the discourse—which at this point would be the witness for the indefinite—and, since the pronoun is in the subject position, promotes it to position 0 ($\langle \pi 0 @ she \rangle$). The object pronoun, ‘her,’ searches for the top-ranked entity in the current attentional state other than the subject—which will be Sue—and, since the pronoun is in the object position, places it at position 1 ($\langle \pi 1 @ she^{x_0} \rangle$). Finally, we require that the subject greeted the object—the aforementioned woman, greeted Sue ($[greeted(x_0, x_1)]$). As expected, the whole discourse is true just in case a woman met Sue, and greeted her.

We have thus met the first challenge. But what about our second challenge, the one reflected in (33), repeated below? What makes John prominent in (33-a) and Tim in (33-b)?

- (33) John was disappointed with Tim.
 a. He fired him.
 b. He disobeyed him.

At this point, discourse coherence enters stage.

The key idea behind discourse coherence is an obvious, yet often overlooked, observation: a discourse is more than a random sequence of grammatical sentences. Successive contributions to the discourse must be linked into a coherent whole by a recognizable flow of interpretive relationships. To make this idea concrete, consider the contrast between (39) and (40), from Hobbs (1979).

- (39) John took the train from Paris to Istanbul. He has family there.
 (40) John took the train from Paris to Istanbul. He likes spinach.

While (39) is perfectly smooth, (40) is odd. What explains the contrast? (40) doesn't just reveal two random facts about John in no particular order: it suggests that the reason he went to Istanbul was to visit his family. Understanding this connection between the two consecutive utterances in (39) is a part of understanding the contribution of the discourse: if one doesn't realize this, one doesn't fully understand the discourse in (39). Just as well, we see the requirement for a coherent connection that ties the two utterances together in (40). The two utterances seem unrelated out of the blue. We are left wondering how the train trip to Istanbul is explained by the preference for spinach. (Is Istanbul known for its spinach? Maybe spinach causes a fear of flying?) The failure to establish the connection results in a felt oddity of (40): we are left with the impression that we haven't fully understood the discourse.

An influential way of modeling this observation has been developed within the Coherence theoretic framework for modeling discourse interpretation, which originated in computational linguistics and artificial intelligence research (Hobbs, 1979, 1990; Knott, 1996; Kehler, 2002; Asher and Lascarides, 2003; Webber et al., 2003, *inter alia*). Coherence theory captures the observation by positing an implicit organization to the discourse that signals how its individual utterances are connected into a coherent whole. A discourse must be organized by recognizable connections—coherence relations—that weave individual utterances into a coherent structure. So, in (39), the discourse is organized around a coherence relation of Explanation, which links the second sentence to the first. This is why we understand that the train trip to Istanbul is explained by the location of John's family. Coherence theory maintains we have to represent this connection explicitly to capture the correct interpretation of (39). Note that the requirement that a discourse must be coherent is strikingly evident in the interpretive effort (40) elicits. Given apparently unrelated facts about John in (40), we search for a connection. We are trying to figure out how the second sentence explains the first. The failure to confirm a relation results in infelicity. Interlocutors can draw on various resources in disambiguating—in determining which relation organizes the discourse.⁶ But, just as clearly, a failure to establish a coherent connection constitutes a failure to understand the discourse.⁷

⁶ Notice, however, that general reasoning *alone* will not explain why we are searching for an explanatory connection in (40), since after all a number of other relations are possible, but not pursued. See Stojnić, Stone, and Lepore (2018) for further discussion.

⁷ It is worth pointing out that representing coherence relations explicitly does not mean an explanatory relation is a part of the truth-conditional content of (39). One can maintain, as I do here, that connections signaled by coherence relations, like Explanation in (39), are not conveyed as a conversational implicature while also maintaining that the relational connection, here one of explanation, is not *at issue*, in the sense of Tonhauser et al. (2013). The formalism in this chapter

What types of coherence relations can we expect to find in discourse? One useful typology is proposed by Kehler (2002) who, channeling David Hume, argues that coherence relations cluster into three qualitatively different categories, reflecting alternative strategies for organizing discourse.⁸ These are illustrated in (41)–(43).

(41) Max spilt a bucket of water. He tripped on his shoelace.

(42) Max spilt a bucket of water. He then mopped the floor.

(43) Max spilt a bucket of water. John dropped a jar of cookies.⁹

(41) illustrates the kind of an explanatory discourse we already saw in (39). For Kehler, Explanation is an instance of a broader class of Cause-effect (or Event-result) relations. (42), in turn, gives an extended description of a sequence unfolding events: it is organized around Narration relation, which for Kehler belongs to a broader class of Contiguity relations. (43) exemplifies what Kehler calls Resemblance relations, organizing a discourse around comparisons and contrasts. Specifically, in (43), the organizing relation is Parallel, as the discourse is drawing a parallel between Max's and John's respective accidents.

Coherence relations thus signal relationships among utterances in discourse. However, these relations also shape how other material in a discourse is interpreted. Of particular importance for our purposes is their effect on context-dependent elements.¹⁰ As already gestured at in Chapter 1, it turns out that the task of establishing coherence and that of pronoun resolution are not independent.¹¹ In (33–b), only resolving 'he' to Tim allows for a plausible explanation

includes only one dimension of semantic content, in the sense of Potts (2005), though we will later see how this can be modified to capture conventionalized not *at issue* content as well.

⁸ See Hume (2007/1738). Kehler'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. In any case, the classification of discourse relations is an area of active ongoing research. Cf. Mann and Thompson (1988); Knott (1996); Asher and Lascarides (2003); Wolf and Gibson (2006). As will become clear in what follows, the important element for us will be the attention-shifting effect of coherence relations. So, the overlaps and distinctions between different typologies on offer will not affect our main conclusions, and the account needn't be explicitly committed to Kehler's typology.

⁹ See Kehler (2002). (41) is adapted by Kehler from Asher and Lascarides (2003).

¹⁰ Discourse coherence has been argued to affect a number of linguistic phenomena including verb phrase ellipsis, gapping, extraction from coordinate structures, pronominal reference resolution, tense, co-verbal gesture interpretation, etc. A part of the task of the present work is to argue it affects context-sensitivity resolution, quite generally and that it does so as a matter of linguistic convention, rather than as a defeasible pragmatic bias.

¹¹ As will become clear, coherence does not merely affect the resolution of demonstrative pronouns, but other context-sensitive elements as well. For instance, as illustrated by (41)–(43), coherence affects the resolution of temporal anaphora in discourse as well. In these examples we infer a temporal relation between the spill described initially and the tripping, mopping, or dropping described next that matches the inferred coherence relation. Reference and coherence relations fit together in such cases. (See Kehler, 1994, and Bittner, 2014, for further discussion.)

of John's disappointment. In (39), only resolving 'there' to Istanbul allows for a plausible explanation of John's trip.¹² (44), from Hobbs (1985) provides another illustration:

- (44) John can open Bill's safe.
 (a) He knows the combination.
 (b) He should change the combination.

In (44a) the organizing relation is naturally understood to be Explanation (John can open the safe because he knows the combination), and so 'he' is resolved to John. In (44b), we understand the discourse as organized by Result (John can open the safe, therefore Bill should change the combination), and the pronoun is resolved to Bill.

Indeed, a large number of studies have experimentally confirmed the interdependence of resolving pronouns and establishing coherence.¹³ We have already seen an illustrative example in (3), from Smyth (1994), repeated below:

- (3) Phil tickled Stanley. Liz poked him.¹⁴

As discussed briefly in Chapter 1, speakers tend to interpret (3) in either of two ways: either assuming Liz's action was *prompted* by Phil's, giving rise to cause-effect interpretation, or assuming Liz's action was *similar* to Phil's in certain respects, giving rise to resemblance interpretation. The first reading understands the discourse as organized by Result relation, giving rise to the understanding that 'him' must refer to Phil. The second reading organizes the discourse around Parallel relation, 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.¹⁵ Pragmatic theories

¹² Note that changing (40) 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 (39) and (40), the pronoun is understood to refer to Istanbul. This further supports the 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 (40), because it is clear how the second sentence provides an explanation of the first, as is required by Explanation: John has a preference for the spinach from Istanbul. But that is beside the point.)

¹³ See, for instance, Kaiser (2009), and references therein.

¹⁴ We assume no demonstrative gestures are accompanying (3).

¹⁵ When (3) is embedded within a larger discourse, we might be able to get other resolutions, but this is because different embeddings can license different coherence relations. Similarly, prosodic marking can affect pronoun resolution, as can demonstrative gestures (discussed earlier). 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

of reference resolution—even standard coherence theoretic accounts—take this to show that reasoning about speaker’s referential intentions (what the intended referent of a pronoun is) and about her intentions in organizing the discourse (which coherence relation was meant), mutually constrain and inform one another. The interpretation involves abductive reasoning with an aim of detecting the overall most plausible interpretation. On this understanding, coherence relations place merely defeasible constraints on pronoun resolution. Yet these studies seem to suggest a tighter connection between coherence relations and pronoun resolution than Coherence theorists have been inclined to posit.

There is a reason to think that it’s a mistake to treat (3) as harboring *two separate* ambiguities, or *two separate* underspecified elements that must be resolved in turn—one involving discourse coherence, another concerning pronoun resolution. The examples suggest that, once a coherence relation is established, a certain pronoun resolution is automatically set up. As we shall see shortly, indeed, there’s good reason to conclude that pronoun resolution is *settled* by whichever coherence relations organize a discourse. Much like a deictic gesture of pointing, coherence relations are another sort of mechanism that linguistically effects changes in the attentional state of a discourse.

If this is right, then we should capture these interpretive effects formally by representing coherence relations and their effects on attentional state of a discourse explicitly in logical form, for the same reasons we should represent the effects of deictic gestures in the logical form.¹⁶ So, coherence relations will be represented in the logical form not only as signaling interpretive relations between bits of discourse, but moreover, as inducing grammatically specified shifts in attention. Departing from standard Coherence theory, on the present account, coherence relations induce shifts in attention as a matter of a linguistic convention, much like noun phrases or pointing gestures do. First, let us look at

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).

¹⁶ It’s in line with standard coherence theoretic approach to maintain that coherence relations themselves are represented in the logical form of a discourse. The literature on discourse coherence points to both syntactic and semantic constraints on coherence relations, that suggests coherence relations are compositionally integrated with the rest of the interpretation. See, in particular Asher and Lascarides (2003); Webber et al. (2003). In addition, we shall see reasons to represent not only coherence relations, but also their effects on the attentional state within the logical form of a discourse.

how this account captures data in (33-a) and (33-b), and then turn to defense of the view that the attention-shifts associated with coherence relations are grammatically encoded, not pragmatically implied.

The difference between resolutions in (33-a) and (33-b) is fixed by the coherence relation that figures in their respective logical forms; these come with different attention-shifting updates, which affect, semantically, the resolution of subsequent pronouns. In (33-a), the coherence relation is Result: the eventuality described by the first sentence prompted the one described by the second. John was disappointed with Tim, so he fired him. In (33-b), the coherence relation is Explanation: the content of the second sentence is taken to explain the content of the first. John was disappointed with Tim, because he disobeyed him. According to Coherence theory, one or the other of these relations surfaces in formal representations of (33-a) and (33-b). The claim is that 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: in addition to establishing an inferential connection, they induce attention-shifts. When a formal representation features Result, 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.¹⁷ Accordingly, (33-a) is represented as (45), and (33-b) as (46):

- (45) $\langle \pi 0j \rangle; \langle \pi 1t \rangle; [was.disappointed.with(x_0, x_1)];$
 $[Result(x_0, x_1)]; \langle \pi 0x_0 \rangle;$
 $\langle \pi 0@he \rangle; \langle \pi 1@he^{x_0} \rangle; [fired(x_0, x_1)]$
- (46) $\langle \pi 0j \rangle; \langle \pi 1t \rangle; [was.disappointed.with(x_0, x_1)];$
 $[Explanation(x_0, x_1)]; \langle \pi 0x_1 \rangle;$
 $\langle \pi 0@he \rangle; \langle \alpha 1 \rangle; [disobeyed(x_0, x_1)]$

Let us work through the key steps in (45) and (46). The first line in (45) represents the contribution of the first sentence. It introduces John into subject position, and Tim into direct object position, requiring that John was disappointed with Tim. The second sentence describes a result of the eventuality described by the first one; we represent the contribution of Result relation as $[Result(x_0, x_1)]; \langle \pi 0x_0 \rangle$, interpreted as an update that requires that the eventuality described by the second sentence describes a result of that described by the

¹⁷ As we shall see below, not all explanatory coherence relations raise the direct object to prominence. It's important to bear in mind that coherence relations encode linguistic inter-dependencies 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.

first, and correspondingly, renders the subject, x_0 , prominent in the attentional state of the discourse.¹⁸ ‘He’ picks out the most prominent candidate that satisfies its character, which is John, while ‘him’ picks out the currently most prominent candidate other than John, which is Tim. Further, it is required that the subject (John) fired the object (Tim).

The first line in (46) is the same as in (45). The crucial difference comes next. This discourse harbors Explanation, represented by the update $[Explanation(x_0, x_1)]; \langle \pi 0 x_1 \rangle$. This update requires that the two bits of discourse stand in an explanatory relation, and promotes the entity in object position to prominence. The update associated with the pronoun, $\langle \pi 0 @ he \rangle$, picks out the top-ranked candidate interpretation, but due to the update associated with Explanation, this is now Tim, rather than John. The formula then proceeds with a further condition that ensures that the subject, Tim, disobeyed the object, which is now the second-ranked entity, John. Predicting the difference in the interpretation of ‘he’ in (33-a) and (33-b), we have now met our second challenge.

On this representation, differences in pronoun resolution *follow from* the coherence relations that organize the discourse and update the attentional state, and *not* from an open-ended pragmatic process of pronoun resolution. Now, just as it is easy to represent the contribution of demonstrative gestures in the logical form, so it is easy to represent the contribution of coherence relations. But why do so? Why not 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 but rather a result of a pragmatic default, feeding into defeasible reasoning that occurs once a hearer has established that a particular coherence relation is organizing the discourse. A hearer reasons that in (33-a), since the second sentence describes the result of the eventuality described by the first, it must describe what John was prompted to do as a result of his disappointment. So, it only makes sense that the speaker intends to continue to talk about John. Similarly, in (36-b), since the speaker is describing what about Tim made John disappointed, it certainly makes sense that ‘he’ should select Tim. On this view, intention recognition affects the re-ranking of the list of prominent referents.

¹⁸ This representation of Result and Explanation, as relating entities that are the agents of respective eventualities, suffices to capture the effects on attention that concern us at present. 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, and Stojnić, Stone, and Lepore (2018) as relations between situations. We will introduce a less simplified way of representing the relation of coherence relations in subsequent chapters.

While such reasoning makes sense, it nevertheless seems that attention-shifting operations prompted by coherence relations are grammatically encoded. Speakers and hearers take a restricted set of cues—discourse internal cues—into account in 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 phenomenon, consider (47), from Kehler (2002):

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

Kehler reports (47) is judged to be infelicitous by his subjects. What explains this? By virtue of following ‘admires’ with ‘absolutely worships’—a stronger term in a scalar relationship—the speaker provides clear evidence (47) 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 prominent, and preferred referent for the pronoun in object position in the second clause.¹⁹ And, indeed, reading (47), it seems as if the speaker has erred, inadvertently referring to Reagan as ‘her.’ Of course, since Thatcher has been mentioned in the previous sentence, in the most prominent, subject position, and given that she 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 abductive reasoning that interpreters use to recover the overall most plausible interpretation; yet, this is not what we find.

The infelicitous interpretation, by contrast, is perfectly expected if we assume that the effect of coherence in (47) is a matter of a linguistic rule. The choice of a coherence relation—in this case Parallel—requires that the antecedent of the pronoun in the object position be introduced in the object position. This is what underwrites the judgment of infelicity. This effect cannot be overridden or trumped by pragmatic reasoning: indeed, it itself trumps pragmatic considerations. Plausibility, relevance, and charity considerations all point to a resolution that selects Thatcher as the antecedent: this antecedent is available and matches the grammatical gender of the pronoun. Moreover, she is a well-known object of Bush’s admiration, so selecting this interpretation would be both relevant and plausible. The interpretation actually recovered, according to which the speaker has mistakenly uttered ‘her’ to refer to Reagan, meanwhile is uncharitable: it ascribes an error to the speaker. That it is nevertheless the naturally recovered interpretation, despite all these pragmatic considerations, suggests that Parallel

¹⁹ Kehler himself is not committed to this preference being a result of a conventionalized, linguistic contribution of coherence.

dictates a particular interpretation of a pronoun as a matter of an underlying linguistic rule, which is hard to override by such pragmatic considerations.

This point is analogous to that about the conventionality of demonstrative gestures. Regardless of how salient Thatcher may be in the speech situation, if the speaker is pointing elsewhere while saying ‘her’ we do not take ‘her’ to refer to Thatcher. The Parallel relation in (47) accomplishes a kind of demonstration, indicating Reagan as the referent in a way that’s difficult for a common-sense inference to override.²⁰

Just as with conventionality of demonstrative gestures, further support for the claim that coherence affects prominence as a matter of convention comes from the variation we find across languages. We have already seen that languages differ in how they signal prominence, and they also differ in how they exploit coherence to do so. Many languages have explicit operations for shifting attention, such as grammatical topic marking, or a distinction between topic and non-topic pronouns. They can be more constrained than English in the extent to which they allow for signaling attention-shifts via coherence relations. As an illustration, consider two possible translations of (33-a) in Serbian, a pro-drop language, which allows (sometimes requires) pronouns to remain unpronounced in certain grammatical positions:

- (48) *Džon je bio razočaran Timom.*
 John–NOM is–PRS–3MS be–PPA–3MS disappointed–ADJ–3MS Tim–INS
 ‘John was disappointed with Tim.’
- a. *Otpustio ga je.*
 Fired–PPA–3MS him is–PRS–3MS
 ‘He fired him.’
- b. *On ga je otpustio.*
 He him is–PRS–3MS fired–PPA–3MS
 ‘He fired him.’

There two ways one can translate (33-a) into Serbian. When translated with its third-person male, singular pronoun “dropped” (as in (48-a)), the interpretation is that John fired Tim. Yet when the pronoun is overt as in (48-b), the interpretation is that Tim fired John. Furthermore, in the first case the natural interpretation is, as in English, that the two sentences connect via Result (John

²⁰ Note that focal stress on ‘her’ makes (47) felicitous, the referent of ‘her’ being Thatcher. But on this interpretation, the organizing relation is no longer Parallel: the discourse explains how Bush follows Thatcher’s opinions, not how conservative politicians feel about Reagan. The crucial point is that holding the relation fixed, the pronoun resolution is fixed as well. In other words, there is no reading where the relation is Parallel and the pronoun resolves to Thatcher. But this should be possible if a particular interpretation of the pronoun were merely suggested by default, but not required by Parallel.

fired Tim because he was disappointed with him), but in the second case, we naturally understand this as an Explanation (Tim was the one who fired John, which is why John is disappointed with him.) Moreover, the pronoun resolution remains the same even if we try to force the Result interpretation, by including the explicit discourse marker signaling the intended coherence relation: ‘so’ (‘pa’):

- (49) *Džon je bio razočaran Timom,*
 John-NOM is-PRS-3MS be-PPA-3MS disappointed-ADJ-3MS Tim-INS
 ‘John was disappointed with Tim’
- a. *pa ga je otpustio.*
 so him is-PRSS-3MS fired-PPA-3MS
 ‘so he fired him.’
- b. *pa ga je on otpustio.*
 so him is-PRSS-3MS he fired-PPA-3MS
 ‘so he fired him.’

As before, in (50a) the interpretation is that John fired Tim, because he was disappointed with him. But, (50b), by contrast can only be interpreted as if somehow John being disappointed with Tim prompted Tim to fire him. It seems that the overt pronoun signals a change in prominence that the covert counterpart does not. English, lacking these explicit means of signaling a shift in attention, is more flexible in how it associates coherence relations with attention-shifting updates.²¹

Note that general reasoning can still have a role to play, but this role, again, is not one of assigning content to the form, but rather one of disambiguating which form has been uttered to begin with. So, for instance, (3) is ambiguous between a form 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’ or who they named when they uttered ‘Betty,’ or, as we have seen before, in disambiguating a particular form of a gesture. To interpret, a hearer must first settle disambiguations.²² This may involve assessing the plausibility of

²¹ Cross-linguistic variation in how coherence is exploited to signal prominence is not limited to pronominal domain. Similar variations are seen in the effects of coherence relations on prominence of candidate resolutions for other context-sensitive items. For instance, there is variation in the domain of temporal anaphora resolution (see Bittner, 2014; for further theoretical discussion of some of the examples from Bittner, see Stojnić, at al., 2017), as well as in the domain of modal anaphora, as we shall see in Part II.

²² Note that I’m not taking this priority of disambiguation as a commitment to a specific processing claim. And obviously, one might take an infelicity as evidence that a different disambiguation is in order.

possible coherence relations that could be operative in a given context. It may involve evaluating whether a particular disambiguation of coherence relations delivers a plausible resolution of pronouns. In short, it can serve in recognizing which form of the available ones that grammar delivers was uttered, but not in determining which meaning a particular form takes on.²³ Once a coherence relation is established, pronoun resolution is determined by grammar, not by general reasoning. And, as I have been arguing throughout, any extra-linguistic parameters such as interpretive reasoning may invoke do not serve to *determine* meaning.

Summing up what I've argued for at this point, attention-shifting effects of coherence relations are governed by linguistic rules, not pragmatic reasoning. These attention-shifting updates guide the evolution of the attentional state of a discourse, 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. Hence, a demonstrative pronoun behaves like a pure indexical: it has a uniform, unambiguous meaning and representation, and a fully specified resolution given a (linguistically established) context.

The present Attention-Coherence account, unlike the Kaplanian account and the standard dynamic semantics account, assigns a uniform linguistic meaning—a character—to each pronoun, and it provides a uniform representation of a pronoun, regardless of where it occurs or whether it is interpreted referentially or not. If the account is on the right track, it captures an important fact about the interpretation of demonstrative pronouns that the previous accounts fail to capture. A pronoun is sensitive to an evolving attentional state of the discourse—it's a prominence-sensitive expression—where the attentional state evolves as dictated by linguistic mechanisms that affect the prominence ranking of candidate interpretations: at any given point, the attentional state is a result of a series of attention-shifting updates, which intuitively re-rank candidate resolutions for pronouns. The updates are contributed by the diverse, yet nevertheless linguistic, mechanisms for structuring discourse, including the grammatically specified preference for discourse entities in specific grammatical roles, demonstrating entities with non-verbal (yet linguistic) gestures, and signaling the direction of discourse through coherence relations 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.

²³ See Stojnić, Stone, and Lepore (2018) for a more detailed discussion of the role of general reasoning in disambiguation as opposed to semantic interpretation.

In modeling the attentional state, we have represented it by an appeal to assignment functions modeled as stacks. Thus, pronouns—on their bound and referential uses alike—behave like variables with dependent interpretations, co-varying with antecedent expressions that pushed a value onto the stack.²⁴ However, since stacks are interpreted as modeling relative prominence of candidate interpretations, pronouns are interpreted relative to a prominence ranking, selecting the most prominent interpretation that satisfies their associated character. Though we interpret demonstrative pronouns relative to an assignment function, they are not treated as true demonstratives, but rather as pure indexicals. They are sensitive to a (dynamically evolving) contextual parameter—the attentional state of the discourse—that governs their interpretation. At any point within a discourse, the attentional state allows the pronoun to automatically select an interpretation in accordance with its character. Importantly, an utterance can contain elements that update the attentional state mid-utterance, even between different occurrences of a demonstrative pronoun; we have seen that each noun phrase updates the attentional state in accordance with its grammatical role, or that deictic gestures, or coherence relations between sub-clauses update the attentional state throughout the interpretation of an utterance. As these shifts are triggered by conventionalized linguistic mechanisms, they are neither *ad hoc* nor mysterious. Thus, the present account, unlike the alternatives, not only allows the context to change between different occurrences of a demonstrative, but explains precisely how it is that the context changes in a specific way, so that it determines the interpretation it does.

The present account of pronoun resolution is still provisional in several respects.²⁵ More work, needs to be done toward providing an exhaustive list of coherence relations. To see what's at stake, consider (50) from Winograd (1972):

²⁴ This is, of course, a bit simplified. We have seen that, on the Attention-Coherence account, a pronoun itself is interpreted as an update that contributes to the attentional state in accordance with the grammatical role within which it is realized. The point is just that a pronoun's semantic value (which the update associated with the pronoun pushes onto the attentional state in accordance with the pronoun's grammatical role) is selected in the same way as is that of a dependent variable: by looking at an assignment function as modified by some prior linguistic expression.

²⁵ For example, the present account doesn't offer a treatment of dependent clauses. In this regard, it is worth pointing to examples such as 'Ann saw him hit him,' where the object pronoun 'him' can simultaneously play the role of the subject of the dependent clause. This means that the complete account of the dynamics of prominence will have to discriminate not just grammatical roles, but further syntactic distinctions to capture the difference between dependent and independent clauses. Furthermore, more needs to be said about coherence relations holding between main and dependent clauses, in particular, in attitude reports. Finally, more work is needed in understanding coherence relations operating at a sub-clausal level (see Hobbs, 2010). That said, the framework can in principle be extended to handle these further complexities. I take the main goal of the formal implementation in the present chapter to be an illustration of an account that captures the key idea, which is that grammar provides a rich repertoire of resources that determine the context relevant for the resolution of a pronoun. Thus, my system here invokes just the mechanisms needed to handle the target examples. In the interest of simplicity, I relegate the details of an analysis of further complex cases to future work.

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

Both (50-a) and (50-b) exhibit an explanatory relation, yet the occurrences of ‘they’ are resolved differently. However, this is not a challenge for the present account so long as the explanatory relation (50-a) displays is qualitatively different from the one in (50-b). And there’s a reason to think this is the case. In (50-b), the council’s decision about the demonstrators can be explained on the basis of the facts about the demonstrators, and the council’s beliefs about them: 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 (50-a), the council’s decision about the demonstrators can be explained based on the facts about the council, and (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, what we might call *Subject-based* vs. *Object-based* Explanation. If so, it’s no surprise they come with different attention-shifting operations—when the explanation is subject-based, the subject is promoted to prominence; when it is object-based, the pattern is reversed.

In this regard it is worth pointing out that we have good reasons to think that verbs signaling implicit causality interact with discourse structure (Kehler et al., 2008). Further we know that implicit causality verbs signal a bias to re-mention the causally implicated referent: where some verbs carry the subject-bias, others signal the object-bias (Garvey and Caramazza, 1974).²⁶

- (51) Sue frightened Mary because she was scary.
 (52) Sue feared Mary because she was scary.

Verbs like ‘frightened’ are subject-biased, signaling the subject as the cause of fear, whereas those like ‘feared’ are object-biased, signaling the object as the cause of fear. It is not unnatural to expect that when the *explanation* features an object-biased verb, as in (50-a), the pronoun in the subject position of the *explanans* will be resolved in accordance with what I called subject-based explanation, one in terms of *earlier* subject’s (in the *explanandum*) beliefs and attitudes. (By contrast, when *what’s to be explained* is featuring an object-biased verb, we expect the pronoun in the subject position of the explanation, as in (52), to be resolved to

²⁶ See also Caramazza et al. (1974); Brown and Fish (1983); Au (1986); McKoon et al. (1993). For an overview on implicit causality in verbs, see Rudolph and Försterling (1997).

the object of the previous sentence, i.e., the cause.) And when the explanation is featuring a subject-biased verb, it's natural to expect it to provide what I call an object-based explanation, one in terms of the features of the object in the *explanandum*. (Again, by contrast, when *what's to be explained* is featuring a subject-biased verb, we expect the pronoun in the subject position of the explanation, as in (51), to be resolved to the subject of the previous sentence, i.e., the cause.)

And we see this clearly in the following contrast:

- (53) The city council denied the demonstrators a permit.
 a. They feared them.
 b. They frightened them.

As one would expect, one naturally understands 'they' in (53-a) to refer to the city council and in (53-b) to demonstrators.

If this is on the right track, then we have clear empirical basis for distinguishing these qualitatively distinct kinds of explanation. Of course, to flesh out this suggestion in full detail, we would need to substantiate systematic differences between these kinds of explanations. Further empirical research is needed for us to gain a firm grasp on all of these issues.

Interlude: Context and Common Ground

I have argued that we should think of the contextual dependencies as determined by the rules of grammar which dictate how the context dynamically evolves with the unfolding discourse. It is thus worth pausing for a moment to compare my account of contextual dynamics with one prominent alternative we have not considered up until this point—that of a context as determined by the *common ground* of a conversation (Stalnaker, 1978, 2002, 2014).

It is worth comparing this alternative to the account I have been defending since, it too, like my account, allows for the context to dynamically change as the discourse unfolds, in a way that affects the interpretation. However, unlike my account, the common ground account construes the context as the interlocutors' mutually shared body of information which changes in response to manifest evidence, as it becomes available and mutually recognized. Furthermore, it incorporates the import of the contextual parameters relevant for interpretation of context-sensitive expressions—for example, the elements of the Kaplanian context—by assimilating the information about these parameters to the mutually shared information that the context comprises. Thus, the account, as I shall explain in more detail in what follows, assimilates the changes in contextual parameters relevant to interpretation of context-sensitive items to other kinds of changes in information resulting from a rational response to incoming evidence. As we shall see, the account does not maintain that linguistic rules play a key role in determining the meaning of a context-sensitive item. Indeed it treats such rules as, at best, bits of evidence on a par with other kinds of extra-linguistic evidence one might have available for a particular interpretation. As a result, while the account incorporates contextual change, the dynamics of context change leaves room for the underspecification in the interpretation of a context-sensitive expression. Indeed, the account has been, and is most naturally, understood as incorporating the extra-linguistic approach to context-sensitivity resolution.¹ This leads to problematic consequences.

Here, I will first briefly outline the details of the common ground account of context and its treatment of contextual dynamics. Then I shall point to what I argue is lacking in the account. By failing to clearly separate the role of linguistic

¹ Cf. Stalnaker (2014).

conventions in governing context-dynamics from other kinds of information a context makes available, the account fails to capture and explain the distinctive and robust effects such conventions have on interpretation. At the same time, by modeling the change in contextual parameters as a change in the common ground the account complicates the explanation of how we can contribute specific and precise content even when we are ignorant about, or fail to track, the values of such parameters.

By contrast with the Kaplanian account of context, the Stalnakerian common ground account proposes to model a context not as a bundle of parameters representing objective features of the real-world speech situation, but rather, as the information determined by a *common ground*, a body of propositions mutually accepted by the interlocutors (at least for the purposes of the conversation). More precisely, the common ground comprises a set of propositions p such that everyone in the conversation accepts that p , and they all believe that they all accept that p , and all believe that all believe that they all accept that p , and so on, *ad infinitum*.² The common ground determines the context: the set of worlds compatible with all the propositions that are in the common ground, that is, the set of worlds in which all the propositions in the common ground hold.³ Intuitively, a context characterizes the space of open possibilities in discourse, those that have not yet been ruled out by the information conversational agents have accepted (for the purposes of a conversation). Specifically, common ground context characterizes the possible ways that the world might be like given what has so far been accepted, that is, given what is common ground.

The common ground model of context, like my own, is in part motivated precisely by the idea that the features on which the interpretation depends can change over time, as the discourse progresses. However, on this model, unlike on mine, the change in interpretive dependencies results straightforwardly from the interlocutors ability to track and accommodate new information and coordinate joint action. Let me explain.

The common ground context is constantly changing because, as the conversation evolves, we acquire more information about the world (or what we mutually agree the world to be like for the purposes of our conversation). As a result of this information agglomeration, the space of open possibilities shrinks. When agents come to accept a new bit of information—when a new proposition is added to the common ground—the possibilities that are incompatible with that proposition are ruled out from the context, and are thus no longer considered

² Common ground is thus closely related to the notion of common belief. See Stalnaker (2002).

³ Formally, 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. This notion of context is sometimes also called *the context-set*.

as live or open ways the world can be. So, as the common ground grows, the context-set shrinks. (For instance, if we start with the common ground that is undecided about whether John is in NYC, and then we come to mutually accept this, perhaps because someone said 'John is in NYC,' we will end up with a context in which all the worlds are such that, in them, John is in NYC (at the time of utterance).)

The common ground, on the Stalnakerian model, can grow as a response to manifest evidence—for example, if all the participants witness that a certain event happens, and believe they all witnessed that it happened, etc., that the event happened will become a part of their common ground. To borrow one of Stalnaker's examples, if a goat walked into the room and stood in front of the conversational participants, so that everyone saw it, and everyone saw that the everyone saw it, etc., it would then become common ground among the group that the goat is in the room. But this mutual acceptance that a manifest event took place then affects the state of the common ground, and induces interpretive effects on the subsequent discourse. The worlds in the context in which the room in question contains no goats are ruled out—one can no longer felicitously presuppose that there is no goat in this room. But the interlocutors can now also felicitously refer to the goat—they can felicitously say things which would have been infelicitous prior to its entrance. For instance, they can say 'It's stinky,' or 'The goat wants some grass.'⁴

But the common ground is also characteristically affected by linguistic moves in a discourse. For instance, it is affected by an act of assertion. When a speaker makes an assertion, if the assertion is successful, the proposition expressed by the utterance is added to the common ground, while the worlds incompatible with it are ruled out from the context-set. So, the characteristic effect of an assertion, on this model, is to add the asserted proposition to the common ground. As the common ground is determined by the group attitude of mutual acceptance among the participants in a conversation, if an assertion of *p* is successful, the interlocutors will coordinate on a particular attitude toward *p*—for example, they will accept for the purposes of the conversation (or perhaps believe, or know) that *p*.⁵ To take an earlier example, if I say 'John is in NYC,' and you don't challenge my assertion, it will become common ground among us that John is in NYC, and the worlds in which he isn't will no longer be live possibilities.

The context then dynamically changes with the evolution of a discourse, as affected by the conversational moves made and the mutual recognition

⁴ I am assuming, for simplicity, that prior to the entrance of the goat, there was no other (salient) goat in the common ground that would satisfy the definite description.

⁵ Precisely, as *p* becomes common ground among a group of participants, they will come to accept that *p*, and believe that they all accept that *p*, and believe that they all believe that they all accept that *p*, and so on.

of manifest events taking place. Because the dynamics of context is understood as information agglomeration, any change in contextual parameters that might govern the context-sensitive items is likewise a byproduct of information agglomeration: for instance, if I speak up, and we all witness this (and witness that we witness this, etc.), it will be manifest that I have become the speaker. This is why ‘I,’ in my mouth, will pick out myself. Similarly, if it becomes manifest, and mutually accepted, that a goat has entered the room, this will make it prominent in a way that allows me to pick it up with a pronoun, say, by saying ‘That’s a smart goat’. In other words, it will now be common ground among us that there is a prominent goat in the room, and I will be able to pick it out with ‘that’. In this way, the change in the contextual parameters relevant for interpreting context-sensitive items, including the parameters of the Kaplanian context, is reflected as the change in common ground.⁶

One might think, then, that this account adequately supplies what was missing in the Kaplanian account—it provides a dynamically changing context that keeps track of how our information, including our information about the contextually relevant parameters, changes as the discourse develops. The model offers an explanation, or so it seems, of why in (10), repeated below, the two occurrences of the pronoun are interpreted the way they are:

(10) He [pointing at Bill] is sad, because he [pointing at Tim] is leaving.

After all, if all the interlocutors recognize that the speaker is pointing at Bill, and then at Tim, it will become mutually accepted that one, and then the other individual is the intended referent. One might think, further, that this alleviates the worry about the ambiguity as well. After all, given that now we have an account according to which a context can change in response to manifest

⁶ One might wonder how to understand the relation between the Kaplanian model of context as a representation of a concrete situation on which interpretation depends, and Stalnakerian common ground. For a detailed discussion, see Stalnaker (2014). Briefly, given that a Kaplanian context is a tuple of parameters representing features of a real-world situation on which interpretation depends, one could think of the Stalnakerian common ground as a parameter of context, a specific body of information that the conversational agents happen to mutually accept at a given point in their conversation. However, this is not the approach Stalnaker takes. He thinks of a Kaplanian context—or more specifically, parameters of a Kaplanian context—as of a feature about which one can be uncertain, and uses the common ground context to model potential uncertainty about the specific conversational situation one is in. The idea is that common ground represents how agents take their real-world conversational situation—the Kaplanian context—to be (given what they accept for the purposes of a conversation). Since their information about the conversational situation can change, it is the common ground, and its dynamics, that is driving the dynamics of contextual parameter change. To model this, in the latest installment of his model, Stalnaker employs a richer notion of context, as a set of Kaplanian contexts, modeled as (multi-)centered worlds (Stalnaker, 2014), representing the interlocutors mutually accepted understanding of the conversational situation which they find themselves in.

evidence, we might think that a pronoun, as a matter of its meaning, selects the most salient entity (satisfying its character) at a given point in the evolution of the context.⁷

However, this is too quick. For one, while this kind of an account allows for the context to change in response to manifest evidence, it is one thing to say that a context can change, and it is quite another to provide an account that explains how exactly the context changes to predict a specific interpretation.⁸ To see the issue, consider the following example, of the sort we have discussed earlier:

- (54) Little Mary is jumping up and down, making herself salient, clearly attracting our joint attention. Amid this, I tell you:
 a. Sue arrived late last night, and checked into a hotel, but she didn't get much sleep. She has been antsy all day today.

(54) is straightforwardly interpreted as conveying that Sue arrived late last night, checked into a hotel, and due to a lack of sleep, has been antsy all day today. Unless I do something—for example, point at Mary, or stress the pronoun—the anaphoric interpretation will be determined for the occurrences of the pronoun 'she.' In, particular, we do not retrieve a reading according to which Mary has been antsy all day (even though she is clearly antsy at the moment, and is thus an available and fitting referent for 'she'). The question is why not?

At the very least, this illustrates that just allowing the context to change in response to manifest evidence doesn't yet tell us how a particular interpretation is determined given a specific state of the context. Second, and more importantly, (54) suggests that not all manifest evidence is on a par. A straightforward explanation of what goes on in (54) is that the speaker starts off by introducing a subject, Sue, and proceeds to develop a narrative about her: about her arrival and subsequent tiredness. The final sentence continues on this thread, describing a result of her lack of sleep. The pronoun is thus straightforwardly interpreted as anchored back to Sue. Of course, all this could be a part of what we mutually recognize. But we also have overwhelming evidence of Mary's salience—we both see her prominent, and we mutually recognize that she fits the description provided. After all, she too is antsy, and she is the most salient entity in our visual field, and is attracting our joint attention.

⁷ Obviously, things get more complicated with respect to the ambiguity between bound and referential uses of pronouns. But it is possible to integrate a broadly Stalnakerian model of context with the dynamic account of binding we have seen in Chapter 3. See, for example, Heim (1982); Groenendijk, Stokhof, and Veltman (1996); *inter alia*.

⁸ Note that the account treats a demonstrative pronoun again as a true demonstrative. A diverse set of potentially non-linguistic cues can work to determine what the interpretation of a demonstrative pronoun on an occasion of use is.

A natural reaction one might have to this is to say that we are simply weighing our overall evidence for particular interpretations, and then selecting the most plausible one. After all, if the speaker is in the middle of a narrative about Sue, it is only natural that she doesn't suddenly change topics, all of a sudden talking about some other individual and events. But the real question to ask is how is it that we recognize that the second sentence in (54) is a continuation of a previous thread, and not a change in topic. Surely, sometimes we do change topics, and we naturally understand that we do when we do—just imagine I made a significant pause before the final sentence in (54), or pointed at Mary. And here we note, again, that the speaker has to do something different—for example, change the intonation contour by stressing the pronoun, or point, or pause—in order to get the interpretation according to which the pronoun is resolved to Mary. The cues signaling such shifts are limited and, to a high degree, arbitrary; for example, there isn't any natural connection between a particular intonation pattern and a shift in attention, or a continuation of a narrative. We could easily imagine practices quite different than the one we actually adhere to.⁹

Of course, one can maintain that such specific conventions are a part of our “manifest” evidence, just as the fact that the word ‘she’ has a specific linguistic meaning can be manifest evidence that in uttering it the speaker is aiming to talk about a female. The question is how such linguistic evidence is weighed against other kinds of evidence, say, against the perceptual salience, general world knowledge, etc. One might expect that such evidence would be on a par with all other; however, we have good reasons to think conventional constraints on interpretation are privileged. For instance, consider the following:¹⁰

(55) John came to the party, and he didn't come.

(55) is infelicitous, unless the pronoun is stressed, or uttered in conjunction with a pointing gesture. This is surprising: after all there is overwhelming evidence that the pronoun should not be co-varying in its interpretation with the antecedent ‘John.’ This yields an inconsistent interpretation. So, charity would require us to discount it. That it doesn't suggests that the conventionalized constraints on pronoun interpretation trump other available, manifest evidence. This is just the same point illustrated earlier in Chapter 5 by (47), repeated here:

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

⁹ Recall our discussion of the conventionality of deictic gestures in Chapter 4.

¹⁰ Compare Stojnić (2019).

So if, as these examples, and the arguments in Chapters 4–5 suggest, conventionalized mechanisms play a distinctive and privileged role in determining interpretation, then we would want this to be reflected in our theory of contextual change, and distinguished from whatever role other kinds of manifest evidence have in interpretation.¹¹

At the same time, we would want to separate the effects of the linguistic rules on the conversational context from our mutual knowledge or understanding of these effects. On the common ground account, an utterance has a characteristic effect on the conversational context, but such an effect is primarily reflected in the transformation of the shared body of information, potentially affecting contextual parameters governing the resolution of context-sensitive items insofar as the information about these parameters is a part of the common ground. But we also often exploit context-sensitivity even when we are completely uncertain about the value—or even potential values—of a particular context-sensitive item, where the specific interpretation of the expression at hand is unaffected by this uncertainty. For instance, suppose we are playing a mystery person game: we don't know who the woman behind the curtain is, but our task is to guess. We can ask questions: 'Is she famous?' We can exchange information: 'She is not an actress.'¹² The game presupposes that there is a chain of reference linking our utterances to a particular individual: we are all the while talking about one and the same person. But the whole point of the game is that we do not know who that is. And while we might be uncertain about which context we are in—what the content of the pronoun is—this does not prevent us from contributing specific, precise truth-conditions, a proposition about a specific individual.

So, an account of context which captures that we publicly communicate such specific contents under uncertainty is desirable. A model of context according to which assertion proceeds by adding asserted content to a shared body of information—where the uncertainty about the values of contextual parameters is reflected in a context which affects the interpretation—complicates this issue. If instead there are linguistic mechanisms that determine specific resolution of context-sensitive items, by manipulating contextual parameters independently of our ability to track them let alone mutually agree on their effects, as I have argued there are, then this kind of communication under ignorance, too, becomes straightforward.

Of course, I'm not offering any of these considerations as knock-down arguments against the Stalnakerian model of context. There are various ways

¹¹ Recall, a diverse set of cues might play a role in disambiguating the form of an utterance, but as I argued, this is not to be confused with the role of fixing the meaning of an expression uttered.

¹² For simplicity, let us assume that we know the person is a female and that, as a matter of fact, there is indeed a female person behind the curtain.

one could try to handle these issues.¹³ I'm merely presenting this as an illustration of the key differences that separate my account of the dynamics of context-change from the familiar alternative. I take it as a surprising, yet welcome result that a model of context and context-sensitivity resolution that isolates the effects on context brought about by linguistic mechanisms from general rational response to shared evidence straightforwardly overcomes these issues.

¹³ Cf. Stalnaker (2014).

PART II

CONTEXTUALIZING CONTENT

Content in Context

The Attention-Coherence account developed in Chapters 4 and 5 posits linguistic mechanisms that change the context by manipulating the attentional state of the discourse. The key insight is that a discourse can harbor elements that not only depend on context for their interpretation, but also *change* the context in a way that can affect the interpretation of subsequent context-sensitive items. We captured this aspect of meaning by interpreting utterances as updates, changing the attentional state of the discourse by updating potential interpretive dependencies.

This points to an important departure from the Kaplanian picture. On the original Kaplanian picture, sentences express content relative to a particular context. The role of the context is to determine the proposition expressed by an utterance. But on the present picture, that cannot be right: contexts constantly change, and do so as governed by the updates associated with the expressions that utterances comprise. Thus, there isn't a single context of utterance that fixes the proposition that the utterance expresses. This raises several important questions and leads to several important ramifications.

First, we must ask, if utterances are assigned dynamic meanings—updates to context—how does such dynamic content relate to the truth-conditional, propositional content, if at all? In what sense, if any, can we say that an utterance of (23), repeated below, expresses a proposition *that some woman came in and Betty sat down*, and an utterance of (17), also repeated below, a proposition *that some woman came in and sat down*? Moreover, if the context, constantly changes, what is the relation between such an evolving context and the propositional meaning that an utterance expresses?

(17) A woman came in. She sat down.

(23) A woman came in. She [pointing at a female cat, Betty] sat down.

The question of whether and how one can extract propositional meanings from our dynamic meanings is important. The notion of propositional, truth-conditional content is central to the standard understanding of the interaction between thought, communication and action. To illustrate, consider the following exchange between A and B:

- (56) a. A: It is raining.
 b. B: Thanks! I should bring an umbrella. (B takes an umbrella.)

A natural way to describe what is going on in (56-a) and (56-b) is as follows: A has a thought about the world—namely, that it is raining (now, at a particular location *l*)—that she wants to convey to B. She utters a string of words of English that (in the given context) expresses this thought—namely, ‘It is raining.’ Upon hearing the utterance, B, being a competent speaker of English, comes to understand A’s utterance. If she thinks that A is sincere and reliable, she might even come to believe what A said. And, furthermore, this belief might impact her action; for example, she might decide to bring an umbrella.

This sort of description has traditionally been cashed out in terms of truth (and truth-conditions), reference, and satisfaction. A’s thought represents the world as being a certain way (namely, as being such that it is raining in it at the location *l*, and time *t*). Her utterance is true (in the given context) just in case the world is this way. It is because B understands what it takes for A’s utterance to be true—what its truth-conditions are—that she comes to believe the world to be a certain way (given that she takes A to be reliable), and acts in it as if it were that way (thus, bringing the umbrella). Further, the notion of propositional content has been either identified with, or taken to determine this truth-conditional, representational content. Propositions are taken to be what we assert and believe.¹

The role of context, on this picture, was to determine the proposition expressed by a given sentence on an occasion of use. If the shifty context could no longer do this job, that would require radical re-thinking of the interaction between thought, communication, and action. At the same time, if truth is defined relative to contexts, we have to appeal to context in our account of validity; and, further, if the context is shifty, we need a way of thinking about validity in the wake of potential shifts in context within the course of an argument. We have already seen, in our discussion in Chapter 4, that a change in context induced by a linguistic item—for example, a pointing gesture—might affect validity, and, further, that the effect of such an item should be reflected in the logical form, if we want to adequately represent arguments which are valid by virtue of their form as, for instance, I argued (27) was. At the same time, given that context can change in a way that affects interpretation, the account of validity should allow us to also keep track of shifts in context that can lead to equivocation as, for instance, was the case in (28).

¹ In this tradition, propositions are thus the primary bearers of truth-values. Utterances, assertions, and beliefs all have truth-values derivatively, by virtue of expressing a truth-value bearing proposition. Compare footnote 21 in Chapter 1.

Luckily, even within our present dynamic, shifty account of context, and given our interpretation of utterances as updates to such context, we can still capture the idea that utterances express propositional, representational content which is contextually determined, and, moreover, we can provide an account of validity that doesn't mandate an arbitrary requirement of a fixed stable context, but that captures intuitively valid patterns of inference and that, importantly, allows us to separate valid patterns of inference from those involving equivocation.²

In Chapters 4 and 5, we interpreted utterances as instructions to update interpretive dependencies, that is, as binary relations on assignment functions. We said that an utterance is true just in case, given an input assignment, it leads to a non-empty output assignment.³ How can one recover propositional, representational meaning from the truth-condition delivered here?

So far our interpretation has been purely extensional. To capture propositional content we need to allow for an interpretation to depend on possible worlds as well. That is, we have to relativize truth to worlds, not just assignments. For concreteness, let us modify our earlier definition of the basic update, repeated in Definition 7.1, replacing it with one defined in Definition 7.2, which relativizes interpretation to a world as well:

Definition 7.1 $\llbracket [C] \rrbracket(g, g')$ iff $g = g'$ and $\llbracket C \rrbracket^g$

Definition 7.2 $\llbracket [C] \rrbracket(w, g, g')$ iff $g = g'$ and $\llbracket C \rrbracket^{w,g}$.

So, while before we said that $\llbracket [woman(x_0)] \rrbracket$ is true on an input assignment just in case there is an identical output assignment on which x_0 is a woman, now we say it is true on an input assignment and *a world*, just in case there is an identical output assignment on which x_0 is a woman in that world.⁴

Similarly, the interpretation of updates that promote new values on the stack are also interpreted relative to worlds. As explained above, we had $\langle an \rangle$ introduce a new (indefinite) value to the n th position of the stack pushing everything else down. This is formally captured in Definition 7.3, which specifies an update that changes the stack by storing a new (unspecified) value at the designated position, k , pushing all values stored at positions k or higher, one position down in the ranking.

² As mentioned before, not everyone will agree with the traditional picture according to which the notion of propositional, representational meaning generally plays the role of the content of propositional attitudes and assertion, nor even that such content should be in general recoverable from compositional semantic content. I turn to the challenge of non-propositionalism in the remainder of Part II. For now, the goal is just to show that the account we developed so far allows us to recover propositional content in the first place.

³ So, $\llbracket [woman(x_0)] \rrbracket$, as before, doesn't affect the ranking: it just checks that whatever the input assignment assigns to x_0 is a woman in w . The basic update is still a partial identity relation.

⁴ Accordingly, predicates are interpreted as denoting properties that are world sensitive—intuitively, entities have properties relative to a world. For formal details, see Appendix A.

Definition 7.3 $\llbracket \langle \alpha n \rangle \rrbracket(g, g')$ iff $g'_i = g_i$ for all $i < k$, $g'_k = o$ for some (indefinite) entity o from the domain of entities supplied by the model, and $g'_i = g_{i-1}$ for all $i > k$.

To capture the idea that the interpretation is sensitive to worlds, we now modify our definition only slightly:

Definition 7.4 $\llbracket \langle \alpha n \rangle \rrbracket(w, g, g')$ iff $g'_i = g_i$ for all $i < k$, $g'_k = o$ for some (indefinite) entity o from the domain of entities supplied by the model, and $g'_i = g_{i-1}$ for all $i > k$.

Definition 7.4 brings no essential change to the effect of the update—it just ensures, again, that the truth-conditions are specified relative to worlds.

Now that we have relativized truth to worlds, this already gives us enough to capture propositional content. In particular, we can now say that an update with an utterance is true relative to a *world* and an assignment just in case there is an output assignment that makes it true. Then, we can recover the propositional content as follows:

Definition 7.5 Given an input assignment g the propositional content expressed by an update K is as follows:

$$\{w \mid \exists g' \text{ such that } \llbracket K \rrbracket(w, g, g')\}.$$

To see how this delivers the desired truth-conditions, let us work through an example. Take the update associated with the first sentence in (17), (35), repeated below:

$$(35) \quad \langle \alpha 0 \rangle; [woman(x_0)]; [came.in(x_0)]$$

Given an initial assignment g , the update will express the following proposition:

$$(57) \quad \{w \mid \exists g' \text{ such that } \llbracket \langle \alpha 0 \rangle; [woman(x_0)]; [came.in(x_0)] \rrbracket(w, g, g')\}$$

The update associated with the first sentence in (17) will be true relative to a world w and an assignment g , just in case there is an output assignment g' , that results from g by promoting a new unspecified value to position 0, pushing all others one position down, and which is such that g'_0 is a woman who came in in w ; that is, just in case some entity (from the domain specified by the model) is a woman who walked in in w .⁵

⁵ See Appendix A for detailed formal definitions.

Following up with the second sentence in (17) places further requirements: taking an output assignment delivered by the first sentence in (17), say g' , as a new input assignment for the second one, the update with the second sentence is true relative to g' and w just in case there is some g'' that results from g' by pushing the highest ranked entity that satisfies the character of 'she' in g' to the top position, and requiring that that entity sat down. Given that g' is some output of the update with the first sentence, this update with the second sentence will be satisfied just in case some woman came in and sat down in w . And this is precisely what we want. So, for the update associated with (17), repeated below, given an input assignment g , we get the proposition expressed in (58) that just reduces to the proposition that some woman came in and sat down in w :

- (37) $\langle \alpha 0 \rangle; [woman(x_0)]; [came.in(x_0)];$
 $\langle \pi 0 @ she \rangle; [sit.down(x_0)]$
- (58) $\{w \mid \exists g' \text{ such that}$
 $\llbracket \langle \alpha 0 \rangle; [woman(x_0)]; [came.in(x_0)]; \langle \pi 0 @ she \rangle; [sit.down(x_0)] \rrbracket (w, g, g')\}$

If we, by contrast followed the first sentence of (17) with the one containing a pointing gesture, as in (23), we would get a different result. The update associated with (23) is the following:

- (59) $\langle \alpha 0 \rangle; [woman(x_0)]; [came.in(x_0)]; \langle \pi 0 b \rangle; \langle \pi 0 @ she \rangle; [sit.down(x_0)]$

The proposition expressed, given an assignment g will be the following:

- (60) $\{w \mid \exists g' \text{ such that}$
 $\llbracket \langle \alpha 0 \rangle; [woman(x_0)]; [came.in(x_0)]; \langle \pi 0 b \rangle; \langle \pi 0 @ she \rangle; [sit.down(x_0)] \rrbracket$
 $(w, g, g')\}$

Now, the update with the first sentence is just as before, and delivers the same content as before—it is true relative to an input assignment g just in case some woman walked in in w ; but taking some output assignment of the update with the first sentence, say g' , as the new input assignment to the second sentence, the update with the second sentence will be true relative to such g' , and a world w , just in case some g'' results from g' by promoting Betty, and demoting everyone else, and requiring that Betty sat down in w . Simplifying, (23) is true relative to g and a world w , just in case some woman came in and Betty sat down in w . This is again the desired result. Notice that, as expected, in (23), the pronoun receives a referential, and rigid reading, whereas in (17), it is interpreted as bound, rather than referential. Yet in both cases, the pronoun is interpreted as co-varying with an antecedent expression: an indefinite in (17) and a pointing gesture in (23).

So, we can recover propositional content. Now, in the traditional, Kaplanian framework, it is the context of utterance that determines the propositional content. The context was understood as a bundle of features of the utterance situation that the utterance interpretation can depend on. At the very least, the context had to include an agent (the speaker), a world of utterance, and a time and location of the utterance. How can we capture the idea that the context of utterance determines the content in our framework that models context-change with the evolving discourse?

So far in our current framework, we have only been focusing on the prominence ranking of candidate referents for the resolution of demonstrative pronouns. But we can factor in other parameters of context. One straightforward way to do this is to simply reserve particular slots in our stacks for further contextual parameters. For instance, we can designate a position on our stack, call it 'w', which stores the world of the utterance: intuitively, this is just our world-parameter of the old Kaplanian context.⁶ The world position of the stack itself is unranked, and it doesn't affect or interfere with the prominence ranking of other elements on the stack; it merely keeps track of the world in which the utterance took place.⁷ We can easily factor in parameters relevant for the interpretation of other indexical expressions in a similar way. Since we only want to show that the context can determine propositional content, we can continue disregarding further parameters.

Now, we can pin down the sense in which context determines propositional content. Let us say that an initial context-assignment⁸ is any stack that stores

⁶ This requires modifying our stacks slightly. Before, a stack was a function from a set of natural numbers plus 0, to a set of entities. Now we will treat it as a function from a set comprising natural numbers, 0 and w, to a set of individuals and worlds.

⁷ We could in principle make the world parameter a ranked element where, initially, the world of utterance—the context-world—is stored in this position. Whether we should do this or not depends on whether we should think of the context-sensitive items dependent on the world of utterance as prominence-sensitive or not. This will partly depend on whether the item in question can receive discourse-bound readings, that is, whether it can co-vary with some world-prominence-affecting item. Similar considerations apply to other indexicals: for instance, to capture the context-sensitivity of 'I', we can simply include the speaker position on the stack, storing the speaker of the utterance. Since 'I' is not prominence-sensitive—as displayed by the inaccessibility of the discourse-bound readings in examples like (i), as well as unavailability of a deictic, prominence-driven selection of an arbitrary referent—this position will be unranked.

(i) The speaker came in. I sat down.

⁸ I'm simplifying somewhat. Recall that an utterance is interpreted as a relation between the assignments that updates interpretive dependencies. This means we cannot identify a context with any single assignment. Rather, a context at any given point has to be a set of such assignments, capturing the constraints on assignment functions encoded at a particular point in a discourse. On this construal, the initial context *C* will be a set of initial context-assignments. And as the discourse progresses, the context will include a set of assignments that are possible outputs given the initial context-assignments.

the actual world—the world of utterance—at the position w .⁹ Given this setup, the content expressed by an utterance relative to a context and a world is the following:

Definition 7.6 (Truth in a context at a world) Where K is an update, and c an initial context-assignment, we say that K is true relative to a world w and the context-assignment c just in case there is some g such that $\llbracket K \rrbracket(w, c, g)$.

And given this notion we can now recover the sense in which the proposition expressed by an utterance is determined by the context, albeit one that changes with the evolution of a discourse. Given an initial context-assignment c , the propositional content expressed by an update K is as follows:

$$\{w \mid \exists g \text{ such that } \llbracket K \rrbracket(w, c, g)\}.$$

The proposition expressed by an update relative to a context-assignment c corresponds to a set of worlds such that the update relates c to a possible output assignment relative to w . And we can now say when an utterance is true relative to a context-assignment *simpliciter*. An utterance is true relative to a context-assignment just in case the proposition expressed by it is true in the world of the context-assignment—the actual world:

Definition 7.7 (Truth in a context) Where K is an update, and c an initial context-assignment, and w_c the world stored at c_w , we say that K is true in c just in case there is some g such that $\llbracket K \rrbracket(w_c, c, g)$.

We have now come full circle. As in the traditional Kaplanian framework, we are able to predict how the context plays the content determining role by supplying the interpretation for context-sensitive vocabulary.

However, while the Kaplanian context was static, not tracking how the parameters change as the discourse unfolds, the context on my account is dynamic, keeping track of the changes in contextual parameters over an ongoing discourse. Consequently, on the current proposal, context dynamically evolves to build a proposition. We can think of this picture of context as a generalization of the Kaplanian framework. For Kaplan, a context is a tuple comprising the parameters that affect the resolution of context-sensitive expressions. But we can generalize this idea, by treating context as a *conversational record*, in the sense

⁹ We can assume that the actual world (the world of utterance, or the context-world) is fixed just like the Kaplanian world parameter. We could also easily have each utterance update the world slot of the context with the world in which the utterance takes place. At this point the choice is insignificant for us.

of Lewis (1979), an abstract record that keeps track of the relevant information about the conversation: who is speaking, at what time, in which world, who is the addressee, and so forth. The scoreboard keeps track of the parameters fixing the meaning of context-sensitive items, but more importantly, unlike the Kaplanian context, it also keeps track of how the relevant contextual parameters change as the discourse evolves. It tracks which conversational moves have been made, what has been said, which propositions have been mentioned, which individuals have been made prominent, and which demoted. In this way, we can think of the prominence ranking of candidate referents as one aspect of the conversational scoreboard. Unlike the static Kaplanian context, the conversational scoreboard keeps track of the dynamic change in relevant contextual parameters as the discourse evolves.

This dynamics of conversational scoreboard, I have argued, is governed by linguistic rules that specify a change to the context. More precisely, we have isolated an important dynamic layer of meaning that serves to build the truth-conditional, propositional content by affecting the change in the parameters of context. This dynamic layer of meaning specifies rules of context-change which build the truth-conditional content as a function of that change. So, on this picture, it is a part of the meaning of an expression that it actively induces a change in the context, in a way that might affect the interpretation of what comes after, thus contributing to the specification of the propositional content expressed.

The emerging account of context I have developed so far departs from the traditional view on context and context-sensitivity resolution which treats most context-sensitive expressions as true demonstratives. On that picture the process of interpretation of context-dependent elements is a defeasible, abductive one, selecting a content from an open-ended array of candidate interpretations through broadly pragmatic mechanisms.¹⁰ By contrast, the present account suggests that, at least as far as demonstrative pronoun resolution goes, the context is not as powerful as traditionally presumed. Even in the paradigmatic case of a true demonstrative—a demonstrative pronoun—the interpretation of a discourse lacks a dependency on *non-linguistic* context. It is the linguistic rules of context-change, encoded in dynamic meaning, rather than abductive pragmatic reasoning, that determine the truth-conditional contribution of demonstrative utterances.

So far, we have seen how the linguistic model of context-change and context-sensitivity resolution allows for a uniform account of the meaning of a

¹⁰ As we have seen, this is to a significant degree true even within standard dynamic accounts—as such accounts typically allow for pragmatic, abductive mechanisms to guide both selection of an antecedent from a prior discourse and selection of a referent from the utterance situation.

demonstrative pronoun, determining truth-conditional content for referential, bound and discourse-bound readings. But the model of context can be extended to handle other cases of apparent contextual variability. We turn to a particularly important and illustrative application in the remainder of Part II and in Part III. The application to the class of context-sensitive expressions I will focus on—modal expressions—is particularly important and illustrative as it directly addresses some of the philosophically provocative arguments challenging the traditional model of content, the interaction between content and context, and the logic such interaction underscores. As briefly mentioned in Chapter 1, puzzling behavior of modal discourse has been argued to evade standard contextualist treatments. While traditionally treated as context-sensitive expressions (Lewis, 1973; Kratzer, 1977, 1981; Kripke, 1980), the growing consensus in the recent literature is that context fails to determine a (non-trivial) propositional, truth-conditional meaning for modal claims. Even relative to a specific context, modal utterances fail to express propositions. The behavior of these expressions has thus been taken to mandate non-representational accounts of meaning. The non-propositional accounts that have been developed deny modal utterances propositional content, and instead propose that the key aspect of meaning of modal discourse is a certain kind of dynamic effect it has on context.¹¹ At the same time, the non-representational treatment has been argued to receive further support from the alleged violations of certain classical patterns of inference in modal discourse as these types of account invalidate some of these patterns.¹²

¹¹ See Veltman (1985); Swanson (2006); Yalcin (2007, 2012b); Gillies (2010); Willer (2014); Bledin (2015); Charlow (2015); Starr (2016); *inter alia*. As cautioned before, this presentation of non-propositionalism glosses over some important distinctions between various versions of non-propositional account. There are proponents of non-propositionalism who will argue that modal expressions do have assertoric content, but will deny that this content is best modeled as traditional representational, propositional content (see, e.g., Kolodny and MacFarlane, 2010; Schroeder, 2013; MacFarlane, 2014; Moss, 2015). These authors will thus typically not think that the notion of assertoric content breaks down for modal discourse, and that instead the characteristic effect of the modal utterances on context is the only notion that transpires in an adequate account of communication with modal discourse. These authors will all, however, argue for non-standard, non-propositional models of assertoric, informational content. I will return to these distinctions in what follows. My arguments against the linguistic basis of non-propositionalism extend to such accounts as well.

¹² See Veltman (1985); Gillies (2004, 2010); Yalcin (2012b); Willer (2014); Bledin (2015); *inter alia*. Note, as mentioned before, I am not arguing that one cannot, in principle, have a non-propositional account that obeys classical logic; nor am I arguing that any propositional account automatically preserves classical logic—it need not. But the brand of non-propositionalism that has been offered in response to the problematic behavior that challenged standard contextualist accounts is typically coupled with non-classical logic; that is, a particular brand of non-propositionalism along with a non-classical semantics has been proposed as a package deal in response to the puzzling behavior of modals in context. I argue that this package deal should be rejected, and that the seeming support for it stems from the faulty assumptions built into the non-linguistic accounts of context and context-content interaction these accounts implicitly assume.

In what follows, I trace the problematic behavior of modals, and the alleged tension with classical logic, that has led the advocates of non-representationalism to reject the propositional, truth-conditional model of meaning, to the presumption of the standard, non-linguistic model of context-sensitivity. A linguistic model of context-change and context-sensitivity resolution, building on the one we developed thus far, delivers an account that, I will argue, is empirically more adequate than the competing non-representational ones. The account captures the problematic behavior that motivates non-representationalism, as well as further phenomena that, I argue, escape non-representational treatment, and does so while preserving classical logic. We shall see that dynamic meaning, which encodes the effects of discourse on context-change, is neither in tension nor somehow incompatible with propositional meaning and with classical logic. The modal discourse, then, turns out to be one domain where the presumption of flexible, pragmatic model of context has led to unwarranted, yet radical, philosophical conclusions: conclusions about the nature of content, and the relation between natural language and logic.

The Challenge: Non-propositionalism

Some of the more philosophically provocative arguments touting dynamic accounts of meaning argue for a dynamic interpretation of fragments of ordinary declarative discourse that is fundamentally incompatible with propositional accounts of meaning. The key motivating arguments for such accounts stem from the behavior of certain bits of natural language discourse—most notably, the behavior of modal vocabulary—that seemingly evades, and is in tension with, the propositional, representational model of content. The sole communicative effect of these bits of discourse, it appears, is to produce a certain kind of effect, irreducible to contributing a bit of information representing a world as being a certain way. An utterance of a modal sentence, on this account, does not represent the world as being a certain way, in the way an ordinary non-modal one does. And, moreover, no (non-trivial) propositional content can be recovered from its dynamic meaning. It is this thesis that I have called non-propositionalism.

Non-propositionalism motivated by the relevant behavior of modal expressions has received different implementations; the most prominent exemplars come in the form of expressivist models of semantics-pragmatics interaction, and in dynamic update semantics.¹ Though these accounts differ importantly,

¹ See, for example, Swanson (2006); Yalcin (2007); Charlow (2015), for the former; and, for example, Veltman (1985), Gillies (2004, 2010), von Fintel and Gillies (2007), Willer (2014), for the latter type of view. A different type of non-propositionalism comes in the form of various versions of relativist semantics (Egan, 2007; Kolodny and MacFarlane, 2010; MacFarlane, 2014; *inter alia*). However, though a version of non-propositionalism, relativism is less straightforwardly described as a non-representational account. Relativists argue that truth has to be relativized not just to contexts and standard indices, including worlds and perhaps times, but that one has to enrich indices with further, “non-standard” parameters. So, in the case of modality, specifically, they argue one has to relativize truth to an information state (where an information state, like a context, is a set of worlds compatible with certain information). Now, we have already seen how one could, in principle, recover propositional content from such compositional semantic content by saturating the values of the parameters of the index with a contextually provided value (recall footnote 21 in Chapter 1). So, even if compositional semantic content receives a truth-value only relative to an index comprising a world and some other, non-standard, parameter, one can recover the propositional content by saturating the value of any non-standard parameters with the value supplied by the context. For example, if indices contain worlds and times, my utterance of ‘John is happy’ in a context expresses the proposition that John is happy at the time of the context (i.e., time of utterance). The relativist, however, will deny that one can recover propositional content in this way—that is, that such propositional content is what is expressed by an utterance in a context—arguing that some relevant “non-standard” parameter (e.g., an information state, in the case of modality) is not fixed by the context of use but by the context of *assessment* (or the *assessor*). Since

they all deny that modal claims express propositional content. Further, they all hold that the primary conversational purpose of a modal utterance is to induce its characteristic effect on a conversational context, what we might characterize as its *dynamic aspect of meaning*, or what's usually called the *context-change potential*.

Non-propositionalism is radical. We have seen in Chapter 7 how a simple, appealing picture of how thought, communication, and action interact is underscored by the notion of propositional, representational content. If I believe that it is raining, the content of my belief—the proposition that it is raining—represents the world as being a certain way. My belief is true just in case the world is this way. I can decide to convey this belief to you, by making an assertion. To do this, I choose the appropriate expression which, in a given context, conveys my thought. For example, I say 'It is raining,' which in a normal context will have expressed that it is raining,² and which will be true (on the given occasion of use) just in case it is raining. My assertion conveys the content of my belief, because both express the same proposition: namely, that it is raining. Upon understanding my utterance, you might come to believe this expressed content, and this belief might guide your actions—you might decide to stay at home and avoid rain. The explanation exploits the idea that my thought and assertion, and your subsequent belief, all represent the world as being the same way.

Denying that modal claims express propositions implies that for a significant subset of natural language claims this simple picture is misguided. First, it implies that assertion of modal claims works in a fundamentally different way than assertion of the propositional discourse. While one might have expected that the same explanation underscores the exchange in the simple dialogue in (56), repeated below, and that in (61), it turns out that the two are fundamentally different. In (56), the speaker, A, is describing a world as being a certain way, by way of contributing a bit of informational content with a specific truth-condition. Meanwhile, in (61), A isn't contributing a bit of truth-conditional, informational content at all. She is not expressing the proposition that represents

for any utterance there are indefinitely many contexts of assessment, there is no one (ordinary) propositional content expressed, that is, asserted, by an utterance. (Indeed, most relativists would thus maintain that the assertoric content is the non-propositional semantic content, which falls short of full propositional content, at least if we understand propositions in the traditional way.) As I will explain below in more detail, expressivists go further in arguing that not only is the relevant "non-standard" parameter not fixed by the context of utterance, but that there is no sense in which any context—be it of utterance or of assessment—fixes such parameters. For ease of exposition, I shall not discuss relativism directly, but I note that my arguments against non-propositionalism carry over against the relativist accounts straightforwardly. For further discussion of the theoretical differences between relativism and expressivism, see MacFarlane (2014).

² Throughout this description, I again omit the sensitivity to times and locations of the context for the ease of the exposition. Officially, one will have expressed that it is raining at a particular contextually specified location and time.

the world as being a certain way, for example, as such that in it it might be raining. What she is doing, by contrast, is expressing that she is in a mental state which does not rule out raining.³ And what the appropriate uptake on the part of B is, even if B takes A to be sincere and reliable, is not to come to believe a proposition that is representing the world as being such that in it, it might be raining, but rather to coordinate her mental state with A's, so that, now, B too is leaving open the possibility of raining.⁴

- (56) a. A: It is raining.
 b. B: Thanks! I should bring an umbrella. (B takes an umbrella.)
- (61) a. A: It might be raining.
 b. B: Thanks! I should bring an umbrella. (B takes an umbrella.)

Assertion and its integration with thought, communication, and action aren't the only aspects of communication that need to be revised once we abandon propositionalism. This theoretical move seems to require a revision of the standard understanding of pragmatic inference as well. The assertoric content of a sentence—its representational, propositional content—is what Grice (1975) called *what is said* by an utterance; this content in turn is construed as the input to further pragmatic reasoning, required for calculating, for example, conversational implicatures.⁵ 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 enter into pragmatic calculations. For instance, under normal circumstances, if I tell you 'Betty might have invited some of her friends to the party,' I implicate that it is not the case that she might have invited all of them. While a Gricean account would explain this by appealing to pragmatic inference that begins with the assertoric content of my utterance—the proposition that Betty might have invited some of her friends—a denial of propositionalism would require a

³ Note, this is to be distinguished from her describing herself as being in such a mental state—the content of her utterance isn't a proposition, and so it is not a proposition describing her as being in a certain mental state.

⁴ This is a characterization of the effect of an assertion of a modal claim we find in standard expressivism. See, for example, Yalcin (2007). A very similar characterization naturally fits the update semantics account as I will explain in more detail below.

⁵ As mentioned before, there is a further debate concerning the need for pragmatic inference in recovering what is said itself—the literal content of an utterance. While this book argues that what is said is determined by the contributions of the rich repertoire of linguistic mechanisms through the course of semantic composition, regardless of how the issue is resolved, the standard account of implicatures and non-literal content maintains that the calculation of an implicature begins with recovering literal truth-conditional content of an utterance, and then using this content as an input to the calculation of the non-literal meaning.

modification of this account to accommodate for the fact that the modal claim in question does not express any proposition to begin with.⁶

This is not to say that a non-propositionalist couldn't in principle provide an account of pragmatic inference that fits in with their account of assertion. After all, non-propositionalists will not deny that logical relations obtain among modal utterances. However, the standard Gricean conversational maxims that one appeals to in the calculation of implicatures are couched in terms of *what is said*, that is, in terms of *informational*, *assertoric* content. Furthermore, even setting this aside, standard pragmatic accounts typically start from the premise that the speaker made a particular kind of a speech act—that they asserted something. It is a thorny and controversial question how exactly one should understand assertion in the non-propositionalist framework which denies that there is, strictly speaking, anything *said by* a modal claim.⁷ And, furthermore, the pragmatic calculations often implicitly or explicitly rely on further background principles concerning assertion, for instance principles maintaining that in making an assertion the speaker typically undertakes a particular kind of commitment (e.g., to the truth of the asserted content) or is constrained by particular kinds of norms (e.g., knowledge norm of assertion). One would, thus, need an account of what such commitments amount to within the non-propositionalist framework as well.⁸

Thus, a rejection of propositionalism calls for a re-thinking of the nature of meaning, communication, assertion, belief, and pragmatic inference.⁹ So,

⁶ There is of course, a question of whether this specific inference is best explained as conversational implicature. However, insofar as *any* pragmatic inference on the basis of an utterance of a modal claim counts as a conversational implicature, in Grice's sense, it would require, on the standard understanding, truth-conditional content as a basis for inference. For a recent sustained criticism of the Gricean account of conversational implicatures in general, see Lepore and Stone (2015).

⁷ Indeed, it might be more natural, in a non-propositionalist framework, to talk of various different types of speech acts, each with its own kind of characteristic effect on the context, which do not necessarily all fit under a single uniform category of assertion. (See, e.g., Yalcin, 2012b,a; see also Cariani, 2019, for a recent discussion.)

⁸ For some discussion of the challenges involved, see Cariani (2019). I note, again, that at least some non-propositionalists will maintain that modal utterances do have assertoric content, but will nevertheless deny that this content is to be understood as propositional, in the traditional sense (see, e.g., Schroeder, 2013 or Moss, 2015, 2018). These views will by design operate with non-traditional understanding of the contents of attitudes, assertion, and knowledge. It is not uncontroversial whether such views can fully answer the challenges I mention in the paragraph above (for some discussion, see, e.g., Cariani, 2016; MacFarlane, 2020). I refrain from discussing the details of such attempts at this point, as I shall argue against some of the key motivations for such views in the first place.

⁹ Of course, I don't take the need for such a revision in itself to be an argument against non-propositionalism. Non-propositionalists will precisely argue that such re-thinking is needed. I will, however, argue that the linguistic arguments for non-propositionalism, which take as a starting point the puzzling data concerning modal expressions discussed below, should be rejected. Thus, insofar as such linguistic arguments were to show that a radical re-thinking is required, they, I argue, fail. Further, once we take into account a broader range of linguistic data, we seem to have

why subscribe to such a radical thesis? To present the key data in support of non/propositionalism, let me start with the traditional understanding of the meaning of modal claims in context.

Modal expressions are traditionally understood as quantifiers over possible worlds (Kripke, 1959, 1963, 1980; Lewis, 1973; Kratzer, 1981). Of course, this quantification—much like any ordinary quantification—is not unrestricted: a modal does not simply quantify over *all* possible worlds. Though *might* is normally understood as an existential quantifier over possible worlds, in uttering ‘It might be raining,’ the speaker is not typically conveying that there is at least one world out of *all possible worlds* in which it is raining. Rather, such an utterance is typically understood to convey something more informative, for instance, that according to the available evidence, raining is an open possibility. The domain restriction is normally assumed to be determined by the context of utterance. This means that, according to the traditional account, a modal is a quantifier over a contextually restricted domain of possible worlds.¹⁰ So, modal claims, on this standard account, express propositions relative to contexts. An utterance of ‘It might be raining,’ for instance, 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 is raining. The utterance relative to

a *prima facie* argument against non-propositionalist accounts. While one might try to modify such accounts to capture the data, at this point, the burden of the argument is on non-propositionalists to motivate the departure from the traditional account. This is not to say that there might not be other arguments for abandoning the propositional paradigm—after all, the original expressivists, such as Ayer, Stevenson, and Hare, were offering a non-linguistic argument for non-propositionalism. More recently, Moss (2015) has argued for a non-propositionalist account of content, and correspondingly of belief and assertion. While Moss appeals to linguistic arguments against propositionalism I discuss below in motivating her account, she also offers foundational, theoretical arguments for non-propositionalism. However, her account, while non-propositional, is somewhat unusual among non-propositionalist accounts insofar as it maintains that modal utterances *do* express informational, assertoric content. Still such content is probabilistic, not propositional. My arguments against the linguistic motivation for non-propositionalism carry over against Moss’ account as well, as do my linguistic arguments in favor of propositionalism. Regarding foundational arguments, in broad outlines, Moss first argues that it is preferable to theorize about doxastic attitudes in terms of a single attitude—belief—toward probabilistic content, than in terms of different attitudes, some of which fall short of full belief, toward propositions; and, second, that given this, and given that a uniform account of content of belief and assertion is to be preferred, we need a non-propositional, probabilistic account of assertion as well. Discussing these foundational arguments in detail would take us too far afield. What is important for our purposes is that, in any case, any satisfactory theory will have to capture the empirical observations I argue are straightforwardly captured by my propositionalist account. That is, if my arguments are on the right track, we are not forced, from the standpoint of linguistic data, to adopt non-propositional accounts. Furthermore, any account, propositional or not, will have to take into account the further data that I point to below. So, my arguments will likewise be of relevance for those who maintain that one should abandon propositional content for foundational reasons.

¹⁰ As is standard, I shall sometimes call a modal’s domain of quantification *the body of information* that the modal quantifies over.

a context expresses the proposition corresponding to the set of worlds in which it is true.¹¹

The traditional account has, however, been increasingly challenged in the recent literature. What fueled the challenge was a battery of data involving the so called *epistemic contradiction*—the construction of the form *not ϕ and might ϕ* —that appears to evade standard truth-conditional, contextualist treatment. The problem is illustrated by the following example, due to Yalcin (2007):¹²

- (62) If it is not raining and it might be raining, then I'm uninformed about the weather.
- (63) If it is not raining and the body of information *i* doesn't rule raining out, then I'm uninformed about the weather.

The standard contextualist account described above cannot explain this stark contrast between (62) and (63): on this account 'It might be raining' just means that the (contextually determined) body of information *i* does not rule out raining. The problem is that there seems to be no body of information that can plausibly serve as a contextually provided domain for the modal in (62). That is, whichever body of information *i* the context were to deliver, (62) remains incoherent. To see this more clearly, suppose that the body of information is the one compatible with what the speaker knows. Then one has to explain why (62) is infelicitous, while (64) is perfectly acceptable, given that the account predicts they are truth-conditionally equivalent:

- (64) If it is not raining and for all I (the speaker) know it is, then I'm uninformed about the weather.¹³

¹¹ The proposition expressed by ϕ in an utterance of *might ϕ* is standardly called the *prejacent* of the modal *might*.

¹² Yalcin's original minimal pair is the following:

- (i) If it is not raining and it might be raining, then...
- (ii) If it is not raining but for all I know it is raining, then...

The minimal pair I have in the main text is more general, as it does not presuppose a particular contextual resolution for the domain: for example, that the contextually determined body of information the modal quantifies over is the one corresponding to what is known by the speaker, or maybe what is known by a group of conversational participants of which the speaker is a part.

¹³ Note that, unembedded, 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, so such utterances are akin to Moore-paradoxical 'It's raining, but I don't believe that it is.' But, as Yalcin (2007) points out, the problem is that only the sentence containing the modal continues to be problematic even when it is not asserted, but instead 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, or Moore-paradoxical. Note also that the problematic knowledge

The problem for the contextualist account is general: there are numerous bodies of information a context could deliver that would yield a plausible interpretation of (62). Yet whichever body of information *i* the context sets, the contrast between (62) and (63) persists. For example, just like (64), all of the following, *inter alia*, are perfectly felicitous, whereas the counterpart (62) remains incoherent:

- (65) If it is not raining and for all our group knows it is, then we are uninformed about the weather.
- (66) If it is not raining and given the salient body of information it is, then that body of information is uninformed.

The contrast between (62) and (63), non-propositionalists maintain, shows that there is no systematic way of determining plausible truth-conditional content for a modal utterance in a context: that incoherence of (62) is persistent across contexts which otherwise make various bodies of information available suggests the embedded modal claim in (62) does not express the proposition that given such body of information, the possibility of rain is open. Thus, the non-propositionalists' diagnosis is simple: the reason why we witness this behavior is because modal utterances do not express such propositional content.

Fueled by this battery of data which they take to challenge proposition-ism, the non-propositionalists further argue that a natural account of the interaction between modal utterances and context yields additional support to non-propositional accounts.¹⁴ Specifically, non-propositionalists start from the Stalnakerian understanding of context and contextual dynamics we have described in Chapter 6.¹⁵ Consider again the characteristic effect of an assertion:

ascriptions are only infelicitous in the first-person. The corresponding third-person ascriptions, for example, 'It's not raining, and for all John knows it is,' are perfectly felicitous, just as one would expect if the original infelicity had a pragmatic explanation. It is difficult, however, to get any reading of the modal in (63) that makes the utterance felicitous, again suggesting that such utterances are not merely pragmatically infelicitous.

¹⁴ While this battery of data has been at the center of the argument against propositionalism, there are further and, we'll see, related data that have been used against propositionalism. In particular it has been argued that standard contextualist propositional accounts cannot accommodate data concerning agreement and disagreement regarding modal discourse, and that they cannot accommodate data concerning iterated modality and higher-order uncertainty. The account I develop handles such cases as well. I will discuss issues arising from disagreement in Part IV, and I will return to the issue of iterated modality in Chapter 9.

¹⁵ I am simplifying somewhat. While the simple form of non-propositionalism, of the expressivist and update semantics flavor alike, is couched in terms of a broadly Stalnakerian picture of context-content interaction, at least some of the non-propositionalists replace the Stalnakerian picture of context with a more complex structure, for example, letting the context be a set of probability spaces, rather than worlds (Yalcin, 2007, 2011a). Such complications have to do with the treatment

when I utter ‘It is raining,’ if my assertion is successful, it will become common ground that it is raining, and the worlds in which it is not will be eliminated. It will thus become established (or accepted for the purposes of our conversation) that it is raining. Now suppose that you assert ‘It is raining’ against a common ground in which it is already accepted that it is raining. Given that the common ground already includes the proposition that it is raining, all the worlds in the initial context-set, call it c , are the worlds in which it is raining. But then, there are no worlds the proposition *that it is raining* will rule out from c —all the worlds in c are already compatible with that proposition. So, adding a proposition that is already in the common ground to the common ground doesn’t change the context. If by contrast the context is such that it includes some non-raining worlds, the assertion will rule those out transitioning into a context in which it is now true that it is raining. So, a successful assertion commits the agents to the truth of the proposition expressed.¹⁶ Finally, if it is true in c that it is *not* raining—if there are no raining worlds in c —then asserting that it is raining will eliminate all the worlds from the context-set leading to an absurd, empty context, $\emptyset$; this reflects the fact that at that point the conversational agents are committed to contradictory contents as part of their common ground.

Thus, an assertion of a declarative sentence brings about a characteristic effect on the context—the characteristic change that asserting the sentence induces. We can call such an effect a *context-change potential*. It is natural, given this setup, to think of the meaning—or at least a key aspect of the meaning—of a sentence as this characteristic effect that asserting it has on context. Further, it is natural given this setup, to say that a sentence is true in a context c just in case adding the proposition it expresses (in c) to the common ground does not change c .¹⁷ The idea is that, if ‘It is raining’ is already true (or accepted) in c , then asserting that it is raining does not change c .

Now, with this in mind, consider what kind of an effect an assertion of ‘It might be raining’ would have on the context—what it would take for it to

of probabilistic vocabulary, and are orthogonal to our main concern, so I omit further discussion in the interest of space. As we shall see in Chapter 10, the account I defend straightforwardly captures the behavior of probabilistic modals as well. Recall, also, that even some propositionalists appeal to richer notions of common ground (Stalnaker, 2014). These complications are, likewise, orthogonal to our main point.

¹⁶ Notice that for the characteristic effect to take place, the assertion has to be successful. At the very least my audience must not object to accepting what I asserted as a part of the common ground. So, if I utter, ‘It is raining’ and you respond with, ‘No, it’s not,’ the proposition I asserted will not become a part of our common ground.

¹⁷ This is the notion of truth in a context that is familiar from dynamic update semantics (Veltman, 1985; von Stechow, 2007; Gillies, 2010). A closely related notion is one that Yalcin (2007) calls *acceptance*, which we will return to below. Note that since the common ground reflects what is mutually accepted by conversational participants, it can happen that the actual world fails to be a part of the context-set, if the agents accept something false of the actual world.

be true in a context. As explained earlier, modal expressions are understood as quantifying over possibilities. So, given a context *c*, 'It might be raining' quantifies over the possibilities in *c*, and requires that there is at least one raining world in the context, that is, that raining is an open possibility.¹⁸ Now, suppose *c* does not rule out that it is raining. Then, there is no proposition to be added to the common ground—that is, no worlds to remove from *c*—so as to make 'It might be raining' true in *c*. It is already true in *c* just by virtue of *c* not ruling out raining. 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 *c* after such change.¹⁹ So, the characteristic effect of a modal utterance is different from that of an assertion of a non-modal claim: it does not add some asserted proposition to the common ground—there's no such proposition to be added—but merely *tests* it for compatibility with the prejacent.²⁰ That is, the goal of a successful assertion of *might p* is not to coordinate on a particular attitude—for example, the attitude of mutual acceptance—toward the proposition that *might p*, for there is no such proposition. Instead, the effect of a successful assertion of such utterance is coordination on a particular attitude toward the prejacent proposition, specifically the attitude of leaving open that *p*. And just as the effect of a non-modal assertion is reflected in the common ground, the effect of the modal claim is reflected in the common ground. To say that *might p* is true in a context *c* is just to confirm that the context does not rule out *p*. And this will be the case just in case the agents do *not* mutually accept that *not p*, that is just in case they coordinate on leaving open that *p*.²¹

¹⁸ Recall, the context outlines the ways the world can be, given what we mutually accept for the purposes of a conversation.

¹⁹ Except for, of course, eliminating all the worlds from the context, leaving us in an absurd state that supports everything. This said, it might be possible to *accommodate*—to repair the context by expanding it to include some previously ruled out raining worlds. But this does not amount to adding a proposition 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. Such accommodation involves revision of what we previously accepted, and it is typically conversationally strained, especially if the possibility was explicitly ruled out: after I tell you that it is not raining, I can't just continue by saying that it might be. Of course, a *hearer* might reject my assertion in the first place by following up with, 'No, it might be raining,' but in that case it did not become a part of our common ground that it is not raining in the first place.

²⁰ Nor can such proposition in general be extracted from the context post-change.

²¹ But if it's already common ground that *might p*, what is the point of uttering it? How can modal claims be informative on this account? For one, it could be that the agents don't actually mutually leave open that *p* prior to the utterance; the utterance can then serve as the proposal to strike this particular attitude toward *p*. (This might require some revision of the prior common ground.) And even if they do, this might not be immediately salient, so the reminder might still not feel redundant, but serve to draw attention to a particular possibility. That said, note that one might need to complicate the current account to allow for the distinction between leaving open that *p*

Non-propositionalist propose that the characteristic effect of a modal utterance on the context—its context-change potential—is the key aspect of modal meaning, and since this effect is irreducible to adding a proposition to the common ground, maintain that modal claims do not express propositional content. The general approach sketched above can be implemented in different ways, and in different frameworks. The two most prominent and influential versions of the account are the dynamic update semantics and expressivism. While they share the general theoretical foundations outlined above, their implementations of these ideas differ. Thus it will be useful to sketch and compare the two approaches.

The key difference between the two implementations lies in the division of semantic and pragmatic labor. While dynamic update semantics views the context-change potential as the semantic content of an utterance, the expressivist account delegates it to pragmatics. More specifically, dynamic update semantics proposes we take the characteristic effect that an utterance has on a context as its *compositional 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. As described above a sentence is true in c , just in case updating with it does not further change c .²²

On the dynamic update semantics account, we can easily extract propositional content of a (declarative) sentence from the dynamic meaning—the context-change potential—provided that the utterance does not contain any modal expressions: if ϕ is a non-modal declarative sentence, updating with ϕ just amounts to adding the proposition expressed by ϕ to the common ground.²³ Formally, this means that the dynamic meanings—context-change potentials—of sentences within 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 . Clearly, there is nothing *essentially*

and simply not having considered p one way or another. See Yalcin (2011b), Willer (2013), Stalnaker (2014) for some discussion. I set this further complication for non-propositionalist accounts aside.

²² Notice, since in the case of a modal claim there is no proposition that can be recovered as the asserted content of the utterance, the dynamic update semanticists will still deny that there is some assertoric, informational content, over and above the context-change potential, that a modal claim contributes. In the case of non-modal assertions as we shall see, such proposition is recoverable, but any remaining context-sensitivity in the sentence asserted, and so, any context-sensitivity that needs to be resolved to determine the propositional assertoric content, is taken to be resolved by potentially non-linguistic features of the input context (as was the case for the Stalnakerian model in general (see Chapter 6)).

²³ Notice, any other context-sensitivity in ϕ has to be resolved by appeal to context, in the usual way, to determine what proposition ϕ expresses in that context. So, while we can recover propositional meaning from the dynamic meaning, even on this account, we still need to appeal to context-sensitivity resolution in bridging the gap between semantic content and assertoric content.

non-propositional about such context-change potentials, and we can always recover propositional contents from them. So, dynamic update semantics for a non-modal fragment is still not a radical departure from the standard propositionalist account.²⁴

However, as anticipated by our informal discussion of the characteristic effect of an utterance of a modal claim, once we turn to modal claims, the update will no longer reduce to an intersection with the propositional content expressed. ‘It might be raining,’ too, is interpreted as an update function that takes a context and returns a context, but as per the characteristic effect it encodes, this function merely tests the context to check if the prejacent is compatible with it: if it is, the context is returned unchanged (the modal claim is true in it); otherwise, the context is reduced to an absurd state, $\emptyset$.²⁵ More generally, the dynamic

²⁴ That a dynamic update semantics for a non-modal fragment is intersective and so reduces to updating by intersecting a context with a proposition is underwritten by a formal result due to van Benthem (1986), which establishes that any such reduction holds just in case all updates within the update system are *eliminative* and *distributive*. Take a propositional fragment of English, closed under negation and conjunction, where atomic sentences are represented as propositional constants. Let $\llbracket \phi \rrbracket$ be the update function associated with a sentence ϕ . In the standard prefix notation, let us write ‘ $c\llbracket \phi \rrbracket$ ’, where the context c is an argument to the function $\llbracket \phi \rrbracket$. Then eliminativity and distributivity are defined as follows:

- *Eliminativity*: An update is eliminative just in case it always (at most) eliminates worlds from the input context c —that is, it always delivers a subset of c :

$$c\llbracket \phi \rrbracket \subseteq c, \text{ for all } c \text{ and } \phi.$$

- *Distributivity*: An update u is distributive just in case 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, that is, just in case the following holds:

$$c\llbracket \phi \rrbracket = \bigcup_{w \in c} \{w\}\llbracket \phi \rrbracket$$

- An update that is *Eliminative* and *Distributive* is also *Intersective*, which just means that it amounts to intersecting the context with a proposition, that is, to adding a proposition to the common ground. Formally, an update is intersective just in case it satisfies the following property:

Intersective Update: $c\llbracket \phi \rrbracket = c \cap W\llbracket \phi \rrbracket$, where W is the set of all possible worlds.

See van Benthem (1986); Groenendijk and Stokhof (1990). For more on van Benthem’s results in relation to truth-conditionality, see von Fintel and Gillies (2007) and Rothschild and Yalcin (2017). Rothschild and Yalcin (2017) provide a generalization of the result which does not rely on the notion of unstructured contents.

²⁵ For concreteness, here is a simple dynamic language that captures this interpretation (see in particular Veltman, 1996; see also van der Does, Groeneveld, and Veltman, 1997, Gillies, 2004, von Fintel and Gillies, 2007, *inter alia*). Take again a basic propositional fragment, a smallest set including atomic propositional letters, closed under negation ($\neg$), conjunction ($\wedge$), and a unary modal operator ($\Diamond$). Let a model be a pair of a set of worlds and an interpretation function, $\langle W, \mathcal{I} \rangle$, where, $\mathcal{I}$ maps propositional letters to subsets of W . We define a function, $\llbracket \cdot \rrbracket$, from well-formed formulae of the language to functions from contexts to contexts, as follows:

Definition 8.1 Basic update semantics

- $c\llbracket p \rrbracket = \{w \mid w \in c \cap \mathcal{I}(p)\}$
- $c\llbracket \neg \phi \rrbracket = c \setminus c\llbracket \phi \rrbracket$

updates associated with modals are merely testing whether the context as a whole has a certain property. This means they are not reducible to adding some proposition to the common ground, that is, to narrowing down the space of open possibilities.²⁶ So, the dynamic meaning of modal claims, and specifically, of *might* ϕ is not merely superficially dynamic. We cannot reduce it to adding some (non-trivial) propositional content to the common ground.²⁷

What about expressivism? According to this type of view, the primary goal of a conversation is to coordinate interlocutors' mental states.²⁸ The Stalnakerian common ground is a reflection of this coordination. The common ground supervenes on the collective attitudes of the conversational participants.²⁹ As

- $c[\phi \wedge \psi] = c[\phi][\psi]$
- $c[\Diamond \phi] = \{w \in c : (c[\phi]) \neq \emptyset\}$

Note that, as described above, the system is eliminative, but not distributive. Note also that since we are dealing with a toy propositional fragment to which we assume expressions of a fragment of English are mapped, any sub-clausal context-sensitivity that might require resolution before we can recover a proposition expressed by an English sentence in context is abstracted away in the toy system.

²⁶ The non-propositionality is a result of the test behavior of modal updates. Being a test, a modal update is sensitive to the properties of the context globally, as a whole. But these properties need not carry over to each singleton subset of the context. In other words, an update with a *might* claim violates distributivity. Here's a quick counterexample from von Fintel 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). Since c contains a p -world, the update with *might* p returns c as is, that is, $\{w_1, w_2\}$. However, updating each singleton subset pointwise with *might* p , and then taking the aggregate of the results, does not return c ; 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. Intuitively, the failure of distributivity reflects the fact that modal claims are sensitive to global properties of contexts or information states—properties of context taken as a whole. Thus, their effect is irreducible to the sum of the effects on the local points in the context.

²⁷ Nor can such content be recovered from the update or the context post-change. Notice that an update of an input context c with *might* ϕ , and that with *must* ϕ (assuming a test-like interpretation for *must*, a dual of *might*), if successful, will return exactly the same result, namely c , and if unsuccessful, it will again return the same result, namely $\emptyset$. Both simply test whether a context satisfies some property. So, the only proposition retrievable from any context after updating with either *might* ϕ or *must* ϕ is either the whole input context, or the empty set.

²⁸ See in particular Yalcin (2007). For related accounts, see, for example, Swanson (2006); Charlow (2015); Starr (2016). See also Moss (2015, 2018). As mentioned before, Moss' account is unique, as it argues for non-propositional, probabilistic content which plays the role of the objects of belief and assertion. She does not hold that modal utterances fail to express any assertoric content. Compare footnote 9 in this chapter.

²⁹ I don't mean to suggest mere supervenience on the mental states of the conversational participants suffices for coordination. For instance, suppose everyone has a credence that the Earth will become uninhabitable by the year 3000. Then there is an average credence in that proposition that supervenes on the attitudes of all the individuals, yet no coordination is involved in drawing out that average. (Thanks to the anonymous reader for the example and for pressing the point.) The thought here is just that there are no changes in the common ground without changes in the relevant attitudes of the group. p is common ground just in case we mutually accept that p ; likewise, *might* p is common ground just in case we mutually accept it, that is just in case we mutually agree to treat p as live possibility for the purposes of the conversation. Note that on the non-propositionalist picture, in both cases our mutual acceptance requires us to coordinate on an attitude toward p —to come to

before, an assertion of a non-modal claim is a proposal to add the proposition expressed by it to the common ground—to make it common knowledge that the interlocutors all accept it for the purposes of a conversation. This is in essence a proposal to strike a particular attitude, that of acceptance, towards the proposition expressed. However, an utterance of a modal claim is not a proposal to add the proposition expressed to the common ground, nor a proposal to strike an attitude towards such a proposition, for there is none. It is a proposal to strike a particular attitude toward 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*. Note, this in turn means that the context will have to have a certain property, namely, that of *being compatible with p*. This just means that the effect of a modal claim is again to test whether the context satisfies a particular property, just as was the case in the dynamic system.

The dynamic update semantics characterizes this test, encoded in the update function, as the compositional semantic content of a modal claim. Expressivism, by contrast, maintains that the compositional semantics determines only a property of bodies of information (in the case of *might*, the property of being compatible with the prejacent). The test, in turn, is just a characteristic pragmatic effect of an assertion of a modal claim on the context, a byproduct of the fact that such an assertion proposes that the context instantiates the property in question. This is the sense in which the two theories differ in terms of the division of semantic labor.

To understand the idea behind expressivism more clearly, recall the standard truth-conditional, contextualist account of content. As described in Chapters 1 and 2, the standard truth-conditional, contextualist 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*. Recall, *w* is an index, or a point of evaluation relative to which the proposition receives a truth-value. The proposition is taken to be what's expressed by an utterance, its assertoric content. Now, if the assertoric content were identical to compositional semantic content, then the point of evaluation would not have to contain anything more than just a possible world. However, as discussed in Chapter 1, it is possible that semantic composition might require compositional content to be sensitive to indices that contain parameters beyond possible

accept it, or to agree to leave it open; in neither case is it an attitude toward some modal proposition *might p*, as there is no such proposition.

worlds: for example, semantic composition might require relativizing truth to times as well.³⁰

Standard contextualists maintain we can recover propositions asserted from this content by letting the context of utterance supply the value of any such further parameters in the index. Now, expressivists, too, propose to further enrich the index, by relativizing truth not just to a context and a world, but also to an information state. 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 . In other words, the content places no constraints on the information state parameter. The relativization to an information state is idle and, so, the propositional content can be easily recovered. But, once we turn to modals, 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 provides the domain of quantification for the modal operator.³¹

How does this non-idleness lead to non-propositionalism? From a formal standpoint, the expressivist kind of compositional semantics is closely related to the standard, truth-conditional contextualist semantics of Kaplan (1989b,a). The compositional semantic content of a sentence in a context, on both of these accounts, is a property of points of evaluation, that is, of indices. From a formal

³⁰ As before, I'm not endorsing the idea that semantic content has to be sensitive to richer indices containing times; I'm just using it for the purposes of illustration.

³¹ For the sake of concreteness, here is a sketch of an expressivist semantics based on Yalcin (2007, 2012b). Take a propositional fragment, a smallest set including atomic propositional letters $(p, p_1, \dots)$, closed under conjunction ($\wedge$), negation ($\neg$), and a one-place propositional operator ($\Diamond$). Let a model be a pair of a set of worlds and an interpretation function, $\langle W, \mathcal{I} \rangle$, where $\mathcal{I}$ maps propositional letters to subsets of W . Given a model $\mathcal{M}$ let $\llbracket \cdot \rrbracket$ be a function assigning a truth value (0 or 1) to each well-formed formula given a context and an index consisting of a pair of a world and an information state. We define truth relative to a context and an index recursively:

Definition 8.2 Expressivist semantics

- $\llbracket p \rrbracket^{c, \langle w, s \rangle} = 1$ iff $w \in \mathcal{I}(p)$
- $\llbracket \neg \phi \rrbracket^{c, \langle w, s \rangle} = 1$ iff $\llbracket \phi \rrbracket^{c, \langle w, s \rangle} = 0$
- $\llbracket \phi \wedge \psi \rrbracket^{c, \langle w, s \rangle} = 1$ iff $\llbracket \phi \rrbracket^{c, \langle w, s \rangle} = 1$ and $\llbracket \psi \rrbracket^{c, \langle w, s \rangle} = 1$
- $\llbracket \Diamond \phi \rrbracket^{c, \langle w, s \rangle} = 1$ iff $\exists w' \in s. \llbracket \phi \rrbracket^{c, \langle w', s \rangle} = 1$

Notice that only the modal operator places a constraint on the information state parameter; for all other clauses, constraints are only placed on the world parameter, and so we get ordinary, propositional, truth-conditional content. Here, too, since we are dealing with a toy propositional fragment assumed to reflect the logical form of a fragment of English at some level of abstraction, any sub-clausal context-sensitivity that might require resolution before we can recover a proposition expressed by an English sentence in context is abstracted away in the toy system. We are also assuming—for simplicity—that no further enrichment of the index, beyond the world and information state parameter, is needed.

point of view, the only difference is that the expressivist compositional semantic content of a sentence in a context is a property of information states, that is, a function that maps a world and an information state to a truth-value, while on the Kaplanian account, the compositional semantic content, as we have seen in Chapter 2, is a function mapping a world, time and location (and possibly other parameters) to truth-values—a property of worlds, times, and locations. What makes the Kaplanian account a standard propositionalist one, recall, is that, on this account, we can recover propositional, assertoric content from these compositional semantic contents by saturating the time and location parameter of the index with a contextually supplied value. So, a sentence in a context expresses a proposition p comprising a set of worlds w such that p maps w and the time and location of the context to truth.³²

Contrary to Kaplan, however, expressivists deny that the context of utterance determines the information state in the index. Otherwise, we could easily recover the propositional content of modal claims in exactly the same way, by simply letting the context saturate the value of the information state parameter. Somewhat less abstractly, by contrast with Kaplan, expressivists deny the propositional content, as the assertoric, informational content of an utterance—the content we assert and believe—can be recovered from the compositional semantic content by saturating the information state parameter by some contextually provided value. Indeed, they maintain that modal utterances lack assertoric content altogether—they are not even in the business of expressing informational content. They merely have a contextual effect, insofar as uttering them places a constraint on a conversational context that it has to satisfy the compositionally assigned property of information states.³³ Adopting expressivist terminology

³² Here we treat a proposition as a function from worlds to truth values, that is, just a property of worlds. Note that, as described in Chapter 2, we can also determine the notion of truth in a context *simpliciter*: a sentence is true in a context *simpliciter*, just in case it expresses a true proposition in that context.

³³ As mentioned earlier (see footnote 1 in this chapter), both expressivist and the standard Kaplanian account are structurally similar to relativism, as developed in, for example, Egan, Hawthorne, and Weatherson (2005); Egan (2007); Kolodny and MacFarlane (2010); MacFarlane (2014). Relativists likewise assign properties of information states as semantic contents of utterances in context, and like expressivists, they deny that this information state parameter can be determined by the *utterance* context. However, as explained earlier, they maintain that the value of the information state parameter is determined by, and varies with, the context of assessment. The primary difference between expressivism and relativism lies in the theoretical role assigned to this compositional semantic content, as relativists take the compositional semantic content of an utterance to be its *assertoric* content. This assertoric content is not true or false relative to a context of utterance *simpliciter*: what we say with our modal utterances can be true relative to the information possessed by some assessors, but false relative to other. Expressivists, by contrast, maintain that modal utterances lack assertoric content altogether—they are not even in the business of expressing informational content, and the relevant information state parameter isn't fixed either by context of utterance or that of assessment. (Exception to this are expressivist accounts such as those developed in Schroeder (2013) or Moss (2015, 2018), which agree that modal utterances have assertoric content, but like relativists, deny that that content is standard propositional content.)

(Yalcin, 2007), we can say that an information state s *accepts* a sentence ϕ just in case $\forall w \in s: \llbracket \phi \rrbracket_{c, \langle w, s \rangle} = 1$.³⁴ Then an utterance of *might* ϕ is a proposal to make *might* ϕ accepted in the context, in this sense; that is, it is the proposal to make the context-set satisfy the property that the semantics assigns to *might* ϕ .

Why do expressivists deny that the context can fix the information state parameter—the contextually relevant body of information? First, the idea is, again, that a modal utterance does not contribute a bit of information to the common ground—there’s no proposition you can add to the common ground to make *might* p accepted if it isn’t already accepted; *might* p merely tests whether common ground leaves open the possibility p or not. Second, the non-propositionalist picture appears to fit the puzzling data in (62) and (63): the data shows, expressivists maintain, that the context cannot play a role of supplying an information state in a way that would deliver plausible truth-conditions.

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 compositional 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 utterance reducible to a proposition, nor can a proposition be recovered from it. So, on neither account is a proposition expressed or asserted by an utterance of a modal claim. And both approaches maintain that the dynamic effect of a modal utterance on the context is the key aspect of its interpretation.

We said that the motivation for non-propositionalism lies in a seeming failure of propositionalist account to explain the contrast between (62) and (63) (repeated below). But how do non-propositionalists explain this contrast? In order to answer this question, we need to specify how both accounts interpret conditionals.

- (62) If it is not raining and it might be raining, then I’m uninformed about the weather.
- (63) If it is not raining and the body of information i doesn’t rule raining out, then I’m uninformed about the weather.

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 .³⁵ Thus, the conditional’s effect

³⁴ Here again $\llbracket \cdot \rrbracket$ is a function that maps a sentence, context, and an index comprising a world and an information state pair to a truth-value. See footnote 31 in this chapter.

³⁵ Formally, we can extend our system in footnote 25 in this chapter to be closed under a conditional ($\rightarrow$). Then, we interpret the conditional as follows:

on the context can be likewise described as *performing a test*: it tests whether updating the context with the antecedent would land one in a state that accepts the consequent. If it does, the context passes the test, and is returned unchanged, and the conditional is thus true; if not, the test fails, returning the absurd, empty context, and the conditional is thus false. The expressivist account endorses a corresponding non-dynamic interpretation: a conditional hypothetically shifts an information state parameter s to the maximal non-empty subset of s that accepts the antecedent, and checks whether this state accepts the consequent as well. The conditional is true, that is, accepted (relative to a context, a world, and the original information state s), just in case it does.³⁶ So, on the expressivist account the conditional likewise serves as a kind of a test on an information state: it hypothetically adds the antecedent to the information state, and then checks whether thus updated information state accepts the consequent or not. If it does, the conditional is true, otherwise, it is false.

Given this, the dynamic and expressivist accounts can provide very similar explanations of the contrast between (62) and (63). 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 we have just eliminated all raining worlds from c , this test fails, and we get the 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 (in the given context) is compatible with raining. Hence, the contrast between (62) and (63).

An expressivist offers a similar explanation. The conditional in (62) 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, the

Definition 8.3 Update semantics for a conditional

- $c[\phi \rightarrow \psi] = \{w \in c : (c[\phi])[\psi] = c[\phi]\}$

³⁶ So, suppose we extend the expressivist system in footnote 31 in this chapter, to be closed under a conditional ($\rightarrow$). We interpret the conditional as follows (Yalcin, 2007):

Definition 8.4 Expressivist semantics for a conditional

- $\llbracket \phi \rightarrow \psi \rrbracket^{c, \langle w, s \rangle} = 1$ iff $\forall w' \in s_\phi. \llbracket \psi \rrbracket^{c, \langle w', s_\phi \rangle} = 1$
where s_ϕ is the largest s' such that $s' \subseteq s$, $s' \neq \emptyset$ and $\forall s'. \llbracket \phi \rrbracket^{c, w, s'} = 1$

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.³⁷ This is why one cannot hypothetically suppose that it isn't raining, but it might be. And this explains why (62) is infelicitous. Meanwhile, there are some non-empty states that accept 'it is not raining, and the relevant body of information i (in the given context) doesn't rule out raining'; hence, the contrast with (63).

So things are looking good for non-propositionalists. Does this mean that there is a compelling reason to abandon representational accounts of content for modal discourse? Is the radical conclusion warranted? Not so fast. I will now proceed to argue, in Chapter 9, that the key assumption which underscores the non-propositionalist argument from (62) and (63) is the traditional assumption of the non-linguistic model of context-sensitivity resolution. Once this assumption is removed, and once we invoke a linguistic account of context and context-sensitivity resolution, we shall see that not only do data like (62) and (63) not tell against the view that modal utterances express propositional content, but moreover, a host of other data that non-propositionalism struggles to predict, are suddenly straightforwardly accounted for.

³⁷ Notice that on this account, the order of conjuncts makes no difference to whether (62) is infelicitous or not. On the dynamic account, an asymmetry is predicted, by contrast, as the update with the 'It's not raining and it might be' leads to an absurd state, but 'It might be raining and it's not,' doesn't.

Dynamic Propositionalism

On the standard contextualist, truth-conditional account, a modal utterance expresses a content—a proposition—relative to a context, and the role of a context is to determine the proposition expressed by supplying the value of the relevant contextual parameter, fixing the domain restriction for the modal. More generally, the context of an utterance determines its semantic content, the proposition expressed, which in turn is, or at least determines, the utterance's truth-conditions

The contextual resolution has been assumed to proceed on the model of a true demonstrative: speaker's intentions and/or general contextual cues fix the interpretation of a modal, in a similar way they fix the interpretation of an utterance of 'He is happy.' And it is the familiar, open-ended abductive process of interpretation that guides the audience in retrieving the relevant contextual interpretation. Ultimately, the aim of interpretation is to recover the overall most plausible interpretation.

The model, while *prima facie* plausible, is misguided for exactly the same reason it was misguided in the case of pronouns. As in the case of pronouns, there is very good evidence that modals are sensitive to the contextual prominence of candidate interpretations: of possibilities, or information states in a discourse. They select the most prominent body of information in the given context as the restriction on their domain of quantification. More importantly, just as in the case of pronouns, prominence is determined by linguistic rules—by discourse conventions. This means that, on the one hand, modals exhibit a more flexible type of context-sensitivity than has been assumed, because prominence dynamically changes as the discourse progresses. So, a modal will not always be anchored to some specific body of information, in some static, fixed context, say, the speaker's or the interlocutors.' But this type of context-sensitivity is also more constrained than what the traditional model assumes, because the changes in prominence are governed by linguistic rules, specifically, by discourse conventions of the sort we have seen in Chapter 5. They are not sensitive to an open-ended array of extra-linguistic resources. In short, I argue that modals exhibit the same type of context-sensitivity I have argued is present in demonstrative pronouns: they are prominence-sensitive, pure indexical type expressions, where the kind of contextual prominence they are sensitive to is one governed by discourse conventions.

That modals are prominence-sensitive expressions is not all that surprising if we approach the issue from the standpoint of the literature on modal subordination, a phenomenon whereby modals receive a coordinated interpretation, as in (67), repeated below, where one modal is interpreted relative to another introduced earlier in the discourse (i.e., relative to a linguistically introduced antecedent):

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

In (67) the modal in the second sentence isn't interpreted as unrestricted: the second sentence doesn't merely convey that a wolf is going to eat the addressee come what may, or even given all we know. The possibility introduced by the modal in the first sentence restricts its interpretation: a wolf would eat the addressee, *given that one walks in*. A natural way of describing this simple instance of modal subordination is by analogy with an anaphoric dependence of a pronoun on a preceding linguistically introduced antecedent, as in (17), repeated below:

(17) A woman came in. She sat down.

As discussed in detail in Part I, in examples like (17), the antecedent 'A woman' makes a particular interpretation prominent, and a pronoun picks it up, thereby receiving a dependent interpretation. Similarly, in (67), a modal in the first sentence makes a certain possibility—the one in which a wolf walks in—prominent, and the modal in the second sentence is in turn restricted by this possibility.

The parallel in interpretive range between pronouns and modals runs deeper. Modals and pronouns exhibit a strikingly similar range of possible interpretations in a discourse. For instance, Stone (1997, 1999) argues that modals allow for referential, bound and discourse-bound readings.¹

First, both classes of expressions can be used referentially—what we would describe as deictic uses—to pick out some entity from the non-linguistic conversational situation. To wit, consider the following:

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

¹ Stone's argument for the anaphoric parallels between pronouns and modals mirrors Partee's argument for the anaphoric parallels between pronouns and tense (Partee, 1984). Thus, the full observation is that pronouns, modals, and tense all allow for a parallel range of interpretations and interpretive dependencies. For other accounts of modal subordination that emphasize this kind of anaphoricity of modal expressions, see Kibble (1994, 1995); Geurts (1995); Frank (1996); Frank and Kamp (1997); Kaufmann (2000); Bittner (2011, 2001); Brasoveanu (2010a,b).

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

Just as (68) can be used to refer to some salient female—made prominent in the speech situation—so too can (69) be used to refer to prominent possibility in the speech situation. (69) doesn't just mean that the neighbors will kill the speaker unconditionally, come what may; they would only do so provided that the speaker buys the stereo in question.

Furthermore, like pronouns, modals can be bound by an operator introduced prior in the discourse. So just like 'she' in (15), repeated below, can co-vary with 'every woman,' so too, in (70), 'will' in the consequent co-varies with different instantiations of the antecedent—different late arrivals will lead to different instances of denials of permission.

- (15) Every woman drives the car she owns.

- (70) If a concert goer arrives late, he or she will not be permitted into the auditorium.² (Stone, 1997)

Furthermore, like pronouns, modals allow for the so-called "donkey anaphoric" readings. Just like 'it' in (71) co-varies with the indefinite noun phrase 'a donkey' even though it is outside of its syntactic scope, so in (72) the modal in the consequent, 'will,' co-varies with the sub-constituent of the antecedent clause 'if an enemy captures it.'

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

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

Stone concludes that the interpretation of modal expressions is "an anaphoric process" analogous to the anaphoric interpretation of pronouns.³

² I cite the original example that Stone provides, but in case you are worried about treating 'will' as a modal, note that 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 (76) and (72)). I shall assume that 'will' indeed is a modal, but nothing hinges on this decision. See, for example, Comrie (1989); von Stechow (1995); Kissine (2008); Portner (2009); Cariani and Santorio (2018) for further discussion concerning the modal status of 'will'.

³ Further support for the interpretive analogy between modals and pronouns comes from cross-linguistic considerations: it has been observed that languages such as Warlpiri or American Sign Language (ASL) allow for a single anaphoric expression to be ambiguous between pronominal and modal interpretation, which calls for a uniform treatment. This strongly suggest that we should

It is worth noting that, just like modals, conditionals also exhibit prominence-sensitivity:

- (73) (Looking at a high-end stereo in an electronics store) If I turned the volume up, my neighbors would kill me.
- (74) If a wolf walks in, it might attack. If John has a tranquilizer gun, he will shoot it. (Stojnić, 2019)

So, cross-discourse dependencies that modals and conditionals exhibit show that these expressions are sensitive to relative contextual prominence of candidate interpretations, in the same way pronouns are. That pronouns allowed for discourse-bound readings, whereby an interpretation depended on some linguistically introduced antecedent, as in (17), suggested to us that the effects induced by the prior discourse have to be reflected in the context relevant for the interpretation of a particular utterance or expression: in order to predict the correct interpretation, we have to take into account how the context changes with the evolving discourse, not just how it represents some isolated real-world conversational situation at a particular static moment. That modals exhibit this type of discourse dependence suggests they are sensitive to the evolving context in a similar way. Whatever else the context is like in which the second sentence in (67) is uttered, it has to at least 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, given that one walks in*. And, just as we have seen in our earlier discussion of pronouns, the context can dynamically shift in this way not only between distinct sentences of a discourse, but even within a single sentence, as is in the following examples:

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

In (75) and (76), an element occurring earlier within the same sentence—the antecedent of the conditional in (75) and *may* in (76)—introduces a possibility that affects the interpretation of the modal that comes later within the same sentence. In (75), we understand that a wolf would eat the addressee, if it comes

in, and in (76), we understand that given that there are other 1961 committee retirements, they will be leaving by the choice of Republican voters.

Adopting the standard treatment of modals as quantifiers over possibilities, we can still maintain that their context-sensitivity is reflected in contextually provided restriction on their domain of quantification. However, what the examples above demonstrate is that the way the restriction on the domain of quantification of a modal is determined is sensitive to prominence: just as the pronoun selects a prominent antecedent from a prior discourse or the situation of utterance, so too a modal selects a prominent possibility either introduced by a prior discourse, or selected from the conversational situation. Since the domain of quantification is just a set of worlds, we can think of the restrictor likewise as a set of worlds—a proposition. The restriction on the domain of quantification then results from intersection with a restrictor proposition.

Notice that these data already pose a challenge to the non-propositionalist accounts. Take, for instance, (67). Non-propositionalist accounts model utterances such as ‘A wolf might walk in,’ as exhibiting test-like behavior, testing whether the context-set is compatible with the prejacent—whether it contains a world in which a wolf walked in. So, the first sentence merely tests the input context, *c*, in this way; if the context passes this test, it is passed down unchanged 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. As a result, the second sentence is interpreted relative to that whole non-restricted context. And this delivers the intuitively wrong, unrestricted reading: (67) conveys 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 these propositions that were introduced earlier in the discourse, and use them to restrict the domain of quantification of subsequent modals. The standard dynamic update, and expressivist, account has no built-in resources that allow for this kind of bookkeeping.

Similarly, the following are problematic from the non-propositionalist standpoint as well:

(77) If a wolf might walk in, we should hide.

(78) If a wolf might walk in, we must be in an area where wolves roam.

Intuitively, ‘should’ in (77), as well as ‘must’ in (78), is restricted by the proposition expressed by the antecedent, that is, the proposition that a wolf might walk in (given some prominent body of information). But on a non-propositionalist

account there is no such proposition. The output of the update with ‘a wolf might walk in’ is not the proposition *that a wolf might walk in*, the proposition we need for the restriction of the subsequent modal. Nor is such a proposition recoverable from the output information state.

Now, it is possible to complicate the simple non-propositionalist semantics I described above to accommodate the examples like (77) and (78).⁴ However, more is needed, since there are examples like the following, which require keeping track of *both* modal and non-modal possibilities, introduced simultaneously:

- (79) If a friend might come to the dinner, we should prepare an extra seat, since they would need to sit somewhere.
- (80) If a wolf might walk in, we should hide. It could attack us.

Further, modals also allow for the so-called exocentric readings, where the modal is interpreted relative to information that *excludes* the speaker’s, that is, as picking out a contextually prominent information state:

- (81) We don’t know whether the new test we are piloting rules out that John has cancer, but we know (by means of independent, reliable tests) that John does not, in fact, have cancer. Before looking at the results of the new test, we say:
- a. If John might have cancer, then it must be a false positive.

Exocentric uses are generally difficult to explain from the perspective of non-propositionalism, as they intuitively seem to describe some contextually prominent information, and so, require a shifting of the information state to the one that’s contextually prominent. But from the standpoint of my account, these are just referential (deictic) uses of modals.⁵ The point is that accommodating full range of data would require a non-propositional system to allow for a complex mechanism of tracking both modal and non-modal possibilities across discourse simultaneously. While this might be possible, the resulting system would in essence mimic the propositional system I develop below.

The standard contextualist truth-conditional account, of the kind developed in Kratzer (1977, 1981), does not automatically capture the anaphoricity of modals (and conditionals) either. This is due to the fact that the way in which the Kratzerian account represents context-sensitivity of modal utterances does

⁴ For instance, see Yalcin (2011a), Willer (2013); or, for the examples involving deontic modality, Cariani (2016).

⁵ The observation that *might* allows for such exocentric uses goes back to Egan, Hawthorne, and Weatherson (2005) and von Stechow and Gillies (2007).

not easily allow for the modal to receive a restriction by a prominent body of information introduced earlier in the discourse.⁶ What we need is a way of keeping track of an anaphorically retrieved restrictor explicitly as a separate contextual element. One simple way to do this is the following. As in the usual quantificational treatment, we take a modal to be a quantifier over possibilities, but requiring that a relation specified by the modal force hold not simply between the prejacent and the domain, but between the prejacent and the body of information that comprises the worlds that belong both to the domain and the contextually provided restrictor. So for instance, since *might* is an existential modal an utterance of *might* ϕ is true (at a world, in a context) just in case there is an accessible world that's both a prejacent world and among the contextually prominent worlds.

Making this a bit more concrete, where ' $M(p, q)$ ' is an abbreviation for the truth-condition expressed by an utterance of *might* ϕ , q is the proposition expressed by the prejacent ϕ of the utterance of *might* ϕ , p the proposition corresponding to an anaphorically retrieved restrictor, and ' R ' denotes a contextually provided accessibility relation, that determines a set of epistemically accessible worlds from a given world w , we define the truth-condition expressed by an utterance of *might* ϕ , ' $M(p, q)$ ', as follows:

Definition 9.1 $M(p, q) = \{w \mid \exists w' : wRw' \ \& \ w' \in p \ \& \ w' \in q\}$

Definition 9.1 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.⁷ Crucially, we factor in the contribution of the restrictor separately, which will allow us to track how it is recovered in a context. (*Mutatis mutandis*, we can specify the truth-conditions for other modals; for example, *must* can be characterized as the universal dual of *might*.)

⁶ As Stone (1997) points out, neither of the contextual parameters characterizing the context-sensitivity of modals on Kratzer's semantics—the modal base or the ordering source—is suitable to provide a semantic parameter that can be contributed anaphorically by prior discourse so as to serve as restriction on a modal domain. For further discussion, see Kratzer (2012, ch. 4), as well as Brasoveanu (2010a) and Bittner (2014).

⁷ Notice that the accessibility relation must be contextually specified. So, there 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, that is, how the accessibility relation is contextually determined. I will address this issue in some detail in Chapter 12.

Since our key examples involve conditionals, we need to define the truth-conditions for the conditional that likewise capture the anaphoricity. We do this as follows:⁸

Definition 9.2 $Cond(p, q, r) := \{w \mid \forall w' : wRw', \text{ if } w' \in p \ \& \ w' \in q, \text{ then } w' \in r\}$

As before, p is the anaphorically retrieved restrictor, q the proposition expressed by the antecedent, and r the one expressed by the consequent. An utterance of a conditional thus 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. That is, 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. Like modals, conditionals, too, are always evaluated against some contextually prominent body of information. Crucially, this body of information need not correspond to the unrestricted set of epistemically accessible worlds discourse initially as was illustrated by examples (73) and (74).

Given these modified truth-conditions, we can now track domain restriction as a separate contextual parameter, just as we did with the antecedent of a pronoun, and so, we have resources to capture the anaphoric dependencies, as well as deictic and exocentric uses of modals displayed in the examples surveyed above. But haven't we already seen that any account on which a modal expresses a proposition partially determined by a contextually supplied body of information struggles to explain the contrast between (62) and (63), repeated below? The key moral of these examples was that there is no body of information—restricted or otherwise—that can predict this contrast. Isn't this a problem?

- (62) If it is not raining and it might be raining, then I'm uninformed about the weather.
- (63) If it is not raining and the body of information i doesn't rule raining out, then I'm uninformed about the weather.

This is where discourse conventions play a crucial role. The contrast between (62) and (63) is only problematic for the propositionalist account of the form just described if we assume something like the intentionalist or extra-linguistic

⁸ I am ignoring here the Kratzerian ordering source, as it will be of no importance for our purposes. On Kratzer's account, the ordering source is the contextually supplied parameter that ranks worlds relative to their proximity to the world of evaluation given some standard or ideal. We could easily factor this parameter in. See Kratzer (1983).

account of context-sensitivity resolution. More specifically, it is only challenging if we assume that, on the contextualist story, (62) should have the overall most plausible interpretation, say, one truth-conditionally equivalent with (63). And indeed, such an assumption is natural on the standard accounts of context-sensitivity resolution, which assume that the resolution of context-sensitivity is an abductive process aiming at the overall most plausible interpretation. And since there are in principle numerous domain restrictions that would make (62) receive a perfectly plausible interpretation in normal circumstances, the account is at loss to explain the apparent infelicity of (62).

However, things seem much different if the problem of context-sensitivity resolution we are faced with is not one of searching for what the overall most plausible interpretation of (62) would be, but rather simply a matter of which interpretation (62) *in fact* receives in its context. I argue that, much as in the case of pronominal resolution, where the interpretation of a discourse seems to lack reliance on extra-linguistic cues, and is instead determined by discourse conventions, so in the case of (62), too, the interpretation is ultimately resolved by discourse-internal considerations that are a reflex of discourse conventions.

(62) is not incoherent because there is no body of information that the context can select that yields a plausible interpretation; rather, it is incoherent because the context determines the body of information that delivers an inconsistent interpretation. That is, discourse conventions that affect the interpretation of modal expressions force the inconsistent interpretation in (62), but not in (63). This is what generates the contrast. Modals are not special in this regard: exactly the same type of a contrast can be recreated with discourses containing demonstrative pronouns, as reflected by the contrast between (55), repeated below and (82):

(55) John came to the party, and he didn't come.

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

While (55) is infelicitous, (82) is not. Further, note that (83), (84), and (85) are all likewise perfectly consistent. Just as with modals, it does not matter what one takes to be the character or the linguistic meaning of 'he'. That is, it does not matter which properties we think the standing linguistic meaning encodes: replacing the occurrence of 'he' with a paraphrase corresponding to whatever we take to be the associated character or the standing linguistic meaning by which the context determines the value of 'he' still produces the contrast, in the same way as replacing a modal with an explicit paraphrase referencing the description of the contextually supplied body of information encoded in the standing meaning of the modal, as in (63), does.

- (83) John came to the party, and the male person I intend to refer to didn't come.
- (84) John came to the party and the demonstrated male didn't come.
- (85) John came to the party and the actual contextually salient male didn't.

The conclusion to draw is not that somehow context fails to determine the referent for 'he,' or that (55) does not express truth-conditions: we know that 'he' is context sensitive, and that it selects some contextually prominent male. Rather, the contrast just reflects the fact that there are rules governing the resolution of context-sensitivity of a demonstrative pronoun, rules that are not bound to only deliver plausible or consistent interpretations. Indeed they determine an inconsistent interpretation in (55). And since the rules force the inconsistent interpretation by fixing a particular interpretation of the pronoun, 'he,' (83)–(85) which lack the pronoun in the same position allow for consistent interpretations.

We have already seen the relevant rules—discourse conventions—that govern the resolution of pronouns in Chapters 4 and 5. Here is how they are operative in (55): the second sentence is naturally understood as elaborating on the first. It provides further information about the event of John coming to a particular party. The discourse is thus understood as organized by Elaboration relation. But, recall, Elaboration relation has an attention-shifting effect: it makes John the most prominent candidate referent, and consequently, 'he' picks John out as a matter of its character. It is this rule affecting prominence signaled by the relation organizing the discourse that forces a particular interpretation of 'he.' And it is because this is a rule, rather than a pragmatic reflex, that the utterance remains infelicitous even though there might be potential values of 'he' which would deliver a felicitous interpretation. In this regard, (55) reflects the same conventionality of the attention-shifts associated with coherence relations that we have seen at work in (47), repeated below:

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

As was the case with (47), under appropriate circumstances, (55) can be made felicitous. For instance, (55) is felicitous if the speaker is pointing at some male other than John while uttering 'he.' This is as expected given our discussion of the attention-shifting effects of demonstrative gestures in Chapter 4. Similarly, (55) can receive a felicitous interpretation if it was uttered as a part of a larger discourse in which another male individual has been mentioned, and so introduced into the attentional state of the discourse, and if it is accompanied with contrastive intonation pattern, and a contrastive stress placed on the

pronoun, 'he.' But in this case one naturally interprets (55) not as organized by the Elaboration relation, but rather Contrast, contrasting John's arrival with someone else's. As before, the change in relation and/or an accompanying demonstration is needed in order for the pronoun to select someone other than John. And unless a different interpretation is signaled explicitly—for example, with a contrastive intonation pattern, or a demonstrative gesture—(55) receives an infelicitous, inconsistent interpretation, whereby 'he' refers to John, and it is asserted of John that he both did and didn't come. This just means that (55) is ambiguous between various different logical forms, differing in organizing coherence relations, or in the presence or absence of a demonstrative gesture. But on any given disambiguation, the pronoun is fixed by the discourse conventions.

The source of contrast between (62) and (63) is exactly analogous to that of the one between (55) and (82). First, as with utterances containing pronouns, modal utterances do not simply passively receive interpretation from context. They induce changes in the context that can affect the interpretation of subsequent discourse. And these changes not only transpire between utterances, but thread throughout the interpretation of a discourse creating effects that can result in different interpretations of different tokens of a single expression even within a single utterance. In Part I, we have seen that it is part of the meaning of an indefinite, or a pointing gesture, to bring about a certain change to prominence ranking of candidate interpretations for pronoun resolution. Similarly, it is part of the meaning of an *if*-clause, or a modal, that it raises a certain possibility to prominence, which can then be picked up as a restrictor by a modal or a conditional in the subsequent discourse. We see this in examples like (67) or (75).⁹ Such changes in contextual prominence, thus, affect the interpretation of modal expressions.

Further, as with pronouns, the prominence of candidate possibilities that can serve as restrictors for modals is likewise dictated by discourse conventions,

⁹ Several accounts of conditionals in the literature, including the dominant Kratzerian restrictor analysis maintain that antecedents of conditionals restrict the interpretation of the consequent, thus essentially making it a part of the meaning of an antecedent that it provides a restrictor for the consequent. For instance, according to the Kratzerian account, the sole function of an *if*-clause is to restrict the domain of an (overt or covert) modal expression in the consequent, while on the update semantics and expressivist accounts of conditionals, as we have seen in Chapter 8, 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 capture the full range of discourse dependencies that conditionals and modals support. As they only capture the dependence between the antecedent and the consequent of a conditional, they fail to account for modal subordination more generally. At the same time, such accounts are in a sense too constrained: they make the dependence of the consequent on the antecedent obligatory, thus making it difficult to account for possible non-restricted interpretations of modals in the consequent (Zvolenszky, 2002). For further discussion, and an attempt to account for this issue within the restrictor-analysis, by positing a covert modal operator that the *if*-clause restricts instead of the overt one, see Kratzer (2012, 1977).

specifically by mechanisms of discourse coherence. We can already see this in (67): the second sentence elaborates on the hypothetical possibility described by the first, the one in which a wolf walks in. Elaboration promotes this possibility, making it the most prominent one at that point in discourse. The modal in the second sentence just selects the most prominent possibility as the restrictor. Thus, it is restricted by the possibility just made prominent by Elaboration, and as a result we get the restricted reading: a wolf would eat you, given that one walks in. Note that the same explanation extends to the conditional in (75): the consequent further elaborates on the possibility described by the antecedent, and the Elaboration relation between the antecedent and the consequent promotes this possibility, so that as a result ‘would’ in the consequent is understood as restricted by this proposition.¹⁰

Note further that 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, as witnessed by the following example:¹¹

- (86) If a wolf walks in, it would eat you. But one probably won’t walk in. (Stojnić, 2017b)

Just as in (75), in (86), the Elaboration relation holds between the antecedent and the consequent of the first sentence. And just as in (75), it promotes the hypothetical scenario introduced by the antecedent; since ‘would’ searches for the most prominent possibility as a restrictor, it ends up being restricted by it. But the second sentence, even though it immediately follows the first one, does not further elaborate upon the scenario described by it. Rather, the two sentences stand in Contrast relation signaled by the discourse marker ‘but,’ contrasting two hypothetical scenarios—one in which a wolf walks in, and one in which a wolf does not.

Contrast relation has a rather different effect on prominence than Elaboration. It requires that the two bits of discourse connected by it provide contrasting information about some information state, regarding some common topic or

¹⁰ In this way the account predicts that, given the appropriate relation between the antecedent and the consequent, a modal in the consequent will be interpreted relative to the possibility described by the antecedent. But since the mechanisms of coherence affect the interpretation of modals generally, not only in consequents of conditionals, my account carries over to other cases of modal subordination, unlike the alternative accounts that hard-wire this dependency of the consequent on the antecedent into the semantics of a conditional. At the same time, since prominence of possibilities is linguistically determined, my account also sharply differs from flexible contextualism of Dowell (2011), which characterizes prominence as a function of complex but defeasible pragmatic factors. We will see that data favor my linguistic notion of prominence.

¹¹ See also Asher and McCready (2007) for a discussion of this point.

dimension of comparison—in our example the topic is what is possible regarding a wolf's entrance.¹² In (86), the first sentence is interpreted as providing information about whatever is the most prominent body of information at the point at which it is uttered. Assuming it was uttered discourse initially, it is understood as providing information about the set of epistemically accessible worlds determined by the context discourse initially. Consequently, the modal in the first sentence is interpreted relative to it, and is understood as describing what might be the case if a wolf walked in, *given this body of information*. Contrast requires, then, that the second sentence provides a contrasting bit of information about this same body of information. The effect of Contrast is to make this body of information prominent, and, hence, the modal in the second sentence is restricted by it. Thus, we get the interpretation that *given the same overall body of information that the first sentence is about*, a wolf probably won't come in. Putting this together, then, (86) conveys that *given the overall body of information, if a wolf walks in*, it would eat the addressee, but *given the same body of information*, one probably won't walk in. A slightly more elaborate example illustrating the same point is the following one, from Asher and McCready (2007):

- (87) John might come to the party. He might drink quite a bit. We would all have fun. But then again, he might not drink anything. And then we wouldn't have fun.

Here we have a contrast between 'He might drink quite a bit' and 'he might not drink anything.' Contrast is centered around the common topic: the possibility of John coming to the party. This possibility is, in turn, made prominent for the interpretation of the two occurrences of 'might.' Hence we get the reading that John might drink quite a bit *if he came to the party*, and that he might not drink anything, *if he came to the party*. 'We would all have fun' elaborates on the sentence preceding it, making the possibility elaborated on prominent for the interpretation of 'would': we understand that the interlocutors would all have fun *if John came and drank quite a bit*. Similarly, 'we wouldn't have fun,' elaborates on the possibility preceding it, making the possibility elaborated on prominent.

¹² The 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 (where I use capitalization to mark focal stress):

- (i) John likes MARY.
- (ii) JOHN likes Mary.

(i) 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 (ii). See Rooth (1992); Roberts (1996).

Hence, we understand that the interlocutors wouldn't have fun *if he came and abstained from drinking*.

So discourse conventions induce context-change reflected in the dynamically evolving prominence of candidate possibilities that affects the interpretation of modal expressions. In this way, they determine the proposition expressed by subsequent utterances. Since the propositional content depends on the resolution of modal anaphora—much like it depends on the resolution of pronominal anaphora—by fixing particular resolutions, discourse conventions fix the propositional content. It is only once we account for the systematic context-change effects induced by discourse conventions, that we can predict the intuitive truth-conditions. So the propositional content of an utterance—a modal utterance in particular—isn't determined by some static context; it is rather *built* by the evolving, dynamically changing context, one that accommodates potential effects of each linguistic item a discourse comprises: a modal, an antecedent of a conditional, a discourse relation.

We can capture the way in which modals depend on the context by continuing to exploit our model of context as a dynamically changing conversational record; in order to account for modal anaphora, we will include among the parameters of context a ranking of possibilities—candidate restrictors, according to their relative prominence, those higher in the ranking preferred as candidate resolutions than those lower in the ranking, where the most prominent possibility is the top-ranked one.¹³ Then, taking the truth-condition of modal utterances to be of the sort defined in Definitions 9.1 and 9.2, 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 by simply letting the modal, as a matter of its standing linguistic meaning, retrieve the most prominent proposition representing an epistemically live possibility as the restrictor for its domain of quantification.

The prominence ranking changes with the evolving discourse: new possibilities are introduced and promoted, demoting others. The dynamics prominence is governed by systematic mechanisms of context-change: by linguistic items that introduce new propositions, and by mechanisms of discourse coherence that

¹³ The idea of exploiting a ranking of possibilities to account for modal subordination is not entirely novel. See, for instance, Kaufmann (2000); Bittner (2011, 2001); Brasoveanu (2010a,b); *inter alia*. As shall become clear, the system I defend differs from these alternative ones in several theoretically important respects. First, according to the account defended here, the ranking of candidate possibilities is entirely determined through linguistic mechanisms—through the effects of discourse conventions. This allows us to make systematic predictions regarding complex discourse embeddings, such as, for instance, the ones in (87). Second, the account defended here allows us to determine propositional content contributed by each clause in the course of semantic composition. This, as I argued, is needed to account for the full range of data. Finally, as we shall see, the account I shall defend provably embeds classical (modal) logic.

promote certain propositions, while demoting others. This allows us to capture the anaphoric behavior of modals manifested in the examples above.

With this at hand, we can also explain the puzzling contrast between (62) and (63), repeated below:

- (62) If it is not raining and it might be raining, then I'm uninformed about the weather.
- (63) If it is not raining and the body of information *i* doesn't rule raining out, then I'm uninformed about the weather.

Let us begin with (62). Its antecedent embeds a conjunction. The first conjunct introduces a hypothetical scenario in which there is an absence of rain. The second conjunct, 'It might be raining,' elaborates on this scenario, describing it further.¹⁴ As before, the Elaboration relation between the two conjuncts makes the proposition elaborated upon—the proposition comprising the epistemically accessible worlds in which it is not raining—prominent. *Might* in the second conjunct selects the most prominent possibility as its restrictor; consequently, it selects the proposition just promoted by Elaboration.

Thus, the modal in (62) is understood as quantifying over the epistemically accessible worlds in which it is not raining. But this delivers an inconsistent interpretation of the antecedent: it predicts that the antecedent expresses 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 leads to a contradiction, and so the conditional is infelicitous.

What, in turn, explains the felicity of (63)? Like (62), (63) likewise embeds a conjunction in the antecedent. The first conjunct, as before, introduces a hypothetical non-raining possibility. The second conjunct still elaborates on this possibility providing more information about it, and Elaboration still makes it prominent. However, since the second conjunct does not contain a dependent

¹⁴ The second conjunct elaborates on the first one just as, for example, in (i) the second conjunct elaborates on a hypothetical scenario in which there are wolves around the house, or in (ii) on one in which the battery is low:

- (i) If there are wolves around our house, and one might walk in, then we are in trouble.
- (ii) If the battery is low, and the phone might stop working, we should charge it.

In (ii), a mere possibility of the phone dying isn't sufficient for recommending the action advocated in the consequent. Indeed, we can even imagine that if the phone might stop working for reasons other than the battery running low, charging would be detrimental (perhaps the phone can stop working if it has been overcharged, in which case charging further would just damage it). The relevant interpretation is thus one according to which the second conjunct in the antecedent is understood as Elaborating on (and so the modal it contains as restricted by), the possibility described by the first one, namely, that the phone's battery is low.

expression anymore—a modal expression requiring the most prominent possibility as the restrictor—we end up with a perfectly consistent reading: the antecedent in (63) expresses the proposition corresponding to the most prominent epistemically live possibility in which *it is not raining, and moreover the body i doesn't rule out raining*. So, we have an explanation of the contrast between the two examples.¹⁵

¹⁵ But what happens if we reverse the order of conjuncts, as in (i)?

- (i) If it might be raining and it's not raining, then we are uninformed about the weather.

My account does not predict that all conditionals of the form *If it might ϕ , and not ϕ , then...* will automatically be bad. This is a good result: as Sorensen (2009) and Dorr and Hawthorne (2014) note, reversing the order of the conjuncts in some cases makes the conditional felicitous. Dorr and Hawthorne (2014) illustrate the point with the contrast between the following:

- (ii) If Jack and Jill fail and they might pass, the preparation is to blame.

- (iii) If Jack and Jill might pass and they fail, then their preparation is to blame.

Second, notice that such examples can be made felicitous more easily with an appropriately set context that provides the relevant body of information against which the modal is to be interpreted and additional cues about the proper disambiguation of the discourse structure. Suppose a marble is selected at random out of an urn with a hundred marbles. I don't know whether any red marbles were included, but I know that whenever some red ones are included then ninety-eight of them are red. I can say:

- (iv) I don't know if it might be red or not. But if it might be, and it's not, I'd be surprised.

Notice also that, while arguably *if ϕ and ψ then γ* is equivalent with *if ϕ then if ψ then γ* , creating a distance between the conjuncts of *might ϕ and not ϕ* , for instance by embedding the second conjunct in the postposed antecedent of the conditional embedded in the consequent, helps as well. So, in the same scenario, we can say:

- (v) I don't know if it might be red or not. But if it might be red, then I'd be surprised if it's not.

- (vi) I don't know if it might be red or not. But if it might be red, then you'd be justified in guessing it's red even if it is not.

Notice, in these examples the speaker signals she doesn't know whether the *might* claim is true or not—the *might* claim is interpreted as commenting on a situation that makes some relevant information available (e.g., about which colors are represented), but that the speaker has incomplete access to. (I'll return to the question of how modals and conditionals are interpreted discourse initially in some more detail below.) What if the speaker endorses the *might* claim? Suppose I found out that there are indeed red marbles in the urn. Then, I could say, commenting on this new bit of information:

- (vii) The marble might be red. But if that's true and yet the marble isn't red, I'll be surprised.

- (viii) The marble might be red. But if so, I'll be surprised if it isn't.

So, the examples improve once the context makes available the relevant body of information and once there are enough resources to establish a coherence relation between the conjuncts (here, the Violation of Expectation relation, another instance of the cause-effect or event-result type relation). But my account does not predict that all cases in which the order of the conjuncts is reversed will automatically be felicitous either. This is again a good result. Whether such conditionals are felicitous or not depends on which coherence relations hold between the two conjuncts in the antecedent, and between the conditional and the prior discourse. All these factors affect the resolution of modal anaphora. Specifically, when the modal occurs in the first conjunct, we have to determine which body of information it is interpreted relative to—which information the conditional is commenting on. Further, we need to establish the relation between the conjuncts that actually makes the discourse coherent. (Notice that the subtle variation in the structure that helps prompt Violation of Expectation interpretation in the examples above helps improve the discourse.) Indeed, with respect to examples

Note, it is crucial for our explanation that the effect of a coherence relation—here Elaboration—is understood as underwritten by linguistic convention, and not as a mere pragmatic default that's operative in trying to recover the most plausible interpretation. In (62) we are precisely not retrieving the most plausible interpretation, nor can such an interpretation be easily recovered by overriding the effect of Elaboration.

If 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 that of the speaker). Indeed, that there is a contrast between (62) and (63) is utterly surprising on a pragmatic account. Clearly, there is at least one possible, plausible interpretation of (62), namely (63), that is available. Indeed, any of (64)–(66), repeated below could be relevant, charitable, plausible interpretations in most typical circumstances:

- (64) If it is not raining and for all I (the speaker) know it is, then I'm uninformed about the weather.
- (65) If it is not raining and for all our group knows it is, then we are uninformed about the weather.
- (66) If it is not raining and given the salient body of information it is, then that body of information is uninformed.¹⁶

If we were guided by general pragmatic considerations that aim at recovering the overall most plausible interpretation, any of these interpretations would be more suited as candidate interpretations of (62). Consider the pragmatic constraints of relevance and plausibility. They would clearly favor any of these more plausible interpretations over the one actually retrieved: by contrast with the incoherent interpretation actually retrieved, they all provide a plausible scenario that the speaker can be hypothetically considering and be plausibly reasoning about. Or consider the constraint of charity. Charity also clearly favors

like the ones above, my account has an advantage over the views that hard-wire the preference for either a strong order effect, or no order effect whatsoever into the semantics. Notice that accounts that hard-wire order insensitivity into the semantics, or into the dynamics of interpretation more generally, have to treat these examples as exceptional, or argue that such order-insensitivity is at best a default which must be overridden in these examples, so as to allow an alternative body of information relevant for the interpretation of the modal to be selected. This selection of the relevant body of information on my account is just a matter of modal anaphora resolution, as in all other cases.

¹⁶ (64) and (65) are particularly good candidates, since epistemic modals are most often (though, as the present account makes clear, not always) interpreted as quantifying over bodies of information that correspond to, or at least include, the information of the speaker.

any of the alternative interpretations over the one actually retrieved: they are all plausible and potentially true, while the actually retrieved one is inconsistent, requiring that we attribute some inconsistency or confusion to the speaker. So, from the standpoint of pragmatic considerations it is mysterious indeed why we are stuck with the interpretation we are actually stuck with. But this robust, persistent infelicity that remains even in the presence of more plausible alternative interpretations is precisely what we would expect if such infelicity were forced as a matter of linguistic constraints on the available interpretation: specifically, if it were a matter of linguistically encoded rule that Elaboration promotes the possibility elaborated on, in this case, the possibility *that it is not raining*, and that modals look for the most prominent possibility as their restrictor, the inconsistent truth-condition would be calculated as a matter of grammar. That it would be the one recovered, even if infelicitous, is precisely predicted.

Thus, I argue, the source of the infelicity of (62) is the same as the one of (47): it lies in the effects of discourse conventions. The reason why the sentence is infelicitous is that a rule that governs its interpretation forces an infelicitous interpretation. So, rather than the context failing to determine a truth-condition for (62), it determines an infelicitous one. It was the presumption of the non-linguistic model of context, and the idea that context-sensitivity resolution aims to deliver the overall most plausible interpretation, that led to the conclusion that the infelicity of (62) is a sign of non-propositionality.

Further data support this conclusion. For instance, consider the contrast between (88) and (89):

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

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

While (88) seems odd, (89) is perfectly fine. The only difference between the two is in the choice of the relation organizing the discourse: (89) is organized around the Contrast relation, signaled by 'but,' while in (88) we naturally expect Elaboration. Contrast promotes the body of information that the two bits of discourse provide contrasting information about, given some common topic—here, regarding the likelihood of John's attendance. This yields the intuitive interpretation of (89)—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, Elaboration would deliver an infelicitous interpretation for (88)—it is compatible with the information available discourse initially that John will come to the party, but (given this same information) he will

probably not come, *given that he comes*.¹⁷ If this effect of Elaboration were a mere byproduct of pragmatic reasoning, there would be no reason, in principle, why 'won't' in (88) couldn't be interpreted as restricted only by what is known discourse initially (as it is in (89)). But, even if one could re-interpret (88) in this way when faced with inconsistency, one could no longer understand (88) as organized by the Elaboration relation—one cannot freely change the resolution of modal anaphora while holding the relation fixed. As in the case of pronouns, a disambiguation of a relation settles the interpretation of a subsequent modal. This again suggests that the effect of the relation is linguistically encoded.¹⁸

That the effects of discourse coherence resist re-interpretation even in the face of infelicity is further illustrated by the examples like the following:

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

'Would' in the final sentence cannot be understood 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 fixed 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. Worse yet for the pragmatic account, (91), which differs from (90) only in the order in which the sentences are uttered, is perfectly felicitous:

¹⁷ Note that, by contrast, the following has a perfectly felicitous interpretation:

- (i) John might come to the party and he might not come.

The reason is that (i) allows for an interpretation where the discourse is organized around Parallel relation: there's a possibility that John will show up, and similarly, there is a possibility that he will not. (88), by contrast, is hard to interpret this way: the possibility of John's attendance, and a low probability of this event, given our information, are not appropriately presented as similar to support Parallel.

¹⁸ One might then wonder, shouldn't changing 'and' to 'but' improve the problematic example (62)? That is, shouldn't 'If it is not raining but it might be raining, then we are uninformed about the weather' be felicitous? Insofar as one succeeds in embedding Contrast within the antecedent of a conditional, the example indeed improves. However, 'but' does not automatically guarantee that the relation between the two conjuncts is Contrast. 'But' can sometimes signal Violation of Expectation relation, (recall footnote 15 of this chapter), which has the same effect on prominence as Elaboration, as witnessed by examples such as 'A wolf might walk in, but it will not eat you.' Note, moreover, that Contrast is indeed somewhat hard to embed in the antecedent of a conditional, as witnessed by the oddity of the following:

- (i) If it might be raining but it might not be raining, then it's not guaranteed that it will rain.

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

In (90) and (91), exactly the same possibilities have been mentioned. It is perfectly clear, insofar as we are reasoning about a rational speaker's intentions, that the speaker can only describe the scenario in which Harvey has been eaten by a wolf if there is a wolf in that scenario, not if the scenario describes an absence of a wolf. A reasonable speaker clearly knows this, and if she utters (90), then if the domain restriction is determined by speaker's intentions, it's reasonable to suppose that the speaker intended the possibility introduced by the first sentence, not the second, as the restriction for the modal. More generally, if we are looking for the overall most plausible interpretation, (90) should mean the same as (91). So, while the linguistic account of context-sensitivity resolution predicts the contrast between (90) and (91) in a straightforward manner, from the standpoint of the non-linguistic account, the contrast is puzzling. It seems, again, that the possible resolutions of a modal anaphora are much more constrained than what one would expect if they were determined by pragmatic principles.

One could perhaps try to explain the contrast between (90) and (91) by invoking a certain kind of pragmatic principle, for example, requiring one shouldn't move from one sub-task to another until the first is complete (cf. Clark, 1996). So we could say, in (90), the speaker is in violation of such a principle, because she introduces a new possibility while still in the middle of describing the previous, contrasting one. Such a principle, while not explained purely by rationality, still describes contingent regularities about human cooperation rather than regularities about language *per se*. However, *even if* we could somehow classify different possibilities as pertaining to different sub-tasks in (90), it is clear that no such explanation of (62) is available. So overall, the linguistic account seems to fit the data better.

Further support for the linguistic approach comes from the behavior of epistemic contradiction under a large number of embeddings. We have seen that the construction of the form *not φ and might φ* is infelicitous not only unembedded, but also once embedded in the antecedents of conditionals, as in (62) or under supposition operators. But the behavior is not merely isolated to such embeddings, but is indeed quite general. For instance, it persists under disjunctions, attitude reports, and quantifiers.¹⁹ Such persistence is expected on the linguistic account of the sort I defended here, as what is generating the infelicity on my account is simply the relation between the two conjuncts, which

¹⁹ For a range of data supporting this claim, see, for example, Dorr and Hawthorne (2014); Mandelkern (2019).

is internal to the construction itself and largely insensitive to the embedding environment.²⁰ That is, on my account, the modal is interpreted relative to its linguistic context as set up by its prior surrounding discourse. Since the relevant features of the linguistic context that determine the resolution of the modal in the second conjunct of the epistemic contradiction are fixed by the linguistically encoded relation between the two conjuncts, the inconsistent interpretation will be predicted regardless of how the conjunction is embedded.

In addition, as in the case of pronominal anaphora resolution, there is also cross-linguistic variation in how languages exploit coherence to signal shifts in prominence that affect modal anaphora. For example, Asher and McCready (2006) point out that the direct truth-conditional translation of (67) in Japanese is infelicitous; however, if an overt discourse marker is inserted signaling an Elaboration relation, the discourse becomes felicitous.²¹ This, again, would be surprising if the interpretation according to which the second modal elaborates on the scenario described by the first one were merely a by-product of general rationality.²²

I conclude that the contextual effects of discourse coherence on the interpretation of modal expressions are genuine discourse conventions. They are specified by the rules of grammar, much like, I have argued, the effects of discourse conventions on pronoun resolution are specified as a part of grammar. Not only do we need to take into account mechanisms of context-change in order to derive the truth-conditions for modal utterances, but moreover, language

²⁰ Provided, of course, that the environment allows for the embedding of the relevant relation. But Elaboration can, in fact, be easily embedded across the full range of the relevant embeddings discussed in, for example, Dorr and Hawthorne (2014) and Mandelkern (2019), and indeed quite generally:

- (i) Suppose that John is at home and that he's drinking tea.
- (ii) Mary believes that John is at home and that he's drinking tea.
- (iii) Either John is at home and he's drinking tea, or he is in his office.
- (iv) Someone who is at home and is drinking tea should be in their office.

So, on my account, the explanation for the infelicity of the epistemic contradiction in such embeddings exactly parallels the explanation of its infelicity in the antecedents of conditionals.

²¹ Asher 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.

²² Just as the effects of coherence on prominence seem 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 preajacent prominent for subsequent modal anaphora, much like it is a part of the meaning of an indefinite or a definite noun phrase to set up a candidate for the resolution of pronominal anaphora. Note that a proposition is made prominent regardless of whether the modal is embedded, say under a negation, as in (ii) and (iii), or in a question, as in (i), or an antecedent of a conditional, as it is in (62):

encodes the effects of these mechanisms along (and beyond) ordinary truth-conditional content.

If this is right, the pessimism toward truth-conditional accounts rests on an overly simplistic conception of the context-sensitivity of modal expressions, and more importantly, it relies on an overly simplistic idea of how context interacts with content. The traditional assumption of the extra-linguistic account of context is what drives the arguments against propositionalism. The linguistic account of context, describing the linguistic mechanisms by which the context changes both within and between utterances precisely predicts the strange behavior, and does so while delivering ordinary propositional content for modal utterances.

While my model and non-propositionalist ones all predict the contrast between (62) and (63), my model allows for a more general account of modal subordination and, I have argued, captures data that are *prima facie* problematic for the non-propositionalists. This is a significant advantage. Non-propositionalists, we have seen, argue that modal utterances do not contribute propositional content, and that their meaning is exhausted by their context-change potential—the effect that an utterance brings about on the conversational context, understood in Stalnakerian terms. On my account, utterances likewise bring about a change in conversational context, understood as a conversational scoreboard. Indeed, these effects are a part of the meaning of the expressions an utterance comprises, much like it was a part of the meaning of, say, an indefinite, to induce a particular attention-shift. In Part I, we have seen that once we ascribed a dynamic layer of content to expressions an utterance comprises, specifying their grammatical effects on the context, we were able to use it to determine the proposition the utterance expresses. Extending this

- (i) Might a wolf walk in? It would probably eat Harvey.
- (ii) No wolf will come in. It would have to get through the security gate.
- (iii) 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 effects on prominence were a matter of a pragmatic import of a modal utterance, it would be surprising that they persist even under embeddings, and even when the possibility that is made prominent is questioned (as in (i)), or explicitly said not to obtain (as in (ii) or (iii)). Moreover, 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. We can still access the far-fetched possibility even once it's been deemed unlikely (as in (iv)), and we cannot simply ignore it (as witnessed by (v), where one naturally expects the second sentence to Elaborate on the preceding conditional):

- (iv) John is late. He might be abducted by the aliens. But that's very unlikely: there are probably no aliens around. Besides, they would have to sedate him somehow.
- (v) (We are waiting for John who should arrive any minute, and I say:) If John is abducted by the aliens, then he might be transported to a distant planet now. He will arrive on time.

idea to modal discourse, we can represent an utterance as endowed with two layers of meaning: a dynamic meaning, which induces a change in context, and a propositional content, which is built as a result of that change.

To be clear, I am not arguing that non-propositionalist in principle could not modify their semantics to account for complex patterns of modal subordination we have seen so far. A non-propositionalist account could in principle adopt a system that dynamically keeps track of the bodies of information, or information states, or local contexts relative to which modals are interpreted, in a way that allows for cross-discourse dependencies.²³ However, if the argument of this chapter is on the right track, to account for the full range of data one has to take into account a full range of mechanisms that affect the interpretation of modal expressions, including the mechanisms of discourse coherence. But if this is right, then the dynamics of context-change affecting the relevant changes in stacks of information states (or contexts) across a discourse will have to closely mimic the system I have defended here, and shall lay out in formal detail in what follows. Indeed, such a system has to allow for referential uses of modal expressions, including (69), but also the exocentric uses where the relevant body of information excludes some of the speakers' information.²⁴ It would also have to allow for anaphoric dependencies not only on non-modal, but also on modal contents (as in, e.g., (77) and (78)). These uses pose a *prima facie* challenge to non-propositionalist accounts. If the non-propositionalist system is modified sufficiently to allow for all such readings, then there might not be anything essentially non-propositional left in the resulting system.

Meanwhile, the propositional account I develop allows for a uniform treatment of all such uses of modal expressions. Further, as we shall see in III, the propositionalist account I develop also allows for a more satisfactory individuation of argument patterns in natural language discourse. Specifically, my account can trace the failures of certain argument patterns back to equivocations on the information expressed by the premises and the conclusion of the relevant argument patterns. By contrast, an account that does not assign informational content cannot account for such failures by appealing to equivocation.²⁵

²³ For instance, see Kaufmann (2000), who likewise employs a stack-based system with the elements of a stack being information states or contexts.

²⁴ See Egan, Hawthorne, and Weatherson (2005) and von Stechow and Gillies (2009) for related examples of exocentric uses of modal expressions. Related data concerning agreement and disagreement involving modal utterances are sometimes argued to pose challenge for contextualist accounts. But my account naturally captures such data as well, since it allows for deictic uses of modal expressions, and recognizes the possibility of ignorance and ambiguity in the interpretation of such deictic uses. I will discuss this issue in more detail in Part IV.

²⁵ There is another line of anti-propositionalist argument that I have not discussed so far. It is sometimes argued that data concerning iterated modality or higher-order uncertainty are

Finally, setting aside the empirical advantages, I pause to make a dialectical point. By having compositional semantic content fully determine propositional assertoric content, my account vindicates a tight connection between assertoric content and semantic content and, by extension, the traditional account of the interaction between belief, communication, and action. Now, as mentioned above, this is not in itself an argument against non-propositionalists, or in favor of my view. But, given the empirical advantages I mentioned, the dialectical burden is on those who advocate the need for departure from the traditional picture, as we no longer have a purely linguistic argument for non-propositionalism.²⁶

I end this chapter by briefly outlining the details of the process whereby such propositional content is built by the dynamic meaning on my account, which I will then spell out in more detail in Part III.

Building on the model of context as conversational scoreboard, as described in Part I, as before, we represent a scoreboard as an abstract record of linguistic information about the discourse, reflecting which moves in the conversation have been made. Earlier we have exploited an attentional state of the discourse as a contextual parameter governing the interpretation of demonstrative pronouns. Now, to capture prominence-sensitivity of modal expressions, our record will have to keep track of ranking of possibilities that serve as candidates for restrictors on the domain of quantification of modal expressions. We will thus include this ranking of possibilities that are candidate restrictors for the domains of quantification of modal expressions as a contextual parameter within our scoreboard. For simplicity, since we are only dealing with modal expressions in the present chapter, we can focus on this single aspect of context. So, abstracting away from other parameters of the scoreboard (including the ranking of indi-

problematic for the propositionalist accounts of meaning. The relevant observation is that, sometimes, epistemic modal utterances can non-vacuously iterate (see Moss, 2015, 2018; c.f. Yalcin, 2007):

- (i) It is probably the case that Jones might be the next hire.

But my account allows non-trivial iterations by design. Each modal is interpreted relative to some prominent possibility (and, further, possibly distinct possibilities can be prominent for distinct modals), not merely interpreted relative to the agent's information states. Suppose that it's likely that research will be an important search criterion, and that, if it is, Jones would be a live candidate (but, otherwise, they would be ruled out). With this background, commenting on the—still unknown—facts about the search criterion, it is natural to say "I don't know whether Jones might be the next hire, but it's probable that they might be." On our present proposal, we interpret this as conveying that it is probable (given the relevant information about possible search criteria), that Jones might be the next hire *given the facts about the search criteria*. This is intuitively right. Moreover, note that the relevant pre-jacent proposition is something that the agent can have non-extremal credence toward. The interlocutors can lack knowledge regarding what the prominent body of information is, and what is compatible with it; and the prominent body of information that is relevant for the interpretation of the modal need not be the agents' information state. Thus, my version of propositionalism avoids the type of worry raised in Charlow (2020).

²⁶ As explained earlier, I am not denying that there might be theoretical reasons for non-propositionalism, but merely that putative linguistic arguments for non-propositionalism fail.

viduals tracked by the attentional state of the discourse), let the context be a ranking of propositions by their prominence, those higher in the ranking being more prominent than those lower.²⁷ Then to capture the fact that utterances can bring about changes to the ranking that affect the truth-conditions expressed, we again interpret them not simply as expressing a truth-conditional, propositional content, but as updates to the conversational scoreboard, that is, as instructions on how to change the state of the scoreboard. Specifically, they will now encode instructions on how to change the prominence ranking of candidate possibilities.

So, an utterance of *might* ϕ is not merely associated with truth-conditional content, specified in Definition 9.1; 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 (67). Suppose that (67) is uttered out of the blue, and that the relevant body of information that it is about is some set of epistemically accessible worlds, say, 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 expressing a propositional content, a modal claim also induces a change in the 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.²⁸

²⁷ Of course, a more complete representation of context would simultaneously model different contextual parameters relevant for the interpretation of distinct context-sensitive items. We have seen, already, in Chapter 7, how we can integrate parameters relevant for the interpretation of prominence-*insensitive*, context-sensitive expressions with the attentional state of the discourse. There are also various ways one can simultaneously model prominence ranking of individuals and possibilities. One would be to keep track of different stacks, tracking different types of entities, and affecting the interpretation of different expression in the discourse. Another would be to keep track of different types of entities within a single ranking. Whichever way we chose, we will have to keep track of the interactions between modal and pronominal anaphora. See Bittner (2001) and Brasoveanu (2010a) for a related semantic framework that allows simultaneous representation of modal and pronominal anaphora.

²⁸ Plausibly, 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, for example, by following up with ‘That is true,’ or a modal anaphora: ‘You will need an umbrella.’ So, more precisely, the modal claim in (67) 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 a live epistemic possibility*. 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 (67) elaborates on the epistemic possibility of a wolf walking in, introduced by the first, as a result of the

The mechanisms of discourse coherence also affect context by changing the prominence ranking. For instance, in (67), the operative relation is Elaboration, which promotes the possibility elaborated on—the epistemic possibility of a wolf walking in—demoting all others one position in the ranking. Thus, much as we did in Part I, 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 (67) offers an elaboration of the epistemic possibility of a wolf walking in). 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 layers of content. One is characterized by the updates which structure and re-arrange the context and contextual parameters as the discourse evolves. This layer 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 they 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, for example, by making a proposition prominent for subsequent modal anaphora. This context-change, in turn, affects the truth-conditions. Given the anaphoric behavior of modals, the proposition expressed by a modal utterance depends on which propositions were made prominent by the prior discourse. And this propositional content can then itself be promoted on the record via an update associated with the utterance expressing it, thus making it available for subsequent anaphora. Now we see that the two types of content are interrelated: an utterance depends on the context to determine the proposition it expresses, but it also dynamically changes the context, thus directly affecting the propositional content of the subsequent discourse. Far from being incompatible with the propositional content, this dynamic meaning is precisely what allows us to systematically compute it.

So, much like the non-propositionalist accounts, my account maintains 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

Elaboration relation, this proposition is again promoted, and, thus, the modal in the second sentence selects it as a restrictor. We will see a more formal rendition of this explanation in Chapter 11.

naturally falls out. Indeed, on my view, we can see the dynamic meaning as an instruction, a recipe, for building propositional content. Thus, though dynamic, my account is in no way in competition with traditional truth-conditional semantics; rather, the dynamic meaning is a crucial element in determining propositional content.

PART III

THE LOGIC OF PROMINENCE

Content, Context, and Logic

Modal discourse is particularly problematic from at least one further perspective. Namely, it seems to violate some of the most entrenched classical rules of inference.¹ Many theorists have come to the conclusion that this reveals a deep tension between the behavior of modal vocabulary and classical logic, with the moral being to revise the semantics for modality in order to invalidate the problematic classical patterns. This is taken as further fuel to non-propositionalism, since non-propositionalist accounts typically invalidate some of these patterns.

I, argue, instead, that this reaction, too, stems from the implicit assumption of the non-linguistic account of context-sensitivity, as well as a simplistic account of validity for inference patterns presented in natural language discourse. As we shall see, the account of modality sketched in Chapter 9 has resources to accommodate the apparent counterexamples, while preserving classical logic.²

Let us then start by looking at some of the problematic cases. I begin with the following counterexample to the pattern *If ϕ , ψ ; $\neg\psi$: $\neg\phi$* , known as *modus tollens* (MT) (Yalcin, 2012b):

Take an urn with a hundred marbles. Ten of them are big and blue, thirty big and red, fifty small and blue, and ten are small and red. One marble is randomly selected and hidden (you do not know which).

Given this setup, (92) and (93) are licensed:

(92) If the marble is big, then it is likely red.

¹ See, for instance, McGee (1985) and Willer (2014) for counterexamples to *modus ponens* (MP), and Veltman (1985) and Yalcin (2012b) for counterexamples to *modus tollens* (MT). See also Kolodny and MacFarlane (2010) who reject MP and MT, 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-propositionalist. An exception to this is McGee (1985), who rejects MP, but offers a framework which, while non-classical, is nevertheless not of non-propositionalist variety. However, his account, too, as shall become clear, rests on a failure to account for discourse conventions.

² To be precise, the underlying logic is a classical modal logic—a simple extension of classical logic with a modal necessity operator (precisely, given the specification of the accessibility relation I will adopt, the system is the modal logic S_4). This logic preserves a classical inference system, and is sound and complete.

(93) The marble is not likely red.

But, surely, (94) does not follow:

(94) So, the marble is not big.

Since the inference from (92) and (93) to (94) looks to be an application of MT, we seem to have a counterexample.³

It is important to note that the non-propositionalists do not deny the obvious fact that, intuitively, the big premise in (92)–(94) is (in some sense) about conditional probability, that is, 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 concedes this. Indeed, given the non-propositionalist account, (of either dynamic or expressivist flavor) the conditional in (92) will be true, relative to an input context *c*, just in case updating *c* hypothetically with the antecedent lands you in a context in which it’s accepted that the marble is likely red. So, they do capture the idea that the conditional in (92) somehow concerns a restricted probability. What the non-propositionalists 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. In the light of our discussion in Chapter 8, it is clear that on such accounts neither (92) nor (93) express any truth-conditional content to begin with.

Yalcin’s own reaction to the counterexample, and one endorsed by non-propositionalists more generally, is to invalidate MT.⁴ Now, (92)–(94) is a terrible argument, so this might be considered a welcome result. After all, we certainly wouldn’t want to say that *that* argument is valid.

Yet, it is important to note that the issue cannot be localized to one single inference pattern, MT: it’s easy to construct 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 a hundred marbles, etc. But then:

(95) Suppose that the marble is big.

(96) Then it is likely red.

³ It 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.

⁴ See, for example, Yalcin (2007, 2012b) and Moss (2015) for expressivist accounts which invalidate MT; and Gillies (2010, 2004) for a dynamic update semantics one. Relativist semantics, of the form developed in, for example, Kolodny and MacFarlane (2010) likewise invalidate MT.

(97) But the marble is not likely red.

(98) So, the marble is not big.

(95)–(98) is also horrible. If (92)–(94) provides grounds for rejecting MT, then by parity (95)–(98) provides grounds for rejecting either MP or *reductio*. And before we are quick to blame MP, here is a slightly modified scenario that can be seen as a counterexample to either *reductio* or disjunctive syllogism (*modus tollendo ponens*):

There's an urn with a hundred marbles. We know that eighty of them are blue, eighteen are red, and two are green. One marble is randomly selected and hidden (you do not know which).

Then, the premises (99)–(102) are all licensed, but the conclusion, (103), doesn't follow.

(99) The marble is either blue or it is likely red.

(100) Suppose that the marble is not blue.

(101) Then, it is likely red.

(102) But the marble is not likely red.

(103) So, it is blue.

So, then the question is, should we invalidate all of these rules? Or can we isolate once culprit that's common to all?

To make things worse, consider the following famous counterexample to *modus ponens* (MP), due to McGee (1985):

Opinion polls taken just before the 1980 election showed the Republican Ronald Reagan decisively ahead of the Democrat Jimmy Carter, with the other Republican in the race, John Anderson, a distant third. Those apprised of the poll results believed (104) and (105), with good reason. Yet they did not have reason to believe (106).

(104) If a Republican wins the election, then if it's not Reagan who wins it will be Anderson.

(105) A Republican will win the election.

(106) If it's not Reagan who wins, it will be Anderson.

As McGee (1985) argues, (104)–(106) seems pretty bad. While we would gladly subscribe to the premises, we wouldn't want to subscribe to the conclusion. After all, if it's not Reagan who will win, the winner will be Carter, the close second.

Yet, by contrast to their reaction to the counterexample to MT, most non-propositionalists validate MP.⁵ Both types of non-propositionalist accounts described in Chapter 8 do. Take any context *c*. If the conditional in (104) is true (accepted) in *c*, then hypothetically updating *c* with the antecedent, which just is (105), leads to a context in which the conditional in the consequent, which just is (106), is true (accepted). If *c* accepts (105), then all the worlds in *c* are ones in which a Republican wins. So, if both premises are accepted in *c*, the conclusion has to be accepted in *c*.⁶

Now, non-propositionalists argue that there's a reason for invalidating MT in light of counterexamples like (92)–(94), but not MP in light of counterexamples like (104)–(106). Specifically, Gillies (2004) argues that while valid, (104)–(106) strikes us as invalid because we tend to revise our commitments to the second premise once we reach the conclusion. So long as we remain committed to (104) and (105), we are thereby automatically committed to (106) too. It is only when we allow the second premise to slip out of sight that we are inclined to reject the conclusion. Meanwhile, MT is genuinely invalid, and (92)–(94) is a genuine counterexample to it.

But given our previous discussion, I want to suggest that the phenomenon has once again been misdiagnosed. The reason why we can seemingly generate counterexamples to classically deductively valid patterns of inference in the presence of modal vocabulary is not that modals somehow have deviant semantics. It lies, once again, in the failure to appreciate the import of discourse conventions on the resolution of context-sensitivity of modal expressions. Capturing the import of these mechanisms in the logical form of an argument, we shall see, is crucial for understanding validity, just as it was crucial, recall, to capture the effects of discourse conventions in the logical form of arguments involving pronouns. We saw earlier that it was only once we made explicit the effect of deictic gestures in the logical form of arguments like (27) and (28), repeated below, that we could explain why the first is valid, and valid by virtue of form, while the second one is not.

⁵ Kolodny and MacFarlane (2010) are an exception. On their semantics, both MP and MT are invalidated.

⁶ MP is also valid, assuming the update semantics in Chapter 8 and the so called update-to-test notion of entailment, according to which ϕ entails ψ just in case, for any context *c*, updating *c* with ϕ leads to a context in which ψ is accepted. Any context *c* updated sequentially with (104) and (105) indeed lands us in a context in which (106) is accepted. For formal details, see Gillies (2004). I will return to this point in more detail in Chapter 11.

- (27) A woman came in. She sat down. Therefore, some woman sat down.
- (28) A woman came in. She [pointing at Betty, the cat] sat down. Therefore some woman sat down.

Indeed, a moment's reflection shows that it is fairly easy to generate counterexamples to classically valid patterns of inference involving pronouns as well. For instance consider the following counterexample to MT:

- (107) If Jane is at the party, then she is having fun.
- (108) She (pointing at Betty) is not having fun.
- (109) So, Jane is not at the party.

Similarly, here is a counterexample to MP, where the pronoun in the conclusion is uttered without any pointing gestures, and without the stress on the pronoun.⁷

- (110) If Jane is at the party, then she [pointing at Betty] is having fun.
- (111) Jane is at the party, so she is having fun.

These are obviously bad arguments. And, of course, no one would invoke them seriously as counterexamples to MP or MT. But why is this? After all, going by the surface sentential form, (107)–(109) fits MT, and (110)–(111) fits MP. Now, the intuitive, simple answer is, of course, that various different occurrences of 'she' have a different, non-uniform interpretation; we know this because we can point to reasons why 'she' is resolved as it is. For instance, we know that (108) does not *really* contradict the consequent of (107), because in (107), the pronoun 'she' occurs in an elaboration of a hypothetical scenario about Jane, and thus, is resolved to Jane, whereas in (108), the pronoun occurs in tandem with a demonstrative gesture—a pointing at Betty—and thus is resolved to Betty. The Elaboration in (107) and the pointing gesture in (108) render a certain referent the most prominent for the subsequent pronoun; the pronoun, as a matter of its meaning, just picks out the most prominent referent.

⁷ Alternatively, here's another counterexample:

- (i) If Jane is at the party, then she is having fun.
- (ii) Jane is at the party.
- (iii) So she [pointing at Betty], is having fun.

Now, how exactly does this difference in interpretation connect to the idea that (107)–(109) isn't a genuine instance of MT? The standard, traditional way of explaining away such pseudo-counterexamples is by indexing on a pronoun. But we have already seen, in Chapter 4, reasons against this strategy. If we want to capture the idea that (107)–(109) does not have a valid form, we should represent the effects of deictic gestures, and of coherence relations, in the logical form of the argument.

So, given the account in Chapters 4 and 5, we would give the following explanation of (107)–(109). The consequent in (107) stands in an Elaboration relation with the antecedent; it is *because* 'she' in (107) occurs in an elaboration of the hypothetical scenario described by the antecedent that the pronoun refers to Jane. And it is *because* in (108) the pronoun occurs in a tandem with a pointing gesture, that it refers to the individual pointed at, namely, Betty. The shift in prominence, reflected in the attentional state of the discourse, renders one or the other referent prominent, affecting the resolution of the subsequent occurrence of the pronoun. As a result of the effects of these mechanisms, (107)–(109) does not express the propositional pattern corresponding to the MT inference.

It is important to remember 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 (107)–(109) are resolved the way they are. But these contextual changes are induced by linguistic mechanisms, reflected in the logical form of a discourse. And *this* is why we can say that (107)–(109) is not associated with a valid logical form.

Had we failed to appreciate how these mechanisms affect the resolution of pronouns within a discourse, we would be misled to interpret examples such as (107)–(109) as "counterexamples." I want to suggest the same goes for the arguments involving modal language. We have seen in Chapter 9 that modals are analogous to pronouns in two crucial respects. First, like pronouns, they are prominence-sensitive: their interpretation is sensitive to contextual prominence of possibilities that can serve as a restrictor on their domains of quantification. And, second, like the resolution of pronouns, the interpretation of modals is guided by discourse conventions, in particular, by mechanisms of discourse coherence. In short, modals are prominence-sensitive, and the relevant notion of prominence guiding their interpretation is responsive to the same mechanisms that pronoun resolution is responsive to. These mechanisms govern the interpretation of both types of expressions, and do so as a matter of linguistic convention. With this in mind, we have all the ingredients we need to provide an explanation of the problematic counterexamples we have encountered.

Let us start with (92)–(94). We can explain what is going on with (92)–(94) as follows: in (92), 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), providing more information about what would be the case if the marble is in fact big. The Elaboration relation renders this possibility the most prominent one for subsequent anaphora. Consequently, the modal 'likely' in the consequent 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 (93) as connected to (92) by Contrast. This is seen even more clearly if we insert an explicit discourse marker 'but' in (93): 'But the marble is not likely red.' Note that some such way of signaling contrast is required for the discourse consisting of (92) followed by (93) to be felicitous. This contrastive interpretation is natural because the initial context sets up a topic—the color of the marble, depending on a certain assumption about its size. (92) and (93) are then understood as providing contrasting bits of information about some prominent body of information regarding this topic. The first sentence already sets the stage in determining the body of information the contrast has to be about—since (92) is uttered discourse initially, it is interpreted relative to the overall body of knowledge available discourse initially, which includes the information about the urn and the marbles, described in the scenario. Thus, one understands (92) and (93) 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. Consequently, Contrast promotes this body of information as the most prominent for the interpretation of the modal in the small premise. So we get the interpretation that 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). Thus, we see that the two occurrences of 'likely' in (92) and (93) are interpreted differently, much like the two occurrences of 'she' are in (107) and (108). So, (92)–(94) is no more an instance of MT than (107)–(109) is.

There are a couple of things we should note here. First, notice that, on the present account, the informational, propositional content expressed by the discourse in (92)–(94) will not correspond to that of MT, just as the propositional content of (107)–(109) doesn't correspond to it. Specifically, where p is the proposition that the marble is big, q that the marble is likely red *given that it is big*, and r that the marble is likely red *given all we know*, we get the following pattern:

(112) If p , then q .

(113) Not r .

(114) So, not p .

It is not a surprise that this is not valid.

Second, as argued in Chapter 9, the propositional content expressed by modal discourse, and, thus, by arguments like the one in (92)–(94), is determined by the linguistic structure of (92)–(94), and the dynamic effect such a structure has on the conversational scoreboard (specifically, on the ranking of candidate possibilities), which is encoded in the dynamic layer of meaning associated with each expression. This means that (92)–(94) is associated with a discourse structure that carries a dynamic effect on context which, in essence, serves as a recipe for building precisely the propositional pattern in (112)–(114). One cannot express another propositional content with this particular discourse.⁸

Let us now turn to the counterexample to MP, the one in (104)–(106). The conditional premise, (104) is uttered against the body of information that includes the description of the elections and the polls. Relative to this information, it introduces a hypothetical possibility, one in which a Republican wins. The consequent further elaborates on this possibility, and Elaboration makes it prominent. This is why the conditional in the consequent is restricted by it. Relative to this restriction, the conditional in the consequent introduces a possibility in which Reagan loses. Its consequent further elaborates on this possibility and requires that it be one in which the winner is Anderson. As per our truth-conditions defined in Definition 9.2, the premise expresses the proposition that's true at w , just in case all the w -accessible worlds in which the Republican other than Reagan wins, are ones in which Anderson wins.

The second premise provides further elaboration on the first one, constraining it to be one in which a Republican will win. Finally the conclusion is related to the premises by relation of Conclusion. Conclusion promotes the body of information which incorporates the information expressed by the premises—the body of information comprising epistemically live worlds in which (104) and (105) hold. Note that throughout this possibility, the winner is a Republican. This body of information, in turn, restricts the conditional in the conclusion: relative to it, the antecedent introduces a hypothetical scenario, one in which Reagan loses. The consequent further elaborates on this scenario, requiring that throughout it the winner is Anderson. The conditional in the conclusion is

⁸ A bit more precisely, one cannot express any other propositional pattern with this discourse, given its natural disambiguation. This is consistent with holding that, as explained earlier, discourses may sometimes be ambiguous with respect to their structure.

true just in case within all the worlds compatible with the scenario, in which a Republican wins, and in which Reagan loses, Anderson is the winner. Note that this is precisely the content expressed by the consequent of the conditional premise, and precisely the content that MP-inference commits us to.

Where does this leave us regarding the status of the counterexample? Note that, given our explanation above, the propositions expressed by the premises entail the one expressed by the conclusion: the propositions expressed by the premises are true of a world, only if the one expressed by the conclusion is. And note that given our setup, this is guaranteed by virtue of the structure of the discourse in (104)–(106). It is just not possible, given our setup, that the discourse in (104)–(106) has the structure it does while expressing any other content.⁹ And indeed, we get exactly the propositional content corresponding to MP inference pattern: (104)–(106) is, indeed, valid, and an instance of MP.

What explains the appearance of a counterexample? Notice that, since modals and conditionals are context-sensitive, utterances containing them express propositions only relative to a context. And, since they are prominence-sensitive, they express propositions only relative to a prominence ranking as it is at a point in discourse at which they occur. It is easy to forget this when one is considering the conclusion, (106), in isolation, and not as a part of the broader discourse. Considered as an isolated sentence, there is no guarantee (106) would express the proposition it does when embedded in the discourse in (104)–(106), just as there is no guarantee that ‘she sat down’ expresses the same content unembedded, and when uttered as a part of the discourse, as in ‘Betty came in. She sat down.’ Indeed, in a complete isolation from context, we cannot really ask which content (106) expresses, just as we cannot ask which proposition is expressed by ‘she is happy’ in a complete isolation from context.¹⁰

Let’s take stock: on my account then, (92)–(94) is not a genuine instance of MT, because the structure it exemplifies—its logical form—determines propositional pattern that does not correspond to MT. (104)–(106) is a (valid) instance of MP, and the structure it exemplifies—its logical form—determines the propositional pattern corresponding to MP. Then, one might wonder, in what sense is the present account advancing over the non-propositionalists, who also render (92)–(94) invalid and (104)–(106) valid? And in what sense have we given a uniform treatment of the phenomenon?

The answer is that, while indeed, my account agrees with non-propositionalists in this regard, it also captures something they cannot. One cannot say, on the

⁹ We have, at this point, only informally described the import of the structure. The formal account of the dynamic meaning is forthcoming in Chapter 11. Consult also Appendix B.

¹⁰ A similar point is made by Gillies (2004), albeit, as we have seen, he draws a different conclusion from it.

non-propositionalist account that (92)–(94) is not only an invalid pattern, but also one that isn't a genuine instance of MT. On my account, we can explain why (92)–(94) is a bad argument while at the same time maintaining that MT is a perfectly valid inference pattern. While on the non-propositionalist account the context dynamically changes throughout an interpretation of an argument, such a change is neither reflected in the logical form of an argument, nor, more importantly, does it have an impact on content expressed by the premises (for there is none). One cannot say that the argument in (92)–(94) involves an equivocation on a non-propositionalist account. On my account, by contrast, the contextual change is both reflected in the logical form of an argument and it impacts the content expressed—it is because of the contextual effect of coherence relations that (93) does not express a proposition that negates the consequent of (92). And notice that exactly the same explanation applies in the case of (104)–(106). It is because of the effects of coherence relations the discourse harbors that the propositional pattern is guaranteed to be a valid one. This is the sense in which we have a uniform explanation of the phenomena, one only relying on general linguistic mechanisms of context-sensitivity resolution, that is, on discourse conventions required to explain argument patterns involving even the most basic type of context-sensitive expressions such as demonstrative pronouns.

We also have resources, at this point, to address the question of validity in the presence of context-sensitivity more generally. The initial Kaplanian reaction to the concerns regarding validity raised by context-sensitivity was to simply require that throughout an argument pattern, the context is held constant. In presenting an argument, we are implicitly committed to a stable context, where the relevant context-sensitive items are interpreted uniformly. But we know, already, that this account is inadequate: one cannot simply impose a ban on context-shifting, because the very expressions an argument comprises affect the context in a way that has interpretive effects on the content of the premises and conclusions. The question we faced was how to account for validity within the dynamically changing context. On the present account we can do so while straightforwardly preserving classical patterns of inference. We can do so because the dynamics of context, which affects the interpretation, is governed by the elements in the logical form of a discourse; so the form determines the underlying propositional, informational, pattern. What remains, then, is to specify formally the dynamic content of modal discourse, and show that it delivers the relevant truth-conditional content. I turn to this presently.

Prominence Semantics for Modality

We take utterances containing modals and conditionals to express propositional content, partly determined by the contextually supplied restriction on their domain of quantification. But at the same time we said that modal discourse carries a layer of content—dynamic content—the role of which is to govern the dynamics of contextual parameters, inducing the shift in prominence ranking of possibilities that can serve as restrictors for subsequent modal expressions. So we specify the dynamic meaning which builds the propositional content expressed by modal utterances. The propositional content so built can, in turn, be made prominent by further prominence-shifting updates, and thus affect the interpretation of downstream expressions, in a similar way as the referent of a pronoun can be made prominent, and thus affect the interpretation of subsequent pronouns.

Just as an attentional state of the discourse, the prominence of candidate possibilities can be seen as a part of the conversational scoreboard, an abstract record keeping track of the values of the contextual parameters as the discourse progresses. For convenience, since we are dealing with modal discourse specifically, it will simplify things if we abstract away from the elements of the conversational scoreboard that do not affect the interpretation of modal vocabulary. In this way, we can focus specifically on the relative prominence of possibilities in a discourse context. That is, we focus on the aspect of the conversational scoreboard which represents a ranking of possibilities—sets of worlds, or propositions. The conversational record tracks propositions put into play in the course of a conversation as well as their relative, and dynamically changing, prominence.

From a formal standpoint, we can model this exploiting the same resources we exploited to model the attentional state of the discourse. That is, we can use the assignment functions, modeled as stacks ordered by prominence, to represent prominence ranking of possibilities in a discourse. The only difference will now be that we will rank possibilities, that is, propositions, and ignore the ranking of individuals. For simplicity, let's call the relevant ranking “context,” abstracting away from the fact that it represent only one aspect of the contextual scoreboard.¹

¹ Formally, just as before, a stack is an assignment function. As we are ranking possibilities, we could think of such a ranking as a function defined on a finite convex subset of $\mathbb{IN}$, mapping onto a set

Just as the ranking in the attentional state of the discourse changes as the discourse unfolds, so does the prominence of possibilities: new ones are promoted, demoting others. We have seen, in Part II, that this dynamics is governed by the updates associated with the expressions the discourse comprises, as the semantic composition unfolds, word-by-word. We can model this, then, just as we did in Chapters 4–7, by interpreting utterances as relations between contexts, that is, as instructions on how to update the contextual dependencies, potentially creating new dependent interpretations for subsequent expressions. As before, we say that an update is true (relative to an input context and a world of evaluation w) just in case it relates the input context to some non-empty output context.

As explained informally in Chapters 9 and 10, utterances make a two-fold contribution: they express truth-conditions, but they also contribute updates that change the context, by updating the ranking of possibilities. The dynamic content serves as a recipe for building propositional content expressed; the propositional content thus built can be promoted by subsequent updates, thus affecting the truth-conditional content expressed by subsequent discourse. In this way, 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.

We shall specify updates in such a way so that they recursively build the proposition expressed by each component clause.² In this way, we will fully represent

of propositions, that is, sets of worlds, together with an undefined value ($\perp$). Then we could interpret our utterances as relations between assignments, as we did in Part I. Given our target examples, this would serve us just fine, and in certain respects would be simpler. But we will not endorse this strategy. Rather, officially, we will treat stacks as orderings of worlds, not propositions, that is, as mappings from a finite convex subset of $\mathbb{N}$ to a set *worlds* together with an undefined value $\perp$. We will in turn take the relevant aspect of the conversational scoreboard representing the ranking to be a set of such assignments—a set of sequences of worlds. Then, relative to G , w_i is a variable that stores a proposition comprising the worlds at the i th position of every stack in G . I shall assume that every stack in G begins with the 0th position, as the top-ranked one, and that for each position n , the position $n + 1$ is one position lower in the ranking. G_i will denote the proposition that comprises the worlds at the i th position of every stack in G , that is, the proposition stored at the i th position in G . G_0 is then, the top-ranked proposition in G , and for each n , G_{n+1} is a proposition one step lower in the ranking. (For precise definitions, consult Appendix B.)

Why model ranking with sets of assignments, rather than single assignments? Because in this way we can keep track not only of propositions (sets of worlds), but also relations between individual elements of different sets (van den Berg, 1996; Brasoveanu, 2010a; Bittner, 2014). This becomes crucial once we want to develop an integrated semantics for modals and pronouns, since at that point we will want to be able to track different individuals across different worlds a proposition comprises. For the moment we are ignoring the ranking of individuals required for pronominal interpretation, but if we wanted to simultaneously model modal and pronominal interpretation, thus representing simultaneously two different parameters of the scoreboard, we would include individuals within the range of our stacks.

² I will only gloss the effects of the relevant contextual updates here. Note that I shall proceed by specifying a translation of a (modal) fragment of English into a formal language with dynamic interpretation. This formal language has atomic expressions—propositional constants (p, q, r) and variables (w_i for $i \in \mathbb{N}$)—conditions (propositional expressions comprising a fixed set of atoms

utterances as having two layers of meaning: propositional content and a dynamic effect on context.³ Take a simple, atomic clause, for example, ‘it is raining.’ This clause expresses propositional content, *that it is raining*. It can also contribute that propositional content to larger expressions it is a part of, as in ‘it might be raining.’ So, to capture this aspect of semantic composition, our simplest update will be one that simply records the propositional content expressed by the clause it corresponds to. This update doesn’t itself affect the ranking, it just makes a proposition available for semantic composition. Correspondingly, we will include a designated, unranked position in our context—what I dub *comp*, for “composition”—which gets updated by the simplest update. Where p is the proposition expressed by the clause, the simplest update is $\langle \text{comp} := p \rangle$; its effect is to update an input context by recording the content p as the new value of the position *comp*. When we have a more complex, non-atomic construction, it will be interpreted as a more complex update. However, we will ensure that an update associated with any clause, no matter its complexity, stores the truth-conditional contribution of the clause as the new value of *comp*, capturing the fact that it can further serve in composing more complex structures. This is the way in which updates recursively compute the proposition expressed by each clause.

A modal, say, ‘might,’ like a pronoun, is a context-dependent element, looking for a prominent possibility to restrict its domain of quantification. At the same time, like a pronoun, it is a prominence-affecting element, promoting the possibility it describes as the new top-ranked one, as showcased by (67), repeated below:

(67) A wolf might walk in. It would eat you first.

An update corresponding to *might* ϕ , reflects this: it selects the top-ranked possibility in the input context as a restrictor, and promotes the possibility comprising the epistemically live worlds in which the restrictor and the prejacent proposition hold. More precisely, where K is the update associated with the prejacent ϕ , and $@E$ denotes the top-ranked epistemically open possibility in the input context G , the update associated with *might* ϕ , $\text{MIGHT}(@E, K)$, first updates G with the prejacent K , storing the proposition it expresses in *comp* position, and then, given the top-ranked epistemic possibility in G , makes the subset of this

closed under conjunction and negation), and update expressions. The official formal definitions are best relegated to an appendix (see Appendix B).

³ For simplicity, we will only look at clause-level composition, thus a minimal update will be associated with an atomic clause, that is, one that has no sentential sub-parts.

possibility in which this new value of *comp* holds prominent, demoting others.⁴ In the end, the truth-conditional content of the modal utterance is stored as a new value of *comp*, thus becoming available for further semantic composition. Given our truth-conditions for *might*-claims, defined in 9.1 this will be the proposition that's true at a world *w*, just in case there's a world accessible from *w* in which the top-ranked possibility in the input context holds, and the prejacent proposition holds as well, that is, just in case there is a prejacent world in the top-ranked epistemically open possibility.

So, for example, where *p* is the proposition *that a wolf walks in*, an update with (67) will be (115):

(115) MIGHT(@E, ⟨*comp* := *p*⟩)

(115) first stores *p* as the new value of *comp*, and then takes the top-ranked epistemically live possibility in the input context, and promotes a subset of it in which *p* holds, thus making the most-prominent epistemically live possibility in which a wolf walks in the new top-ranked one, demoting others. Finally, it stores the proposition expressed by (67), which is true at a world *w* just in case there is some epistemically accessible world among the most prominent ones in the input context, in which a wolf walks in—that is just in case a wolf might walk in, given the prominent body of information. This is just what we want: (67) expresses the proposition that it is compatible with the prominent epistemic possibility in the input context that a wolf walks in, and promotes a possibility comprising the top-ranked epistemically live worlds in which a wolf walks in.

Updates with modal expressions of different forces can be defined in exactly the same way, varying only in the truth-conditional contribution as needed to capture the difference in force. What about the conditional? Its effect is similar, except that we have to sequentially process first the update associated with the antecedent, and then the one with the consequent. A conditional, like a modal (or a pronoun), is context-sensitive, looking for a restrictor, and it's prominence-affecting—making certain possibilities prominent. Where *K*₁ and *K*₂ are updates associated with the antecedent and the consequent, the conditional update, IF(@E, *K*₁, *K*₂), first updates the input context with the antecedent, storing the propositional content it expresses as the value of *comp*. Then it takes the most prominent body of information in the input context, and promotes as the new top-ranked proposition the one corresponding to the set of worlds in this body of information in which the antecedent proposition

⁴ Formally, extending the basic vocabulary of our dynamic language, we include 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 the property denoted by *E* and delivering the top-ranked proposition satisfying it. For a precise formal specification, please see Appendix B.

holds. Further, relative to the state of the context in which *this newly promoted possibility introduced by the antecedent* is prominent, it updates with the consequent, storing its contribution as the new value of *comp*, and promotes as the new top-ranked proposition the one comprising the set of worlds in the currently top-ranked body of information—the one already updated with the antecedent proposition—in which the consequent holds.⁵ And finally it stores the proposition expressed by the whole conditional as the final value of *comp*. Given Definition 9.2, this will be the proposition true at a world *w*, just in case, in all the worlds accessible from *w*, which are in the top-ranked body of information in the input context, and in which the antecedent proposition holds, the consequent holds as well.

In this way, then, we get exactly what we were aiming for: utterances affect the context expression-by-expression, in a way that sequentially builds the propositional content of each component clause. Modals and conditionals, specifically, update the context by affecting the prominence ranking of possibilities that are candidate restrictors for subsequent modals and conditionals. At the same time they are context-sensitive: the restrictor prominent at a point in context at which a modal expression occurs affects the proposition that the clause to which it belongs expresses. In this way, the updates associated with utterances drive the evolution of context as the discourse progresses; the context in turn determines the truth-conditions expressed by the utterances. Thus, we capture the two-way interaction between context and content.

Given that the updates specify how the context will evolve, they are essentially building the truth-conditional content expressed. As explained, the update associated with each clause stores its corresponding truth-conditional content as the value of *comp* at some point during semantic composition. This captures what it takes to express a certain content. 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, as it can be picked up, for instance, by a propositional anaphor:

- (116) a. It is raining.
b. That is true.

⁵ In this way, we predict that the epistemically live possibility in which the antecedent and consequent both hold is available as a potential restriction for subsequent modals. This prediction is correct, as witnessed by examples like (i):

(i) If a wolf walks in it would eat you. It would then eat Harvey, too.

To capture this we ensure that the effect of an assertion is reflected in our context, by introducing an assertion update that achieves this effect. Where K is an update corresponding to an utterance, the assertion update, $\text{ASSERT}(K)$, promotes the set of top-ranked epistemically live worlds in which the asserted content—the propositional content expressed by the utterance asserted—holds, as the novel top-ranked possibility, and requires that the actual world be within that set.

So far we have specified the updates associated with individual clauses and sentences, but we have not yet said anything about the updates associated with coherence relations holding between those clauses. We have seen that those, however, are crucial for the dynamics of prominence governing the resolution of context-sensitivity of modal and conditional expressions. For instance, recall (92)–(94), repeated here:

(92) If the marble is big, then it is likely red.

(93) The marble is not likely red.

(94) So, the marble is not big.

I argued (92)–(94) 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 of information that both premises are about—that is, the whole set of epistemically accessible worlds discourse initially. To capture this, we have to say something about how coherence relations update the context.

As was the case in Part I, we represent coherence relations as contributing prominence-affecting updates. More precisely, they have a two-fold contribution: a *relational* contribution, requiring that a specific relation hold between the relevant bits of discourse, and a *prominence-affecting* contribution, making certain possibilities prominent.

Consider Elaboration first. As we have seen, Elaboration promotes the possibility that is elaborated upon. We capture this idea as follows: where ϕ is a possibility elaborated on, and K an update associated with the utterance that provides an elaboration on ϕ , the update contributed by Elaboration relation, $\text{ELAB}(\phi, K)$, makes a two-fold contribution: it makes the possibility elaborated upon, ϕ , prominent, demoting others, and it requires that the propositional

content expressed by the utterance that provides elaboration (which will be built by K), stands in the Elaboration relation to ϕ .⁶

The Contrast relation between two bits of discourse, as was explained in Chapters 9 and 10, requires that they provide contrasting information about some body of information regarding some common topic. In turn, this body of information is made prominent. So, where K_1 and K_2 are updates associated with the bits of discourse related by the Contrast relation, the update associated with the relation, $\text{CONTRAST}(K_1, K_2)$, first updates the input context with K_1 , computing its propositional content relative to the top-ranked body of information in the input context, then it again promotes the top-ranked body of information in the original input context, and updates with K_2 , thus computing the propositional content it expresses relative to this body of information. Finally, it requires that K_1, K_2 stand in Contrast relation to each other.⁷

Let's look at how all this comes together by focusing on one of our target examples. Where K is an update, let $\text{NOT}(K)$ be an update associated with a negation: it simply stores the complement of the truth-condition expressed by K , as the value of *comp* of the output context.⁸ Then where p stands for "the marble is big" and q for "the marble is red," we represent (92) and (93), the premises of the argument in (92)–(94), as follows:

$$(117) \text{CONTRAST}(\text{ASSERT}(\text{IF}(@E, \langle \text{comp} := p \rangle, \\ \text{ELAB}(\mathbf{w}_0, \text{LIKELY}(@E, \langle \text{comp} := q \rangle))), \\ \text{ASSERT}(\text{NOT}(\text{LIKELY}(@E, \langle \text{comp} := q \rangle))))^9$$

⁶ We can maintain a provisional characterization of the Elaboration relation: we characterize an elaboration relation between propositions ϕ and ψ , $\text{Elab}(\phi, \psi)$, as holding just in case ϕ and ψ are centered around the same event or entity, that is, just in case the event or scenario described by ψ is a part of the event or scenario described by ϕ (cf. Hobbs, 1979; Asher and Lascarides, 2003). This suffices because we primarily care about the effect on prominence.

⁷ Note that CONTRAST is operating on two updates, and ELAB 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).

⁸ So, where ϕ is the proposition expressed by K , $\text{NOT}(K)$ makes the complement of ϕ the new value of *comp*.

⁹ Note that given this formalization, CONTRAST will not contribute any asserted content on its own, or more precisely, it will not be a part of the asserted content that the two bits of discourse contrast each other. This is by design, as the contrasting dimension of meaning does behave as part of backgrounded, rather than asserted content. 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. So, we see now how we can exploit the resources of the present formal framework to represent not just asserted but also backgrounded content.

(117) is our logical form associated with the premises of the argument in (92)–(94). It makes explicit the changes in context, which in turn determine the underlying propositional pattern. The updates associated with the two premises are (i) and (ii), respectively. They are organized by Contrast relation, and thus, the prominent possibility for the interpretation of both premises will be the top-ranked body of information in the input context.

(i) ASSERT(IF(@E, <comp := p>, ELAB(w₀, LIKELY(@E, <comp := q>))))

(ii) ASSERT(NOT(LIKELY(@E, <comp := q>)))

Here are the key steps in (117). The conditional update in the first premise, IF(@E, <comp := p>, ELAB(w₀, LIKELY(@E, <comp := q>))), first introduces the possibility corresponding to the set of top-ranked epistemically accessible worlds in which the proposition expressed by the antecedent holds, that is, the set of epistemically live worlds in which the marble is big. The consequent, ELAB(w₀, LIKELY(@E, <comp := q>)), provides an elaboration of this possibility, which at this point is ranked at the top-ranked position 0 (and, so, denoted by w₀ in (117)); as a result, the possibility is made prominent. Furthermore, the update requires that the possibility introduced by the consequent stands in Elaboration relation to the possibility introduced by the antecedent. Since the consequent contains an occurrence of the modal *likely*, which is prominence-sensitive, 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, given the prominent information (discourse initially), the marble is likely red, given that it is big. Then by Definition 9.2, the whole conditional expresses the proposition that given the prominent information, if the marble is big, then it is likely red, given that it is big. The assertion update wrapped around the conditional update ensures that the premise is asserted, so that it is required that the conditional holds of the actual world. Finally, the set of epistemically live worlds in which it holds is promoted.

Due to the effect of Contrast, which connects (i) and (ii), the body of information that the first utterance was about (which is the top-ranked body of information in the initial context, which, by assumption, is the one describing the situation concerning the urn and the marbles) is promoted. Then the update associated with the small premise, ASSERT(NOT(LIKELY(@E, <comp := q>))), is interpreted relative to this possibility, and the modal in the small premise will select it as a restrictor. This provides the interpretation that, given all that

is known, the marble is likely red. Given negation, then, the small premise expresses the complement of this possibility; that is, the proposition that it is not the case that given all that is known the marble is likely red. The assertion update wrapped around the small premise ensures that this proposition holds of the actual world, and promotes the set of epistemically accessible worlds in which the content holds. In addition, the propositions expressed by the two premises are required to stand in the Contrast relation—that is, to provide contrasting information regarding a common topic, about the body of information they are about, that is, the set of epistemically accessible worlds discourse initially.

(117) 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, given that it's big, but, given all that is known, the marble is not likely red*. This of course, will not yield the pattern corresponding to the premises of MT. So, *a fortiori*, (92)–(94) is not a counterexample to MT. So, on my account, (92)–(94) is infelicitous for exactly the same reason (62) was inconsistent: just as the internal discourse structure of the antecedent of (62) determines an inconsistent truth-condition, so in (92)–(94) it determines an invalid propositional pattern. The context-change threading through the argument is critical for the correct interpretation, and is forced by the dynamic layer of meaning.

We can also capture our other target argument, the alleged counterexample to MP, in (104)–(106), repeated below:

- (104) If a Republican wins the election, then if it's not Reagan who wins it will be Anderson.
- (105) A Republican will win the election.
- (106) If it's not Reagan who wins, it will be Anderson.

To represent the argument in our current system, we just need to specify the contribution of the Conclusion relation. Like other coherence relations, Conclusion has a two-fold contribution: it promotes the proposition corresponding to the set of worlds in the top-ranked body of information in the input context within which the premises hold as the new prominent body of information for the interpretation of the conclusion, and it requires that the content expressed by the premises entails that of the conclusion.¹⁰ So where K_1 is an update corresponding to the premises, and K_2 the one corresponding to the conclusion, an update

¹⁰ Where I assume that proposition ϕ entails ψ iff $\phi \subseteq \psi$.

with the Conclusion relation, $\text{CONCLUSION}(K_1, K_2)$, will first update the context with the premises, then promote the subset of the most prominent body of information in the input context in which the premises hold, and then update with the conclusion; finally, it will require that the propositions expressed by the premises entail the one expressed by the conclusion.

We can then represent (104)–(106) as follows:

- (118) $\text{CONCLUSION}(\text{ASSERT}(\text{IF}(@E, \langle \text{comp} := r \rangle, \text{ELAB}(\mathbf{w}_0, \text{IF}(@E, \text{NOT}(\langle \text{comp} := n \rangle), \text{ELAB}(\mathbf{w}_0, \langle \text{comp} := a \rangle))))); \text{ASSERT}(\text{ELAB}(\mathbf{w}_0, \langle \text{comp} := r \rangle)), \text{ASSERT}(\text{IF}(@E, \text{NOT}(\langle \text{comp} := n \rangle), \text{ELAB}(\mathbf{w}_0, \langle \text{comp} := a \rangle)))$

As before, the conditional update in the first premise first introduces a hypothetical scenario in which a Republican wins relative to the prominent body of information in the initial input context, which includes the information about the polls. The conditional in the consequent elaborates on this newly introduced possibility—the one in which a Republican is the winner, and the information about the polls holds. The update with Elaboration, $\text{ELAB}(\mathbf{w}_0, \text{IF}(@E, \text{NOT}(\langle \text{comp} := n \rangle), \text{ELAB}(\mathbf{w}_0, \langle \text{comp} := a \rangle)))$, first makes this possibility elaborated on prominent, and then proceeds to update with the consequent. This ensures that the most prominent possibility in the context in which the consequent is interpreted will be the one promoted by the Elaboration. The conditional in the consequent, too, first updates this context with the antecedent, introducing the possibility that Reagan lost, relative to the possibility made prominent by Elaboration. Finally, the consequent of the embedded conditional elaborates on this possibility, and requires it is one in which Anderson wins. So, the propositional content will be one true of a world w , just in case all epistemically accessible worlds in which the prominent information in the initial context holds, and in which a Republican wins and Reagan loses, are worlds in which Anderson wins. Finally, the conditional is asserted, which ensures that it must hold of the actual world, and that the proposition corresponding to the set of worlds in the most prominent body of information in the initial context, in which it holds, is made prominent.

Now, the small premise elaborates on the possibility made prominent by the assertion of the conditional premise, further asserting that in it, the winner will be a Republican. The assertion promotes the subset of the possibility made prominent by the first premise comprising the worlds in it in which the winner is a Republican, and requires that it holds of the actual world. Finally, the Conclusion update, relating the premises and the conclusion, promotes the

body of information prominent in the initial context and updated with the information in the premises. This in turn becomes the most prominent possibility in the context in which the conditional in the conclusion is interpreted. The conditional in the conclusion, then, selects this possibility as a restrictor, and, relative to it, it introduces a possibility in which Reagan loses. The consequent elaborates on this possibility, and finally requires that in it Anderson is the winner. The conclusion expresses the proposition that's true at a world w just in case all the worlds accessible from w in which the information prominent in the original input context holds, and in which the premises hold and in which Reagan loses, are ones in which Anderson wins. Because throughout the prominent body of information in which the premises hold a Republican is the winner, the proposition expressed by the conclusion will be entailed by the one expressed by the premises. And the argument is valid in our dynamic system as well.¹¹

Our discussion raises an interesting question about argument individuation that we already touched upon in Chapter 10, but that now we can give a fully precise answer to.¹² The question is, how do we individuate discourses that instantiate genuine form of MT, from discourses like (92)–(94), that do so only superficially? The key to answering this question is to appreciate that forms like the ones in (118) or (117) unequivocally determine a specific propositional pattern, and do so fully as a matter of their internal structure. Thus, saying that (92)–(94) is a counterexample to MT is like saying that the following is a counterexample to the $\phi.\phi$ argument pattern:

(119) He (pointing at Bill) is happy. Therefore, he (pointing at Jim) is happy.

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 dependence (i.e., the value of $@E$) of the modal in the big premise and the small one is resolved. And this, as we have seen, is a matter of discourse-internal linguistic structure—discourse conventions—which is reflected in the logical form of the discourse.

¹¹ An inference pattern is dynamically valid just in case, if the sequential update with the premises leads to a non-empty context, so does the sequential update with the premises followed by the conclusion. (See §B.3.1 for a precise definition.) Indeed, we can prove that the dynamic system described here 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. For formal details consult Appendix B and see Stojnić (2017b).

¹² For a related discussion of some issues concerning argument individuation, see also Stojnić (forthcoming).

This is because the updates that constitute a discourse determine which truth-conditions are expressed by each clause within the discourse. Whether or not a discourse will be an instance of MT (partly) depends on whether the truth-conditional content expressed by the consequent of the big premise in the context obtained after updating with its antecedent, and the truth-conditional content expressed by the small premise in the context obtained after updating with the big premise, indeed do negate each other. But since a full-blown logical form of an argument determines when this is the case, we can capture the idea of validity as a matter of logical form: in particular, a sequence of updates expressing classically valid truth-conditional pattern will be valid in our dynamic system as well.¹³

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. And a discourse structured in such a way so that it determines this particular propositional pattern will instantiate MT and it will be valid. So, while in a sense the superficial form of (92)–(94) is misleading insofar as the string in (92) and the one in (93), taken together, might seem like they correspond to the premises of a MT argument, the logical form of (92)–(94) is determined directly as a matter of its internal discourse structure which is richer than its overt superficial form. Still, on my account, unlike on the standard contextualist account, pragmatic inference does not play a role in determining form or affecting validity.

One might worry, given the logical form of the argument in (117), that there might be no actual English arguments that are instances of MT. But this is not right. There are many arguments in natural language that clearly do express MT pattern. For instance, the following is a good candidate:

- (120) A random marble is selected from an urn. Suppose we know that all big marbles are red. So we know (120-a). But then we learn that the selected marble is not red. So, we know (120-b). Then, we conclude (120-c).
- a. If the marble is big, then it's red.
 - b. The marble is not red.
 - c. So, it's not big.

The point is that whether an argument expresses a particular pattern of inference depends, in part, on the truth-conditional content its premises and conclusion

¹³ What one can show is that, for any way of constructing a dynamic formula, there is a classical translation (i.e., a translation in a classical (modal) system). That is, the dynamic system I develop embeds classical (modal) logic. See Appendix B.

express.¹⁴ But note that on my view there is still a tight connection between the logical form of an argument and the propositional pattern it expresses, since the propositional pattern expressed is entirely determined by the logical form—the discourse structure—the argument exemplifies. No argument exemplifying the form in (117) is valid.

But does this mean that no argument with a modal in the consequent is an instance of MT? Notice, in order for a form to express an MT pattern, the modal in the big (conditional) premise, and that in the small premise would have to be resolved in the same way.¹⁵ I don't predict that no such argument is an instance of MT. However, I do predict that there are challenges to expressing this form in English discourse. Specifically given that modals are anaphoric, insofar as the occurrence of the modal in the conditional premise of a MT argument is anaphoric on the antecedent, the one in the small premise will likewise have to be thus anaphoric for the argument to be an instance of MT. But this is at odds with the contrastive structure of a MT discourse, which organizes the conditional and non-conditional premise around Contrast relation. In other words, if it is possible to have a structure that guarantees that both occurrences of the modal are resolved to the same antecedent, we will indeed get genuine

¹⁴ Note, in characterizing the argument pattern in this way we are characterizing it partly in terms of its underlying truth-conditional content, not merely in terms of its superficial sentential form, or a purely syntactic form. One might be worried about this. For instance, Yalcin (2012b, pp. 1007–9) writes that since standard Kratzerian restrictor analysis of conditionals does not recognize the English conditional as a binary operator, and, hence, assuming this analysis is correct, 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 is supposedly a generalization about this alleged dyadic operator, is ill-conceived from the start (as is, according to this construal, MP as well). There simply is no stable notion of “antecedent” and “consequent” for English conditional, so there is no MT (or MP). But this argument is too quick. Even if something like the Kratzerian analysis of a conditional is correct, that will hardly constitute a demonstration that MT and MP actually do not exist. At best, it shows that 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. Concluding that the absence of a particular unique syntactic structure entails the invalidity or non-existence of MP and MT would be like arguing that, say, a formalization of the propositional calculus by means of the Sheffer stroke invalidates or dissolves MP and MT.

¹⁵ Essentially, this would have the same effect as if Contrast “scoped low,” that is, if Contrast relation held between the *consequent* of the conditional premise and the small premise, as in ‘If a wolf walks in, it might eat you. But it might not.’ If Contrast were to attach low in our target case, the propositional content expressed would be that *if the marble is big then it's likely red, given that it's big and the marble is not likely red, given that it is big*, in which case it would follow from these premises that the marble is not big. In essence, we would get the same result as: *if ϕ then ψ , if ϕ then $\neg\psi \therefore \neg\phi$* , which is, of course, valid. So, we do, on my account, validate *if ϕ , likely ψ ; but not likely ψ ; $\therefore \neg\phi$* when the contrast attaches low in this way, as this, in essence, expresses the propositional form *if ϕ , then ψ ; if ϕ , then $\neg\psi$, $\therefore \neg\phi$* , which is valid. But this form is also arguably dispreferred on pragmatic grounds: in uttering this discourse I'm introducing a possibility, describing it in a roundabout way, only to tell you it does not obtain. If I told you ‘If I won the lottery, I'm rich. But I'm not rich if I won the lottery,’ that would be a strange way to inform you that I did not win the lottery.

instances of MT.¹⁶ But even if it did turn out that it is hard, or even impossible, to express genuine instances of MT the conditional premise of which contains (otherwise unembedded) modals in the consequent in English, that in itself would not be a problem for my view, nor would it entail that MT is invalid in the resulting semantics.¹⁷

The traditional approach to validity within a standard contextualist framework, recall, guards against illicit shifts in interpretation by imposing a ban on context-shifting. But we have seen that this is unsatisfactory if we want to capture intuitively valid arguments involving occurrences of context-sensitive expressions. On my account, however, no such ban is needed to explain validity. Indeed, as is clear given my account, it is impossible to impose a ban on context-shifting: this is because discourses are composed of linguistic items which govern the evolution of the conversational scoreboard as a part of their meaning. This precisely explains why a discourse like (92)–(94) cannot be felicitously interpreted as an instance of a valid MT argument.

As briefly and informally explained in Chapter 10, when it comes to argument individuation, my account also has an advantage over the non-propositionalist approaches. We can now see this more precisely as well. We have seen, in

¹⁶ Notice that arguments such as the following will correctly be predicted to be valid:

- (i) a. If the marble is big, then it must be red.
- b. But the marble cannot be red. (It's not the case that it might be red.)
- c. So, the marble is not big.

Here, the premises would have the following structure:

$$\begin{aligned} &\text{CONTRAST}(\text{ASSERT}(\text{IF}(@E, \langle \text{comp} := p \rangle, \\ &\quad \text{ELAB}(\mathbf{w}_0, \text{MIGHT}(@E, \langle \text{comp} := q \rangle)))), \\ &\quad \text{ASSERT}(\text{NOT}(\text{MIGHT}(@E, \langle \text{comp} := q \rangle)))) \end{aligned}$$

In such a structure, the propositional content expressed by the small premise is guaranteed to entail the negation of the propositional content expressed by the consequent of the big premise.

¹⁷ Notice, similarly, that on my account, the so-called import/export pattern—the argument pattern of the form *if* ($\phi \wedge \psi$), *then* γ : *if* ϕ , *then* (*if* ψ , *then* γ) is invalid (as expected). Yet, it is also the case, given my account, that—as intuition suggests—arguments like ‘If the butler is innocent and the gardener is innocent, then the cook is guilty; so, if the butler is innocent, then if the gardener is innocent the cook is guilty’ will be valid. This is because conditionals are anaphoric, and so long as the relation between the antecedent and the consequent of the embedding conditional in the conclusion makes the antecedent possibility prominent (as does Elaboration, in this particular example, and as would any other relation that has the same effect on prominence as Elaboration), the right-embedded conditional in the conclusion is bound to be restricted by the possibility described by the antecedent of the embedding conditional. But failures of exportation arise precisely when such restriction is lifted—that is, when the right-embedded conditional is allowed to be interpreted as unrestricted by the antecedent of the embedding conditional. So, my semantics predicts that, just as intuition suggests, in English, such conditionals are mutually entailing, while at the same time avoiding the result that exportation is valid. This is a desirable result. As shown by Gibbard (1981), under minimal plausible assumptions, a semantics that validates MP, exportation, and deduction theorem collapses a conditional into the material conditional. Our semantics invalidates exportation, while at the same time predicting why it is difficult to express counterexamples to exportation in ordinary English discourse. The moral is that, when individuating natural language inference patterns, we have to take into account potential equivocations due to contextual effects. We have to do this for arguments involving modals, just as we have to do it for those involving pronouns, or indexicals. For further discussion, see Stojnić (forthcoming).

Chapter 8, that these accounts maintain that the main aspect of the meaning of modal discourse is its context-change potential—it's effect on the context understood according to the Stalnakerian model, as a shared body of information. On this account, an utterance was true just in case it doesn't further change the context. And this notion of truth is typically coupled with a corresponding notion of entailment: we say that a sentence ϕ entails ψ just in case updating with ϕ leads us to a context in which ψ is true (accepted). Intuitively, a set of premises entails the conclusion just in case updating with the premises sequentially lands us in a context in which the conclusion is accepted.¹⁸ A bit more formally, here is how we can characterize this within the dynamic update semantics.¹⁹

Definition 11.1 Truth

ϕ is true (supported) in c , $c \models \phi$, just in case $c \llbracket \phi \rrbracket = c$.

Then, we define entailment as follows:

Definition 11.2 Entailment

$\phi_1, \dots, \phi_n$ entails ψ just in case, for any c , $c \llbracket \phi_1 \rrbracket \dots \llbracket \phi_n \rrbracket \models \psi$

The idea is that an argument is valid just in case any context that accepts the information encoded in the premises also accepts that of the conclusion. It is easy to see, then, that on this account, (92)–(94) would be invalid (given the appropriate test-like interpretation of the modal):²⁰ if the maximal subset of the context within which the marble is big supports the update with the consequent (so that, given that subset, the marble is likely red), and the context as a whole supports the update with the claim that the marble is not likely red, it doesn't follow that the context as a whole accepts that the marble is big. Now, this in itself is not a bad result. But such an account misses something important about (92)–(94). It does not register an implicit shift in *interpretation* that occurs in (92)–(94)—one quite akin to the one we see in (107)–(109) or (119). As a result, it concludes that (92)–(94) is a genuine instance of MT, and thus a counterexample to the rule.

¹⁸ A family of related dynamic notions of entailment is available. We focus on update-to-test entailment (instead of, e.g., test-to-test entailment), as we are interested in both the agglomeration and preservation of information in a context.

¹⁹ The corresponding notion of entailment as acceptance preservation can be characterized within the expressivist framework, as well, and my point will equally apply. I omit the details for simplicity.

²⁰ Note, on non-propositionalist interpretation *likely* will have a test like behavior, just as *might* does. So, we are assuming here a test-like interpretation of the modal *likely* analogous to that of *might*, introduced in Chapter 8. (See Yalcin (2012).) Note that analogous putative counterexamples can be constructed with other modals, e.g., *must*, or *have to* all of which are assumed to have a test-like interpretation on the type of account under consideration.

The reason why the non-propositionalist semantics fails to capture this shift is rooted in the way it characterizes the dynamic meaning. It treats modal and conditional sentences as sensitive to a background body of information they are interpreted against, but does not really treat them as context-sensitive in an important sense: the truth of an update with a sentence varies depending on the input context, but since the account doesn't associate modal utterances with propositional, informational content, it does not generally predict that the information carried by a sentence is affected by the potential interpretive dependencies encoded in the input context. That is, we cannot say that the change in context brought about by the antecedent changes the proposition expressed by the consequent. Yet, I have argued, this information is crucial for individuating the logical form of an argument.

The theory I defended, by contrast, provides an account of context-change which 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. They are reflected in the logical form of a discourse, allowing us to restore a tight connection between validity and logical form despite the presence of context-sensitivity, and in a way which preserves classical patterns of inference. Unlike the non-propositionalist accounts, the account I developed allows us to keep track of changes in information brought about by the changing context, because we can read those changes off of the logical forms of a discourse. This, I have argued, allows for a more plausible way of individuating argument patterns—and evaluating argument validity.

As is clear from what I have been arguing throughout the book, my account exploits resources that affect context-sensitivity resolution more broadly, and thus are needed to account for argument validity quite generally. Recall, for instance, that it was representing the effects of discourse conventions explicitly in the logical form that allowed us to capture the fact that (27), but not (28), is valid (I repeat both discourses below).

(27) A woman came in. She sat down. Therefore, some woman sat down.

(28) A woman came in. She [pointing at Betty, the cat] sat down. Therefore, some woman sat down.

The account developed allows us to maintain that the forms of the arguments in (27) and (28) depend on their underlying discourse structure. It thus provides us with resources to capture the effects of the changing context, and at the same time, it individuates argument forms finely enough that differences in the information expressed will transpire as differences in form.

PART IV

CONCLUSIONS AND DIRECTIONS

The Grammar of Prominence

We began by asking how it is that what we say when we speak depends not solely on the words we use, but also on the context in which these words are uttered. What is it that fixes the meaning of context-sensitive expressions, and how is it possible that we can understand such varying meanings effortlessly, “on the fly”?

I have argued against philosophical and linguistic orthodoxy that relegates the mechanisms of meaning-fixing of context-sensitive expressions to extra-linguistic contextual factors—speaker’s intentions, and other non-linguistic facts about the context. I argued that context-sensitive expressions are best thought of as pure indexicals in the sense that their character—their linguistic meaning—determines their content given a context, where the context, too, is maintained and manipulated by linguistic rules, encoded in the dynamic layer of meaning of a discourse. On the account I defended, then, extra-linguistic features of context (e.g., world-knowledge and speaker’s mental states) do not play a role as determinants of semantic content. So, it is the linguistic meaning of an expression, together with linguistic mechanisms governing contextual parameters that determine the referent; no extra-linguistic supplementation is required.

The model of context I defended departs from the standard Kaplanian model. It departs from this model not only with respect to how parameters of context affect the interpretation at a particular point in discourse, but also with respect to how these parameters change with an evolving discourse. This, I have argued, is crucial for correctly determining truth-conditions. My account also differs from the familiar Stalnakerian model of context, understood as a body of shared information, mutually accepted for the purposes of a conversation, which supervenes on interlocutors’ mental states, and changes as the interlocutors’ attitudes are adjusted in rational response to manifest evidence. By contrast with this account, I have argued that a conversational context relevant for interpretation of context-sensitive expressions is best conceived of as an objective conversational record, keeping track of how the context changes with the unfolding discourse, as governed by discourse conventions, encoded in the dynamic layer of utterance meaning.

Though my account treats context-sensitive expressions as pure indexicals, it preserves an important insight behind the original Kaplanian distinction between pure indexicals and true demonstratives. While on my account the

interpretation of context-sensitive elements is not determined by appeal to extra-linguistic features of context, it is still governed by prominence, linguistically determined. So, while on my account context-sensitive items do not behave like true demonstratives, many context-sensitive expressions are *prominence-sensitive*. However, by contrast with the tradition which maintains that prominence is largely a matter of non-linguistic features of utterance, I argued that the dynamics of contextual prominence is governed by linguistic mechanisms.¹

I have argued that the account I developed is both empirically more adequate than its competitors and that it allows for a theoretically preferable account of the interaction between context and content, the nature of semantic content, and that of validity and entailment for natural language discourses.

Still there are a number of open questions we have to address. For one, while the model of context-sensitivity presented here was developed as a general model of context-sensitivity resolution, and while the framework allows for natural extensions to different context-sensitive items, I have focused on a relatively small number of theoretically interesting expressions, particularly, on demonstrative pronouns and on modal operators. This choice was underscored by their theoretical interest. Demonstrative pronouns represent a paradigm case of context-sensitivity, and have been taken as a model that has shaped theorizing about context-sensitivity resolution quite generally. Modal expressions, on the other hand, have motivated the key challenge to the truth-conditional paradigm of meaning, and to classical logic as an underlying logic of natural language discourse. One might wonder whether the same kinds of discourse conventions govern other kinds of context-sensitive expressions; which context-sensitive items are prominence-sensitive, and which pattern with the first-person indexical; and which further specific theoretical considerations might different kinds of context-sensitive items raise for my model.

Second, while I have argued for a linguistic model of contextual prominence, maintained through the effects of discourse conventions, I have not said much

¹ 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., ϕ -features of a pronoun). The character of prominence *insensitive* expressions, like 'I', does not reference prominence in this way. (Though this 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.) 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. 'I', by contrast, does not allow for such readings. Interestingly, 'here' and 'now', though often classified as pure indexicals, 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'). For an account of the prominence-sensitivity of 'now' along the lines of the theory I defended in this book, see Stojnić and Altshuler (2020).

about the role of prominence, and the mechanisms governing it, in the resolution a prominence-sensitive expressions that occur within an utterance that does not follow upon a prior discourse. So, for instance, one might wonder how a demonstrative pronoun receives a resolution when the utterance containing it is not a continuation of a prior discourse, and when a pointing gesture is absent. In what follows, I will address these issues, tackling them in reverse order. That is, I first sketch an account of discourse initial uses of prominence-sensitive expressions, and then discuss extensions of the model developed in this book to other cases of context-sensitivity, and specific challenges and insights that they might bring about.

Let us begin, again, with the basic case of a demonstrative pronoun.² Demonstrative pronouns, we have seen, can be interpreted as co-varying with a prior linguistic expression, or as selecting an object introduced by a demonstrative gesture, as when I point at an object, saying: ‘That is lovely.’ We took this as evidence that demonstrative pronouns require an antecedently set context, one in which a particular interpretation has been made prominent—prominent, I have argued, as a result of an effect of a discourse convention, such as a coherence relation, or a pointing gesture. And indeed, it is a familiar observation that demonstrative pronouns are limited in terms of their acceptability out of the blue. If I just state: ‘She is happy,’ out of the blue, with no proper setup, there is a sense in which my utterance has diminished felicity. This is not surprising—the audience has no idea whom I might be talking about.³

However, this is not to say that a demonstrative pronoun can never be felicitously used discourse initially. For instance, Stojnić, Stone, and Lepore (2018) describe a scenario, based on an actual case from Julie Child’s cooking show, *The French Chef*, where at a particular instance there is a close up of the skillet from above, an omelette’s forming inside, and Child’s voice is heard saying:⁴

² I have defended an account of deictic uses of demonstrative pronouns along the lines sketched here in Stojnić, Stone, and Lepore (2018). Here I expand upon that discussion, and argue for its extension to further cases of discourse initial uses of context-sensitive expressions.

³ The requirement that the context should provide a specific antecedent value for an expression is sometimes called ‘strong contextual felicity constraint’ (Tonhauser et al., 2013). Note that one can sometimes exploit this constraint to dramatic effect, as when a book opens with a pronoun, only revealing the referent several pages on.

⁴ You can see the relevant episode of the French Chef here: <https://www.youtube.com/watch?v=4RoLavF2ZLU> The actual utterance Child performs (occurring at 5:19) is:

(i) There’s your omelet, turning over on itself, forming itself in the bottom of the pan.

The slightly idealized (121) is more convenient to work with since it isolates the problem we are dealing with directly, that of resolution of (discourse initial) demonstrative pronouns unaccompanied by pointing gestures.

(121) That's your omelet.

(121) is a *situated utterance*, used to describe or comment on a specific real-world situation. Understanding Child requires recognizing the fact that she's describing the ongoing events that we see on the screen as she speaks. Part of the interpretive task in determining the meaning of (121) is understanding what 'that' refers to. And, clearly, (121) doesn't require a pointing gesture in order to be interpreted: we readily understand it, just by looking at the screen. And one might think, this is precisely where the kind of pragmatic open-ended reasoning about speaker's intentions drives the interpretation. The omelet is the referent of 'that' because that interpretation makes most sense: after all, surely, she isn't saying that the pan handle or the spatula she is using to flip the egg-mixture is an omelet. But this initial semblance of open-ended reasoning is misleading. To see this, note that the key task in interpreting an utterance like (121) is understanding what the utterance is doing—what kind of a contribution it is making. It matters that you are being presented with a *how-to* cooking video in which a particular recipe for making an omelet is being described. The utterance occurs embedded within this task and is understood as making a comment about that perceptually present situation, summarizing what's visible on the screen at that particular point in time.

Drawing on this observation, Stojnić, Stone, and Lepore (2018) argue that interpreting (121) involves recognizing a coherence relation connecting the utterance of (121) to the perceptually present, real-world situation it is about. Specifically, (121) connects to the ongoing situation on the screen via *Summary* relation. *Summary* relation will hold between an utterance and the situation visible on the screen, only if the description 'omelet' fits the relevant object visible there. In addition, Stojnić, Stone, and Lepore (2018) argue that *Summary* makes the central entity in the situation described prominent.⁵ The demonstrative 'that' then, as always, selects this entity as its referent. The key point of departure from the account that posits open-ended reasoning as a key interpretive mechanism is that in interpreting (121) one is not presented with two independent tasks—establishing a coherent connection with an ongoing perceptually present situation, and, in addition, resolving a demonstrative pronoun. As soon as one recognizes that the speaker is making a comment about that situation, one knows that the referent of 'that' is the most prominent entity in that situation. To this end, it is perhaps worth pointing out that, while in the scenario the audience

⁵ Stojnić, Stone, and Lepore (2018) provide a formal extension of the framework we used in Chapters 4 and 5 to accommodate situated utterances, and specifically to accommodate the idea that a coherence relation can hold not just between utterances in a discourse, but also between the utterances and real-world situations those utterances are about.

is aware that the description ‘omelet’ fits a specific entity on the screen, even if one uttered an unfamiliar description expression, one would still resolve ‘that’ to the food in the pan, not a pan handle or the stovetop. This is partly because how the situation is described itself guides its conceptualization. As mentioned earlier it matters that we are being instructed on how to cook an omelet.⁶ That’s what the video is about: the situation visible at one point on the screen, that the specific comment in (121) is about, is embedded within this larger narrative, imposed by the task at hand. This is why the omelet rather than the stovetop, or the pan-handle is prominent. And this is why it would be somewhat bizarre to utter ‘That’s a pan handle’ at this point, even though, strictly speaking, there is a pan handle on the screen, and you can see it clearly.⁷

Of course, since not all perceptually present situations we can comment on come within neatly narrated sequences, as in a TV show, there can be significant ambiguity as to which situation the speaker is making a comment about, and, moreover, what type of a relation her utterance bears to it, just as there can be ambiguity as to what relation an utterance bears to the prior discourse, and which prior bit(s) of discourse it relates to.⁸ But the point is, again, that one isn’t presented with two independent tasks: understanding the relation and resolving the pronoun. Rather, fixing the relation fixes the interpretation of the pronoun.

To see this more clearly, it is useful to consider the following example, from Stojnić, Stone, and Lepore (2018). Suppose there is a conjunction of Venus and Jupiter in the sky: you cannot tell them apart by the naked eye. The speaker has set a telescope facing a window over the western sky. She adjusts the telescope, and says to you, without any further demonstration:

(122) That’s Jupiter. You can even see four moons.

Given the astronomical conjunction, one cannot know, without looking through the telescope, which of the two objects the speaker is referring to—indeed, the speaker herself cannot by her naked eye tell the two objects apart; yet once the audience recognizes that the speaker is making a summary of a scene viewed

⁶ See Stone and Stojnić (2015) for some discussion of coherence constraints relating utterances not merely to the discourse and real-world situations, but also the ongoing activity.

⁷ Of course, a reference to a pan handle could be perfectly appropriate had the situation been a part of an advertisement for non-stick pens with super strong handles. In that case, saying ‘That’s an omelet’ while shaking the pan would be bizarre. And that’s precisely the point.

⁸ Consider the M&M “Trailer” Public Service Announcement, in which an agent hands out a “classified” envelope to Yellow and Red M&M saying ‘This is your mission,’ to which Yellow responds by holding up the envelope and saying ‘No, this is an envelope.’ Here, the ambiguity in understanding the utterance as a summary of the content of the envelope, or as one about the situation in which an object is handed over is exploited for a comedic effect (https://www.youtube.com/watch?v=j_eabL16a5w).

through the telescope, there is no further ambiguity of the pronoun left to resolve. The referent is whatever the most prominent entity in the situation seen through the telescope is. Once it is recognized that Summary relation relates the utterance to the situation seen through the telescope, one is not left with a further task of disambiguating the pronoun. Summary makes whatever entity is central in the situation seen through the lenses of the telescope prominent, and that's what 'that' selects as its referent. And at the same time, unless the audience recognizes that the speaker is summarizing this situation, they will have failed to understand the utterance.

Of course, as before, there can be ambiguity in which situation the speaker's utterance relates to. If there were a microscope set up next to the telescope, and the speaker uttered (123), instead of (122), it might be ambiguous whether my utterance is a summary of the situation seen through the microscope, or the one seen through the telescope. But once again, resolving *this* ambiguity automatically settles the resolution of the pronoun as well.

(123) That is α -complex. It has just been discovered.

Note, moreover, that in order for the speaker to successfully refer to Jupiter with (122), and for you to understand this utterance, it is not necessary for either you or the speaker to actually look through the telescope. You understand that (122) has a precise truth-condition, and you understand what it would take to identify the referent: you would have to look through the telescope.

This account extends to discourse initial uses of other types of context-sensitive expressions. For instance, an utterance of a modal expression can be related to a real-world situation, or a specific perspective or body of information it makes prominent. And, just as the interlocutors don't know what 'that' in (122) is until they look through the telescope—even though they perfectly well know that whatever 'that' picks out is the most prominent entity seen through the telescope—they might not know, or might even disagree about what the prominent possibility is in a given situation. Consider a case, from von Stechow and Gillies (2007), where two agents are searching for the missing keys, and the following dialogue ensues:

- (124) a. A: The keys might be in the drawer.
b. B: No, they can't be! I looked there already.

Given the exchange, A could react to B's remark by saying either (125), or (126):

(125) A: I guess I was wrong. Let's look elsewhere then.

(126) A: I only said they *might have* been. And they might have.

It's natural to explain this as an ambiguity in the prominent body of information A's utterance is commenting on. That is, there might be ambiguity about which body of information she has made the comment about. The disagreement can result from the fact that the speaker thinks that the prominent possibility is the one that involves what the evidence suggests, or what is compatible with the knowledge of those who have examined the evidence, while the hearer might think that it involves what the agents best informed about the case know. And even if it is settled that the prominent possibility is, say, the one that is compatible with the information of those who have examined the evidence available in the relevant situation, the agents might still disagree about which information this is, since they might have different conceptions of what the totality of the available evidence is, or who the relevant agents that have examined it are.

However, often, especially in the case of epistemic modal utterances, it might not matter for our practical purposes to exactly identify a specific resolution, as it typically matters in the case of a pronoun, like 'that' in (121): if I ask you where the keys are and you say "They might be in the drawer," if I don't know whether they are in the drawer, I might not care whether the prominent possibility that your 'might' statement quantifies over is the one involving my information or only yours. My primary goal is to find the keys, and your utterance suffices to help my search in this case, regardless of which possibility is selected. I'll look in the drawer all the same.⁹

It is worth pointing out that, while pronouns are associated with strong contextual felicity constraint, modals do not seem to be subject to this constraint. Indeed, it is perfectly felicitous to begin a discourse with a modal utterance. One might wonder what explains such discrepancy on the present account. Our previous considerations point to an answer to this question as well. In uttering a modal out of the blue, while there will normally be ambiguity of the sort described as to which body of information is prominent, resolving this ambiguity might typically be unimportant for the main purpose of the utterance: if all I want is to find the keys, and I haven't ruled out the drawer, my understanding of the utterance will suffice, even if I haven't fully identified which body of information is prominent.¹⁰ By contrast, often, it will matter for

⁹ Thanks to Matthew Stone and Gala Stojnić (p.c.) for discussion concerning these types of example.

¹⁰ This is a good point to address a further worry Yalcin (2007) raises for the contextualist treatment of epistemic modality. Yalcin argues that contextualist accounts of modality give intuitively wrong predictions about the subject matter for belief and assertion in assigning propositional content to epistemic modal claims. If Mary believes that Betty might be in the library, a contextualist—Yalcin argues—would ascribe Mary a second order belief about her information state, not a first order mental state that is about Betty. That is, her belief would end up being about what is compatible with her beliefs, and not about Betty. But this is too quick: first of all, the contextualist account of the sort I defended does not maintain that modal propositions are about the speaker's or conversational

the main point of the utterance resolving who the pronoun refers to.¹¹ If my goal is to inform you that a certain person is happy, by saying ‘She is happy,’ I won’t succeed unless you actually identify the person in question. Similarly, if my goal is to get you to give me a specific object, it matters a great deal that you can identify what ‘that’ refers to when I say ‘give me that.’

In sum, the account I defended has resources to accommodate discourse initial uses of context-sensitive expressions, and does so by appeal to the same discourse conventions that govern context-sensitivity resolution more generally. The account predicts not that context-sensitive expressions cannot be used when not following on a prior discourse, but that they can be successfully used in this way when linguistic mechanisms connect the utterance to some real-world situation or an ongoing activity in a coherent way, thereby setting the relevant parameters of the contextual scoreboard.

Now, let us turn to potential extensions of the framework to other kinds of context-sensitive items. While our discussion largely focused on modals and pronouns, we should expect the same kinds of discourse conventions to affect resolution of context-sensitivity more generally. A complete defense of this position would require looking at each case independently, and so our discussion must remain provisional in certain respects. However, there are good reasons to think that the resolution of context-sensitive items quite generally is driven by discourse-internal conventions.

Given our discussion of modal expressions, it might not come as a great surprise that we see prominence-sensitivity in nominal and adverbial quantifiers as well. Like modal expressions, nominal and adverbial quantifiers require a contextually supplied restriction, and this restriction can be provided either by the real-world situation or by the prior discourse. So, for instance, just as one can use a pronoun to refer to a prominent individual in a non-linguistic

participants’—or indeed, anybody else’s—mental states. A modal quantifies over the prominent body of information, whatever that might be in the given context; further, this body of information is provided by the context directly. So, where *i* is the prominent body of information that restricts the relevant occurrence of the modal, the content is simply *that given i Betty being in the library is not ruled out*. The speaker might believe this without knowing which body of information is prominent, and without being able to pick it out in any way more substantive than simply “the prominent body of information.” They might even have further false beliefs about what *i* is. (Just as I might believe, upon hearing someone speak, that they—the prominent individual—have a deep voice, even if I have false beliefs about who they are, perhaps falsely believing that they are my friend, John.) They simply believe that the prominent body of information (whatever it is) leaves this possibility open. Nor is this belief cognitively complex: it just expresses that the relevant information (whichever it is) does not rule out that Betty is in the library. One can reason about this without having any beliefs about mental states at all. Intuitively, the content expressed is about Betty’s whereabouts, as far as they are revealed by some evidence. Thanks to an anonymous reader for the OUP for pressing this point.

¹¹ Often, but not always. Recall (122).

context, so too one can use (127) in a context to talk about a prominent group of individuals.¹²

- (127) (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:

- (128) Many students came to the party. Some brought cookies.

Here the second quantifier is understood as quantifying over a subdomain of the first one, thus receiving a more restricted reading—some *of the students who came to the party* brought cookies. In such cases, prominence of potential restrictors on the domain of quantification is governed by mechanisms of discourse coherence, just as the resolution of pronominal and modal anaphora is. The first sentence describes a situation about the group of students at a particular party. The second one 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 (128) to mean that many of the students came to the party, and some of those students who came to the party brought cookies. While this example might seem uninteresting, as before, the import of coherence seems crucial:

- (129) All students graduated eighth grade. Some students didn’t graduate eighth grade.

It is hard to understand (129) as anything but contradictory, even if we know that there are two cohorts of students, one from the seventh and one from eighth grade, both of which the speaker has been teaching for the past year. But with such background knowledge it would clearly make the most sense that ‘some students’ in the last sentence quantifies over the set of students in the lower grade. Yet, this is not the interpretation we find.

We see similar kinds of coherence constraints with adverbial quantifiers as well. Consider the following example:¹³

- (130) If there is a fire, the alarm always rings.

¹² Whether the domain restriction comes in terms of sets, or functions from objects to sets, or properties or perhaps situations is a matter of debate (von Stechow, 1994; Stanley and Szabó, 2000). The details do not matter for our purposes.

¹³ The example is from (Yalcin, 2012b), though he uses it for different purposes, and does not exploit it in the context of domain restriction.

(131) But the alarm doesn't always ring.

In (130), the consequent of the conditional elaborates 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 (131), however, 'always' in the second sentence is *not* thusly restricted—we cannot conclude from (130) and (131) 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 (130) and (131). 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. 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 there is fire. So (130) conveys 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:

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

(133) The firefighters usually arrive within five minutes.

Superficially, (132) and (133) seem to have exactly the same shape as (130) and (131). However, they are not naturally understood as standing in the Contrast relation. Rather, (133) is understood as Elaborating on the situation described by the consequent of (132) (which in turn is itself understood as elaborating on the antecedent). Elaboration makes the situation described by the consequent—the one in which there is fire, and the alarm rings—prominent, and the quantifier in (133) 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.

We can even generate examples akin to the problematic example we have seen exhibited by modal discourse. For instance, consider the following:

(134) Sometimes it's raining and it's not raining.

(135) Suppose sometimes it's raining and it's not raining.

(134) and (135) are both odd or awkward. This is as expected if we understand *and* to signal an organization around Elaboration relation. In that case, the

second conjunct is understood as elaborating further on the situation described by the first, and hence we witness infelicity.¹⁴ Of course, if we force a contrastive interpretation, the infelicity disappears:

(136) Suppose sometimes it's raining but it's not raining *now*.

Similar behavior is exhibited by other context-sensitive expressions. For instance, we see this with the so called incomplete definite descriptions:

(137) 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*).¹⁵ But even though there are many tables in the world, we can imagine (137) uttered in a room with exactly one table, successfully conveying that the table in the room is covered with books.¹⁶ So, 'the table' in (137) 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. As before, the restriction can be provided by the non-linguistic context, as in (137), or by the prior discourse, as in (138):

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

Once again, prominence appears to be affected by discourse conventions. In (138), we understand the second sentence as elaborating on the situation described by the first. Similarly as 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

¹⁴ Notice that, here, the order clearly improves the interpretation:

(i) Suppose it is not raining and sometimes it is raining.

This is because what's driving the infelicity is the preference for the narrow scope of 'it's not raining,' in the original example, which is in line with Elaboration, but not with Contrast. Such narrow scope is not possible in (i), resulting in improved interpretation.

¹⁵ This observation goes back to at least Frege (1892) and Russell (1905).

¹⁶ As Roberts (2011) points out, even a weak requirement that the referent be unique merely within the common ground won't help us get around the problem. A typical common ground will normally entail the existence of many tables.

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.¹⁷

As with other expressions we considered, here, too, coherence effects cannot be overridden, even in the face of infelicity:

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

In (139), 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, and the bedroom as providing background circumstances to the event of entering the bedroom, and thus as restricted to the situation describing the bedroom.¹⁸ Since the descriptions are incompatible—there can be no window in a windowless room—(146-b) is infelicitous.

Next, consider gradable adjectives, like ‘tall.’ Gradable adjectives are typically treated as expressing relations between individuals and degrees, where an individual satisfies the property just in case it possesses the property to at least the contextually specified degree.¹⁹ Thus, gradable adjectives are looking for a contextually prominent degree or threshold. And, examples (140)–(141) suggest that the mechanisms of discourse coherence affect the possible values of this contextual parameter, making certain resolutions prominent.

- (140) John is tall! His height is 2m.

- (141) We should put John on our basketball team. He’s tall.

In (140), the second sentence is related to the first by Specification, providing more detail about how one should understand what counts as tall.²⁰ In particular, the second sentence makes it clear that being two meters tall suffices to count as tall in the given context. In (141), the organizing relation is Explanation.

¹⁷ There is a long-standing tradition of treating domain restriction by exploiting possible situations (Barwise and Etchemendy, 1987; Recanati, 1996; Kratzer, 1998, 2014; 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 Stechow, 1994; Stanley and Szabo, 2000).

¹⁸ Background relation is similar to Elaboration. While the latter provides more information about an eventuality, the former describes circumstances that lead to an event.

¹⁹ See, for example, Hamann (1991); Klein (1991); Barker (2002); Kennedy (2007). We need not concern ourselves with specific formal details of the semantics of gradable adjectives for our present purposes.

²⁰ For a similar point involving Specification, see Stone and Stojnić (2015).

It likewise specifies the relevant degree of height—John has to be at least as tall as an average basketball player.

The same point extends to relational expressions, like ‘local’, or ‘friend’ that require a contextually specified ‘implicit argument.’ An utterance of (142) can convey that John went to a bar local to him, or local to some other, contextually specified, person or place. Relational expressions require an implicit argument prominent in the non-linguistic context or in the prior discourse.

(142) John went to a local bar.

Once again, we witness sensitivity to coherence relations:

(143) John was in Paris last week. He visited a nice local restaurant.

Here, ‘local’ occurs within a Narrative about the Paris trip, and is thus understood as *local-to-Paris*. But, as with our other examples, ‘local’ likewise cannot just freely be used to denote a location local to just any place or person even if that place or person is made salient in the context, and even if that would lead to an overall more plausible interpretation. Consider the following:

(144) John was in Paris in May, but then he came to Belgrade. He checked into a hotel, and proceeded to explore the city. He particularly enjoyed a local restaurant near the Eiffel Tower.

Clearly, resolving ‘local’ in such a way so as to mean *local-to-Paris* would make the most sense in (144), 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 (144) as a part of the chain of elaborations about the Belgrade trip, ‘local’ is interpreted as *local-to-Belgrade*. Assuming that prominence of a particular resolution of ‘local’ is dictated by discourse coherence, we can predict this straightforwardly: the final sentence in (144) 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. The same kinds of considerations extend to other kinds of examples we have considered so far.

These examples illustrate that discourse conventions 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 cases—pronouns, quantifiers, modal expressions, definite descriptions, gradable adjectives, and relational expressions. This

suggests that these kinds of effects on prominence are not isolated to a small sample of expressions, but are rather general and pervasive. And, more importantly, in all such cases, as we have seen, the effects seem to cross-cut common-sense reasoning about what makes the most sense.

This is not to say that specific examples might not require specific treatment. For instance, a full-blown treatment of nominal domain restriction, definite description or gradable adjectives might require allowing for shifts in prominence of a specific value of the relevant contextual parameters to occur at a sub-clausal level. For instance, consider (145), from Stanley and Szabó (2000), which can be uttered to describe a scenario in which every sailor on the ship waves at every sailor on the shore:

(145) Every sailor waived at every sailor.

Assuming that ‘every sailor’ quantifies over a set of sailors in a contextually provided situation, in order to predict the relevant reading of (145), one would have to allow the situation parameter to shift mid-clause, between the two occurrences of the quantifier phrase. It is clear that such readings do involve similar coherence constraints as those that are operative between distinct clauses: the relevant interpretation draws on the parallelism of the two situations, one describing the sailors on the ship, and the other those on the shore. Note that it is harder to get this reading if we force an interpretation to be organized around a single situation—(146-b) is better than (146-a):

(146) Four students were alone in the classroom:

- a. Every student smiled at every student.
- b. The students smiled at one another.

Of course, in order to model such examples properly, we need to say more about coherence constraints at sub-clausal level.²¹ This is an important task for future research.

Summing up, I have argued the interpretation of context-sensitive items relies on discourse-internal cues—discourse conventions—and the interpreters consequently have to rely on knowledge of these linguistic conventions, not reducible to common-sense inference that draws holistically on the epistemic cues reflecting speaker intentions. My model of context and context-sensitivity resolution constrains and informs the appeal to context-sensitivity in philosophical theorizing in ways quite different than what has been customary.

²¹ For some discussion of, and motivation behind, clause-internal coherence, see Hobbs (2010).

We have seen how the traditional assumption of the non-linguistic model of context-sensitivity resolution can lead to philosophically radical conclusions. In particular, we have traced the prominent, yet radical, conclusions about the nature of semantic content, its relation to context, and its underlying logic to this assumption of the underlying non-linguistic, flexible model of context-sensitivity. This suggests that an appeal to context in philosophical arguments has to proceed with caution: such arguments must take into account the underlying conventions of language that play the key role in fixing the meaning in context.

APPENDICES

APPENDIX A

Formal Definitions for the Attention-Coherence Approach to Pronoun Resolution

Appendix: Formal definitions

This appendix provides complete definitions for a logical language that formalizes 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 .

A.1 Models

The interpretation of constants is set up in terms of frames and models in the usual way.

- **A Frame** is a tuple $\mathcal{F} = \langle D_w, R, D_e, D_t \rangle$ 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 $\mathcal{M} = \langle \mathcal{F}, \mathcal{I} \rangle$, where $\mathcal{F}$ is a frame and $\mathcal{I}$ 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 $\mathcal{D}_e$.

A.1.1 Operations on Stacks

The Attention-Coherence approach captures interpretive dependencies across formulas using dynamic semantics. Operations on sequences of individuals from the model (that is, stacks) play a key role in the semantics. A sequence is a function defined on a finite convex subset of $\mathbb{N}$, mapping onto a set of entities defined by a Frame $\mathcal{F}$ together with an undefined value $\perp$. Given the entities defined by a Frame $\mathcal{F}$, we use the following notation to specify these operations.

- i_m
If $m \in \mathbb{N}$, and i is a stack, then i_m is the m th element of i if i is defined on m , otherwise $i_m = \perp$.

- $i_{m,n}$
If $m, n \in \mathbb{N}$, and i is a stack, then $i_{m,n}$ is a stack j , defined on the set $\{0, \dots, ((n - m) - 1)\}$, such that for $k \in \mathbb{N}$, $j_k = i_{(m+k)}$ if j is defined on k .
- $i_{m\dots}$
If $m \in \mathbb{N}$, and i is a stack, then $i_{m\dots}$ is a stack j defined on a set $\{k \in \mathbb{N} \mid i \text{ is defined on } m + k\}$ such that for $k \in \mathbb{N}$, $j_k = i_{(m+k)}$.
- $i + j$
If i is a stack with a finite domain of definition with maximal element $k - 1$, then for a sequence j , $i + j$ is a sequence h , such that for $x \in \mathbb{N}$, $h_x = i_x$, if i is defined at x , and $h_x = j_{(x-k)}$ otherwise.
Note then that $i = i_{0,k} + i_{k\dots}$.
- $u.i$
If $u \in D_t \cup D_w \cup D_e$, and i is a stack, then $u.i$ is the sequence 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.
- $w(i)$
If i is a sequence, then $w(i)$ is the first element u of i such that $u \in D_w$ (if any).

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 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 .
(Constants from model.)
- $\llbracket x_m \rrbracket i, w = i_m$.
(Variables. x_0 will correspond to the subject, x_1 to the direct object, x_2 to the indirect object, and so forth.)
- $\llbracket @p \rrbracket i, w = \perp$ if $i_0 = \perp$, $\llbracket @p \rrbracket i, w = i_0$, if $\langle w, i_0 \rangle \in \mathcal{I}(p)$, and $\llbracket @p \rrbracket i, w = \llbracket @p \rrbracket i_{1\dots}, w$ otherwise.
- $\llbracket @p^t \rrbracket i, w = \perp$ if $i_0 = \perp$, $\llbracket @p \rrbracket i, w = i_0$, if $\langle w, i_0 \rangle \in \mathcal{I}(p)$ and $i_0 \neq \llbracket t \rrbracket i, w$, and $\llbracket @p^t \rrbracket i, w = \llbracket @p^t \rrbracket i_{1\dots}, w$ otherwise.

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 $\langle w, \langle \llbracket t_1 \rrbracket i, w, \dots, \llbracket t_n \rrbracket i, w \rangle \rangle \in I(r)$
(Observe that the condition is false if the interpretation of constituent terms is undefined)

Dynamic updates:

- $\langle \alpha k \rangle$ is an update, for $k \in \mathbb{N}$.
(push new indefinite assignment for variable x_k)
- $\langle \pi kt \rangle$ 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.
(sequencing—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 $j = i_{0,k} + o.i_{k\dots}$ 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 $j = i_{0,k} + o.i_{k\dots}$ 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 $\llbracket K \rrbracket (v, i, k)$.

A.2 Truth, Validity, and Entailment

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 .

A.3 Worked Out Examples

- A woman met Sue. She greeted her.

formula	gloss	output
$\langle \alpha 0 \rangle$; [woman (x_0)];	"A woman (is the subject)"	(w , ...) where w is a woman
$\langle \pi 1 s \rangle$;	"Sue (is the object)"	(w , s , ...)
[met (x_0 , x_1)];	"(the subject) met (the object)"	(w , s , ...) where w met s
$\langle \pi 0 @ \mathbf{she} \rangle$;	"She (is the subject)"	(w , w , s , ...) since w is a she
$\langle \pi 1 @ \mathbf{she}^{x_0} \rangle$;	"her (is the object)"	(w , s , w , s , ...) since w , s are she but $w = x_0$
[greeted (x_0 , x_1)]	"(the subject) greeted (the object)"	(w , s , w , s , ...) where w 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
$\langle \pi 0 j \rangle$;	"John (is the subject)"	(j , ...)
$\langle \pi 1 t \rangle$;	"Tim (is the object)"	(j , t , ...)
[disapp.with (x_0 , x_1)];	"(the subject) was disappointed with (the object)"	(j , t , ...) where j was disappointed with t
[Result (x_0 , x_1); $\langle \pi 0 x_0 \rangle$;	"describe the result of the subject's disappointment"	(j , j , t , ...) where j was disappointed by t
$\langle \pi 0 @ \mathbf{he} \rangle$;	"He (is the subject)"	(j , j , j , t , ...) since j is a he
$\langle \pi 1 @ \mathbf{he}^{x_0} \rangle$;	"him (is the object)"	(j , t , j , j , t , ...) since j , t are he , but $j = x_0$
[fired (x_0 , x_1)]	"(the subject) fired (the object)"	(j , t , j , j , t , ...) where j fired t

- John was disappointed with Tim. He disobeyed him.

formula	gloss	output
$\langle \pi 0 j \rangle$;	"John (is the subject)"	(j , ...)
$\langle \pi 1 t \rangle$;	"Tim (is the object)"	(j , t , ...)
[disapp.with (x_0 , x_1)];	"(the subject) was disappointed with (the object)"	(j , t , ...) where j was disappointed with t
[Explanation (x_0 , x_1); $\langle \pi 0 x_1 \rangle$;	"explain appealing to the object"	(t , j , t , ...) where j was disappointed by t
$\langle \pi 0 @ \mathbf{he} \rangle$;	"He (is the subject)"	(t , t , j , j , ...) since t is a he
$\langle \pi \mathbf{he} \rangle$;	"him (is the object)"	(t , w , t , j , t , ...) since j , t are he , but $t = x_0$
[disobeyed (x_0 , x_1)];	"(the subject) disobeyed (the object)"	(t , j , t , j , t , ...) where t disobeyed j .

We can reduce the difference in the minimal pair to an attentional shift associated with two different coherence relations, Result and Explanation, represented explicitly in logical form.

APPENDIX B

A Formal Language for Modality with Coherence

This appendix provides a dynamic language that formalizes my approach to modality with anaphora, and a translation from a fragment of English into this language.

B.1 Syntax

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 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 an 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.
- $K; K_1$ is an update expression, if K is an update expression and K_1 is an update expression.
- $\text{MIGHT}(\phi, K)$ is an update expression, if ϕ is a condition and K an update expression.

- $\text{MUST}(\phi, K)$ is an update expression, if ϕ is a condition and K an update expression.
- $\text{LIKELY}(\phi, K)$ is an update expression, if ϕ is a condition and K an update expression.
- $\text{IF}(\phi, K_1, K_2)$ is an update expression, if ϕ is a condition and K_1 and K_2 are update expressions.
- $\text{AND}(K_1, K_2)$ is an update expression, if K_1 and K_2 are update expressions.
- $\text{NOT}(K)$ is an update expression, if K an update expression.
- $\text{ASSERT}(K)$ is an update expression, if K is an update expression.
- $\text{ELAB}(\phi, K)$ is an update expression, if ϕ is a condition and K an update expression.
- $\text{CONTRAST}(K_1, K_2)$ is an update expression, if K_1 and K_2 are update expressions.

B.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:
 - $\mathcal{P}(\mathcal{D}_w) = 1$.
 - $\mathcal{P}(p \cup q) = \mathcal{P}(p) + \mathcal{P}(q)$, when p and q are disjoint subsets of $\mathcal{D}_w$.
 - $\mathcal{P}$ is a regular probability measure: if $p \neq \emptyset$ then $\mathcal{P}(p) > 0$.
- A **Model** is a pair $\mathcal{M} = \langle \mathcal{F}, \mathcal{I} \rangle$, where $\mathcal{F}$ is a frame and $\mathcal{I}$ 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.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 Definition 9.1 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.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.3

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

Definition B.4 Definition 9.2 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.2.2 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 a function defined on a finite convex subset of $\mathbb{N}$ plus *comp* mapping to a set of worlds plus $\perp$, 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.)

- 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 definition of i , and $i_m = \perp$ otherwise. (i_{comp} is the member of the stack stored at the designated position *comp*.)
 - 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 $m = comp$.
- 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 , $i_m \dots$ 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, $G_{m \dots} = \bigcup_{g \in G} \{j \mid j = g_{m \dots}\}$ and for $H = G_{m \dots}$, $H_{comp} = G_{comp}$.
- If i is a stack with a finite domain of definition 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 $h_{comp} = i_{comp}$).
- 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 $j_{comp} = i_{comp}$.

¹ A 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 G is a context, u is a world, and g, j are stacks, $G_{u\dots} = \bigcup_{g \in G} \{j | j = u, g\}$ and for $H = G_{u\dots}$, $H_{comp} = G_{comp}$.
- $g[n]g'$ 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 $\{g_{0,n} + g_{n+1\dots} | g \in G'\} = G$ and $G_{comp} = G'_{comp}$.

B.3 Semantics

The Interpretation of Atoms: We interpret an expression e , relative to the interpretation function $\mathcal{I}$ a context G , and a world w :

- $\llbracket p \rrbracket^{G,w} = \mathcal{I}(p)$, 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 $G_0 \in \mathcal{I}(P)$, and $\llbracket @P \rrbracket^{G,w} = \llbracket @P \rrbracket^{G_{1\dots},w}$ 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 $G \sim_{comp} H$ & $H_{comp} = \llbracket p \rrbracket^{G,w}$
- $\llbracket [\phi] \rrbracket(w, G, H)$ if and only if $H = G$ and $w \in \llbracket \phi \rrbracket^{G,w}$
- $\llbracket [K; K_1] \rrbracket(w, G, H)$ iff $\exists G' : \llbracket K \rrbracket(w, G, G')$ and $\llbracket K_1 \rrbracket(w, G', H)$

The following are updates that describe how propositional content (§B.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 $\llbracket K \rrbracket(w, G, G') \& G' \approx_0 G''$ & $G'_0 = G'_{comp} \cap \llbracket @E \rrbracket^{G,w} \& G'' \sim_{comp} H \& H_{comp} = M(\llbracket \phi \rrbracket^{G,w}, G'_{comp})$
- $\llbracket \text{MUST}(\phi, K) \rrbracket(w, G, H)$ iff there is a G' and G'' such that $\llbracket K \rrbracket(w, G, G') \& G' \approx_0 G''$ & $G'_0 = G'_{comp} \cap \llbracket @E \rrbracket^{G,w} \& G'' \sim_{comp} H \& H_{comp} = N(\llbracket \phi \rrbracket^{G,w}, G'_{comp})$
- $\llbracket \text{LIKELY}(\phi, K) \rrbracket(w, G, H)$ if and only if there is a G' and G'' such that $\llbracket K \rrbracket(w, G, G') \& G' \approx_0 G''$ & $G'_0 = G'_{comp} \cap \llbracket @E \rrbracket^{G,w} \& G'' \sim_{comp} H \& H_{comp} = P(\llbracket \phi \rrbracket^{G,w}, G'_{comp})$
- $\llbracket \text{IF}(\phi, K_1, K_2) \rrbracket(w, G, H)$ iff there is a G', G'', G''' and G'''' such that $\llbracket K_1 \rrbracket(w, G, G') \& G' \approx_0 G''$ & $G'_0 = G'_{comp} \cap \llbracket @E \rrbracket^{G,w} \& \llbracket K_2 \rrbracket(w, G'', G''') \& G''' \approx_0 G''''$ & $G''''_0 = G'''_{comp} \cap \llbracket @E \rrbracket^{G'',w} \& G'''' \sim_{comp} H \& H_{comp} = \text{Cond}(\llbracket \phi \rrbracket^{G,w}, G'_{comp}, G'''_{comp})$
- $\llbracket \text{AND}(K_1, K_2) \rrbracket(w, G, H)$ iff there is a G', G'', G''' and G'''' such that $\llbracket K_1 \rrbracket(w, G, G') \& G' \approx_0 G''$ & $G'_0 = G'_{comp} \cap \llbracket @E \rrbracket^{G,w} \& \llbracket K_2 \rrbracket(w, G'', G''') \& G''' \approx_0 G''''$ & $G''''_0 = G'''_{comp} \cap \llbracket @E \rrbracket^{G'',w} \& G'''' \sim_{arg_0} H \& H_{comp} = G'_{comp} \cap G'''_{comp}$
- $\llbracket \text{NOT}(K) \rrbracket(w, G, H)$ iff there is a G' such that $\llbracket K \rrbracket(w, G, G') \& G' \sim_{comp} H \& H_{comp} = \llbracket \neg comp \rrbracket^{G',w}$
- $\llbracket \text{ASSERT}(K) \rrbracket(w, G, H)$ iff there is a G' such that $\llbracket K \rrbracket(w, G, G') \& G' \approx_0 H \& H_0 = G'_{comp} \cap \llbracket @E \rrbracket^{G,w} \& w \in H_0$

In order to define the truth-conditions for updates associated with coherence relations, we assume the following abbreviations:

Definition B.5 *Elab*(ϕ, ψ) iff ϕ and ψ are centered around the same event or entity, that is, iff the event or scenario described by ψ is a part of the event or scenario described by ϕ .

Definition B.6 A formula, ϕ , is about 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.7 *Contrast*(ϕ, ψ) iff ϕ and ψ describe contrasting information about some body of information regarding a common topic.

² For generality, I let the restrictor in the definition be any proposition. However, as argued above, epistemic modals and conditionals select the top-ranked possibility in a given context (@E) as their restrictor.

Definition B.8 *Conclusion*(ϕ, ψ) iff $\llbracket \phi \rrbracket^{w,G} \subseteq \llbracket \psi \rrbracket^{w,G}$ for all G, w .

- $\llbracket \text{ELAB}(\phi, K) \rrbracket(w, G, H)$ iff there are G' and G'' such that $G \approx_0 G'$ & $G'_0 = \llbracket \phi \rrbracket^{G,w}$ & $\llbracket K \rrbracket(w, G', G'') \& G'' \approx_0 H$ & $H_0 = G''_{\text{comp}} \& \text{Elab}(\llbracket \phi \rrbracket^{G,w}, H_0)$.
- $\llbracket \text{CONTRAST}(K_1, K_2) \rrbracket(w, G, H)$ iff there is a G' and G'' such that $\llbracket K_1 \rrbracket(w, G, G')$ & $G' \approx_0 G''$ & $G'_0 = \llbracket \theta_{K_1} \rrbracket^{G,w}$ & $\llbracket K_2 \rrbracket(w, G'', H)$ & $\llbracket \theta_{K_1} \rrbracket^{G,w} = \llbracket \theta_{K_2} \rrbracket^{G'',w}$ & $\text{Contrast}(G''_{\text{comp}}, H_{\text{comp}})$.
- $\llbracket \text{CONCLUSION}(K_1, K_2) \rrbracket(w, G, H)$ iff there is a G' such that $\llbracket K_1 \rrbracket(w, G, G')$ & $\llbracket K_2 \rrbracket(w, G', H)$ & $\text{Conclusion}(G'_0, H_0)$.

B.3.1 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.4 Relation Between Dynamic and Classical Translation

I first give a dynamic translation for a fragment of English, specifying the updates associated with utterances containing modals and conditionals, and then a classical translation for the same fragment. Finally, we show that the former preserves the latter. I shall avail myself of abstract level of logical forms (ALFs) for the relevant fragment of English.

B.4.1 Abstract Logical Forms for a Fragment of English

Terms:

- Propositional constants from our base language in §B.1 (set $\mathcal{C}$).

Atoms:

- All terms are atoms, and nothing else is an atom.

ALFs:

- Atoms are ALFs.
- If ϕ and ψ are ALFs, then *might*(ϕ, ψ) is an ALF. (Stands for “it might be the case that ψ , given the restrictor ϕ .”)
- If ϕ and ψ are ALFs, then *must*(ϕ, ψ) 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 ALF. (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.4.2 Dynamic Interpretation:

Dynamic Translations

In this section, I provide a translation of the relevant fragment of English, into our dynamically interpreted language defined in §§B.1–B.3. I’ll assume the ALFs for the relevant fragment of English defined in §B.4.1 (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)

- $T_d(might(\phi, \psi)) = \text{MIGHT}(\phi, T_d(\psi))$
- $T_d(must(\phi, \psi)) = \text{MUST}(\phi, T_d(\psi))$
- $T_d(likely(\phi, \psi)) = \text{LIKELY}(\phi, T_d(\psi))$
- $T_d(if(\phi, \psi, \gamma)) = \text{IF}(\phi, T_d(\psi), T_d(\gamma))$
- $T_d(and(\phi, \psi)) = \text{AND}(T_d(\phi), T_d(\psi))$
- $T_d(not(\phi)) = \text{NOT}(T_d(\phi))$
- $T_d(Assert(\phi)) = \text{ASSERT}(T_d(\phi))$

Next, I shall introduce the classical interpretation of modals and conditionals in §B.4.3, and then go on to prove that our dynamic interpretation of modals and conditionals §(B.4.2) preserves the classical interpretation in §(B.4.3).

B.4.3 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 $\langle p, q, r \rangle$.

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.

Classical Semantics

I now define classical semantics for the simple modal language. I assume models are defined as in §B.2. (Assuming the definition of models in §B.2, and sets of sequences in §B.2.2), where R is an accessibility relation provided by the model, and ϕ a restriction on the domain of quantification of a modal:

- $\llbracket p \rrbracket^{G,w} = w \in \mathcal{I}(p)$
- $\llbracket \Diamond(\phi, \psi) \rrbracket^{G,w} = \{w \mid \exists w': wRw' \ \& \ w' \in \llbracket \phi \rrbracket^{G,w}, w' \in \llbracket \psi \rrbracket^{G,w}\}$
- $\llbracket \Box(\phi, \psi) \rrbracket^{G,w} = \{w \mid \forall w': \text{if } wRw' \ \& \ w' \in \llbracket \phi \rrbracket^{G,w} \text{ then } w' \in \llbracket \psi \rrbracket^{G,w}\}$
- $\llbracket \phi \wedge \psi \rrbracket^{G,w} = \llbracket \phi \rrbracket^{G,w} \cap \llbracket \psi \rrbracket^{G,w}$
- $\llbracket (\phi \wedge \psi) \rightarrow \gamma \rrbracket^{G,w} = \{w \mid \forall w': \text{if } wRw' \ \& \ w' \in \llbracket \phi \rrbracket^{G,w} \cap \llbracket \psi \rrbracket^{G,w} \text{ then } w' \in \llbracket \gamma \rrbracket^{G,w}\}$

Classical Translation

Assuming the same ALFs for the relevant fragment of English as in §B.4.1, 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 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.

Proof

- We prove that our dynamic interpretation (§B.4.2) preserves the classical interpretation (in §B.4.3). 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

1. By (§B.4.2), $T_d(p) = \langle \text{comp} := p \rangle$.
 2. By (§B.3), we have $\llbracket \langle \text{comp} := p \rangle \rrbracket(w, G, H)$ iff $G \underset{\text{comp}}{\sim} H$ & $H_{\text{comp}} = \llbracket p \rrbracket^{G,w}$.
 3. Suppose $\llbracket \langle \text{comp} := p \rangle \rrbracket(w, G, H)$.
 4. By (1)–(3), $H_{\text{comp}} = \llbracket p \rrbracket^{G,w}$, and $\llbracket p \rrbracket^{G,w}$ iff $w \in \mathcal{I}(p)$. Thus, $H_{\text{comp}} = \llbracket T_c(p) \rrbracket^{G,w}$, which we were set to prove.
- * 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 $H_{\text{comp}} = \llbracket \phi \rrbracket^w$, 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.4.2, $T_d(\phi) = \text{MIGHT}(\chi, T_d(\psi))$. Suppose that $\llbracket \text{MIGHT}(\chi, T_d(\psi)) \rrbracket(w, G, H)$. We know by (§B.3) that $\llbracket \text{MIGHT}(\chi, T_d(\psi)) \rrbracket(w, G, H)$ iff there is a G' and G'' such that $\llbracket T_d(\psi) \rrbracket(w, G, G') \& G' \approx_0 G'' \& G''_0 = G'_{\text{comp}} \cap \llbracket @E \rrbracket^{G,w}$ & $G'' \underset{\text{comp}}{\sim} H$ & $H_{\text{comp}} = M(\llbracket \chi \rrbracket^{G,w}, G'_{\text{comp}})$. Take such a G' and G'' . We have that $\llbracket T_d(\psi) \rrbracket(w, G, G')$; thus, 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) and (§B.4.3), 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.4.2, $T_d(\phi) = \text{MUST}(\chi, T_d(\psi))$. Suppose that $\llbracket \text{MUST}(\chi, T_d(\psi)) \rrbracket(w, G, H)$. So, by (§B.3), there is a G' and G'' such that $\llbracket T_d(\psi) \rrbracket(w, G, G') \& G' \approx_0 G'' \& G''_0 = G'_{\text{comp}} \cap \llbracket @E \rrbracket^{G,w}$ & $G'' \underset{\text{comp}}{\sim} H$ & $H_{\text{comp}} = N(\llbracket \chi \rrbracket^{G,w}, G'_{\text{comp}})$. 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.2) and (§B.4.3), 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.4.2, $T_d(\chi) = \text{IF}(\chi, T_d(\psi), T_d(\gamma))$. Suppose that $\llbracket \text{IF}(\chi, T_d(\psi), T_d(\gamma)) \rrbracket(w, G, H)$. So, by (§B.3), we know that there is a G', G'', G''' and G'''' such that $\llbracket T_d(\psi) \rrbracket(w, G, G') \& G' \approx_0 G'' \& G''_0 = G'_{\text{comp}} \cap \llbracket @E \rrbracket^{G,w}$ & $\llbracket T_d(\gamma) \rrbracket(w, G'', G''') \& G'' \approx_0 G''' \& G'''_0 = G''_{\text{comp}} \cap \llbracket @E \rrbracket^{G'',w}$ & $G''' \underset{\text{comp}}{\sim} H$ & $H_{\text{comp}} = \text{Cond}(\llbracket \chi \rrbracket^{G,w}, G'_{\text{comp}}, G''_{\text{comp}})$. 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.4) and (§B.4.3), H_{comp} is equivalent to $\llbracket T_c(\text{if}(\chi, \psi, \gamma)) \rrbracket^{G,w}$, by simple math.

Bibliography

- Adams, Ernest. 1975. *The Logic of Conditionals*. Dordrecht: D. Reidel.
- 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.
- 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.
- Asher, Nicholas and McCready, Eric. 2007. "Were, Would, Might and a Compositional Account of Counterfactuals." *Journal of Semantics* 24:93–129.
- Au, Terry K. 1986. "A Verb is Worth a Thousand Words: The Causes and Consequences of Interpersonal Events Implicit in Language." *Journal of Memory and Language* 25: 104–122.
- Ayer, A. J. 1936. *Language, Truth and Logic*. London: Gollancz.
- Bach, Kent. 1992a. "Intentions and Demonstrations." *Analysis* 52:140–146.
- Bach, Kent. 1992b. "Paving the Road to Reference." *Philosophical Studies* 67:295–300.
- Bach, Kent. 1994. "Conversational Implicature." *Mind and Language* 9:124–162.
- Barker, Christopher. 2002. "The Dynamics of Vagueness." *Linguistics and Philosophy* 25:1–36.
- Barwise, Jon and Etchemendy, John. 1987. *The Liar*. Oxford University Press.
- van Benthem, Johan. 1986. *Essays in Logical Semantics*, Studies in Linguistics and Philosophy. vol. 29. 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, NY.
- Bittner, Maria. 2011. "Time and Modality Without Tenses and Modals." In Renate Musan and Monika Rathert (eds.), *Tense Across Languages*, 147–88. Tübingen: Niemeyer.
- Bittner, Maria. 2014. *Temporality: Universals and Variation*. Oxford: Wiley-Blackwell.
- Bledin, Justin. 2015. "Modus Ponens Defended." *Journal of Philosophy* 112:57–83.
- Brasoveanu, Adrian. 2010a. "Decomposing Modal Quantification." *Journal of Semantics* 27:437–527.
- Brasoveanu, Adrian. 2010b. "Structured Anaphora to Quantifier Domains." *Information and Computation* 208:450–43.
- Braun, David. 1996. "Demonstratives and Their Linguistic Meaning." *Nous* 30:145–73.
- Brown, Roger and Fish, Deborah. 1983. "The Psychological Causality Implicit in Language." *Cognition* 14:237–273.
- Cantwell, John. 2008. "Changing the Modal Context." *Theoria* 74:331–351.
- Cappelen, Herman and Hawthorne, John. 2009. *Relativism and Monadic Truth*. Oxford: Oxford University Press.

- Cappelen, Herman and Lepore, Ernest. 2005. *Insensitive Semantics: A Defense of Semantic Minimalism and Speech Act Pluralism*. Oxford: Blackwell Publishing.
- Cappelen, Herman and Lepore, Ernest. 2007. *Language Turned on Itself: The Semantics and Pragmatics of Metalinguistic Discourse*. Oxford: Oxford University Press.
- Caramazza, Alfonso, Grober, Ellen, Garvey, Catherine, and Yates, Jack. 1974. "Comprehension of Anaphoric Pronouns." *Journal of Verbal Learning and Verbal Behaviour* 16:601–609.
- Cariani, Fabrizio. 2016. "Deontic Modals and Probabilities: One Theory to Rule Them All?" In Nate Charlow and Matthew Chrisman (eds.), *Deontic Modality*, 11–47.
- Cariani, Fabrizio. 2019. "Assertion and Modality." In Sanford Goldberg (ed.), *The Oxford Handbook of Assertion*. Oxford: Oxford University Press.
- Cariani, Fabrizio and Santorio, Paolo. 2018. "Will Done Better: Selection Semantics, Future Credence, and Indeterminacy." *Mind* 127:129–165.
- Carston, Robin. 2002. *Thoguhts and Utterances*. Oxford: Blackwell.
- Carston, Robyn. 1988. "Implicature, Explicature and Truth-theoretic Semantics." In Ruth Kempson (ed.), *Mental Representations: The Interface between Language and Reality*. Cambridge: Cambridge University Press.
- Charlow, Nate. 2014. "Logic and Semantics for Imperatives." *Journal of Philosophical Logic* 43:617–664.
- Charlow, Nate. 2015. "Prospects for an Expressivist Theory of Meaning." *Philosophers' Imprint* 15:1–43.
- Charlow, Nate. 2020. "Grading Modal Judgement." *Mind* 129:769–807. doi:10.1093/mind/fzz028.
- 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.)
- Ciardelli, Ivano, Groenendijk, Jeroen, and Roelofsen, Floris. 2019. *Inquisitive Semantics*. Oxford: Oxford University Press.
- Clark, Herbert H. 1996. *Using Language*. Cambridge: Cambridge University Press.
- Cohen, Stewart. 1998. "Contextualist Solutions to Epistemological Problems: Scepticism, Gettier, and the Lottery." *Australasian Journal of Philosophy* 76:289–306.
- Comrie, Bernard. 1989. "On Identifying Future Tenses." In Abraham Werner and Janssen Theo (eds.), *Tempus-Aspekt-Modus: die lexicalischen und grammatischen Formen in den germanischen Sprache*, 51–63. Tübingen: Max Niemeyer Verlag.
- Condoravdi, Cleo and Lauer, Sven. 2012. "Imperatives: Meaning and Illocutionary Force." In Christopher Pinón (ed.), *Empirical Issues in Syntax and Semantics*, volume 9, 37–58. Paris: Colloque de Syntaxe et Sémantique à Paris.
- Dekker, Paul. 1994. "Predicate Logic With Anaphora." In L. Santelmann and M. Harvey (eds.), *Proceedings From SALT IV*, 79–95. DMLL Publication, Cornell University, Ithaca, NY.
- 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.
- DeRose, Keith. 2009. *The Case for Contextualism: Knowledge, Skepticism, and Context*, volume 1. Oxford: Oxford University Press.
- van der Does, Jaap, Groeneveld, Willem, and Veltman, Frank. 1997. "An Update on 'Might.'" *Journal of Logic, Language and Information* 6:361–380.

- Dorr, Cian and Hawthorne, John. 2014. "Embedding Epistemic Modals." *Mind* 122: 867–913.
- Dowell, J.L. 2011. "A Flexible Contextualist Account of Epistemic Modals." *Philosophers' Imprint* 11:1–25.
- Dummett, Michael. 1993. *The Logical Basis of Metaphysics*. Cambridge, MA: Harvard University Press.
- Egan, Andy. 2007. "Epistemic Modals, Relativism and Assertion." *Philosophical Studies* 133:1–22.
- Egan, Andy. 2010. "Disputing About Taste." In Richard Feldman and Ted A. Warfield (eds.), *Disagreement*, 247–286. 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: Oxford University Press.
- van Eijck, Jan and de Vries, Fer-Jan. 1992. "Dynamic Interpretation and Hoare Deduction." *Journal of Language, Logic and Information* 1:1–44.
- van Eijck, Jan and Unger, Christina. 2010. *Computational Semantics with Functional Programming*. Cambridge: Cambridge University Press.
- Elbourne, Paul D. 2005. *Situations and Individuals*. Cambridge, MA: MIT Press.
- Elbourne, Paul. 2008. "Implicit Content and the Argument from Binding." In *Proceedings of Semantics and Linguistic Theory XVIII (SALT 18)*, 284–301. CLC Publications, Ithaca, NY.
- Evans, Gareth. 1979. "Reference and Contingency." *The Monist* 62:161–189.
- Fiengo, Robert and May, Robert. 1994. *Indices and Identity*. Cambridge, MA: MIT Press.
- Fiengo, Robert and May, Robert. 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: Oxford University Press.
- von Fintel, Kai and Gillies, Anthony S. 2008. "CIA Leaks." *Philosophical Review* 117: 77–98.
- von Fintel, Kai and Gillies, Anthony S. 2009. "'Might' Made Right." In Andy Egan and Brian Weatherson (eds.), *Epistemic Modality*. New York: Oxford University Press.
- Frank, Anette. 1996. *Context-dependence in Modal Constructions*. Ph.D. thesis, University of Stuttgart.
- Frank, Anette and Kamp, Hans. 1997. "On Context Dependence in Modal Constructions." In Aaron Lawson (ed.), *Proceedings from SALT VII*, 151–168. CLC Publications, Cornell University, Ithaca, NY.
- Frege, Gottlob. 1892. "Über Sinn und Bedeutung." *Zeitschrift für Philosophie und philosophische Kritik* 100:25–50.
- Frege, Gottlob. 1918. "Der Gedanke." *Beiträge zur Philosophie des deutschen Idealismus* 1:58–77.
- Garvey, Catherine and Caramazza, Alfonso. 1974. "Implicit Causality in Verbs." *Linguistic Inquiry* 5:459–464.
- Geurts, Bart. 1995. *Presupposing*. Ph.D. thesis, University of Stuttgart.
- Gibbard, Alan. 1981. "Two Recent Theories of Conditionals." In Robert Stalnaker William Harper and Glen Pearce (eds.), *Ifs*, 211–247. Amsterdam: Reidel.

- Gillies, Anthony S. 2004. "Epistemic Conditionals and Conditional Epistemics." *Noûs* 38:585–616.
- Gillies, Anthony S. 2010. "Iffiness." *Semantics and Pragmatics* 3:1–42.
- Ginzburg, Jonathan. 1994. "An Update Semantics for Dialogue." In Harry Bunt (ed.), *Proceedings of the First International Workshop on Computational Semantics*, Tilburg, Netherlands.
- Glanzberg, Michael. 2007. "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.
- Grice, H.P. 1989. "Indicative Conditionals." In *Studies in the Way of Words*, 58–85. Cambridge, MA: Harvard University Press.
- Groenendijk, Jeroen and Stokhof, Martin. 1990. "Two Theories of Dynamic Semantics." In Jan van Eijck (ed.), *Logics in AI*, 55–64. Berlin: Springer.
- Groenendijk, Jeroen, Stokhof, Martin, and Veltman, Frank. 1996. "Coreference and Modality." In Shalom Lappin (ed.), *Handbook of Contemporary Semantic Theory*, 179–216. Oxford: Willey Blackwell.
- 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.
- Hamann, Cornelia. 1991. "Adjectival Semantics." In Arnim von Stechow and Dieter Wunderlich (eds.), *Semantics: An International Handbook of Contemporary Research*, 657–673. New York: de Gruyter.
- Hare, Richard M. 1952. *The Language of Morals*. Oxford: Clarendon Press.
- Haug, Dag Trygve Truslew. 2014. "Partial Dynamic Semantics for Anaphora: Compositionality without Syntactic Coindexation." *Journal of Semantics* 31:457–511.
- Hawthorne, John. 2004. *Knowledge and Lotteries*. Oxford: Oxford University Press.
- Heim, Irene. 1982. *The Semantics of Definite and Indefinite Noun Phrases*. Ph.D. thesis, University of Massachusetts.
- Heim, Irene and Kratzer, Angelika. 1998. *Semantics in Generative Grammar*. Oxford: Blackwell Publishers.
- Hobbs, Jerry R. 1979. "Coherence and Coreference." *Cognitive Science* 3:67–90.
- Hobbs, Jerry R. 1990. *Literature and Cognition* (CSLI Lecture Notes 21). Chicago: University of Chicago Press.
- Hobbs, Jerry R. 2010. "Clause-internal Coherence." In Peter Kühnlein, Anton Benz, and Candace L. Sidner (eds.), *Constraints in Discourse*, 15–34. Amsterdam: John Benjamins.
- Hume, David. 2007/1738. *A Treatise of Human Nature*. Oxford: Clarendon Press.
- Kaiser, Elsi. 2009. "Effects of Anaphoric Dependencies and Semantic Representations on Pronoun Interpretation." In *Anaphora Processing and Applications*, 121–130. Heidelberg: Springer.
- Kameyama, Megumi. 1996. "Indefeasible Semantics and Defeasible Pragmatics." In M. Kanazawa, C. Pinon, and H. de Swart (eds.), *Quantifiers, Deduction and Context*, 111–138. Stanford: CSLI Publications.
- Kamp, Hans. 1971. "Formal Properties of 'Now.'" *Theoria* 37:227–273.
- Kaplan, David. 1970. "Dthat." *Syntax and Semantics* 9:221–43.
- Kaplan, David. 1989a. "Afterthoughts." In Almog Joseph, Perry John, and Wettstein Howard (eds.), *Themes From Kaplan*, 565–614. Oxford: Oxford University Press.

- Kaplan, David. 1989b. "Demonstratives." In Almog Joseph, Perry John, and Wettstein Howard (eds.), *Themes From Kaplan*, 481–563. Oxford: Oxford University Press.
- Kaplan, David. 1990. "Words." *Proceedings of the Aristotelian Society, Supplementary Volumes*, 64: 93–119.
- Kaufmann, Magdalena Schwager. 2012. *Interpreting Imperatives* (Studies in Linguistics and Philosophy). New York: Springer.
- Kaufmann, Stefan. 2000. "Dynamic Context Management." In Martina Faller, Stefan Kaufmann, and Marc Pauly (eds.), *Formalizing the Dynamics of Information*, 171–188. Stanford: CSLI.
- Kehler, Andrew. 1994. "Temporal Relations: Reference or Discourse Coherence?" In *Proceedings of the 32nd Annual Meeting on Association for Computational Linguistics*, 319–321, Las Cruces, New Mexico: Association for Computational Linguistics.
- Kehler, Andrew. 2002. *Coherence, Reference and the Theory of Grammar*. Stanford: CSLI Publications.
- Kehler, Andrew, Kertz, Laura, Rohde, Hanna, and Elman, Jeffrey. 2008. "Coherence and Coreference Revisited." *Journal of Semantics* 25:1–44.
- Kendon, Adam. 1988. *Sign Languages of Aboriginal Australia: Cultural, Semiotic and Communicative Perspectives*. Cambridge: Cambridge University Press.
- Kendon, Adam. 2004. *Gesture: Visible Action as Utterance*. Cambridge: Cambridge University Press.
- Kennedy, Christopher. 2007. "Vagueness and Grammar: The Semantics of Relative and Absolute Gradable Adjectives." *Linguistics and Philosophy* 30:1–45.
- Khoo, Justin. 2013. "A Note on Gibbard's Proof." *Philosophical Studies* 166:153–164.
- Kibble, Rodger. 1994. "Dynamics of Epistemic Modality and Anaphora." In Harry Bunt, Reinhard Muskens, and Gerrit Rentier (eds.), *International Workshop on Computational Semantics (IWCS)*, Tilburg, Netherlands.
- Kibble, Rodger. 1995. "Modal Subordination, Focus and Complement Anaphora." In *Proceedings of the Tbilisi Symposium on Language, Logic and Computation*. Stanford, CA: CSLI.
- King, Jeffrey C. 2003. "Tense, Modality and Semantic Values." *Philosophical Perspectives* 17:195–246.
- King, Jeffrey C. 2014a. "The Metasemantics of Contextual Sensitivity." In Alexis Burgess and Bret Sherman (eds.), *Metasemantics: New Essays on the Foundations of Meaning*, 97–118. Oxford: Oxford University Press.
- King, Jeffrey C. 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: Oxford University Press.
- Kissine, Mikhail. 2008. "Why Will Is Not a Modal." *Natural Language Semantics* 16: 129–155.
- Klein, Ewan. 1991. "Comparatives." In Arnim von Stechow and Dieter Wunderlich (eds.), *Semantics: An International Handbook of Contemporary Research*, 673–691. New York: de Gruyter.
- 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.

- Kratzer, Angelika. 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.
- Kratzer, Angelika. 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.
- Kratzer, Angelika. 1998. "An Investigation of the Lumps of Thought." *Linguistics and Philosophy* 12:607–653.
- Kratzer, Angelika. 2012. *Modals and Conditionals*. Oxford: Oxford University Press.
- Kratzer, Angelika. 2014. "Situations in Natural Language Semantics." <http://plato.stanford.edu/entries/situations-semantics/>
- Kripke, Saul. 1959. "A Completeness Theorem in Modal Logic." *Journal of Symbolic Logic* 24:1–14.
- Kripke, Saul. 1963. "Semantical Analysis of Modal Logic I: Normal Modal Propositional Calculi." *Zeitschrift für mathematische Logik und Grundlagen der Mathematik* 9:67–96.
- Kripke, Saul. 1980. *Naming and Necessity*. Cambridge, MA: Harvard University Press.
- Lascarides, Alex and Stone, Matthew. 2006. "Formal Semantics for Iconic Gesture." In *Proceedings of the 10th Workshop on the Semantics and Pragmatics of Dialogue (Brandial)*, Potsdam.
- Lascarides, Alex and Stone, Matthew. 2009. "A Formal Semantic Analysis of Gesture." *Journal of Semantics* 26:393–449.
- Lepore, Ernest and Stone, Matthew. 2015. *Imagination and Convention: Distinguishing Grammar and Inference in Language*. Oxford: Oxford University Press.
- Lewis, David K. 1969. *Convention: A Philosophical Study*. Cambridge, MA: Harvard University Press.
- Lewis, David K. 1973. *Counterfactuals*. Oxford: Wiley-Blackwell.
- Lewis, David K. 1979. "Scorekeeping in a Language Game." *Journal of Philosophical Logic* 8:339–359.
- Lewis, David K. 1980. "Index, Context and Content." In Stig Kanger and Sven Öhman (eds.), *Philosophy and Grammar*, 79–100. Amsterdam: D. Reidel.
- Lewis, Karen. 2019. "The Speaker Authority Problem for Context-sensitivity (or: You Can't Always Mean What You Want)." *Erkenntnis*. doi:10.1007/s10670-018-0089-2.
- Locke, John. 1689/1975. *An Essay Concerning Human Understanding*, ed. Peter H. Nidditch. (ed.) Oxford: Calderon Press.
- MacFarlane, John. 2014. *Assessment Sensitivity: Relative Truth and Its Applications*. Oxford: Oxford University Press.
- MacFarlane, John. 2020. "On Probabilistic Knowledge." *Res Philosophica* 97:97–108.
- Mandelkern, Matthew. 2019. "Bounded Modality." *Philosophical Review* 1:1–61.
- Mann, William C. and Thompson, Sandra A. 1988. "Rhetorical Structure Theory: Toward a Functional Theory of Text Organization." *Text* 8:243–281.
- McGee, Vann. 1985. "A Counterexample to Modus Ponens." *Journal of Philosophy* 92: 462–471.
- McKoon, Gail, Greene, Steven B., and Ratcliff, Roger. 1993. "Discourse Models, Pronoun Resolution, and The Implicit Causality of Verbs." *Journal of Experimental Psychology* 19:1040–1052.
- 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.
- Moss, Sarah. 2018. *Probabilistic Knowledge*. Oxford: Oxford University Press.
- Murray, Sarah E. and Starr, William B. 2018. "Force and Conversational States." In Daniel Fogal, Daniel W. Harris, and Matt Moss (eds.), *New Work on Speech Acts*, 202–236. Oxford: Oxford University Press.
- Neale, Stephen. 1990. *Descriptions*. Cambridge, MA: MIT Press.
- Neale, Stephen. 2004. "This, That and the Other." In Marga Reimer and Anna Bezuidenhout (eds.), *Descriptions and Beyond*, 68–182. New York: Oxford University Press.
- Ninan, Dilip. 2010. "Semantics and the Objects of Assertion." *Linguistics and Philosophy* 33:355–380.
- Nowak, Ethan, and Michaelson, Eliot. 2020. "Discourse and Method." *Linguistics and Philosophy* 43:119–138.
- Nunberg, Geoffrey. 1993. "Indexicality and Deixis." *Linguistics and Philosophy* 16:1–43.
- Partee, Barbara. 1984. "Nominal and Temporal Anaphora." *Linguistics and Philosophy* 7:243–286.
- Portner, Paul. 2007. "Imperatives and Modals." *Natural Language Semantics* 15:351–383.
- Portner, Paul. 2009. *Modality*. Oxford: Oxford University Press.
- Portner, Paul. 2012. "Permission and Choice." In *Discourse and Grammar: From Sentence Types to Lexical Categories* (Studies in Generative Grammar). Berlin: Mouton de Gruyter.
- Potts, Christopher. 2005. *The Logic of Conventional Implicatures*. Oxford: Oxford University Press.
- Rabern, Brian. 2012. "Against the Identification of Assertoric Content with Compositional Value." *Synthese* 189:75–96.
- Recanati, François. 1996. "Domains of Discourse." *Linguistics and Philosophy* 19:445–475.
- Recanati, François. 2004. *Literal Meaning*. Cambridge: Cambridge University Press.
- 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.
- Roberts, Craige. 1996. "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.
- Roberts, Craige. 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.
- Roberts, Craige. 2011. "Solving for Interpretation." Manuscript. Ohio State University.
- Roberts, Craige. 2012. "Information Structure in Discourse: Towards an Integrated Formal Theory of Pragmatics." *Semantics and Pragmatics* 5:1–69. doi:10.3765/sp.5.6
- Rooth, Mats. 1992. "A Theory of Focus Interpretation." *Natural Language Semantics* 1:75–116.
- Rothschild, Daniel and Yalcin, Seth. 2017. "On the Dynamics of Conversation." *Noûs* 51:24–48. doi: 10.1111/nous.12121.
- Rudolph, Udo and Försterling, Friedrich. 1997. "The Psychological Causality Implicit in Verbs: A Review." *Psychological Bulletin* 121:192–218.
- 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.
- Schiffer, Stephen. 1972. *Meaning*. Oxford: Clarendon Press.
- Schiffer, Stephen. 1981. "Indexicals and the Theory of Reference." *Synthese* 49:43–100.
- Schiffer, Stephen. 2005. "Russell's Theory of Definite Descriptions." *Mind* 114:1135–1183.
- Schlenker, Philippe. 2013. "Temporal and Modal Anaphora in Sign Language (ASL)." *Natural Language and Linguistic Theory* 31:207–234.
- Schlenker, Philippe. 2018. "Gesture Projection and Cosuppositions." *Linguistics and Philosophy* 41:295–365.
- Schoubye, Anders J. 2017. "Type-ambiguous Names." *Mind* 126:715–767.
- Schroeder, Mark. 2013. "Two Roles for Propositions: Cause for Divorce?" *Nous* 47: 409–430.
- Sidner, Candice. 1983. "Focusing in the Comprehension of Definite Anaphora." In Michael Brady and Robert C. Berwick (eds.), *Computational Models of Discourse*, 267–330. Cambridge, MA: MIT Press.
- Smyth, Ron. 1994. "Grammatical Determinants of Ambiguous Pronoun Resolution." *Journal of Psycholinguistic Research* 23:197–229.
- Soames, Scott. 2002. *Beyond Rigidity*. Oxford University Press.
- Sorensen, Roy. 2009. "Meta-agnosticism: Higher Order Epistemic Possibility." *Mind* 118:777–784.
- Sperber, Dan and Wilson, Deirdre. 1986. *Relevance*. Oxford: Blackwell.
- Stainton, Robert J. 1994. "Using Non-sentences: An Application of Relevance Theory." *Pragmatics & Cognition* 2:269–284.
- Stalnaker, Robert. 1978. "Assertion." *Syntax and Semantics: Pragmatics* 9:315–332.
- Stalnaker, Robert. 2002. "Common Ground." *Linguistics and Philosophy* 25:701–721.
- Stalnaker, Robert. 2014. *Context*. Oxford: Oxford University Press.
- Stanley, Jason. 1997a. "Names and Rigid Designation." In Bob Hale and Crispin Wright (eds.), *A Companion to the Philosophy of Language*, 555–585. Oxford: Blackwell.
- Stanley, Jason. 1997b. "Rigidity and Content." In Richard G. Heck Jr. (ed.), *Language, Thought, and Logic: Essays in Honor of Michael Dummett*, 131–156. Oxford: Oxford University Press.
- Stanley, Jason. 2004. "The Linguistic Basis of Contextualism." *Philosophical Studies* 119:119–146.
- Stanley, Jason. 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 B. 2016. "Dynamic Expressivism about Deontic Modality." In Nate Charlow and Matthew Chrisman (eds.), *Deontic Modality*, 355–394. Oxford: Oxford University Press.
- Starr, William B. 2020. "A Preference Semantics for Imperatives." *Semantics and Pragmatics*. 13. doi: 10.3765/sp.13.6.
- von Stechow, Arnim. 1995. "On the Proper Treatment of Tense." In Mandy Simons and Teresa Galloway (eds.), *Semantics and linguistic Theory (SALT)* 5: 362–386. Ithaca, NY: CLC Publications, Cornell University, Ithaca.
- Stevenson, Charles L. 1937. "The Emotive Meaning of Ethical Terms." *Mind* 46:14–31.
- Stojnić, Una. 2016. *Context-Sensitivity in A Coherent Discourse*. Ph.D. thesis, Rutgers University.

- Stojnić, Una. 2017a. "On the Connection between Semantic Content and the Objects of Assertion." *Philosophical Topics* 45:163–179.
- Stojnić, Una. 2017b. "One's *Modus Ponens*. . . : Modality, Coherence and Logic." *Philosophy and Phenomenological Research* 95:167–214.
- Stojnić, Una. 2018. "Discourse, Context and Coherence: The Grammar of Prominence." In *Beyond Semantics and Pragmatics*, 97–124. Oxford: Oxford University Press.
- Stojnić, Una. 2019. "Content in a Dynamic Context." *Noûs* 52:394–432.
- Stojnić, Una. forthcoming. "The Anatomy of Arguments in Natural Language Discourse." In *Oxford Studies in Philosophy of Language*, volume 2. Oxford: Oxford University Press.
- Stojnić, Una and Altshuler, Daniel. 2020. "Formal Properties of 'Now,' Revisited." Manuscript.
- Stojnić, Una and Lepore, Ernest. 2013. "What's *What's Said*." In *What Is Said and What Is Not: The Semantics/Pragmatics Interface*, 17–36. Stanford: CSLI Publications.
- Stojnić, Una, Stone, Matthew, and Lepore, Ernest. 2018. "Deixis (Even Without Pointing)." *Philosophical Perspectives* 27:502–525.
- Stojnić, Una, Stone, Matthew, and Lepore, Ernest. 2017. "Discourse and Logical Form: Pronouns, Attention and Coherence." *Linguistics and Philosophy*. 40:519–547.
- Stojnić, Una, Stone, Matthew, and Lepore, Ernest. 2018. "Distinguishing Ambiguity from Underspecificity." In Ken Turner and Lawrence Horn (eds.), *Pragmatics, Truth and Underspecification: Towards an Atlas of Meaning* (Current Research in the Semantic-s/Pragmatics Interface). Leiden: Brill.
- Stojnić, Una, Stone, Matthew, and Lepore, Ernest. 2020. "Pointing Things Out: In Defense of Attention and Coherence." *Linguistics and Philosophy* 43: 139–148.
- Stone, Matthew. 1997. "The Anaphoric Parallel Between Modality and Tense." *IRCS Report* 97–06, University of Pennsylvania, Philadelphia.
- Stone, Matthew. 1999. "Reference to Possible Worlds." Technical Report 49, Rutgers University, RuCCS, New Brunswick.
- 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.
- Taylor, Kenneth A. 2001. "Sex, Breakfast, And Descriptus Interruptus." *Synthese* 128: 45–61.
- Tonhauser, Judith, Beaver, David, Roberts, Craig, 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.
- Veltman, Frank. 1985. *Logics for Conditionals*. Ph.D. thesis, University of Amsterdam.
- Veltman, Frank. 1996. "Defaults in Update Semantics." *Journal of Philosophical Logic* 25:221–261.
- Vermeulen, C.F.M. 1993. "Sequence Semantics for Dynamic Predicate Logic." *Journal of Logic, Language and Information* 2:217–54.
- Walker, Marilyn A., Masayo, Iida, and Cote, Sharon. 1994. "Japanese Discourse and the Process of Centering." *Computational Linguistics* 20:193–232.
- Webber, Bonnie L. 1988. "Tense as Discourse Anaphor." *Computational Linguistics* 14:61–73.
- Webber, Bonnie L., Stone, Matthew, Joshi, Aravind K., 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*. Abingdon: Lawrence Erlbaum Associates.
- Willer, Malte. 2013. "Dynamics of Epistemic Modality." *Philosophical Review* 122:45–92.
- Willer, Malte. 2014. "Dynamic Thoughts on Ifs and Oughts." *Philosophers' Imprint* 14: 1–30.
- Winograd, Terry. 1972. *Understanding Natural Language*. New York: Academic Press.
- Wolf, Florian and Gibson, Edward. 2006. *Coherence in Natural Language: Data Structures and Applications*. Cambridge, MA: The MIT 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.
- Yalcin, Seth. 2011a. "Context Probabilism." *Proceedings of the 18th Amsterdam Colloquium* 12–21.
- Yalcin, Seth. 2011b. "Nonfactualism about Epistemic Modality." In Brian Weatherson and Andy Egan (eds.), *Epistemic Modality*, 295–332. New York: Oxford University Press.
- Yalcin, Seth. 2012a. "Bayesian Expressivism." *Proceedings of the Aristotelian Society* 112:123–160.
- Yalcin, Seth. 2012b. "A Counterexample to *Modus Tollens*." *Journal of Philosophical Logic* 41:1001–1024.
- Yalcin, Seth. 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: Oxford University Press.
- Zeevat, Henk. 1992. "Presupposition and Accommodation in Update Semantics." *Journal of Semantics* 9:379–412.
- Zvolenszky, Zsófia. 2002. "Is a Possible-Worlds Semantics of Modality Possible? A Problem for Kratzer's Semantics." In Brandon Jackson (ed.), *Semantics and Linguistic Theory XII*, 339–358. Ithaca, NY: Cornell University Press.

Index of Names

A

Adams, Ernest 6
Alahverdzhieva, Katya 46
Altshuler, Daniel viii, 172
Asher, Nicholas ix, 5, 12, 13, 62, 63, 65, 67,
124, 125, 133, 159
Au, Terry K. 73
Ayer, A. J. 6, 99

B

Bach, Kent 15, 16, 26
Barker, Christopher ix, 182
Barwise, Jon 182
Beaver, David ix, 62, 173
van Benthem, Johan viii, 105
Bittner, Maria ix, 40, 41, 58, 59, 63, 70, 114,
116, 119, 126, 137, 154
Bledin, Justin ix, 16, 18, 93
Brasoveanu, Adrian 114, 119, 126, 137, 154
Braun, David ix, 26, 27, 30
Brown, Roger 73

C

Cantwell, John 143
Cappelen, Herman 3, 53
Caramazza, Alfonso 73
Cariani, Fabrizio ix, 98, 115, 118
Carston, Robin 16
Charlow, Nate ix, 6, 8, 16, 18, 93, 95, 106, 136
Chomsky, Noam 29, 56
Ciardelli, Ivano 8
Clark, Herbert H. 132
Comrie, Bernard 115
Condoravdi, Cleo 8
Cote, Sharon 58, 59

D

Dekker, Paul 35, 41
DeRose, Keith 3
Desmet, Timothy 6
van der Does, Jaap 105
Dorr, Cian viii, 128, 132, 133

Dowell, J.L. ix, 3, 124
Dummett, Michael 15

E

Egan, Andy viii, 3, 95, 109, 118, 135
van Eijck, Jan 41
Elbourne, Paul D. 35, 182
Elman, Jeffrey 6, 73
Etchemendy, John 182
Evans, Gareth 15

F

Fiengo, Robert 29
von Fintel, Kai 3, 6, 95, 102, 105, 106, 118, 135,
176, 179, 182
Fish, Deborah 73
Försterling, Friedrich 73
Frank, Anette 114
Frege, Gottlob 6, 181

G

Garvey, Catherine 73
Geurts, Bart 114
Gibbard, Alan 6, 166
Gibson, Edward 6, 63
Gillies, Anthony S. viii, 3, 6, 18, 93, 95, 102,
105, 106, 118, 135, 144, 146, 151, 176
Ginzburg, Jonathan 8
Glanzberg, Michael viii, 4
Greene, Steven B. 73
Grice, H.P. 3, 4, 97, 98
Grober, Ellen 73
Groenendijk, Jeroen 8, 79, 105
Groeneveld, Willem 105
Grosz, Barbara J. 40

H

Hamann, Cornelia 182
Hare, Richard M. 6, 99
Hawthorne, John viii, 3, 109, 118, 128, 132,
133, 135
Heim, Irene 33, 35, 79

Hobbs, Jerry R. 5, 12, 13, 61, 62, 64, 65, 67, 72,
159, 184
Hume, David 63

J

Joshi, Aravind K. 13, 40, 62, 65, 159

K

Kaiser, Elsi 6, 64
Kameyama, Megumi 58, 59
Kamp, Hans 24, 35, 114
Kaplan, David 1, 2, 3, 4, 15, 23, 24, 25, 26, 27,
29, 31, 34, 37, 39, 45, 50, 51, 52, 54, 55, 71,
75, 76, 78, 85, 90, 91, 92, 108, 109, 152,
171, 191
Kaufmann, Magdalena Schwager ix, 8
Kaufmann, Stefan ix, 41, 114, 126, 135
Kehler, Andrew ix, 5, 6, 12, 13, 59, 62, 63,
68, 73
Kendon, Adam 47, 48, 49
Kennedy, Christopher 182
Kertz, Laura 6, 73
Khoo, Justin 19
Kibble, Rodger 114
King, Jeffrey C. viii, 3, 4, 15, 16, 26, 50
Kissine, Mikhail 115
Knott, Alistair 13, 62, 63, 65, 159
Kratzer, Angelika ix, 33, 93, 99, 118, 119, 120,
123, 165, 182
Kripke, Saul 54, 56, 93, 99

L

Lascarides, Alex viii, 5, 12, 13, 46, 62, 63, 65,
67, 159
Lauer, Sven 8
Lepore, Ernest viii, 3, 49, 50, 53, 62, 67, 71, 98,
173, 174, 175
Lewis, David K. 9, 10, 15, 23, 48, 92,
93, 99
Lewis, Karen ix, 4
Locke, John 3, 4

M

MacFarlane, John 3, 16, 18, 93, 95, 96, 98, 109,
143, 144, 146
Mandelkern, Matthew 132, 133
Mann, William C. 63
Masayo, Iida 58, 59
May, Robert viii, 29
McCready, Eric 124, 125, 133
McGee, Vann 18, 19, 143, 145, 146
McKoon, Gail 73
Michaelson, Eliot ix, 54

Moss, Sarah 6, 7, 16, 18, 93, 98, 99, 106, 109,
136, 144
Murray, Sarah E. ix, 8

N

Neale, Stephen viii, 3, 4
Ninan, Dilip 15
Nowak, Ethan ix, 54
Nunberg, Geoffrey 2

P

Partee, Barbara ix, 114
Portner, Paul 8, 115
Potts, Christopher 62

R

Rabern, Brian ix, 15
Ratcliff, Roger 73
Recanati, François 15, 16, 182
Reimer, Marga 4, 26, 50
Roberts, Craige ix, 8, 62, 114, 125,
173, 181
Roelofsen, Floris 8
Rohde, Hanna ix, 6, 73
Rooth, Mats 125
Rothschild, Daniel ix, 105
Rudolph, Udo 73
Russell, Bertrand 181

S

Schaffer, Jonathan ix, 3
Schiffer, Stephen viii, 3, 4
Schlenker, Philippe ix, 46, 48, 116
Schoubye, Anders J. 54
Schroeder, Mark 7, 16, 93, 98, 109
Sidner, Candice 40
Simons, Mandy ix, 62, 173
Smyth, Ron 5, 64
Soames, Scott 54
Sorensen, Roy 128
Sperber, Dan 16
Stainton, Robert J. 16
Stalnaker, Robert 19, 75, 76, 77, 78, 79, 81, 82,
101, 102, 104, 106, 134, 167, 171
Stanley, Jason 3, 15, 16, 179, 182, 184
Starr, William B. ix, 6, 8, 16, 93, 106
von Stechow, Arnim 115
Stevenson, Charles L. 6, 99
Stojnić, Gala viii, 177
Stojnić, Una vii, viii, 14, 15, 49, 50, 54, 62, 65,
67, 70, 71, 80, 116, 124, 163, 166, 172, 173,
174, 175, 182
Stokhof, Martin 79, 105

Stone, Matthew viii, 13, 46, 49, 50, 62, 65, 67,
71, 98, 114, 115, 116, 119, 159, 173, 174, 175,
177, 182

Swanson, Eric 6, 93, 95, 106

Szabó, Zoltán Gendler ix, 3, 16, 179,
182, 184

T

Taylor, Kenneth A. 16

Thompson, Sandra A. 63

Tonhauser, Judith 62, 173

Travis, Charles 3

U

Unger, Christina 41

V

Veltman, Frank ix, 6, 18, 19, 79, 93, 95, 102,
105, 143, 144

Vermeulen, C.F.M. 41

de Vries, Fer-Jan 41

W

Walker, Marilyn A. 58, 59

Weatherson, Brian ix, 3, 109, 118, 135

Webber, Bonnie L. 13, 62, 65, 159

Weinstein, Scott 40

Wettstein, Howard K. 4

Wilkins, David 48

Willer, Malte 6, 16, 18, 93, 95, 104, 118, 143

Wilson, Deirdre 16

Winograd, Terry 72

Wolf, Florian 6, 63

Y

Yalcin, Seth ix, 3, 6, 15, 16, 17, 18, 93, 95, 97,
98, 100, 101, 102, 104, 105, 106, 108, 110,
111, 118, 136, 143, 144, 165, 167, 177, 179

Yates, Jack 73

Z

Zeevat, Henk 41

Zvolenszky, Zsófia 123

Index of Terms

A

abductive reasoning 2, 9, 10, 31, 50, 53, 65, 67, 68, 92, 113, 121
 Aboriginal 48
 ambiguity 11, 26–29, 33–39, 44, 45, 51, 54, 55, 65, 78, 79, 135, 175–77
 exotic 26, 29
 American Sign Language (ASL) 115, 116
 anaphora 36–39, 41–44, 49, 51, 56, 58, 63, 70, 79, 114–120, 126–138, 149, 157, 159, 163, 165–166, 172, 179, 190, 193, 197, 199, 200
 modal 70, 79, 114–121, 126, 128, 129, 131, 132, 133, 137, 138, 179
 pronominal 63, 70, 114–121, 126, 133, 137, 179
 Arrente 48
 artificial intelligence 62
 assertion 8, 15, 16, 77, 81, 87, 96, 97, 98, 99, 100, 101, 102, 103, 104, 106, 107, 157, 158, 160, 161, 162, 177, 199
 knowledge norm of 98
 assignment function 34–38, 40–45, 52, 56, 57, 72, 87–91, 153, 154, 189–190, 195–196
 Attention-Coherence account 71, 72, 85, 189
 attention-shifting 42, 44, 46, 49, 53, 55, 65–73, 80, 122, 134, 192
 attentional state of the discourse 40, 41, 45, 53–61, 65–67, 71, 72, 85, 122, 136, 137, 148, 153, 154

B

Background relation 182
 Binding theory 56

C

Campania 47
 center of attention 32, 40, 42, 43, 44, 45, 49, 55, 57
 Centering theory 40
 character 1, 12, 24, 26, 27, 30, 39, 40, 52, 55, 56, 57, 59, 60, 61, 67, 71, 72, 79, 89, 121, 122, 171, 172
 Chibchan languages 48
 classical logic 7, 18, 19, 93, 94, 126, 143, 152, 163, 164, 168, 172, 198–201

co-indexing 28, 29, 37, 38, 51
 coherence relations 11, 12, 62, 63, 64, 65, 66, 67, 68, 69, 70, 71, 72, 73, 122, 123, 128, 129, 138, 148, 152, 158, 159, 161, 173, 174, 183, 192, 197
 Coherence theory 5, 6, 62, 65, 66, 68
 common ground 64, 75, 76, 77, 78, 81, 102, 103, 104, 105, 106, 107, 110, 181
 computational linguistics 62
 Conclusion relation 150, 161, 198
 conditional 6, 15, 19, 100, 110, 111, 116, 118, 120, 123, 124, 126, 127, 128, 131, 132, 133, 134, 143, 144, 146, 148, 150, 151, 153, 156, 157, 158, 160, 162, 163, 164, 165, 166, 168, 180, 194, 197, 198, 199, 200, 201
 conjunction elimination 28
 content
 assertoric 14, 15, 16, 17, 18, 93, 96, 97, 98, 99, 104, 106, 107, 109, 136
 compositional semantic 14, 15, 16, 17, 24, 87, 95, 104, 107, 108, 109, 110, 123, 136
 informational 6, 15, 16, 17, 93, 96, 97, 98, 99, 104, 109, 135, 149, 168
 probabilistic 6, 99, 106
 context
 of assessment 95, 96, 109
 of utterance 8, 37, 85, 90, 96, 99, 108, 109
 -change 7, 24, 31, 32, 33, 41, 82, 90, 92, 94, 96, 102, 104, 105, 126, 133, 134, 135, 138, 161, 167, 168
 -shifting 18, 30, 31, 52, 152, 166
 Kaplanian 4, 23, 24, 25, 27, 39, 71, 75, 76, 78, 85, 90, 91, 92, 171
 linguistic 11, 12, 13, 14, 37, 46, 92, 133, 179
 local 12, 123, 135
 Stalnakerian 76, 77, 78, 79, 81, 101, 104, 106, 134, 167, 171
 context change potential 7, 96, 102, 104, 105, 134, 167
 context-set 76, 77, 102, 110, 117
 conversational implicature 62, 97, 98
 conversational record 10, 12, 13, 16, 17, 91, 92, 126, 134, 136, 137, 138, 150, 153, 154, 166, 171, 178

cross-linguistic 39, 48, 70, 115, 133
Cuna 48

D

definite 10, 45, 56, 77, 133, 181, 183, 184
deictic gesture 11, 47, 48, 49, 50, 53, 54, 57, 65,
72, 80, 146, 148
deixis 48, 49
demonstration 25, 26, 27, 29, 30, 48, 53, 54, 55,
69, 123, 165, 175
demonstrative pronouns 3, 5, 6, 18, 19, 27, 28,
31, 32, 33, 34, 37, 39, 40, 45, 46, 56, 57, 63,
71, 72, 79, 90, 92, 93, 113, 121, 122, 136,
152, 172, 173, 174
demonstratives 2, 3, 19, 23, 24, 25, 26, 27, 28,
29, 30, 31, 32, 33, 38, 45, 46, 72, 79, 92,
113, 171, 172
true 2, 3, 19, 23, 25, 26, 27, 29, 30, 31, 32,
33, 45, 46, 72, 79, 92, 113, 171, 172
disagreement 101, 135, 177
discourse coherence 5, 6, 12, 13, 61, 63, 65,
124, 126, 131, 133, 135, 138, 148, 179,
180, 182, 183
discourse conventions 5, 8, 10, 13, 16, 17, 18,
19, 20, 40, 41, 57, 58, 113, 120, 121, 122,
123, 126, 130, 133, 143, 146, 148, 152,
163, 168, 171, 172, 173, 178, 181, 183, 184
discourse relations 6, 12, 63, 126, 159
Discourse Representation Theory 41
discourse structure 11, 65, 73, 128, 131, 133,
150, 151, 161, 163, 164, 165, 168
discourse-bound 34, 35, 36, 90, 93, 114,
116, 172
disjunctive syllogism 145
domain restriction 99, 113, 117–121, 126, 132,
150, 153, 157, 166, 178, 179, 180, 181,
182, 184, 200
dthat 25
dynamic meaning 8, 17, 18, 41, 85, 92, 94, 95,
96, 104, 106, 134, 135, 136, 138, 139, 150,
151, 152, 153, 154, 161, 168, 171
dynamic semantics 8, 19, 35, 36, 40, 41, 51, 71,
79, 95, 102, 104, 105, 107, 110, 117, 144,
167, 189, 198, 199

E

E-type 34
Elaboration relation 122, 123, 124, 127, 129,
130, 131, 133, 138, 147, 148, 149, 150,
158, 159, 160, 162, 166, 180, 181, 182, 183
entailment 32, 146, 167, 172, 191, 198
test-to-test 167
update-to-test 146, 167
epistemic contradiction 100, 127, 128, 129,
132, 133

Explanation relation 62, 63, 64, 65, 66, 67, 70,
73, 74, 182, 192
Object-based 73, 74
Subject-based 73
expressivism 6, 95, 96, 97, 99, 101, 104, 106,
107, 108, 109, 110, 111, 117, 123,
144, 167
extra-linguistic vii, 2, 3, 4, 5, 6, 7, 8, 10, 11, 12,
14, 16, 17, 18, 19, 23, 25, 27, 28, 30, 31, 37,
38, 39, 40, 45, 46, 49, 50, 53, 54, 71, 75,
113, 120, 121, 134, 171, 172

F

free enrichment 16

G

gradable adjectives 182, 183, 184

I

identification thesis 14–18
imperative 8
implicit arguments 183
implicit causality 73
incomplete descriptions 181
indefinite 33, 35, 36, 37, 38, 42, 43, 44, 45, 46,
49, 51, 52, 55, 57, 58, 60, 61, 87, 89, 96,
115, 123, 133, 134
index 24, 95, 107, 108, 109, 110
indexicals 2, 3, 23, 30, 31, 32, 55, 71, 72, 90,
113, 166, 171, 172, 191
pure 2, 3, 30, 31, 32, 55, 71, 72, 113,
171, 172
indexing 27, 28, 29, 34, 37, 38, 39, 41, 44, 45,
46, 51, 53, 54, 60, 148
information state 95, 106, 108, 109, 110, 111,
112, 113, 118, 124, 135, 136, 177
intentionalism 54
intentionalist 4, 25, 26, 45, 46, 53, 120
interrogative 8

J

Japanese 133

L

linguistic convention 5, 6, 17, 46, 48, 53, 59,
63, 65, 69, 72, 75, 76, 80, 81, 129, 148,
184, 185
logical form 16, 18, 27, 28, 29, 30, 31, 38, 39,
45, 46, 49, 51, 52, 53, 54, 55, 65, 66, 67, 71,
86, 108, 123, 146, 148, 151, 152, 160, 163,
164, 165, 168, 192, 198

M

modal subordination 114, 123, 124, 126,
134, 135

modality 19, 20, 95, 99, 101, 118, 135, 143, 153, 177, 193

modals 3, 15, 19, 93, 94, 95, 98–112, 113, 114, 115, 116, 117, 118, 119, 120, 121, 123, 124, 126, 127, 128, 129, 130, 132–138, 143, 146, 148, 149, 151, 153–163, 165, 166, 167, 168, 172, 176, 177, 178, 181, 183, 194–200

modus ponens 19, 143, 145–147, 150, 151, 161, 165, 166

modus tollendo ponens 145

modus tollens 143–149, 151, 152, 158, 161, 163–167

Moore-paradoxical 100

N

Narration relation 63

narrative 11, 43, 79, 80, 175, 183
structure 11

non-propositionalism 6, 7, 8, 16, 17, 18, 19, 87, 93, 95, 96, 98, 99, 101, 104, 106, 108, 110, 112, 117, 118, 130, 134, 135, 136, 138, 143, 144, 146, 151, 152, 166, 167, 168

Northamptonshire 47

O

object-bias 73

objectivist 4, 45, 46

P

Parallel relation 12, 63, 64, 68, 69, 70, 131

point of evaluation 16, 24, 107, 108

pointing 1, 2, 3, 11, 13, 16, 25, 27, 28, 30, 33, 37, 40, 42, 44, 45, 46, 47, 48, 49, 50, 51, 52, 53, 54, 55, 65, 69, 78, 80, 85, 86, 89, 122, 123, 147, 148, 163, 168, 173, 174

pragmatic inference 97, 98, 164

prejacent 100, 103, 105, 107, 117, 119, 133, 136, 137, 155, 156

Principle B 29, 56, 58, 59, 60

pro-drop 69

prominence 13, 19, 20, 40, 41, 42, 43, 45, 49, 55, 57, 58, 59, 60, 65, 66, 67, 68, 69, 70, 71, 72, 73, 90, 92, 113, 114, 116, 117, 122, 123, 124, 126, 129, 131, 133, 134, 136, 137, 138, 148, 151, 153, 154, 155, 156, 157, 158, 159, 160, 166, 171, 172, 173, 178, 179, 180, 181, 183, 184, 195

ranking 40, 43, 45, 49, 55, 59, 67, 71, 72, 87, 90, 92, 123, 126, 136, 137, 138, 150, 151, 153, 154, 155, 157, 172, 195

-sensitive 19, 20, 42, 55, 71, 90, 113, 114, 116, 117, 136, 148, 151, 160, 172, 173, 178, 183

proposition 6, 8, 14, 15, 16, 17, 19, 23, 24, 67, 76, 77, 81, 85, 86, 88, 89, 91, 92, 93, 95, 96, 97, 98, 99, 100, 101, 102, 103, 104, 105, 106, 107, 108, 109, 110, 113, 117, 118, 119, 120, 124, 126, 127, 128, 131, 133, 134, 136, 137, 138, 149, 150, 151, 152, 153, 154, 155, 156, 157, 159, 160, 161, 162, 163, 168, 177, 194, 197, 199, 200

propositional content 7, 8, 14, 15, 16, 17, 18, 24, 85, 86, 87, 88, 90, 91, 92, 93, 94, 95, 96, 97, 98, 99, 101, 104, 105, 106, 107, 108, 109, 112, 126, 134, 135, 136, 137, 138, 139, 149, 150, 151, 153, 154, 155, 156, 158, 159, 162, 165, 166, 168, 177, 197

Q

quantification 99, 108, 113, 117, 119, 120, 126, 136, 148, 149, 153, 155, 179, 200

quantifier 15, 33, 34, 35, 36, 42, 43, 44, 70, 99, 117, 119, 132, 178, 179, 180, 181, 183, 184

adverbial 178, 179

nominal 178, 184

R

reductio ad absurdum 145

reference 4, 30, 56, 63, 65, 71, 81, 86, 175, 182, 190

relational expressions 183

relativism 95, 96, 109, 144

representational content 6, 7, 8, 14, 16, 18, 41, 86, 87, 93, 94, 95, 96, 97, 112

restrictor 117, 119, 120, 123, 124, 126, 127, 128, 129, 130, 131, 136, 137, 138, 148, 149, 153, 155, 156, 157, 160, 163, 179, 180, 197, 198, 199, 200

restrictor-analysis 123, 165

Result relation 12, 64, 65, 66, 67, 69, 70, 192

S

salience 3, 4, 9, 10, 13, 33, 79, 80

Serbian 69

Sheffer stroke 165

situated utterance 174

speaker intention 1, 2, 3, 4, 5, 6, 8, 9, 25, 26, 29, 45, 46, 51, 53, 54, 55, 65, 67, 113, 132, 171, 174, 184

Specification relation 182

stack 41, 42, 43, 44, 56, 58, 59, 72, 87, 90, 135,
137, 153, 154, 189, 190, 195, 196
standing linguistic meaning 1, 2, 3, 10, 11, 12,
24, 25, 26, 27, 30, 31, 32, 33, 38, 39, 40, 45,
52, 56, 57, 71, 80, 121, 126, 171
subject-bias 73, 74
subject-object preference 59, 61
subject-prominent language 59
Summary relation 174, 176

T

Tarskian semantics 34, 35, 36
test 103, 105, 106, 107, 110, 111, 117,
146, 167
topic-prominent language 59
truth 15, 16, 24, 28, 32, 52, 86, 87, 88, 91, 95,
98, 102, 108, 109, 167, 168, 191, 198
truth-conditional content 6, 7, 8, 12, 14, 17,
18, 19, 23, 35, 41, 53, 62, 81, 85, 86, 87,
88, 92, 93, 94, 96, 97, 98, 100, 101, 105,
107, 108, 110, 113, 118, 119, 120, 121,
122, 126, 130, 133, 134, 137, 138, 139,
144, 150, 152, 154, 155, 156, 157, 159,
161, 163, 164, 165, 171, 172, 176, 194, 197

U

underspecification 27, 29, 30, 31, 38, 44, 45,
51, 53, 54, 65, 75
update semantics 95, 97, 101, 102, 104, 105,
107, 110, 111, 117, 123, 144, 146, 167
update system 104, 105, 106
distributive 105, 106
eliminative 105, 106
intersective 104, 105, 106

V

validity 24, 26, 28, 31, 32, 51, 52, 86, 87, 143,
146, 152, 164, 165, 166, 168, 172,
191, 198
variable 34, 35, 36, 37, 38, 39, 41, 42, 43, 44,
46, 51, 60, 72, 154, 155, 189, 190, 191,
193, 196
bound 34, 35, 36, 37, 38
free 34, 37, 38, 51

W

Warlpiri 48, 115, 116
Warumungu 48
what is said 97, 98