The Mereology of Attitudes

A Dissertation Presented
by
Robert Stern Pasternak
to
The Graduate School
in Partial Fulfillment of the
Requirements
for the Degree of
Doctor of Philosophy
in
Linguistics

Stony Brook University

August 2018

Stony Brook University

The Graduate School

Robert Stern Pasternak

We, the dissertation committee for the above candidate for the
Doctor of Philosophy degree, hereby recommend
acceptance of this dissertation.

Richard K. Larson — Dissertation Advisor
Professor, Department of Linguistics

Thomas Graf — Chairperson of Defense
Assistant Professor, Department of Linguistics

Jiwon Yun
Assistant Professor, Department of Linguistics

Angelika Kratzer
Professor Emerita, Department of Linguistics
University of Massachusetts Amherst

This dissertation is accepted by the Graduate School

Dean of the Graduate School

Abstract of the Dissertation

The Mereology of Attitudes

by
Robert Stern Pasternak
Doctor of Philosophy
in
Linguistics

Stony Brook University
2018

This dissertation explores the intersection of attitude semantics—the semantics of lexical items expressing concepts like desire and belief—and event semantics. More specifically, this work constitutes an initial foray into the nature of the part-whole relations of attitude states, and the impact these mereological structures have on semantic interpretation. The two main topics covered are (i) the mereological basis of attitude intensity and (ii) non-distributive ascriptions of belief.

After a cursory overview (Chapter 1) and a discussion of prior theories of the semantics of attitudes (Chapter 2), in Chapter 3 I argue that in the model used for semantic interpretation, the intensity of mental states in general, and attitudes in particular, is encoded in the part-whole structure of such states. Put simply, a more intense desire state is "bigger" in a particular dimension than another, less intense desire state. The crux of the argument is that there are several measurement-related constructions in English that impose requirements relating the measure function to the part-whole structure of the measured domain, and all of these constructions can be used to measure the intensity of mental states. While some data from Chinese initially complicate the picture, I show that the similarities and differences between Chinese and English are best accounted for by positing that intensity does indeed track the part-whole structure of attitude states.

Having established the mereological basis for attitude intensity, in Chapter 4 I provide a natural language metaphysics of desire states that meets the conditions set forth in the previous chapter. This requires intertwining the traditional ordering and quantification over possible worlds with the part-whole structure of attitude states, as well as imposing requirements about how these world-orderings and part-whole structures must relate to each other.

In Chapter 5 I shift gears, discussing cases in which beliefs are ascribed to a plurality that cannot be ascribed to the individuals that make up that plurality (i.e.,

non-distributive belief ascriptions). To account for these cases, I offer principles relating the beliefs of plural individuals to those of their atomic parts, focusing on how (dis)agreements between individual epistemic agents are negotiated. The account is then revised so that the mechanisms at play can differentiate between relevant and irrelevant disagreements between individual experiencers.

Contents

Acknowledgments viii

1 Introduction **1**

1.1 The panoramic view 1

1.2 Roadmap 4

1.3 On citation 5

2 A micro-history of attitude semantics **7**

2.1 Flat quantification 8

2.1.1 The classical theory of modals 8

2.1.2 Translation to belief and desire 9

2.2 Best-worlds quantification 11

2.2.1 Background: Kratzer's modal semantics 12

2.2.2 von Fintel 1999 18

2.3 Comparison to alternatives 22

2.3.1 Heim 1992 22

2.3.2 Villalta 2008 24

2.4 Bayesian 36

2.4.1 Probabilistic belief 37

2.4.2 Utility-based desire 39

2.5 (Not) picking sides: Is *want* (Strawson) upward-entailing? 43

2.6 From *want* to *wish* and *regret* 46

2.7 Where's the quantification? 49

2.8 Conclusion 54

3 Intensity is monotonic **55**

3.1 What is monotonicity? 57

3.2 Constructions with monotonicity requirements 59

3.2.1 Pseudopartitives 59

3.2.2 *Out the wazoo* and *in spades* 62

3.2.3 Adverbial measure phrases 63

3.2.4 Nominal comparatives . . . 63
3.2.5 Verbal comparatives . . . 64
3.2.6 Summary . . . 65
3.3 A compositional semantics for (verbal) comparatives . . . 65
3.3.1 Adjectival comparatives . . . 66
3.3.2 Adverbial comparatives . . . 70
3.3.3 Verbal comparatives . . . 75
3.4 Back to intensity comparatives . . . 82
3.5 An alternative hypothesis: Lexical gradability . . . 86
3.6 Against the LGH . . . 92
3.6.1 The evidence from English . . . 92
3.6.2 The evidence from Chinese . . . 94
3.7 Conclusion . . . 98

4 Two-dimensional attitudes **99**
4.1 Non-attitude mental states: The case of *hate* . . . 99
4.1.1 The ontology of intensity: Going vertical . . . 100
4.1.2 Deriving *hate* comparatives . . . 103
4.2 Transitioning to *want* . . . 107
4.2.1 Semanticizing homogeneity . . . 107
4.2.2 Two options for WANT . . . 109
4.2.3 Verbal comparatives with *want* . . . 111
4.3 Desires of varying intensity . . . 113
4.3.1 The problem . . . 113
4.3.2 Lessons from extension gap theories of adjectives . . . 114
4.3.3 The nitpicker's guide to graded intensionality . . . 116
4.3.4 Integration with the mereology . . . 120
4.3.5 Summary . . . 121
4.4 Extension to *wish* and *regret* . . . 122
4.5 Consequences and comparisons . . . 123
4.5.1 Some consequences . . . 123
4.5.2 Some comparisons . . . 131
4.5.3 Summary . . . 136
4.6 Conclusion . . . 136
Appendix: Implementation in premise semantics . . . 138

5 Non-distributive ascriptions of belief **150**
5.1 Introduction . . . 150
5.2 Basic assumptions about plurals . . . 152

5.3 Evidence of non-distributivity and compositional semantics of the two interpretations . . . 154
5.3.1 Non-distributive ascriptions of belief . . . 154
5.3.2 Compositionality of distributive interpretation . . . 157
5.3.3 Compositionality of non-distributive interpretation . . . 160
5.3.4 Defining the problem . . . 162
5.4 Premise semantics and belief-summing . . . 163
5.4.1 A Kratzerian analogy . . . 163
5.4.2 Extension to belief-summing . . . 166
5.5 (Ir)relevance and aboutness . . . 169
5.5.1 The problem of irrelevant disagreement . . . 169
5.5.2 Beliefs are about situations . . . 170
5.5.3 Fixing the irrelevant disagreement problem . . . 170
5.5.4 Context-dependency . . . 172
5.6 Conclusion . . . 175

6 Conclusion **177**

Bibliography **179**

Acknowledgments

While there may be no such thing as the perfect advisor, in my opinion Richard Larson comes about as close as it gets. I have benefited enormously from Richard's infectious (and seemingly infinite) enthusiasm, as well as his genuinely extraordinary ability to distill complex ideas into their simplest intuition and justification. He has always been extremely generous with his time and his extensive knowledge, both about linguistics in particular and about navigating grad school and academia in general. I wish every student could have an advisor like Richard, and I will always be grateful for having had that opportunity myself.

Thomas Graf has taught me a great deal about what it means to be a formally-minded linguist. Among other things, he has helped me appreciate the difference between a clever formalism and the insight that that formalism generates, as well as the instrumental value formal devices have in understanding the human language apparatus. I am grateful to Thomas for such insights, as well as for the many fun non-linguistic conversations we've had over the years, often hashing out petty disagreements about pop culture minutiae. You'd be surprised at how much an argument about *Seinfeld* can help when you're freaking out about a qualifying paper.

Jiwon Yun taught the first two classes in formal semantics that I ever took: an introductory semantics class in the fall of my first year, and a wonderful seminar on indefinites the following spring. Thus, given that I came to Stony Brook with no background in linguistics whatsoever, if I left grad school a halfway competent semanticist then she deserves a substantial portion of the credit. (If I remain incompetent, I'll take the blame.) In addition to teaching me the tools of the trade, Jiwon was a big part of helping me understand in a more fundamental sense what it is that semanticists actually do and care about. It is this, more so than any lesson on lambdas or quantifier raising, that I thank Jiwon for.

As a quick skim of this dissertation or a glance at the bibliography will make clear, Angelika Kratzer's work has massively influenced my own research. But her influence on this dissertation extends well beyond her writings. When I spent a semester visiting UMass in Fall 2016, Angelika and I would meet on a weekly basis, often discussing whatever ill-thought-out ideas were rattling around my skull at the time. Those conversations were crucial in the development of many of the ideas in this dissertation, especially in the early stages. In addition, she always made me feel like what I was doing was interesting, a feeling that gave me much-needed confidence as I was exploring some of the quirkier nooks and crannies of my research. For this and much else, I am thankful.

Philippe Schlenker was not technically on my committee, but I owe him a special degree of gratitude for the million ways in which he has helped me along the way, as an instructor, collaborator, colleague, and friend. In addition to fostering my interest in co-speech supplements and formal pragmatics, he was a constant source of helpful advice on all aspects of navigating the world of academia. But most importantly, Philippe represents to me the model of the innovative spirit of research, always willing to leap straight into completely uncharted territory for the sake of greater understanding. I can only hope that at least a little bit of this spirit has made its way into my own work.

Stony Brook is one heck of a place to study linguistics, and I will always look back fondly on my time here. Many thanks to all of the Stony Brook faculty who were kind enough to teach me: Mark Aronoff, John Bailyn, Michael Becker, Christina Bethin, Ellen Broselow, Dan Finer, Thomas Graf, Bob Hoberman, Richard Larson, Gary Mar, and Jiwon Yun. Sandra Brennan and Michelle Carbone were of constant and enormous help—and often showed near-infinite patience—with all things administrative.

I would also like to thank all the faculty at other institutions who were kind enough to let me sit in on their classes: Philippe Schlenker at NYU; Lyn Frazier, Jeremy Hartman, Vincent Homer, Troy Messick, and Tom Roeper at UMass; Kai von Fintel, Danny Fox, Irene Heim, Sabine Iatridou, Loes Koring, and Roger Schwarzschild at MIT; and Gennaro Chierchia at Harvard. Additional thanks to everyone who I spoke with about matters related to this dissertation, beyond those individuals already mentioned: Rajesh Bhatt, Seth Cable, Lucas Champollion, Masha Esipova, Valentine Hacquard, Jeff Lidz, Maša Močnik, Friederike Moltmann, Roumi Pancheva, Milo Phillips-Brown, Barry Schein, Viola Schmitt, Ildikó Emese Szabó, Alexis Wellwood, and Alexander Williams. Earlier incarnations of this work and predecessors thereto were presented at the NYU Semantics Reading Group, the Workshop on Modality Across Categories (Universitat Pompeu Fabra), SALT 26 (UT Austin), the UMass Semantics Workshop, the MIT Syntax-Semantics Reading Group (LFRG), the 2018 LSA annual meeting in Salt Lake City, and SALT 28 (MIT). My thanks to these audiences for their stimulating questions and comments.

I've been fortunate enough to have many incredible friends in my life—from Williams, Stony Brook, the New York-St. Petersburg Institute, NYU, Institut Jean Nicod, UMass, MIT, Harvard Law School, the Jazz Loft jam sessions, and elsewhere—during my time in grad school. I wouldn't have survived the long transition from Mr. Pasternak to Dr. Pasternak were it not for their unwavering support. While there are too many people to name individually, I hope you each know just how important you have been and still are to me. That being said, Will Piereson, Ben Coleman, and Pete Mertz deserve special mention. Our mixture of dumb jokes and real talk got me through some tough times these past five years (and for years before), and I will always be grateful for your friendship, in all the many forms it has taken.

Lastly, and most importantly, I would like to thank my family: my parents Reuven and Nancy, my brothers Andy and Jordan, my sister-in-law Kate, my grandparents, aunts, uncles, and cousins, and my partner, Ildi Szabó. You all mean the world to me, and while most of you will probably never understand any of the gibberish in this dissertation—I never did get very good at that elevator pitch—this gibberish wouldn't exist if it weren't for you. Thanks especially to Mom and Dad, whose endless and unconditional love and support made failures feel small and successes feel enormous; and to Ildi, who talked me through a million minor freak-outs and made me feel like I could do just about anything. I love you all.

This dissertation is dedicated to the memory of my late grandfather, Richard Pasternak ("Peepa"), who passed away shortly after the dissertation was defended.

Chapter 1: Introduction

1.1 The panoramic view

Davidson (1967) famously argues that the denotations of (certain) sentences contain a variable over *events*, which is then (usually) existentially quantified over. Thus, whereas *Mary hugged Ann* would be assigned a denotation along the lines of (1a) in an event-free semantics, Davidson proposes that it really has a denotation like (1b):

(1) a. $\mathrm{hug}(\mathrm{mary}, \mathrm{ann})$
b. $\exists e[\mathrm{hug}(e, \mathrm{mary}, \mathrm{ann})]$

In the five decades since Davidson's paper, linguists and philosophers have tirelessly poked and prodded at this event variable, with a variety of questions arising in the process:

- How are events individuated? That is, what are the criteria for identity or distinctness of events? (Davidson 1969; Carlson 1984, 1998)
- The domain of entities has been argued to be structured in a way that encodes part-whole relations (Link 1983). Furthermore, the nature of the part-whole relations in the domains of noun phrases has been shown to be reflected in the grammar in a variety of ways. To what extent do parallel facts hold for verb phrases and events? (Bach 1986a, Krifka 1989, a.m.o.)
- As first observed by Castañeda (1967), incorporating an event argument allows other arguments of the verb to be separately introduced, both in the logical representation of the denotation of a sentence and in the syntactic-semantic derivation itself. For example, in addition to Davidson's (1b), denotations like (2a) and (2b) have been proposed, directly injecting thematic relations as a means of introducing some or all of the arguments of the verb (and/or functional verbal heads):

(2) a. $\exists e[\mathrm{Agt}(e) = \mathrm{mary} \wedge \mathrm{hug}(e, \mathrm{ann})]$ (Kratzer 1996)
b. $\exists e[\mathrm{Agt}(e) = \mathrm{mary} \wedge \mathrm{hug}(e) \wedge \mathrm{Thm}(e) = \mathrm{ann}]$ (Carlson 1984)

Which arguments are (not) separated, how are they divvied, and how are they introduced in the syntax and semantics?

- How do events relate to their participants? How do readings of collectivity, cumulativity, distributivity, etc., arise and relate to each other?

One foot of this dissertation will be planted in this tradition of inquiry, and in the proposals and results that have sprung therefrom. The other foot will be planted in a semantic tradition that until recently led a relatively independent life: the interpretation of attitude verbs like *believe*, *know*, *want*, *wish*, and *regret*. Hintikka (1969) was the first to propose the quantificational treatment of attitude verbs, in which attitudes universally quantify over sets of possible worlds. On such an approach, attributions of belief and desire have denotations along the lines of (3a) and (3b), respectively.

(3) a. $[\![\text{Jo believes } p]\!] = 1$ iff p holds in all worlds compatible with Jo's beliefs.
b. $[\![\text{Jo wants } p]\!] = 1$ iff p holds in all worlds compatible with Jo's desires.

Of course, a variety of proposals for the semantics of attitudes have been proferred in subsequent years, with varying degrees of faithfulness to Hintikka's proposal.

More recently, there have been attempts to fuse these two previously separate semantic traditions. For example, Hacquard (2006, 2010) and Anand & Hacquard (2008) use a Davidsonian semantics for attitude verbs to account for facts having to do with epistemic modals embedded under attitudes. To illustrate the problem, consider (4), with its unembedded use of epistemic *might*:

(4) Joanna might be in New York.

A simplistic analysis of (4) would be to say that it is true iff Joanna's being in New York is compatible with the cumulative beliefs of the conversational participants. Now consider (5), in which (4) is embedded under the attitude verb *think*:

(5) Chip thinks that Joanna might be in New York.

The strongly preferred default interpretation of (5) can be summarized as claiming that Joanna's being in New York is compatible with Chip's beliefs. At first glance, this is puzzling for two reasons. First, whereas the translation of unembedded *might* in (4) made reference to the cumulative beliefs of the conversational participants, somehow in (5) this discourse dependence is nixed: all that matters is Chip's beliefs. Second, since there are two intensional operators in (5)—*think* and *might*—one might reasonably but erroneously expect (4) to have a "double-modalized" meaning. That is, one would expect (5) to mean that Chip merely *thinks* that Joanna's being in New York is compatible with his (or someone else's) beliefs. But by all appearances, the interpretation of (5) does not involve such higher-order belief.

Hacquard and Anand & Hacquard note that these puzzles can be jointly solved by adopting a Davidsonian theory of both modals and attitudes. In short, modals are interpreted relative to events, with modals' event arguments being bound by higher operators. In the case of an unembedded epistemic modal, this event argument is bound by a speech act operator in the left periphery. Hence, the modal will be evaluated relative to the information contained in the speech act event, such as the beliefs of the conversation participants. But when the modal is embedded under *think*, the event argument is identified with the event argument of *think*; put informally, the *think*-ing and the *might*-ing are the same event. This explains both the lack of sensitivity to other conversational participants—in a thinking event, the only person whose thoughts matter is the experiencer's—and the absence of a doubly-modal reading, since the attitude and the modal are collapsed together into a single event.[1]

For the most part, work like this at the intersection of attitudes and event semantics has focused on what benefits an event argument confers with respect to compositional factors like argument structure and variable binding. However, one area of inquiry in event semantics that has borne a great deal of fruit there, but that has generally been put aside in Davidsonian treatments of attitudes, is how the part-whole structure of events can affect entailments and acceptability in a variety of constructions. This dissertation is an attempt to start filling this gap, shedding light on the rich inner structure of states of desire, belief, regret, etc. The research in the following chapters thus constitutes part of the tradition of what Bach (1986b) famously refers to as *natural language metaphysics*: the study of the nature of the models used for semantic interpretation, as well as their semantic repercussions.

This dissertation focuses on two main topics in the mereology of attitude states: (i) the way in which intensity is manifested in the part-whole relations of attitude states; and (ii) non-distributive belief ascriptions, i.e., the ascription of beliefs to pluralities that cannot be ascribed to the individuals of which they are constituted. Many times throughout this dissertation, the running thread will be that on a compositional level, attitudes are in many respects unexceptional: they exhibit similar sorts of properties as seemingly simpler verbs like *run* and *eat*. More so than my own particular results and proposals, it is the utility and viability of this general research program that I hope will be conveyed to the reader over the course of this dissertation. In short: while the lexical semantics of attitudes is of unique importance for certain linguistic and philosophical reasons, it is crucial to bear in mind that as far as the compositional semantics is concerned, there is a large extent to which attitude verbs are simply verbs, and behave syntactically and semantically as such.

[1] For event-free analyses of these and similar facts, see Yalcin 2007, Anand & Hacquard 2013.

1.2 Roadmap

The structure of the rest of this dissertation is as follows. Chapter 2 ("A micro-history of attitude semantics") serves as a brief introduction to the history of the semantics of attitudes, starting with the seminal work of Hintikka (1969) and continuing to more contemporary work. The discussion begins with an overview of four of the major types of proposals for the semantics of *believe* and (especially) *want*: flat quantification (Hintikka 1969), best-worlds quantification (von Fintel 1999), comparison to alternatives (Heim 1992, Villalta 2008), and Bayesian (Levinson 2003, Yalcin 2007, Lassiter 2011a,b). Differences between the proposals with respect to predicted entailments in the clausal complement of *want* are discussed, with particular emphasis on whether or not this environment is predicted to be *(Strawson) upward-entailing*. While I remain non-committal on this issue, the formal proposals I make in subsequent chapters require the commitment to a semantics for *want*. With this in mind, I formalize my analysis using the theories of Heim (1992) and von Fintel (1999), as their similarities and differences provide a particularly compelling lens through which to observe what commitments my own proposal does and does not make. The chapter then concludes with the discussion of two more areas of work in attitude semantics: the relation between the semantics of *want* and the semantics of *wish* and *regret*, and theories that propose exporting some of the semantic work of attitudes to other, nearby syntactic heads.

Chapters 3 ("Intensity is monotonic") and 4 ("Two-dimensional attitudes") focus on the issue of the mereological basis of mental state intensity in general, and attitude intensity in particular. In Chapter 3, I argue that intensity tracks the part-whole relations of mental states, meaning that it is, in Schwarzschild's (2002, 2006) terms, a *monotonic* measure function. In other words, for two mental states to have differing intensities is for them to have differing sizes, at least along a particular dimension. In arguing for this claim, I go through a variety of measurement constructions that impose monotonicity requirements on the chosen measure function, showing that the intensity of mental states can be measured using each construction. I also argue against an alternative proposal, rendered initially plausible by evidence from Chinese, which pushes back against claims of monotonicity by positing important differences in lexical semantics and syntactic structure. In fact, I show that when we look closer at the Chinese data, we see that not only does it not constitute sufficient evidence *against* a monotonic view of mental state intensity, but it actually constitutes evidence *in favor* of such a natural language metaphysics.

Given the evidence in Chapter 3 that intensity tracks the part-whole structure of mental states, in Chapter 4 I turn to the question of what a natural language metaphysics meeting this condition might look like. In short, I propose that mental states extend in two dimensions: "horizontally" through time, and "vertically" in the dimension along which intensity is measured. Starting with the simpler case of non-attitude mental states like hatred, I explore the relationship between the semantics of a verb like *hate* and the part-whole structure of two-dimensional states of hatred. I then turn my attention to attitudes (and in particular *want*), showing how the additional complication of ordering and quantification over possible worlds can be smoothly integrated with the part-whole structure of desire states. I also show that my proposal makes certain semantic predictions independent of the mereological prediction of monotonicity, and that these predictions are in fact correct.

In Chapter 5 ("Non-distributive ascriptions of belief") I turn to a different phenomenon in the mereology of attitudes: namely, the interpretation of plural subjects of belief reports. Typically, attitude ascriptions with plural subjects are interpreted distributively. Hence, (6) is usually interpreted as entailing both (7a) and (7b):

(6) Alex and Bertha think that Marty left.

(7) a. Alex thinks that Marty left.
    b. Bertha thinks that Marty left.

In spite of this tendency, I provide evidence suggesting that *non*-distributive belief ascriptions are in fact possible, meaning that beliefs can be attributed to pluralities that cannot also be attributed to the individuals that make up that plurality. The next question to address is what the relationship is between the beliefs of individuals and the beliefs of the pluralities formed from them. Or, in Davidsonian terms, how does the content of two belief states relate to the content of their sum? I show that the ways in which individuals' (dis)agreements are negotiated in forming the beliefs of the plurality are just those one would expect from a Lewis-Kratzer premise semantics (Lewis 1981, Kratzer 1981a). I therefore posit a natural language metaphysical principle of belief summing built on this formal mechanism. This acount is then refined in order to allow for a distinction between relevant disagreements, which have an impact on the way belief states are combined, and irrelevant disagreements, which have no such impact.

Chapter 6 ("Conclusion") ties a bow on the dissertation, offering some concluding remarks and additional areas for future study, beyond those offered in preceding chapters.

1.3 On citation

Before moving on, it is worth noting that various aspects of the work in this dissertation will be appearing in separate publications in the near (or near-ish) future. Some remarks on citation are thus in order.

First, some portions of the material in Chapters 3 and 4 are taken, often verbatim, from drafts of a paper of mine recently accepted for publication in *Linguistics and Philosophy* (Pasternak in revision). Wherever there is overlap, I ask that the reader please cite that paper instead of, or in addition to, this dissertation.

Second, the data and analyses discussed in Chapter 5 are also addressed in a forthcoming proceedings paper of mine (Pasternak in prep). Once again, where there is overlap, please cite that paper (once it has been published).

Chapter 2: A micro-history of attitude semantics

In this chapter, I will go over—in a very cursory fashion—some of the more prominent modern work on the semantics of attitudes. The goal of this chapter is not so much to attain conclusive definitions, or even to give a full explanation of the motivations for each theory, but rather to take a guided tour of the hypothesis space and have a quick glance at what kinds of issues have arisen in the history of model-theoretic research on attitudes.

In Sections 2.1–2.4 I will introduce in order of formal complexity what I take to be the four most popular kinds of theories of the semantics of attitudes, and in particular of *think*/*believe* and *want*. The simplest of these analyses, discussed in Section 2.1, is Hintikka's (1969) classic theory of *believe* and *want*. Next up in Section 2.2 is von Fintel's (1999) "best worlds" theory for the semantics of *want*, which is built on the modal semantics of Kratzer (1981a, 1991, 2012). In Section 2.3, I will discuss theories of *want* involving comparison of a proposition to certain alternatives; this includes Heim's (1992) influential theory, as well as the focus-sensitive analysis of Villalta (2008). Section 2.4 is devoted to a type of theory that differs in more fundamental ways from the ones before: namely, proposals that make use of probabilities in the semantics of *believe*, and the decision-theoretic metric of expected utility in the semantics of *want* (e.g., Levinson 2003; Yalcin 2007; Lassiter 2011a,b).

During the discussion of these four types of theories, frequent reference will be made to predicted entailments in the clausal complements of these attitudes. In Section 2.5 I will discuss one particularly controversial question in this area: namely, whether and to what extent the complement of *want* is an upward-entailing environment. Finally, in Sections 2.6 and 2.7, I will go over some other developments in the semantics of attitudes that are more or less independent of the particular choice of proposals in the first four sections. In 2.6 I discuss the case of "counterfactual" attitudes like *wish* and *regret*, and in 2.7 recent proposals that a significant proportion of the semantic work of attitudes actually comes from the semantics of the embedded clause, and not from the semantics of the verb itself. A brief conclusion follows.

Before starting the tour, it's worth noting that many of the issues that will arise in this chapter are orthogonal to the ones addressed in the rest of this dissertation, so in certain respects the choice of starting theory will not matter. However, the formal implementation of my proposals in Chapters 4 and 5 will require choosing a particular theory of attitude semantics, so in those chapters I will build off of Hintikka's (1969) theory for the semantics of *believe*, and the theories of von Fintel (1999) and Heim (1992) for *want*. My choices in this respect are mostly a matter of convenience. Hintikka's semantics for *believe* is among the simpler options on the market; von Fintel and Heim's theories are formally similar enough that both can be implemented in more or less the same way in the proposal in Chapter 4, while being different enough to illustrate some important dimensions along which my proposal is agnostic.

2.1 Flat quantification

Hintikka's (1969) semantics for attitudes is on a formal level the simplest of the proposals discussed in this chapter. For Hintikka, attitudes have denotations parallel to what was then the mainstream view of modal auxiliaries like *must* and *can*: universal and existential quantification (respectively) over a set of accessible worlds. I will first go over this classical view of modality, and then I will discuss how Hintikka extends it to attitudes.

2.1.1 The classical theory of modals

On a classical account, modal auxiliaries take as arguments a proposition p (often called the *prejacent*) and a world w (the world of evaluation) and return a truth value. In the case of *must*, the result is true if and only if all possible worlds accessible from w are worlds such that p is true. For *can* or *may*, the result is true if and only if *some* possible world accessible from w is a world such that p is true. These definitions can be seen in (1), where $R(w)$ is the set of worlds accessible from w:

(1) a. $[\![\text{must}]\!]_{\text{classical}} = \lambda p \lambda w.\ \forall w' \in R(w)[p(w')]$
b. $[\![\text{can}]\!]_{\text{classical}} = \lambda p \lambda w.\ \exists w' \in R(w)[p(w')]$

While the classical account has some critical faults discussed in the next section, it does generate some nice results as far as predicted entailments between necessity and possibility modals. For example, the apparent semantic equivalence of (2a) and (2b) is rightly predicted:

(2) a. You are not required to go to the party.
b. You are allowed not to go to the party.

To see this, consider the predicted interpretations:

(3) a. $[\![\text{(2a)}]\!]_{\text{classical}} = \lambda w.\ \neg\forall w' \in R(w)[\text{party}(w')]$
b. $[\![\text{(2b)}]\!]_{\text{classical}} = \lambda w.\ \exists w' \in R(w)[\neg\text{party}(w')]$

(3a) is true of a world w iff not all worlds accessible from w are worlds in which you go to the party. Because of the dual relationship between universal and existential quantification, this is equivalent to saying that (3a) is true of a world w iff there is some accessible world in which you don't go to the party. But that is exactly what (3b) says.

Another prediction made by the traditional theory of modality is that if there is at least one accessible world, then *must* p should entail *can* p. After all, if *every* accessible world is a p world, then *some* accessible world is a p world.[1] This prediction does in fact seem to be borne out:

(4) a. You are required to go to the party.
    b. You are allowed to go to the party.

So much for the classical theory of modality. Next, we will see how Hintikka extends this conception of modality to attitudes.

2.1.2 Translation to belief and desire

The Hintikkan definition of *believe* will look like that for *must*, but will take, in addition to a proposition p and world w, an experiencer x denoted by the subject, and will return true iff p holds in every world compatible with x's beliefs in w (x's *belief worlds*).[2] Similarly, *want* will assert that in all worlds compatible with x's desires in w (x's *desire worlds*), p holds. This can be seen in (5), where $\text{Dox}(x, w)$ is x's belief worlds in w, and $\text{Boul}(x, w)$ is x's desire worlds in w.

(5) a. $[\![\text{believe}]\!]_{\text{Hintikka}} = \lambda p \lambda x \lambda w.\ \forall w' \in \text{Dox}(x, w)[p(w')]$
    b. $[\![\text{want}]\!]_{\text{Hintikka}} = \lambda p \lambda x \lambda w.\ \forall w' \in \text{Boul}(x, w)[p(w')]$

Let's hold off on desire for the time being and ask ourselves what this analysis gets us with respect to the semantics of *believe*. The most important prediction of this theory is that the complement of *believe* should be an upward entailing environment. After all, if p holds in all belief worlds, and p entails q, then q holds in all belief worlds. As a result, we rightly predict examples like those in (6) to be contradictory:

[1] The reason for the caveat that there be at least one accessible world is the following: if there are no accessible worlds at all, then the definition of *must* in (1a) will return true for any proposition, due to vacuous quantification over the empty set of worlds. The definition of *can* in (1b), meanwhile, will return false for any proposition.

[2] Lewis (1979a) and Chierchia (1989) famously argue that propositions do not carry enough information to serve as the denotations of the complements of attitudes, and that instead the embedded clause denotes a (self-ascribed) property. For our purposes this distinction is irrelevant, as talk of propositions (i.e., sets of worlds) can easily be replaced with talk of, say, world-individual pairs. I will stick to propositions.

(6) a. # Zelda believes that I ate at least three pies, but she doesn't believe that I ate at least two.

b. # Zelda believes that I wore a red sweater, but she doesn't believe that I wore a sweater.

c. # Zelda believes that I broke the vase, but she doesn't believe that the vase broke.

Since my eating at least three pies entails my eating at least two pies, any world in which the former is true is a world in which the latter is true. Thus, Zelda's believing that I ate at least three pies is expected to entail her believing that I ate at least two. Similar facts hold for my wearing a red sweater and my wearing a sweater, and for my breaking the vase and the vase breaking. Thus, the examples in (6) are contradictory.

However, there is a well-known downside to this approach when it comes to mathematical or other logically or metaphysically necessary truths. A common view in the philosophical literature is that mathematical truths don't change across possible worlds: there is no possible world in which two plus two adds up to anything other than four. Thus, the propositions denoted by the two sentences in (7) are equivalent, as each is simply the set W of all possible worlds.

(7) a. Two plus two equals four.

b. The sum of the angles of a quadrilateral in a Euclidean space is equal to three hundred and sixty degrees.

This leads to two related problems. First, by definition, any proposition p is such that $p \subseteq W$. So given that belief is predicted to be upward-entailing, anyone who believes anything is predicted to also believe any and all mathematical truths, since any proposition entails them. Second, since the denotations of the clauses in (7) are equivalent, believing one should entail believing the other. Both of these predictions are plainly false. Notice, by the way, that the latter problem is not specific to Hintikka's proposal: so long as the referent of the embedded clause is just a set of possible worlds, and two necessary propositions denote the same set of possible worlds, believing one is predicted to entail believing the other, regardless of what the direction of entailment is in the clausal complement of *want*.

There are at least three general approaches to trying to resolve these problems. The first is to bite the metaphysical bullet: perhaps there actually *are* "impossible worlds" where two plus two does not equal four, and likewise for other mathematical truths (see, e.g., Hintikka 1975). In this case, many propositions will not entail a given mathematical theorem, and the sentences in (7) will have non-equivalent denotations. The second option is to say that mathematical truths are indeed true in all worlds, but that clausal complements denote something with more structure than a proposition—perhaps something that incorporates the syntactic structure of the complement (e.g., Larson & Ludlow 1993). One can then restrict the upward entailment in the complement of *believe* by imposing an additional syntactic requirement: while (7a) entails (7b), the two do not stand in the right syntactic relation for a belief in (7a) to entail a belief in (7b). A third option is to accept that mathematical truths hold in all worlds, as well as that a strict Hintikkan semantics is the right semantics, and to use an alternative semantic mechanism to derive these non-entailments. Cresswell & von Stechow (1982), for example, offer an approach utilizing *de re* construal. For them, a belief in (7a) may be a belief *de re* about the number two, the addition operation, the equality relation, and the number four, while a belief in (7b) may be a belief *de re* about quadrilaterals, Euclidean spaces, etc. On views of the *de dicto*/*de re* distinction commonly adopted for independent reasons, attributing such beliefs *de re* is enough to avoid the undesirable entailments without changing either the metaphysics or the denotations of attitude complements.

A similar problem to those discussed above is what is often known as *logical omniscience*: the Hintikkan account predicts that we are perfect reasoners, in that one's beliefs are expected to be closed under logical deduction. Put differently, given all of my beliefs, I am also predicted to believe all of their logical entailments. We are, of course, not so reasonable. Note that on our current assumptions, this problem is in essence a generalization of the first problem of mathematical truths: mathematical truths are the consequences of all propositions, so we should always believe them. Thus, it may be that whatever successfully resolves the mathematical truths problem will equally successfully resolve the more general problem of logical omniscience.[3]

I raise these problems here only to cast them aside for the remainder of this dissertation. To the extent that they are problematic, they are problematic for every account I am familiar with, or at least those that also make the right predictions about (6). I will thus continue without accounting for mathematical truths or the logical omniscience problem, with the hope that whatever the right semantics of belief is, it is an extension of whatever the right account is for facts like (6).

2.2 Best-worlds quantification

Let's put *believe* on hold for a while and focus on the perhaps more contentious case of *want*. The next approach we will look at are what might be called "best worlds" approaches to *want*, as exemplified by the theory of von Fintel (1999). In the same way that Hintikka's theory was based on the standard view of modality at the time, von Fintel's analysis is built on Kratzer's (1981a, 1991, 2012) influential theory of modality. With this in mind, I will start with Kratzer's theory, and will then move on

[3] For extensive discussion of the problem of logical omniscience, as well as an interesting potential solution along Hintikkan lines, see Yalcin 2016.

to von Fintel's extension to *want*.

2.2.1 Background: Kratzer's modal semantics

Kratzer's semantics for modals famously attempts to resolve three issues faced by (certain) prior approaches to the semantics of modality. I will first introduce the puzzles, then show what Kratzer's theory is and how it aims to resolve those puzzles.

2.2.1.1 Three puzzles

Fine-grainedness The first puzzle has to do with the fine-grainedness of the context-sensitivity of modal auxiliaries. Traditional treatments of modality tended to put the possible interpretations of modals into a small class of categories, such as epistemic (roughly, in view of the collective knowledge of the conversation participants), deontic (in view of one's obligations), bouletic (desires), teleological (goals), ability, etc. Examples of various modals with these interpretations can be seen in (8):

(8) a. **Epistemic:**
Bob's not in the kitchen? Well then he **must** be in the living room.

b. **Deontic:**
You **must** pay your taxes by April 15^{th}.

c. **Teleological:**
To capture the queen you **should** take out your opponent's rook.

d. **Bouletic:**
If you want to go to Harlem, you **have to** take the A train.

e. **Ability:**
Magnus **can** deadlift five hundred pounds.

Keeping the set of possible interpretations small renders it at least plausible, though perhaps a bit unsatisfying on a theoretical level, that these interpretations arise due to a simple lexical ambiguity. According to such a view, English really has multiple homophonous *must*s, one for each type of modality expressable by *must*.

However, Kratzer (1977) observes that distinctions between modal interpretations (what she calls *modal flavors*) are far too fine-grained to be reasonably accounted for by means of lexical ambiguity. In fact, by making use of *in view of* phrases, one can practically set the terms of modal evaluation to whatever one likes, as illustrated in (9):

(9) a. In view of the laws of Alabama, citizens must not carry ice cream in their pockets.

b. In view of the rules of chess, bishops must move diagonally.

If the space of possible modal interpretations is exclusively accounted for by means of lexical ambiguity, then accounting for (9a) requires positing a separate *must* for an interpretation in view of the laws of Alabama, and (9b) requires a separate *must* from the vantage point of the rules of chess. Obviously, if a linguistic theory with half a dozen *must*s, *can*s, and *might*s is a bit ungainly, a theory with a (near-)infinite number of each is simply untenable, so the flexibility in interpretation of modals must at least in part be due to something other than lexical ambiguity.

Conditionals The second puzzle has to do with the relationship between modals and conditionals. Perhaps the best-known example of this is what has come to be known as the *Samaritan Paradox*. To use Kratzer's (1991) example, say that a country has two laws: first, there shall be no murder; and second, if there is murder, the murderer shall go to prison. A "flat" quantificational semantics for modals such as (1), in which worlds are simply divided into those that are ideal/acceptable and those that are not, combined with a naïve view of conditionals based on material implication ($\rightarrow$), leads to some strange results in such a scenario. After all, there are two possibilities: either all deontically acceptable worlds are worlds in which there is no murder, or not. If the latter is the case, then we wrongly predict (10) to be false:

(10) One must not commit murder.

If all acceptable worlds are murder-free, then (10) is rightly predicted to be true, but capturing (11) becomes tricky:

(11) If Kathy commits murder, she must go to prison.

In our simplistic analysis, there are two plausible scope arrangements for (11), shown in (12), where R_D is the deontic accessibility relation:

(12) a. $\lambda w.\ \overbrace{\text{murder}(w) \rightarrow \underbrace{\forall w' \in R_D(w)[\text{prison}(w')]}_{\textit{must}}}^{\text{conditional}}$

b. $\lambda w.\ \overbrace{\forall w' \in R_D(w)[\underbrace{(\text{murder}(w') \rightarrow \text{prison}(w'))}_{\text{conditional}}]}^{\textit{must}}$

Because falsehood of the antecedent entails truth of the material implication, an analysis like (12a) leads to the prediction that if Kathy does not commit murder in the actual world, then any conditional like (11) is true regardless of what the consequent is. That is, if Kathy never commits murder, then not only is (11) true, but so is (13):

(13) If Kathy commits murder, she must be given a thousand dollars.

(12b), meanwhile, predicts (13) to be true regardless of what Kathy does in the actual world. After all, by assumption all deontically acceptable worlds are murder-free worlds. Therefore, all deontically accessible worlds w' are such that murder$(w') \rightarrow q(w')$ for any q, again because a material implication is automatically true if the antecedent is false.

So in summary, a straightforward classical semantics for modals, in conjunction with a material implication-based view of conditionals, will not correctly predict the right truth values for all of (10), (11), and (13), regardless of whether or not the set of deontically accessible worlds includes worlds with murder in them, and regardless of the scope position of the modal.

Gradability The third puzzle is that many expressions with seemingly modal or modal-like interpretations are gradable, such as the modal adjectives *likely*, *important*, and *permissible*. The gradability of these adjectives is illustrated in (14), in which all three adjectives appear in felicitous comparative constructions:

(14) a. It is **more likely** that Jonah left than it is that Jordanne left.
b. It is **more important** that Jordanne stay than it is that Jonah stay.
c. It is **more permissible** for Jonah to leave than it is for Jordanne to leave.

What's more, Portner & Rubinstein (2016) note that at least on some interpretations, weak necessity modals like *should* can be gradable as well:

(15) Joelle should do her homework more than she should practice piano.

(15), in addition to its more mundane meaning about the recommended frequency of homework and piano practice, also has a reading in which she has a stronger obligation towards her homework than towards her piano-playing.

Naturally, the simple quantificational analysis of modality does not offer a satisfying account of gradable modality, because everything in the simple quantificational analysis is categorical: a world is either accessible from the world of evaluation or it is not, and either all accessible worlds have a certain property or not all of them do.

2.2.1.2 Kratzer's proposal

In attempting to resolve these three puzzles, Kratzer posits that modals are interpreted relative to two *conversational backgrounds*, contextually determined partial functions from worlds to sets of propositions. The first conversational background, the *modal base* (f), generates a set of propositions used to restrict the domain of possible worlds to those that meet certain basic criteria. In the case of an epistemic interpretation, the propositions in $f^c(w)$—where c is the context parameter and w

is the world of evaluation—will restrict the set of worlds to those compatible with the beliefs or knowledge of some individual or collection of individuals. For other sorts of modals, the restriction will be to those worlds that are circumstantially accessible from the world of evaluation, i.e., those possible states of affairs that are in some sense "achievable." The way that the modal base effects this restriction on the set of worlds is simple: in order for a given world to be included, each of the propositions in $f^c(w)$ must be true in that world. Thus, the restricted set of worlds, which I will refer to as the *modal domain*, will be $\bigcap f^c(w)$.

The second conversational background is what Kratzer refers to as the *ordering source* (g). The propositions in $g^c(w)$ are used to generate an ordering over worlds, with the basis of the ordering depending on the flavor of modality: obeying obligations in the case of deontics, obtaining desires for bouletics, achieving goals for teleologicals, etc. As for how $g^c(w)$ effects an ordering over worlds, Kratzer (1981a) adopts a method first proposed by Lewis (1981) to reconcile two seemingly competing views on the semantics of counterfactual conditionals such as (16).

(16) If Sloan had packed his suitcase properly, he would have gotten the job.

According to the *premise-semantic* view of counterfactuals, (16) is true iff when we conjoin Sloan's packing his suitcase properly with a variety of propositions characterizing relevant natural laws, tendencies, and contingent facts in the actual world, the result entails that Sloan got the job (see, e.g., Goodman 1947; Kratzer 1981b, 1989). The *world-ordering* view, meanwhile, states that (16) is true iff in those possible worlds most similar to the actual world in which Sloan packs his suitcase properly, he gets the job (Stalnaker 1968, 1975; Lewis 1973). The way Lewis bridges the apparent gap between these two proposals is to generate an ordering over worlds based on a set of premises as in (17), where $w_1 \precsim_Q w_2$ iff w_1 is at least as ideal as w_2 with respect to the set Q of propositions:

(17) $w_1 \precsim_Q w_2$ iff $\{p \in Q \mid p(w_1)\} \supseteq \{p \in Q \mid p(w_2)\}$

In short, $w_1 \precsim_Q w_2$ is true iff every proposition in Q that is true in w_2 is also true in w_1. This results in a *preorder* over worlds: an ordering that is reflexive and transitive, but not necessarily antisymmetric or connected, meaning that two distinct worlds can be deemed equally ranked or incomparable. Equal ranking ($\sim_Q$), strict ranking ($\prec_Q$), and incomparability ($\|_Q$) can then be defined as follows:

(18) a. $w_1 \sim_Q w_2$ iff $w_1 \precsim_Q w_2$ and $w_2 \precsim_Q w_1$
b. $w_1 \prec_Q w_2$ iff $w_1 \precsim_Q w_2$ and $w_2 \not\precsim_Q w_1$
c. $w_1 \|_Q w_2$ iff $w_1 \not\precsim_Q w_2$ and $w_2 \not\precsim_Q w_1$

Naturally, for Lewis, the "idealness" relation meant similarity to the actual world, so that $w_1 \precsim_Q w_2$ iff w_1 is at least as similar to the actual world as w_2 is. Kratzer extends this formal mechanism to other forms of idealness rankings, and in particular those notions of preferability and likelihood relevant to the interpretation of modals. Thus, the mechanism in (17) is the means by which $g^c(w)$ will rank worlds, with the criteria for the ordering of worlds partially determining the flavor of modality. Kratzer then defines *must* as being true of a proposition p in a world w iff all of the best worlds in $\bigcap f^c(w)$ with respect to the ordering $\precsim_{g^c(w)}$ are worlds in which p is true. Similarly, *can* p is true in w iff at least one of the best worlds in $\bigcap f^c(w)$ with respect to $\precsim_{g^c(w)}$ is a p world. More formal versions of these definitions can be seen in (19).[4]

(19) If $\textsc{Best}(A, \precsim) = \{w \in A \mid \neg\exists w' \in A[w' \prec w]\}$, then:
  a. $[\![\text{must}]\!]^c_{\text{Kratzer}} = \lambda p \lambda w.\ \forall w' \in \textsc{Best}(\bigcap f^c(w), \precsim_{g^c(w)})[p(w')]$
  b. $[\![\text{can}]\!]^c_{\text{Kratzer}} = \lambda p \lambda w.\ \exists w' \in \textsc{Best}(\bigcap f^c(w), \precsim_{g^c(w)})[p(w')]$

As an example, consider (20), in a context in which Dmitri has parked illegally and gotten a ticket.

(20) In view of the laws of the United States, Dmitri must pay his fines.

In this context, the modal base will restrict the modal domain to those worlds circumstantially accessible from the possible world, while the ordering source will order worlds with respect to obedience to United States law, so that $w_1 \prec_{g^c(w)} w_2$ iff US law is better obeyed in w_1 than in w_2. Thus, (20) constitutes a claim that in all of the legally best circumstantially accessible worlds, Dmitri pays his fines. Notice that even though worlds in which Dmitri never parks illegally in the first place are arguably legally better than worlds in which he parks illegally and pays the fine, worlds of the former sort are absent from the modal domain because of Dmitri's lack of a time machine, so the best currently available possibility is for Dmitri to pay his fines.

How does this view fare with respect to the three puzzles? As for the first, the context-sensitivity of the modal base and the ordering source provides plenty of elbow room for a sufficiently wide range of modal flavors. That being said, replacing a small class of "flat" accessibility relations with a singular, contextually determined

[4] This presupposes what Lewis (1973) calls the *Limit Assumption*, which states that there is no "infinite regress", so that there is always a non-empty set of ideal worlds (i.e., $\textsc{Best}(\bigcap f^c(w), \precsim_{g^c(w)}) \neq \varnothing$). Kratzer defines both *must* and *can* in a way that avoids the Limit Assumption, as follows:

(i) $[\![\text{must}]\!]^c_{\text{No Limit}} = \lambda p \lambda w.\ \forall w' \in \bigcap f^c(w)[\exists w'' \in \bigcap f^c(w)[w'' \precsim_{g^c(w)} w' \wedge \forall w''' \in \bigcap f^c(w)[w''' \precsim_{g^c(w)} w'' \rightarrow p(w''')]]]$

(ii) $[\![\text{can}]\!]^c_{\text{No Limit}} = \lambda p \lambda w.\ \neg[\![\text{must}]\!]^c_{\text{NoLimit}}(-p)(w)$

In a nutshell, (i) states that *must* p is true iff for every world in the modal domain, we can start from that world on a path through more ideal worlds and reach a point where all the worlds ahead of us are p worlds. (ii) defines *can* in a way that preserves the traditional dual relationship between *must* and *can*. That being said, I will adopt the Limit Assumption throughout this dissertation.

accessibility relation would accomplish this just as well. So while Kratzer's proposal resolves the issue of fine-grained context sensitivity, this first puzzle does not serve as direct justification for Kratzer's view in particular.

With respect to the second problem of conditionals, Kratzer's theory makes the right predictions about (10), (11), and (13) if one adopts the view that the antecedents of conditionals serve to restrict the modal domain of a modal in the consequent. A templatic version of this analysis is stated formally in (21) for cases where the modal in the consequent is *must*.

(21) $[\![\text{If } p\text{, then must } q]\!]^c = \lambda w.\ \forall w' \in \textsc{Best}(\bigcap f^c(w) \cap p, \precsim_{g^c(w)})[q(w')]$

Let's say that in the Samaritan scenario, the deontically ideal worlds in the modal domain are those where no one commits murder, the worst worlds are those where there is murder that goes unpunished, and in between are worlds where there is murder and the perpetrator goes to prison. In this case, the overtly unrestricted modal construction (10) is rightly predicted to be true, since all of the deontically ideal worlds are murder-free. Furthermore, (11) is also rightly predicted to be true in an analysis of conditionals like (21). In this case, *if Kathy commits murder* restricts the modal domain to those worlds in which Kathy commits murder, meaning that there are no more murder-free worlds in the modal domain. As a result, the best worlds in this restricted modal domain are those in which Kathy goes to prison. Finally, (13) is predicted to be false, because in the ideal worlds in the restricted domain Kathy is not given a thousand dollars. Thus Kratzer's theory of modals, conjoined with a restrictor view of conditional antecedents, gets the right results because the ordering of worlds allows us to make the best of a bad situation, in a way that flat quantificational theories of modals do not.

As far as gradability is concerned, I will hold off on a more complete discussion until later, as the issue of gradable intensionality in Kratzerian frameworks will become highly relevant in the discussion of attitude comparatives in Chapter 4. But on an abstract level, the fact that worlds are ordered opens the way for an analysis in which sentences like *p is more important than q* and *p is more desirable than q* are true iff *p* worlds in some sense outrank *q* worlds with respect to the relevant world-ordering. The question for such an analysis is how one formally defines what "*p* worlds outrank *q* worlds" means, as well as how this comparison of worlds should be integrated with other elements in the grammar and ontology.

Of course, Kratzer's analysis has been through the ringer in the years since it was first introduced, and while her theory of modals and conditionals has become somewhat of a default theory in the linguistic literature, her account of the facts above is far from universally accepted. But regardless of one's thoughts about the particular details of Kratzer's proposal, one relevant feature thereof has remained more or less constant across the various revisions and wholesale replacements of her account:

for deontic, teleological, bouletic, and other similar types of readings—what Portner (2009) calls *priority* modals—one's obligations and permissions lie somewhere at the intersection of what is circumstantially achievable and what is normatively preferable. For Kratzer, this intersection is cast in terms of the relationship between the modal base, which defines what is achievable, and the ordering source, which defines preference relations. Other proposals encode this relationship differently. Even so, the theories of desiderative attitudes like *want* discussed in the rest of this chapter will similarly cast desires in terms of the relationship between (believed) achievability and preference.

2.2.2 von Fintel 1999

von Fintel (1999) offers a theory of desiderative attitudes that is for the most part faithful to Kratzer's theory of modality, adding a few minor changes in order to account for features that differentiate between attitudes and modals. The first revision is parallel to Hintikka's revision of the classical modal semantics: whereas modals are raising verbs with no external thematic role, attitudes of course have experiencers, and thus require an extra entity argument.[5] Naturally, both the ordering source and the modal domain will need to take into account the new entity argument. The ordering source for desiderative attitudes will therefore have as its domain the set of entity-world pairs, with $\precsim_{g(x,w)}$ ordering worlds based on their preferability to x in w. The modal domain for *want* will be $\mathrm{Dox}(x, w)$, x's belief worlds in w.[6]

Beyond the aforementioned intuition underlying the modal domain's being $\mathrm{Dox}(x, w)$—that desires lie at the intersection of preferences and beliefs—additional support for the view that the modal domain is the experiencer's set of belief worlds comes from facts having to do with presupposition projection. Since presupposition projection is not the primary focus of this dissertation, I will only go over the evidence and background at surface level; the reader is referred to the cited literature for more in-depth discussion.

Presuppositions are typically formalized as a requirement that some proposition

[5] It has been argued at various points that there is a distinction between epistemic and root modals with respect to argument structure, with (some) root modals having an external argument where epistemic modals lack one (Ross 1969, Jackendoff 1972, Brennan 1993). However, see Bhatt 1998, Wurmbrand 1999 for syntactic arguments that all modals are in fact raising verbs, thereby lacking an external argument; see also Hackl 1998 for parallel semantic arguments.

[6] In actuality, von Fintel follows Heim (1992) in using as the modal domain for *want* not the set of worlds compatible with what the experiencer believes, but rather the set of worlds compatible with what the experiencer believes will be true *regardless of the experiencer's future actions*, the latter being a superset of the former. For our purposes this difference will not be of great importance, and to my knowledge a revised modal domain is fully compatible with the proposals in this dissertation.

hold throughout some set of worlds.[7] In unembedded contexts, this set of worlds is based on the *common ground*, which is a set of propositions that the conversational participants take for granted as being true. The intersection of the common ground is the *context set*, the set of all worlds compatible with what the interlocutors take for granted. An unembedded sentence (or an utterance thereof) presupposing p then amounts to a requirement that p hold throughout the context set, i.e., that p be entailed by the prior collective commitments of the conversational participants.[8] Thus, (22) imposes a requirement that the context set entail that Gretchen has a cello, since the definite possessive DP *her cello* triggers an existential presupposition.

(22) Gretchen brought her cello to the party.

But when presupposition triggers are embedded under attitudes, something interesting happens. When embedded under *think*, as in (23), the existential presupposition of *her cello* seems to project: by default, we retain the inference that Gretchen has a cello.

(23) Gretchen thinks that her cello is in Berlin.

However, as (24) makes clear, this needn't necessarily be the case.

(24) Gretchen wrongly thinks that she owns a cello, and she thinks that her cello is in Berlin.

(24) shows that by updating the common ground with the information that Gretchen merely *thinks* that she owns a cello—and just as importantly, not that she *actually* owns a cello—the ensuing context set satisfies the presupposition in the second clause. Thus, what (23) and (24) cumulatively show is that when *her cello* is embedded under *think*, we can end up either with a presupposition that Gretchen actually owns a cello, or a presupposition that she thinks she owns one. The same game can be played with *want*: while (25a) by default carries an inference that Gretchen owns a cello, (25b) shows that it is enough for her to merely *want* to own a cello.

(25) a. Gretchen wants to burn her cello.
b. Gretchen wants to own a cello, and she wants to (then) burn her cello.

The natural way to (informally) account for these facts is to say that when a presupposition trigger is embedded under an attitude, the presupposition can impose a requirement either on the context set (i.e., Gretchen must actually own a cello) or on the set of worlds quantified over by the attitude (i.e., Gretchen must believe

[7]See van der Sandt 1992 for an alternative, anaphora-based view in which presuppositions are discourse referents that need to be bound by elements in the preceding discourse.

[8]While this basic idea is ubiquitous nowadays, the notions of common ground and context set are most intimately connected with the pioneering work of Robert Stalnaker (1970, 1973, 1977, 1978).

she owns a cello/want to own one). While the grammatical or pragmatic source of this optionality is not obvious, a convenient way to think of these facts is in terms of scope: a presuppositional element can take "high" scope (enforced with respect to the context set) or "low" scope (the set of worlds quantified over by the attitude verb) (cf. Heim 1992).

But while both *think* and *want* are similar in allowing for either high or low scope of presuppositional elements, there is an an intriguing asymmetry between the two. Consider (26), where *her cello* is again embedded under *want*:

(26) Gretchen wrongly thinks that she owns a cello, and she wants to burn her cello.

(26) shows that ascribing to Gretchen a *belief* that she owns a cello is enough to satisfy the presupposition under *want*: in (26) we infer neither that Gretchen actually owns a cello, nor that she wants to own one. In other words, in the case of desire ascriptions there are (at least) *three* scope possibilities for embedded presuppositions: matrix scope (evaluated with respect to the context set), low scope (evaluated with respect to desire worlds), and—terminologically jumping the gun a bit—*intermediate* scope (evaluated with respect to belief worlds).[9]

Belief, on the other hand, lacks a third scope reading: ascribing to Gretchen a desire to own a cello does not suffice to satisfy presuppositions embedded under *think*. Consider (27):

(27) Gretchen wants to own a cello, and she thinks that she will burn her cello.

Even on a reading of (27) where the presupposition does not take matrix scope—i.e., where we do not infer that Gretchen already owns a cello—we still retain an inference that Gretchen *thinks* she will own a cello. That is, the resulting interpretation is something along the lines of (28):

(28) Gretchen wants to own a cello, and she thinks <u>that she will own a cello and</u> that she will burn her cello.

This, of course, is classic presupposition accommodation, in the sense of Lewis (1979b): the presupposition in the embedded clause has not previously been established, so in order to retain pragmatic acceptability the listener adjusts by accepting it into the common ground alongside the trigger of that presupposition.[10] The presence of this

[9] Intermediate projection was first observed by Karttunen (1974), though he does not note the additional possibilities of low- and high-scope projection. To my knowledge, Heim (1992) was the first to observe and attempt to analyze the full range of projection possibilities.

[10] More specifically, this is what Heim (1983) calls *local* accommodation: the presupposition is accommodated below the scope of the attitude (i.e., Gretchen *thinks* she'll own a cello). This is in contrast to *global* accommodation, where the presupposition is accommodated at maximally high scope; in this case, we would infer that Gretchen will *actually* own a cello.

accommodation suggests that adding the desire ascription to the common ground was not enough to satisfy the presupposition of the second clause, in contrast to example (26).[11]

This asymmetry suggests that there is an implicational hierarchy of the sort diagrammed in (29), where any construction that allows a presupposition to take a given scope allows it to take any scope above it. Thus, in the case of *want*, all three types of projection are available, while for *think*, only the two highest are available.

(29) Context Set > Belief Worlds > Desire Worlds

As noted by Heim (1992), Geurts (1998), and Maier (2015), this hierarchy can be readily accounted for if one's desire worlds are, to use Maier's word, "parasitic" on one's belief worlds, while one's belief worlds are not similarly parasitic on one's desire worlds. One way of enforcing this is to make the belief worlds the modal domain, utilized in conjunction with a bouletic ordering over worlds. We can then say that the presupposition in a desire ascription can be enforced with respect to the bouletically ideal worlds (i.e., the desire worlds), the modal domain (the belief worlds), or the context set. Meanwhile, bouletic orderings don't enter the picture in belief ascriptions, so only belief worlds and the context set are available.

Going back to von Fintel's definition of *want*, the most direct Kratzerian translation would be as in (30), which states that x wants p in w iff all of x's bouletically ideal belief worlds are p-worlds.

(30) $[\![\text{want}]\!]_{\text{almost}} = \lambda p \lambda x \lambda w.\ \forall w' \in \textsc{Best}(\text{Dox}(x,w), \precsim_{g(x,w)})[p(w')]$

But there's a problem. The set of one's bouletically ideal belief worlds is, of course, a subset of the set of one's belief worlds. And given the presupposition facts discussed above, this is a good thing. But this means that any proposition that holds throughout one's belief worlds will by necessity also hold throughout one's bouletically ideal belief worlds. As noted by Stalnaker (1984), this leads to the prediction that believing p entails wanting p. This, of course, is not the case: Jen's believing that Ben was murdered does not entail that she wants him to have been murdered.

To fix this, von Fintel adopts a proposal from Heim (1992) that *want* presupposes that its propositional argument is compatible with, but is not entailed by, the experiencer's beliefs. That is, $[\![\text{want}]\!](p)(x)(w)$ presupposes that there is at least one

[11] A second strategy for avoiding infelicity in (27) is worth noting: we can covertly restrict the domain of the modal *will* to those worlds in which Gretchen eventually owns a cello:

(1) Gretchen wants to own a cello, and she thinks that <u>if she (eventually) owns a cello,</u> she will burn her cello.

Restricting *will* in this manner thereby satisfies the (low scope) presupposition, since in all of the remaining belief worlds Gretchen owns a cello.

world in $\mathrm{Dox}(x, w)$ in which p is true, and at least one in which p is false. (I will call this presupposition the **diversity condition** or **diversity presupposition**, after a similar constraint proposed by Condoravdi (2002) for root modals.) Naturally, this avoids the problem discussed above: if x believes that p, then none of x's belief worlds are worlds in which p is false, and so x cannot want p. As a result, von Fintel's final denotation for *want* is as in (31):

(31) $$[\![\text{want}]\!]_{\text{von Fintel}} = \lambda p \lambda x \lambda w : \exists w', w'' \in \mathrm{Dox}(x, w)[p(w') \wedge \neg p(w'')]. \\ \forall w' \in \textsc{Best}(\mathrm{Dox}(x, w), \lesssim_{g(x,w)})[p(w')]$$

2.3 Comparison to alternatives

For both Hintikka's theory and von Fintel's, *want* is defined in terms of universal quantification over some set of bouletically ideal worlds. While von Fintel's theory includes the diversity presupposition and adopts more formal machinery in determining the set of ideal worlds, the basic notion underlying the two theories is more or less the same. In this section, we will consider two theories that still utilize quantification over possible worlds, but that take a somewhat different approach: to want p is to consider p more desirable than certain alternatives to p. The first theory we will consider is Heim's (1992), in which p worlds are compared to not-p worlds; the next theory considered will be Villalta's (2008), in which p is compared to propositions in a potentially more robust, focus-derived set of alternatives.

2.3.1 Heim 1992

I will let Heim introduce the basic idea of her proposal herself:

> The analysis of desire verbs I want to pursue here is sketched in Stalnaker (1984: 89): 'wanting something is preferring it to certain relevant alternatives, the relevant alternatives being those possibilities that the agent believes will be realized if he does not get what he wants.' An important feature of this analysis is that it sees a hidden conditional in every desire report. A little more explicitly, the leading intuition is that *John wants you to leave* means that John thinks that if you leave he will be in a more desirable world than if you don't leave.
>
> Heim 1992: 193

Heim's "hidden conditional" is what differentiates between her theory and von Fintel's. As discussed above, theories of conditionals dating back to the work of Stalnaker (1968, 1975) and Lewis (1973) state that a conditional *if p, then q* is true in a world w iff in all of the p worlds most similar to w, q is true. Thus, whereas von Fintel's theory only involved one ordering over worlds (the bouletic ordering), Heim's theory will introduce a second ordering of worlds in terms of similarity. The resulting denotation can then be characterized as follows: x *wants* p is true iff for each world w compatible with x's beliefs, those belief worlds most similar to w in which p is true are more preferable to x than those in which p is false. This is stated more formally in (32–33), which also includes the diversity condition:[12]

(32) Preliminary definitions for (33):
  a. $\text{Sim}_w(p) =_{\text{def}}$ the set of p worlds most similar to w.
  b. $w_1 \prec_{x,w} w_2$ iff w_1 is more preferable to x in w than w_2 is.
  c. $A_1 \prec_{x,w} A_2$ iff for all $w_1 \in A_1$ and $w_2 \in A_2$, $w_1 \prec_{x,w} w_2$.

(33) $[\![\text{want}]\!]_{\text{Heim}} = \lambda p \lambda x \lambda w : \exists w', w'' \in \text{Dox}(x,w)[p(w') \wedge \neg p(w'')].$
$\forall w' \in \text{Dox}(x,w)[\text{Sim}_{w'}(\text{Dox}(x,w) \cap p) \prec_{x,w} \text{Sim}_{w'}(\text{Dox}(x,w) - p)]$

We can simplify this definition by adopting the principle that each world is more similar to itself than any other world is to it. (This is what Lewis (1973) refers to as *strong centering*.) This entails that for each world w, if p is true in w, then $\text{Sim}_w(p) = \{w\}$. As a result, (33) (minus the diversity presupposition) is equivalent to the following: for each belief world w, (I) if p is true in w, then w is preferable to all of those belief worlds most similar to it in which p is false; and (II) if p is false in w, then w is worse than all of those belief worlds most similar to it in which p is true.

There is a notable difference between Heim's proposal on the one hand, and Hintikka's and von Fintel's on the other. As discussed previously, for Hintikka, the propositional argument of *want* and other attitudes is upward entailing: if p entails q (i.e., $p \subseteq q$), then *Ann wants p* is predicted to entail *Ann wants q*. For von Fintel, similar but not identical conditions hold, with the source of the difference being the diversity condition. Recall that according to the diversity condition, $[\![\text{want}]\!]$ requires that its propositional argument be compatible with, but not entailed by, the experiencer's beliefs. Now let's say that Ann wants q. There will be some set of worlds q' such that $q \subset q'$, and q' contains all of Ann's belief worlds. Since Ann's beliefs entail q'—all of Ann's belief worlds are q' worlds—q' violates the diversity condition. Thus, there are propositions q and q' such that Ann wants q, $q \subseteq q'$, and it is not the case that Ann wants q'. In other words, von Fintel's *want* is not, strictly speaking, upward

[12] This is a static translation of Heim's (1992) dynamic definition of *want*. Heim precedes this dynamic definition with a static definition of her own, but one with slightly different truth conditions: the arguments of $\text{Sim}_{w'}$ are not $\text{Dox}(x,w) \cap p$ and $\text{Dox}(x,w) - p$, but simply p and $-p$. The result is that instead of going to each belief world and comparing the bouletic idealness of the nearest *belief* worlds in which p and $-p$ hold, it's simply the nearest p and $-p$ worlds *simpliciter*. While this difference has some interesting implications, so far as I can tell it has no direct bearing on the cases at hand.

entailing. But it does meet a slightly weaker condition: adopting von Fintel's term, it is *Strawson upward entailing*, meaning that if Ann wants q, $q \subseteq q'$, *and* q' *satisfies all relevant presuppositions*, then Ann wants q'.

Heim's semantics for *want*, meanwhile, is neither straightforwardly upward entailing, nor Strawson upward entailing. I provide a constructive proof below; those who don't wish to slog through it can feel free to take me at my word and skip to the next paragraph.

> Let $\{w_1, w_2, w_3\}$ be the set of belief worlds. Furthermore, say that we can rank the similarity of these worlds in terms of (the absolute value of) the numerical difference between their indices: w_1 is more similar to w_2 than it is to w_3, w_2 is equally similar to w_1 and w_3, etc. Finally, say that w_1 is preferable to w_3, which is preferable to w_2.
>
> **First step:** *Show that Heim's definition of* want *is true of the proposition* $\{w_1\}$. First, w_1, which is obviously a world in which $\{w_1\}$ is true, is more ideal than w_2, the closest world to it in which $\{w_1\}$ does not hold. w_2 and w_3, which are worlds where $\{w_1\}$ does not hold, are each worse than w_1, which is the closest world to each in which the proposition does hold. Thus, each $\{w_1\}$ world is preferable to all nearest $-\{w_1\}$ worlds, and each $-\{w_1\}$ world is worse than all nearest $\{w_1\}$ worlds, as desired.
>
> **Second step:** *Show that Heim's definition of* want *is* not *true of* $\{w_1, w_2\}$. w_3 is a world in which $\{w_1, w_2\}$ is false. The nearest world to w_3 in which $\{w_1, w_2\}$ is true is w_2. w_3 is preferable to w_2. Thus, w_3 is not worse than the most similar world(s) in which $\{w_1, w_2\}$ is true. *A fortiori*, Heim's *want* is not true of $\{w_1, w_2\}$.
>
> Since Heim's *want* is true of $\{w_1\}$, but not of $\{w_1, w_2\}$, it is not upward entailing. Furthermore, since both $\{w_1\}$ and $\{w_1, w_2\}$ satisfy the diversity condition, it is not Strawson upward entailing either. □

I mention this distinction because whether or not the complement of *want* is Strawson upward-entailing is a topic of ongoing debate. In Section 2.5, we will go into the details of this debate a little further.

2.3.2 Villalta 2008

In Heim's theory, p worlds are compared to $-p$ worlds. Villalta (2008), on the other hand, proposes a theory in which p is compared with other members of a set of propositions, with the members of this set being the focus alternatives of the clausal complement of *want*. Villalta proposes such an account because of certain alleged problems for Heim's theory (and von Fintel's). I will first go over Villalta's three primary qualms, and then I will discuss her proposed solution. I next discuss Rubinstein's (2012) demonstration that the first of Villalta's qualms is due to a misunderstanding on her part, as well as Phillips-Brown's (2016) attempt to account for the second problem within what is essentially von Fintel's (1999) theory of *want*. Finally, I sketch a means of accounting for Villalta's third apparent problem.

2.3.2.1 "Problems" and problems

Problem 1: The picnic To illustrate the first purported problem, imagine that Sofía is bringing dessert to the company picnic, and that Victoria knows this. As far as Victoria is aware, Sofía will bring one of three desserts: chocolate cake, apple pie, or ice cream. Victoria considers chocolate cake to be by far the most preferred option, but also by far the least likely. She thinks that apple pie is significantly more likely than chocolate cake, but it's also not nearly as good. Finally, ice cream is by a wide margin the most likely thing, but is by a similarly wide margin the least desirable.

Intuitively, in this scenario, (34) is false:

(34) Victoria wants Sofía to bring apple pie.

However, Villalta argues that Heim would inaccurately predict (34) to be true. After all, for each of Victoria's belief worlds in which Sofía brings an apple pie, the nearest worlds in which she *doesn't* bring an apple pie will presumably be worlds in which she brings ice cream instead, since this is the most likely alternative to apple pie. Since the apple pie worlds are better than these ice cream worlds, the truth of (34) is thereby (allegedly) predicted.

Problem 2: Believed equivalence Suppose that Mary wants to teach on Tuesdays and Thursdays this semester, rather than, say, on Mondays and Wednesdays. Furthermore, suppose that Mary thinks that she will teach on Tuesdays and Thursdays if and only if she works very hard. As per both Heim's and von Fintel's semantics for *want*, we seem to wrongly predict (35) to be true:

(35) Mary wants to work hard.

The reason for this is that for both Heim and von Fintel, the only worlds that enter into the picture when evaluating desires are those worlds compatible with the experiencer's beliefs. But if Mary believes that she'll teach on Tuesdays and Thursdays if and only if she works hard, then the set of Mary's belief worlds in which the former happens is identical to the set of her belief worlds in which the latter happens. That is, a desire for one should be equivalent to a desire for the other.

Problem 3: Focus sensitivity The third issue that Villalta raises is that *want*, unlike some other attitudes like *believe* and *know*, seems to be sensitive to focus in its embedded clause. For example, consider a scenario in which Lisa, John, and Lara are faculty in the linguistics department. There is a single syntax class, and Lisa would prefer that Lara teach it, although she knows that John will be picked to do so. She also would prefer for the class to be on Tuesdays and Thursdays, rather than Mondays and Wednesdays. In this scenario, (36a) seems true, and (36b) false:

(36) a. Lisa wants John to teach syntax ON TUESDAYS AND THURSDAYS.
b. Lisa wants JOHN to teach syntax on Tuesdays and Thursdays.

With attitudes like *believe*, this difference doesn't seem to arise. That is, both (37a) and (37b) seem to mean roughly the same thing, modulo independent discourse impacts due to focus assignment:

(37) a. Lisa believes that John will teach syntax ON TUESDAYS AND THURSDAYS.
b. Lisa believes that JOHN will teach syntax on Tuesdays and Thursdays.

Since none of the previous definitions involves any sort of sensitivity to focus, this naturally seems like a problem for them.

2.3.2.2 Villalta's proposal

To solve these three problems, Villalta proposes a semantics for *want* that is, in essence, a superlative semantics: to want p is to view it as more desirable than any of its focus-derived alternatives. She formulates this idea in several ways throughout her paper, but for the most part these have to do more with how the meaning of a desire ascription is composed, and less with the end result. (A notational variant of) Villalta's proposed semantics for α *wants* p can be seen in (38), where $\mathrm{Alt}^c(p)$ is the set of focus alternatives to p. For our purposes, it doesn't really matter how exactly $p <^{\mathrm{DES}}_{\alpha,w} q$ is defined. All that matters is that it is meant as a formal implementation of the idea that p worlds outrank q worlds with respect to the bouletic world-ordering $<_{\alpha,w}$.[13]

(38) $[\![\alpha \text{ wants } p]\!]^c$ is defined in w iff $\forall q \in \mathrm{Alt}^c(p)[\exists w' \in \mathrm{Dox}(\alpha, w)[q(w')]]$.[14]
Where defined, $[\![\alpha \text{ wants } p]\!]^c = 1$ iff $\forall q \in \mathrm{Alt}^c(p)[(q \neq p) \rightarrow (p <^{\mathrm{DES}}_{\alpha,w} q)]$

[13] For those seeking the technical details, Villalta defines $<^{\mathrm{DES}}_{\alpha,w}$ similarly to Kratzer's (1981a, 1991) "better possibility" relation: $p <^{\mathrm{DES}}_{\alpha,w} q$ iff for all $w' \in q$ there exists some $w'' \in p$ such that $w'' <_{\alpha,w} w'$, and it is not the case that for all $w' \in p$ there exists some $w'' \in q$ such that $w'' <_{\alpha,w} w'$.

[14] This is essentially the diversity condition. Since p is itself in the alternative set, the fact that $\mathrm{Dox}(\alpha, w)$ must contain some p world(s) is automatically derived. The existence of non-p worlds in $\mathrm{Dox}(\alpha, w)$ is guaranteed if at least one of the focus alternatives to p is disjoint from p in $\mathrm{Dox}(\alpha, w)$.

How does this semantics address the three problems discussed above? Let's discuss them in reverse order. As for the focus sensitivity, the definition of *want* in (38) utilizes $\text{Alt}^c(p)$, the set of focus-derived alternatives to p. In the case of (36a), the set of alternatives will look something like (39a), while for (36b) it will look like (39b):

(39) a. {John teaches Tuesdays and Thursdays, John teaches Mondays and Wednesdays, John teaches Mondays and Fridays…}

b. {John teaches Tuesdays and Thursdays, Lara teaches Tuesdays and Thursdays, Lisa teaches Tuesdays and Thursdays…}

So (36a) is predicted to claim that Lisa views worlds in which John teaches on Tuesdays and Thursdays as better than worlds in which John teaches on some other schedule. Since this claim is true, (36a) is predicted to be true. Meanwhile, (36b) is predicted to claim that Lisa prefers John teaching Tuesdays and Thursdays to Lara teaching Tuesdays and Thursdays, which is false. Hence, Villalta's semantics gets the facts for (36a) and (36b) right.

What about the problem of believed equivalence? In (38), the experiencer's beliefs are used to determine the definedness conditions, since each alternative must be a live option according to the experiencer's beliefs. But once definedness is determined, the worlds that are compared are *all* of the worlds in which p and its alternatives are true, not just those compatible with the experiencer's beliefs. So even if Mary believes that she will teach on Tuesdays and Thursdays if and only if she works hard, a desire for the former is not predicted to entail a desire for the latter. In fact, this entailment may actually doubly avoided, since the embedded clauses *(PRO) to teach Tuesdays and Thursdays* and *(PRO) to work hard* will presumably have different focus alternatives, and thus different comparisons will take place.

Finally, with respect to the apparent picnic problem, Villalta indeed avoids an inference that Victoria wants Sofía to bring apple pie. After all, in contrast to Heim, Villalta makes no reference to similarity relations. Thus, if we assume that chocolate cake is included as an alternative to apple pie in (34), and that Sofía ranks chocolate worlds as better than pie worlds, then the falsehood of (34) is rightly predicted.

2.3.2.3 Responses

Next, I will go over some responses to Villalta's proposal. I will first discuss Rubinstein's (2012) response to the picnic problem, then I will move on to Phillips-Brown's (2016) discussion of the problem of believed equivalence. Finally, I will discuss why Villalta's arguments for focus-sensitivity should be taken with a grain of salt, since her data may fall out from a focus-insensitive semantics for *want* in conjunction with a sufficiently robust view of the semantics and pragmatics of focus.

The picnic non-problem As Rubinstein (2012) points out, Villalta's alleged picnic problem is actually not a problem at all for Heim, as her theory makes all the right predictions about the scenario at hand. Victoria's belief worlds can be divided into three classes: those worlds in which Sofía brings ice cream, those in which she brings apple pie, and those (unlikely but still not entirely ruled out) worlds in which she brings chocolate cake. What we want to see is whether for each belief world w, the closest worlds to w in which Sofía brings apple pie are better than the closest worlds to w in which Sofía does not bring apple pie. If this is so, we wrongly predict (34) to be true; otherwise, not.

Let's start with the ice cream worlds. For each ice cream world w, the closest world in which Sofía does not bring apple pie is simply w itself due to strong centering, and w will indeed be worse than the closest worlds to w in which Sofía brings apple pie. Shifting to apple pie worlds, the trend continues: if w is an apple pie world, the closest worlds to w in which Sofía does not bring apple pie will presumably be ice cream worlds, given the comparative likelihood of ice cream and chocolate cake. Thus, for all apple pie and ice cream worlds, Villalta's diagnosis looks right, as it is indeed the case that the nearest non-apple pie worlds are worse than the nearest apple pie worlds. However, this is not so for worlds in which Sofía brings chocolate cake. If w is a chocolate cake world, then the nearest non-apple pie world is w itself, and since chocolate cake worlds are better than apple pie worlds, w will in fact be better than its neighboring apple pie worlds. So what we get is that for non-chocolate cake worlds, it is indeed the case that the nearest apple pie worlds are better than the nearest non-apple pie worlds, but for the chocolate cake worlds, this is not the case. As a result, Heim correctly predicts falsehood for (34).

(40), meanwhile, Heim correctly predicts to be true:

(40) Victoria wants Sofía to bring chocolate cake.

After all, each chocolate cake world is better than its neighboring non-chocolate cake worlds, and vice versa.

In fact, given the scenario at hand, the only way that we can force (34) to be true is if we eliminate all of Victoria's belief worlds in which Sofía brings chocolate cake. That is, if Victoria has entirely ruled out the possibility that Sofía will bring chocolate cake, then Heim will predict (34) to be true because of the absence of the counterexemplifying chocolate cake worlds. But so far as I can tell, this isn't a bad thing: if Victoria's only believed options are apple pie and ice cream, and she prefers apple pie to ice cream, then it seems like (34) *is* true.

Finally, it's worth noting that von Fintel's theory, like Heim's, makes all the right predictions in this scenario. If all of Victoria's bouletically ideal belief worlds are chocolate cake worlds, then we rightly predict the truth of (40) and the falsehood of (34). But if Victoria has ruled out the possibility of chocolate cake, then the bouletically ideal worlds will presumably be worlds in which Sofía brings apple pie, meaning that (34) is predicted to be true. Thus, we cannot reasonably award points to Villalta's theory for solving the picnic problem, as there was no actual picnic problem to be solved.

Believed equivalence Unlike the picnic problem, believed equivalence seems to be a genuine problem for von Fintel, Heim, and anyone else for whom desires are in some sense constrained by beliefs. Phillips-Brown (2016), building on prior work by Cariani et al. (2013) on modals and Yalcin (2016) on belief, offers an interesting attempt to resolve this problem in von Fintel's (1999) framework for *want*, though by all appearances the core notions in his proposal are extendable to other theories of *want* as well.

In short, the idea is that in the general Kratzerian apparatus, we swap out *individual* worlds for *sets* of worlds, called *coarse worlds*. An intuitive motivation for the use of coarse worlds is that it provides a representation of the limits of granularity in human perception and decision-making. While the space of possibilities is infinite, we as humans obviously cannot differentiate between possibilities at the level of individual worlds, and thus we lump together worlds that don't differ in any important, relevant, or noticeable respects. Hence, we divide the logical space into sets of worlds whose mutual differences are irrelevant, leading to a partition Π of worlds. Observing that a partition of worlds is also the denotation of a question on certain theories of question semantics (e.g., Groenendijk & Stokhof 1984), a set of coarse worlds could equally well be thought of as the delineation of a question whose possible answers constitute our space of options, with all determinations of belief and desirability being relative only to the options in this space.

Bearing this in mind, let's start with the modal domain, which will follow the previously established trend of being built on the beliefs of the experiencer. As before, $\mathrm{Dox}(x, w)$ returns the set of all possible worlds that x has not eliminated as being viable candidates for the actual world. But given $\mathrm{Dox}(x, w)$ and some set Π of coarse worlds, we can alternatively determine *which options in* Π have not been eliminated by x, which I will call $\mathrm{Dox}^{\Pi}(x, w)$:

(41) $\mathrm{Dox}^{\Pi}(x, w) = \{p \in \Pi \mid p \cap \mathrm{Dox}(x, w) \neq \varnothing\}$

Note that $\bigcup(\mathrm{Dox}^{\Pi}(x, w))$, the disjunction of the non-eliminated options, will always be a superset of $\mathrm{Dox}(x, w)$, and usually a proper superset. For an extreme example, say that $\Pi = \{p, -p\}$, and that x is undecided as to whether p or $-p$ holds. This means that there are belief worlds in $\mathrm{Dox}(x, w)$ in which p holds, and likewise for $-p$. But this in turn means that neither of the options in Π is eliminated, so $\mathrm{Dox}^{\Pi}(x, w)$ is simply Π itself. Thus, while there are many individual worlds that x

has eliminated from consideration, no member of this particular set of options has been eliminated.

As another example, let's say that Dee wonders whether Dennis is happy and/or healthy. In this case, Π might be the four-member set $\{\text{happy} \cap \text{healthy}, \text{happy} - \text{healthy}, \text{healthy} - \text{happy}, -\text{happy} \cap -\text{healthy}\}$. Now let's say that Dee believes that Dennis is happy if and only if he is healthy, meaning that in every world in $\text{Dox}(\text{dee}, w)$, Dennis is either both happy and healthy, or both unhappy and unhealthy. In this case, $\text{Dox}^{\Pi}(\text{dee}, w) = \{\text{happy} \cap \text{healthy}, -\text{happy} \cap -\text{healthy}\}$, since Dee has ruled out worlds in which Dennis is happy and unhealthy, or healthy and unhappy. Notice once again that if Dee has some orthogonal belief, such as that Dennis is six feet tall, this will not be reflected in $\text{Dox}^{\Pi}(\text{dee}, w)$, even though it is obviously reflected in $\text{Dox}(\text{dee}, w)$, since Dee's beliefs about Dennis's height do not bear directly on the question of whether he is healthy and/or happy.

Naturally, in Phillips-Brown's semantics for *want*, the set of belief coarse worlds ($\text{Dox}^{\Pi}(x, w)$) is the modal domain. And perhaps unsurprisingly, Phillips-Brown's semantics will also incorporate a Kratzer-style ordering over coarse worlds. For this we will have to briefly go back to premise semantics, in which sets of propositions are used to order worlds. Recall the Lewis-Kratzer definition of $\precsim_Q$ in (17) for a set Q of propositions, repeated below:

(17) $w_1 \precsim_Q w_2$ iff $\{p \in Q \mid p(w_1)\} \supseteq \{p \in Q \mid p(w_2)\}$

Phillips-Brown proposes repurposing the propositions in Q and using them to generate an ordering over coarse worlds, as follows:

(42) $q \precsim_Q r$ iff $\{p \in Q \mid p \subseteq q\} \supseteq \{p \in Q \mid p \subseteq r\}$

Thus, whereas Lewis and Kratzer rank individual worlds in terms of which premises are true in them, Phillips-Brown ranks coarse worlds in terms of which premises hold *throughout* them (i.e., in all of their worlds). Hence, $q \precsim_Q r$ iff every proposition in Q that holds throughout r also holds throughout q. Much like in the Kratzerian system, and following von Fintel, I will use g as the ordering source with $g(x, w)$ returning the set of propositions used to order coarse worlds in terms of their preferability to x in w.

So now we need a definition for *want* that uses all of these coarse worlds. For this, Phillips-Brown adopts the definition in (43), where Π^c is the contextually determined set of coarse worlds:

(43) $[\![\text{want}]\!]_{\text{P-B}}(p)(x)(w)$ is defined iff there is some $q \in \Pi^c$ such that $q \subseteq p$ or $q \cap p = \varnothing$. Where defined,

$$[\![\text{want}]\!]_{\text{P-B}}(p)(x)(w) = 1 \text{ iff } \forall q \in \textsc{Best}(\text{Dox}^{\Pi^c}(x, w), \precsim_{g(x,w)})[q \subseteq p]$$

The presupposition of *want* is that there is an answer to the question of Π^c that either entails p or entails $-p$. In other words, answering Π^c must somehow be relevant to determining whether p. The assertion, parallel to von Fintel's *want*, is that each of the best options in Π^c as determined by $\precsim_{g(x,w)}$ entails p.

Before moving on, it is worth noting that in this definition, we don't quite capture the primary intuition that was previously laid out for using coarse worlds. After all, in (43) the partition Π^c is pragmatically imposed: how the logical space is divided in determining (beliefs and) desires is stipulated by the conversational participants. But the primary motivation for coarse worlds was stated in terms of the limits of granularity *for the experiencer*. So if we really want to stay true to this intuition, the set of coarse worlds—or set of sets of coarse worlds, since a single experiencer can carve the possibility space up in multiple ways—should be relative to the experiencer, not the context. But while this switch is easily done (see Yalcin 2016 for relevant discussion about belief), it's not necessary for showing how Phillips-Brown's semantics avoids the problem of believed equivalence. With this in mind, we'll keep the definition as is.

Let's see how this all plays out in the case of teaching and working hard. Say that Mary's schedule on Mondays, Wednesdays, and Fridays is already packed with meetings, and thus that she believes that she will have a healthy sleep schedule only if her teaching is on Tuesdays and Thursdays. Furthermore, let's say that Mary has some casual reading she wants to do, and (she believes that) she can only do it if she doesn't work terribly hard. Finally, as before, assume that Mary believes that she will teach on Tuesdays and Thursdays if and only if she works hard. An illustration of Mary's belief worlds can be seen in Figure 2.1, where TR is true in worlds where she teaches Tuesdays and Thursdays.

A TR, workhard, sleep, −read	**B** TR, workhard, −sleep, −read
C −TR, −workhard, −sleep, −read	**D** −TR, −workhard, −sleep, read

Figure 2.1: Mary's belief worlds

As stated above, we intuitively want to say that Mary wants to teach on Tuesdays and Thursdays, but not that she wants to work hard. It turns out that we can get the results if we keep $\text{Dox}(\text{mary}, w)$ and $g(\text{mary}, w)$ fixed, but play with Π^c: the beliefs and priorities stay the same, but the option space relative to which these are evaluated will change. As for teaching on Tuesday and Thursday, say that Π^c is as in (44):

(44) $\{\text{TR} \cap \text{sleep}, \text{TR} \cap -\text{sleep}, -\text{TR} \cap \text{sleep}, -\text{TR} \cap -\text{sleep}\}$

Now we need to find $\text{Dox}^{\Pi^c}(\text{mary}, w)$. As per our characterization above, none of Mary's belief worlds are such that she sleeps well and doesn't teach on Tuesdays and Thursdays. The other three options are manifested: $\text{TR} \cap \text{sleep}$ in region A, $\text{TR} \cap -\text{sleep}$ in B, and $-\text{TR} \cap -\text{sleep}$ in C and D. Thus, $\text{Dox}^{\Pi^c}(\text{mary}, w)$ will be as in (45):

(45) $\{\text{TR} \cap \text{sleep}, \text{TR} \cap -\text{sleep}, -\text{TR} \cap -\text{sleep}\}$

Next, $g(\text{mary}, w)$ gets its turn. As hinted at above, Mary has two (relevant) priorities: sleep and read. So let's just say that $g(\text{mary}, w) = \{\text{sleep}, \text{read}\}$. As stated above, some option in (45) is at least as good as another iff any member of $g(\text{mary}, w)$ that is true throughout the second is true throughout the first. But Mary's reading does not hold throughout any of the options in (45): whether Mary reads is logically independent of whether she works on Tuesdays and Thursdays or gets enough sleep. Thus, it will have no effect on the ordering of options, and only sleep matters, leading to the ordering in (46):

(46) $\text{TR} \cap \text{sleep} <_{g(\text{mary}, w)} \text{TR} \cap -\text{sleep}, -\text{TR} \cap \text{sleep}$

Thus, the single best option entails that Mary works on Tuesday and Thursday. In conjunction with the fact that TR satisfies the presupposition of *want*, what we predict is that *Mary wants to teach on Tuesdays and Thursdays* is indeed true in this context.

What about *Mary wants to work hard*? In the context above, with Π^c as in (44), this is undefined, as the presupposition of *want* is not satisfied: none of the propositions in (44) entails either that Mary reads or that Mary doesn't read. We thus need to partition our worlds differently, "asking" Mary a different question. Our new set of coarse worlds will be as in (47):

(47) $\{\text{workhard} \cap \text{read}, \text{workhard} \cap -\text{read}, -\text{workhard} \cap \text{read}, -\text{workhard} \cap -\text{read}\}$

Since we are keeping $\text{Dox}(\text{mary}, w)$ fixed as in Figure 2.1, $\text{Dox}^{\Pi^c}(\text{mary}, w)$ will look like (48), since the only option in (47) not represented in Mary's belief worlds is $\text{workhard} \cap \text{read}$:

(48) $\{\text{workhard} \cap -\text{read}, -\text{workhard} \cap \text{read}, -\text{workhard} \cap -\text{read}\}$

As stated earlier, we are keeping $g(\text{mary}, w)$ fixed as $\{\text{sleep}, \text{read}\}$. But this time, it is sleep that has no effect on the ordering of options, as it entails none of the options in (48). Thus, options are ordered only with respect to reading, and we get the ordering in (49):

(49) $-\text{workhard} \cap \text{read} <_{g(\text{mary},w)} \text{workhard} \cap -\text{read}, -\text{workhard} \cap -\text{read}$

In this case, the single best option entails that Mary doesn't work hard. Thus, not only do we predict that *Mary wants to work hard* is false, but we predict that *Mary wants to not work hard* is true. We therefore successfully predict that in spite of the fact that Mary believes that she'll teach on Tuesdays and Thursdays if and only if she works hard, she wants the former, but not the latter.

As mentioned earlier, Phillips-Brown notes that the use of coarse worlds in this manner is not inherently tied to von Fintel's (1999) semantics for *want*, and that it could in principle extend to other types of analyses. Either way, for the rest of this dissertation I will not make use of coarse worlds, for the simple reason that I wish to avoid the additional formal complications that they bring along with them.

Is *want* really focus-sensitive? We are now left with the apparent focus-sensitivity of *want*. One way in which *want* could be made semantically focus-sensitive is by adopting Phillips-Brown's coarse worlds approach, and making the choice of coarse worlds sensitive to the set of focus alternatives of the embedded clause. Since both the set of coarse worlds and the set of alternatives are sets of propositions, a requirement could be imposed relating the two, thereby potentially deriving the results in conjunction with the right base theory for the semantics of *want*.

However, I would instead like to push back against the claim that the facts Villalta cites are genuine evidence that *want* is semantically focus-sensitive. I will offer a (very sketchy) counterproposal according to which *want* has a pedestrian, non-focus-sensitive semantics, and Villalta's observations stem simply from the interaction between this semantics and the interpretation of focus. I will offer this analysis in von Fintel's theory, but so far as I can tell the basics of the analysis are more or less framework-independent.

There is a well-known connection, perhaps most famously explored by Roberts (1996), between the semantics and pragmatics of focus and that of questions. This is most plainly seen in the case of question-answer congruence. For the question posed in (50), the reply in (50a) is felicitous, while that in (50b) is odd:

(50) What did Sue eat?
  a. Sue ate BEANS.
  b. # SUE ate beans.

The reason for this, Roberts suggests, is that the set of focus alternatives for the answer has to be the same as the set of alternatives that is the denotation of the question to which it is a response. (This naturally requires certain assumptions about the nature of focus alternatives and the semantics of questions.) Thus, the interpretation of the question in (50), as well as the set of focus alternatives for the reply (50a), will be something like what is informally sketched in (51a), while the set of focus alternatives for the reply (50b) will look like (51b):

(51) a. $\{$Sue ate $x \mid x$ is a thing$\}$
b. $\{x$ ate beans $\mid x$ is a person/thing$\}$

The match in the case of (50a) and the mismatch for (50b) explain the judgments of (in)felicitousness for (50), as well as the inverse case in (52):

(52) Who ate beans?
a. # Sue ate BEANS.
b. SUE ate beans.

Roberts takes this notion of question-answer congruence and extends it to *all* instances of focus interpretation in a discourse. Thus, every assertion over the course of a discourse is an answer to some question, either explicitly stated or tacitly adopted, which she calls the *question under discussion* (QUD). And just as in the case of overt question-answer congruence, the focus alternatives of an assertion must be identical to the focus alternatives of the QUD. In her formalization of this view, Roberts defines contexts using considerably more structure than that adopted in our previous discussion of presuppositions, where they were just sets of possible worlds. In keeping with the sketchy nature of my proposal, I will keep contexts as they are, but I will take seriously Roberts's proposed relationship between focus and questions.

Of importance to our analysis will be two presuppositions involved in the posing of questions. Take (53) as an example:

(53) Where did Stacy do jumping jacks yesterday?

One requirement in order for (53) to be felicitously asked is that the answer not already be established in the common ground. If you have already told me where Stacy did jumping jacks, then it is infelicitous for me to subsequently ask (53). A second, weaker presupposition of (53) is that Stacy did in fact do jumping jacks yesterday. If the set of alternatives in the denotation of (53) looks like (54), then this boils down to a (weak) presupposition that the disjunction of (54) holds, sometimes called an *existential presupposition* (see Abusch 2010 for extensive discussion).[15]

(54) $\{$Stacy did jumping jacks at l yesterday $\mid l$ is a location$\}$

A formalization of these two requirements for question-asking can be seen in (55). (55a) imposes the requirement that the context set not already entail a particular answer to the question, while (55b) requires that the context set entail the disjunction of the question's alternatives.

[15] This is not to be confused with the existential presupposition triggered by definite DPs like *the king of France*, though it is of course possible that the two share a common semantic-pragmatic source.

(55) Given a context set C, the felicitous asking of a question whose denotation is a set Q of alternatives presupposes the following:
a. There is no $q \in Q$ such that $C \subseteq q$, and
b. $C \subseteq \bigcup Q$.

Given the pragmatic relationship between focus and questions, we can impose similar presuppositions on assertions as well. Thus, since (56) would be an answer to a (potentially tacit) QUD like (53), its assertion would similarly presuppose (i) that none of its focus alternatives is entailed by the context set—that is, that where Stacy did jumping jacks was not already established—and (ii) that Stacy in fact did jumping jacks yesterday (the existential presupposition).

(56) Stacy did jumping jacks IN HER LIVING ROOM yesterday.

Let us treat this requirement as a single conjoined presupposition, which I will call **informativity**.

Now let's turn back to Villalta's example sentences, repeated below:

(36) a. Lisa wants John to teach syntax ON TUESDAYS AND THURSDAYS.
b. Lisa wants JOHN to teach syntax on Tuesdays and Thursdays.

Recall from the context provided that Lisa believes that John will teach syntax, even though she wishes that someone else could. Given that John is teaching syntax, a Tuesday/Thursday syntax schedule is preferred to any other schedule. What I will now show is that combining von Fintel's semantics for *want* with the view of focus interpretation discussed above generates the correct prediction that (36a) is true and (36b) false, so long as we say that the focus alternatives are initially generated in the embedded clause, with the concomitant informativity presupposition projecting from there.

For (36a) we start with the embedded clause (57a), which generates the set of alternatives (57b):

(57) a. John to teach syntax ON TUESDAYS AND THURSDAYS
b. {John teaches syntax on k | k is a collection of days of the week}

The informativity presupposition of (57a)—that the context not entail a particular teaching schedule for John, and that it entail that he teaches syntax on some schedule—starts in the scope of *want*. But as we know, this presupposition can project in at least three ways: it can take matrix scope, in which case the presupposition is checked relative to the context set; intermediate scope, in which case it is checked relative to Lisa's belief worlds; or low scope, in which case it is checked relative to Lisa's desire worlds (i.e., her bouletically ideal belief worlds). If the informativity presupposition takes low scope, then it is not satisfied: there is a particular alternative in (57b) that is entailed by Lisa's set of desire worlds, namely, the alternative according to which John teaches on Tuesdays and Thursdays. But if the presupposition takes intermediate scope, then it is indeed satisfied: Mary's belief worlds do not entail a particular alternative in (57b) (she doesn't know when John will teach), and her beliefs entail that John will indeed teach, satisfying the existential presupposition. Thus, with intermediate projection of the informativity presupposition, (36a) has an interpretation along the lines of (58):

(58) Lisa wants John to teach syntax on Tuesdays and Thursdays, she believes that John will teach syntax, and she believes that his schedule is as yet undetermined.

For (36b), on the other hand, there is no level at which the informativity presupposition can take scope and be satisfied. The set of focus alternatives generated in the embedded clause in (36b) is as in (59):

(59) $\{x$ teaches syntax on Tuesdays and Thursdays $\mid x$ is a person$\}$

The informativity presupposition of the embedded clause in (36b) will thus be that who will be teaching on Tuesdays and Thursdays is not determined, and that *someone* will be teaching on Tuesdays and Thursdays. As with (36a), if this informativity presupposition takes low scope, the requirement that the choice of alternative is undetermined will not be satisfied, as John is the instructor in all of Lisa's ideal belief worlds (since he is the instructor in all of her belief worlds *simpliciter*). But unlike with (36a), switching to intermediate scope doesn't save the day. After all, on an intermediate scope reading, (36b) would presuppose that Lisa's beliefs do not entail that a particular person is teaching syntax on Tuesdays and Thursdays, but do entail that *someone* is teaching syntax on Tuesdays and Thursdays. But both of these are false in the scenario at hand: Lisa's beliefs *do* entail that a particular person is teaching syntax (namely, John), but they *don't* entail a particular schedule. We thus get presupposition failure or, if the presupposition is accommodated, falsehood.

This is obviously a very rough analysis of Villalta's data allegedly showing focus-sensitivity of *want*. But I hope that it provides at least for an outline of how these data might be analyzed in a theory in which *want* is not focus-sensitive.

2.4 Bayesian

The final type of approach to the semantics of modals and attitudes that we will review differs greatly from those already discussed in that existential and universal quantification over worlds, as well as direct comparison of worlds in terms of preferability, are partly or wholly replaced with probabilities and utility values. Typically, though not of necessity, this is done in a two-pronged fashion. First, for epistemic modals and belief ascriptions, a more structured notion of belief/knowledge is defined using the language of probability theory (Yalcin 2007, 2010; Lassiter 2011a, 2017). Second, obligations and desires are cast in terms of the decision-theoretic notion of *expected utility*, which combines probabilities with measures of utility (Levinson 2003; Lassiter 2011a,b, 2017).

2.4.1 Probabilistic belief

If belief is to be analyzed in terms of probabilities, then we have to add some new structure to the contents of belief states (and, for epistemic modals, epistemic states or information states), as a simple set of worlds will no longer do. While there are a variety of ways of doing this, I will demonstrate the basic ideas using a relatively simple formalism. For simplicity's sake, in this formulation (and its extension to *want*) I will assume that the set of possible worlds is finite, though this is not an in-and-of-itself requirement for the approaches discussed here.

On a basic Bayesian view, belief states, as well as the information states used for the evaluation of epistemic modals, have the structure of *probability spaces*, as defined in (60).

(60) A **(finite) probability space** is an ordered triple $\langle \Omega, F, \mathrm{Pr} \rangle$, where Ω is a set of possible outcomes, and:

a. $F \subseteq \mathrm{Pow}(\Omega)$ is such that:
   i. $\Omega \in F$;
   ii. if $A \in F$, then $\Omega - A \in F$; and
   iii. for all $G \subseteq F$, $\bigcup G \in F$.

b. $\mathrm{Pr} : F \to [0,1]$ is such that:
   i. $\mathrm{Pr}(\Omega) = 1$; and
   ii. for all $G \subseteq F$, if the members of G are pairwise disjoint (i.e., for all $A, B \in G$, $A \cap B = \varnothing$), then $\mathrm{Pr}(\bigcup G) = \sum_{A \in G} \mathrm{Pr}(A)$.

For an individual x and world of evaluation w, we can define a *doxastic probability space* $\mathrm{PS}_{x,w}$ as a probability space $\langle W, F_{x,w}, \mathrm{Pr}_{x,w} \rangle$. I will assume $F_{x,w}$ to be $\mathrm{Pow}(W)$, the power set of W, so that every proposition is assigned a probability; however, we could just as well assume a less fine-grained probability space in which some propositions are not assigned probabilities. $\mathrm{Pr}_{x,w}$ is x's probability function in w, i.e., a function from propositions to the probability that x assigns to that proposition (in w).

With this in mind, one way in which belief can be formulated in a Bayesian system is as in (61), where n is some threshold probability for qualification as belief:

(61) $[\![\text{believe}]\!]_{\text{prob}} = \lambda p \lambda x \lambda w.\ \Pr_{x,w}(p) \geq n$

An obvious choice for n is 1: to believe p is to believe that there is a 100% likelihood of p. Notice that this is, in essence, a Hintikkan view of *believe*: if we define $\text{Dox}(x, w)$ as the set of worlds w' such that $\Pr_{x,w}(\{w'\}) \neq 0$, then for any proposition p, $\Pr_{x,w}(p) = 1$ iff $\text{Dox}(x, w) \subseteq p$. As an alternative, n could be assigned a lower value, so that to believe p is to believe that it's *nearly* certain that p.

Even if we set n to 1 and thereby simply reconstruct a Hintikkan denotation for *believe*, there may be independent evidence suggesting that some sort of additional structure for belief states and information states is required beyond a mere set of worlds. Consider the case of gradable modal adjectives like *likely*. It is difficult to see how (62) can be analyzed using only a categorical distinction between those worlds compatible with the epistemic states of the interlocutors, and those not:

(62) This coin is as likely to come up heads as it is tails.

If *likely* is an epistemic modal, then (62) shows that whatever is included in an information state must have sufficient structure to encode gradable notions of likelihood.

Similar facts hold for belief as well. Recall (63) from Chapter 1:

(63) Chip thinks that Joanna might be in New York.

As discussed there, in examples like (63) the epistemic modal is by default evaluated relative to the belief state of the matrix experiencer. Thus, (63) is true iff Joanna's being in New York is compatible with Chip's beliefs. But Yalcin (2007) argues that the exact same thing is going on in examples like (64):

(64) Dara thinks that this coin is as likely to come up heads as it is tails.

The interpretation of (64) seems to be parallel to (63), in that it appears to assert that the subjective likelihood that Dara assigns to the coin coming up heads is at least as high as that she assigns to the coin coming up tails. But this suggests that (gradable) likelihood has to be encoded in belief states as well, indicating a structure to belief states that extends beyond Hintikka's flat set of worlds.

It is worth noting, however, that this argument hinges on the assumption that *likely* and similar adjectives like *certain*, *probable*, *possible* etc., are true epistemic modals to begin with. If we allow for the possibility that they aren't, then it might be possible for (62) and (64) to receive a more pedestrian explanation. Adopting a possibly extreme version of this view, it could be that the likelihood of a proposition is not a feature of information states and belief states *per se*, but is instead a feature of a world itself, or of a world-at-a-time, or of a situation within a given world. That is, the likelihood of something or other being the case is an objective fact about a given situation. As a result, to believe that p has such-and-such likelihood is simply to be in a doxastic state such that in all (relevant) situations compatible with that state,

p has such-and-such likelihood. No extra structure for belief states is needed. Note that this is completely compatible with *likely* still being a modal, and even with likelihood being measured probabilistically; the difference is not in what the structure of likelihood looks like, but in the relationship between structures of likelihood and the structure of a belief state or information state.

Such an objectivist view of likelihood leaves unexplained how apparent entailments between adjectives of likelihood and epistemic modals—such as that from (65a) to (65b)—are to be captured.

(65) a. It's 90% likely that Alexis is at the party.
b. Alexis might be at the party.

If an objectivist account is right, (65a) is a straightforward claim about the world (or about some situation in it), while (65b) is a statement that is in some sense about my epistemic state, or about that of myself and my interlocutors. There is thus no obvious way of bridging the semantic divide enough to generate the apparent entailment from (65a) to (65b). While I do not know of a strong objectivist response to this, one possibility is worth considering: (65a) does not *semantically* entail (65b), but an utterance of (65a), along with subsequent acceptance into the common ground, commits the interlocutors to an information state according to which (65b) is true. In other words, the conjunction of (65a) and the negation of (65b) is not *contradictory*, but *unassertable*. However, I must leave a fuller exploration of this possibility for future work.

But regardless of whether or not facts like (62) and (64) are construed as evidence in favor of a non-flat structure for belief states, it is clear that gradable notions of likelihood are things that we have beliefs about, and that they inform our decision-making processes in some manner. With this in mind, let's soldier on with a probabilistic view of belief for the time being, and move to a definition of *want* building on this probabilistic structure.

2.4.2 Utility-based desire

Bayesian definitions of *want* use the decision-theoretic metric of *expected utility*, which (as the name suggests) is a metric for determining the expected value of p, given what is likely to be the case if p is the case. As hinted at by the *if*-phrase in this paraphrase, the definition of expected utility incorporates *conditional probabilities*, defined in (66) (where $\Pr(A|B)$ is the probability of A given B):

(66) $\Pr(A|B) = \frac{\Pr(A \cap B)}{\Pr(B)}$

Say, for example, that we have a non-loaded six-sided die. If we label our outcomes based on the side of the die that comes up, this means that $\Pr(\{1\}) = \frac{1}{6} = .1\overline{6}$, and likewise for each of two through six. One question we can ask about this probability space is what the probability of $\{1,2,3\}$ is, i.e., the likelihood that the die will come up as a one, two or three. In our scenario, this comes out to .5: this set of outcomes includes half of the possible outcomes, and our outcomes are weighted equally. But another question we can ask is the following: what is the probability of $\{1,2,3\}$ *given that five and six are ruled out*? Or put another way, if the die has been rolled and we somehow are sure that the result is between one and four (inclusive), what is the likelihood that it is a one, two, or three? Intuitively, it should be $\frac{3}{4}$ (= .75), since there are four equally likely remaining sides, and three of them are among the outcomes we're looking for. And this is indeed what we get using the formula in (66):

(67) $\Pr(\{1,2,3\} \mid \{1,2,3,4\}) = \frac{\Pr(\{1,2,3\} \cap \{1,2,3,4\})}{\Pr(\{1,2,3,4\})} = \frac{\Pr(\{1,2,3\})}{\Pr(\{1,2,3,4\})} = \frac{\frac{1}{2}}{\frac{2}{3}} = \frac{3}{4}$

(To see the equivalence of the penultimate fraction and the final fraction, multiply both the top and bottom by six.) Another way we can think about conditional probability is as follows: the probability *simpliciter* of a proposition is its likelihood relative to the full space of possible outcomes, while the probability of p given q is the likelihood of p relative only to the space of possible outcomes in which q holds.

The way that expected utility makes use of conditional probabilities is as follows. In determining the expected utility of p, we go to each world w in which p holds and ask two questions: first, how good of a world is w, and second, given p, how likely is it that w will be the world in which we end up? Answering the second question is where conditional probability comes into play, since it simply asks for the probability of $\{w\}$ given p, i.e., $\Pr(\{w\}|p)$. Answering the first question requires a *utility function* Ut, which takes a world and returns a utility value. In general, what the codomain of Ut is doesn't matter much. However, it is convenient to treat the codomain as the set of all real numbers, with worlds with positive utility values being good, worlds with negative utility values being bad, and worlds with utility values of zero being neutral.

Next, the utility value of w and the conditional probability of $\{w\}$ (given p) need to be combined. This is done by multiplication, giving us $\mathrm{Ut}(w) \times \Pr(\{w\}|p)$. The expected utility of p is then the result of summing all of these values for all possible worlds in which p is true:

(68) Given utility function Ut and probability function Pr,
$\mathrm{EU}(p) = \sum_{w \in p}(\mathrm{Ut}(w) \times \Pr(\{w\}|p))$

To illustrate, let's go back to our example of the roll of the die. Say that you and I are playing a simple gambling game. I roll a (fair) die. If the result is anything between one and four (inclusive), you have to pay me a dollar. If the result is a five or a six, I have to pay you three dollars. Thus, if we assign the values of outcomes from my perspective in terms of the amount of money I gain, then for all n from one to four, $\mathrm{Ut}(n) = 1$, and if n is a five or six, $\mathrm{Ut}(n) = -3$. Once again, for each n, $\Pr(\{n\}) = \frac{1}{6}$.

In this case, what is the expected utility of $\{2,3,5\}$, the set of outcomes in which I roll a prime number? The calculation of this expected utility is shown in (69). (I leave it to the reader to prove to his or her own satisfaction that for each $n \in \{2,3,5\}$, $\Pr(\{n\}|\{2,3,5\}) = \frac{1}{3}$.)

(69) $\mathrm{EU}(\{2,3,5\}) = [\mathrm{Ut}(2) \times \Pr(\{2\}|\{2,3,5\})] +$
$[\mathrm{Ut}(3) \times \Pr(\{3\}|\{2,3,5\})] + [\mathrm{Ut}(5) \times \Pr(\{2\}|\{2,3,5\})]$
$= [1 \times \frac{1}{3}] + [1 \times \frac{1}{3}] + [-3 \times \frac{1}{3}]$
$= \frac{1}{3} + \frac{1}{3} - 1$
$= -\frac{1}{3}$

What this means is that if I roll a prime number, I can expect on average to lose about a third of a dollar. We can also use this to see if either of us has the upper hand in the contest as a whole, by checking $\mathrm{EU}(\{1,2,3,4,5,6\})$. This also turns out to be $-\frac{1}{3}$, meaning that on average, this contest will result in you getting a third of a dollar. We can thus say that the game is slightly rigged in your favor.

In generating an expected utility-based definition of *want*, we naturally require a subjective expected utility function $\mathrm{EU}_{x,w}$, using a subjective utility function $\mathrm{Ut}_{x,w}$ and the aforementioned subjective probability function $\Pr_{x,w}$. In (70) we see two possibilities for how to inject $\mathrm{EU}_{x,w}$ into the semantics of *want*. (70a) predicts that x wants p in w iff $\mathrm{EU}_{x,w}(p)$ exceeds some threshold n. (While the choice of n doesn't matter, zero seems like as good a choice as any.) Meanwhile, (70b), proposed by Levinson (2003), requires that the expected value of p exceed that of $-p$.[16]

(70) a. $[\![\text{want}]\!]_{\text{threshold}} = \lambda p \lambda x \lambda w.\ \mathrm{EU}_{x,w}(p) > n$
b. $[\![\text{want}]\!]_{\text{Levinson}} = \lambda p \lambda x \lambda w.\ \mathrm{EU}_{x,w}(p) > \mathrm{EU}_{x,w}(-p)$

(Note that while neither of these definitions includes the diversity condition, introducing it is a simple matter: we just add a presupposition that $0 < \Pr_{x,w}(p) < 1$.)

That these get us different results can be seen in cases where all of the possible outcomes are good, or where all of the outcomes are bad. For an example of the former, say that I have won a lottery. To see how much I win, I have to roll a six-sided die: if I roll a one, I get one million dollars; if I roll a two, I get two million; etc. Now consider the sentences in (71):

(71) a. I want to roll a six.
b. I want to roll a one.

[16] Note that I exclude a third analysis, namely, a translation of Villalta's (2008) focus-sensitive theory into an expected utility framework. I exclude this possibility because there is little new to discuss: the result of the combination is essentially just the sum of its parts. In other words, the benefits and drawbacks of such a theory would simply be a mixture of the ups and downs of its two sources.

On reasonable assumptions, both (70a) and (70b) predict (71a) to be true: rolling a six and getting six million dollars is both an exceedingly good outcome in its own right, and better than the predicted outcome if I don't roll a six (in which case I get less than six million dollars). However, the two proposals will disagree about (71b). Since getting a million dollars is still great, rolling a one still leads to a great outcome. Hence, the expected utility of rolling a one presumably will—or at least could—exceed the threshold of (70a), leading to a prediction that (71b) is true. On the other hand, rolling a one gets me the worst possible outcome, meaning that the expected utility of rolling a one is less than the expected utility of not rolling a one. As a result, Levinson would predict (71b) to be false.

The tables are turned when we consider cases where all possible outcomes are bad. For example, say that in settling a lawsuit, I have to roll a die. This time, the result determines how many millions of dollars I *lose* in paying the settlement. In this case, Levinson's definition predicts (71b) to be true and (71a) to be false, since rolling a one represents the best of this bad set of possible outcomes. Meanwhile, the threshold definition in (70a) predicts both sentences in (71) to be false, since either outcome is very bad for me.

These differences between (70a) and (70b) are illustrated in table form in Figure 2.2 below.

Higher roll ⇒ <u>better</u> outcome, all outcomes <u>good</u>

	$[\![\text{want}]\!]_{\text{threshold}}$	$[\![\text{want}]\!]_{\text{Levinson}}$
(71a)	True	True
(71b)	True	False

Higher roll ⇒ <u>worse</u> outcome, all outcomes <u>bad</u>

	$[\![\text{want}]\!]_{\text{threshold}}$	$[\![\text{want}]\!]_{\text{Levinson}}$
(71a)	False	False
(71b)	False	True

Figure 2.2: Predictions of the decision-theoretic definitions of *want*.

But as important as these differences are, more important for our purposes is a trait shared in common between these two analyses, proved formally in the box below: that neither definition of *want* is Strawson upward-entailing.[17]

[17] In the box, I prove this for $[\![\text{want}]\!]_{\text{threshold}}$ on the assumption that the threshold is zero. Changing the threshold to some different arbitrary number does not change the basic nature of the proof.

Let $\{w_1, w_2, w_3\}$ be the set of belief worlds. Furthermore, let $\text{Pr}_{x,w}(\{w_1\}) = .25$, $\text{Pr}_{x,w}(\{w_2\}) = .5$, and $\text{Pr}_{x,w}(\{w_3\}) = .25$. Finally, let $\text{Ut}_{x,w}(w_1) = 1$, $\text{Ut}_{x,w}(w_2) = -3$, and $\text{Ut}_{x,w}(w_3) = 0$.

First step: *Show that the EU definition of* want *is true of the proposition* $\{w_1\}$. Since $\text{EU}_{x,w}(\{w_1\}) = 1$, while $\text{EU}_{x,w}(\{w_2, w_3\}) = -2$, $\{w_1\}$ has a greater-than-zero EU, as well as a greater EU than $-\{w_1\}$ (which also has an EU of -2). Thus, *want* is true of $\{w_1\}$.

Second step: *Show that the EU definition of* want *is* not *true of* $\{w_1, w_2\}$. Since $\text{EU}_{x,w}(\{w_1, w_2\}) = -1.\overline{66}$, while $\text{EU}_{x,w}(\{w_3\}) = 0$, $\{w_1, w_2\}$ has a sub-zero EU, as well as a lesser EU than $-\{w_1, w_2\}$ (also 0). Thus, *want* is false of $\{w_1, w_2\}$.

Since both EU definitions of *want* are true of $\{w_1\}$, but not of $\{w_1, w_2\}$, neither definition is upward entailing. Furthermore, since both $\{w_1\}$ and $\{w_1, w_2\}$ satisfy the diversity condition, they are not Strawson upward entailing either. □

2.5 (Not) picking sides: Is *want* (Strawson) upward-entailing?

So far, we have seen theories for the semantics of *want* that make three different predictions with regard to upward entailment. Hintikka's (1969) prediction is straightforward upward-entailment: if α wants p and $p \subseteq q$, then α is predicted to want q. von Fintel's (1999) theory is not straightforwardly upward-entailing, but it's close, as it is Strawson upward entailing: if α wants p, $p \subseteq q$, and q satisfies the diversity condition, then α is predicted to want q. Finally, Heim (1992), Villalta (2008), and the decision theorists, whose theories make quite different predictions in other areas, are in full agreement that the complement of *want* is not even Strawson upward-entailing: even if the presupposition of *want* is satisfied, α's wanting p still does not entail that α wants q if $p \subseteq q$.

So who's right? As it turns out, we can rule out straightforward upward entailment fairly quickly. Consider (72):

(72) a. Lana wants to die peacefully in her sleep.
     b. Lana wants to die.

The set of worlds in which Lana dies peacefully in her sleep is clearly a proper subset of the set of worlds in which Lana dies. Therefore, in a purely Hintikkan semantics, (72a) is expected to entail (72b), which it plainly does not. von Fintel's theory, meanwhile, fares better on this front. After all, Lana presumably believes that she will die at some point in the future. Thus, Lana has no belief worlds in which she does not die, meaning that (72b) is a violation of the diversity condition. We thereby successfully derive a lack of entailment from (72a) to (72b).

At this point, the reader might complain that I've pulled something of a bait-and-switch, as this is not what the embedded clause of (72b) is actually interpreted as meaning. Instead, (72b) as a whole is interpreted as meaning that Lana specifically wants to die *now*, or as soon as possible, or something along these lines. (72a), meanwhile, is more likely to be interpreted as claiming that Lana wants her death—whenever it may be—to be peaceful. But since Lana's dying a peaceful death does not entail her dying now or soon, the embedded clause as it is interpreted in (72a) does not actually entail the embedded clause as it is interpreted in (72b), so this pair of sentences does not constitute a true test for straightforward upward entailment.

We can avoid this problem by including a time frame adverbial like *this afternoon* in the embedded clause, as in (73):

(73) Lana wants to die peacefully this afternoon.

If (73) is accepted as true, it seems we can infer that either Lana believes that she is fated to die this afternoon, or she in fact wants to die this afternoon. von Fintel's theory predicts exactly this disjunctive inference, while Hintikka's theory forces the inference that Lana wants to die this afternoon.

It seems, then, that straightforward upward entailment can be safely eliminated, leaving us with either Strawson upward entailment or the lack thereof. One argument that has been levied against Strawson upward entailment is that adding disjuncts in the complement does not appear to preserve truth. For example, say that Isabella is the winner of a competition, and the (somewhat bizarre) prize is her choice of one of the following: receiving a million dollars, getting kicked in the shins, or getting nothing at all. Assuming Isabella wants the money and would hate to get kicked in the shins, (74a) is true, while (74b) seems false.

(74) a. Isabella wants to get a million dollars.
b. Isabella wants to either get a million dollars or get kicked in the shins.

Notice that this time, resorting to the diversity presupposition won't help. Since Isabella believes that getting nothing is also an option, the propositions denoted by the embedded clauses in (74) are each compatible with, but not entailed by, Isabella's beliefs, meaning that the presupposition of *want* is satisfied for both.

However, Crnič (2011) observes that (74) gets a ready explanation within von Fintel's theory, so long as scalar implicatures are properly taken into account. While Crnič operates under a theory in which implicatures are calculated via the semantics of a syntactic head Exh (Fox 2007, Chierchia et al. 2012), the basic reasoning can be understood within a (somewhat informal) Gricean framework. In short, when interpreting an utterance of (74b), we also consider certain structurally-defined alternatives to that sentence. When a sentence includes a disjunction, among its alternatives will be (i) the sentence itself, (ii) a version of the sentence where *or* is replaced by *and*, and (iii) each version of the sentence where the disjunction is replaced with just one of the disjuncts (Sauerland 2004, Alonso-Ovalle 2006, Fox 2007). As a result, the set of alternatives to (74b) looks like (75):

(75) {Isabella wants money or kick, Isabella wants money and kick, Isabella wants money, Isabella wants kick}

Scalar implicatures are calculated by negating those alternatives that, to use Fox's (2007) term, are *innocently excludable*. For this example, we can simply say that an alternative is innocently excludable iff it is not entailed by the main assertion of the sentence. Importantly, by von Fintel's definition (74b) entails none of the latter three alternatives in (75), since Isabella could hypothetically have no preference between getting the money and getting kicked, but know that she wants one of them. (For example, this will be the case if her bouletically ideal worlds are a mixture of worlds where she gets kicked in the shin and receives no money, and worlds where she gets money and no kick in the shins.) Because these three alternatives are innocently excludable, the negation of each is generated as a scalar implicature, so that in effect (74b) is interpreted as (76):

(76) Isabella wants to either get a million dollars or get kicked in the shins, she doesn't want both, and she is apathetic between the two.

As Crnič notes, it may be this implicature of apathy that leads to an interpretation of (74b) as false, since in reality Isabella is very much opinionated about which option she gets. Thus, even though (74a) is interpreted as not entailing (74b), this is not necessarily because *want* is not Strawson upward-entailing.

Another example that has been argued to be problematic for a Strawson upward-entailing view of *want* is (77), adapted from Asher (1987):

(77) a. Nicholas wants to get a free trip on the Concorde.
b. Nicholas wants to get a trip on the Concorde.

If Nicholas refuses to pay the exorbitant price for a ticket to ride the Concorde, but would gladly take a free trip on the off chance that it's offered, then by all appearances (77a) is true, while (77b) is false. However, von Fintel (1999) notes that things are not so simple. If (77a) genuinely does not entail (77b), then we would expect (78) to be acceptable and potentially true, when in reality it sounds plainly contradictory.

(78) # Nicholas wants to get a free trip on the Concorde, but he doesn't want to get a trip on the Concorde.

von Fintel takes the oddness of (78) to indicate that the apparent lack of entailment from (77a) to (77b) is not semantic in nature, but is instead due to discourse factors that alter the domain of world-quantification when shifting from one sentence to the other. The contradictory nature of (78) can then be attributed to the fact that there is not enough "room" for such contextual shifts in between the first and second clause.

To be clear, the dust is far from settled on this issue, and debate continues over whether the clausal complement of *want* obeys something like Strawson upward entailment (see Levinson 2003; Villalta 2008; Crnič 2011; Lassiter 2011a,b; Rubinstein 2012, 2017; Condoravdi & Lauer 2016; Phillips-Brown 2016). I don't plan to pick a side in this dissertation, since as far as I can tell, the issues that arise in the discussion of *want* (and *wish* and *regret*) in Chapters 3 and 4 are issues that have not been addressed in any of the frameworks discussed above. But while I wish to avoid hitching myself to one particular theory, the formulation of my proposal in Chapter 4 will require the adoption of concrete definitions for the attitudes at hand. In doing so, I will implement my proposals using the theories of von Fintel (1999) and Heim (1992). There are two main reasons why I have opted for these two theories in particular. First, they diverge with respect to Strawson upward entailment, so showing that my proposal is compatible with both indicates that it does not commit me to a particular stance with respect to this inference pattern. Second, it happens to be the case that the exact same basic ideas and stipulations that work for von Fintel's theory work just as well for Heim's. While similar ideas can likely be implemented within the other frameworks discussed up to this point, this convenient result allows us to avoid a certain degree of back-tracking, reformulation, and redundancy that would do little to advance the actual content of this dissertation. I can only hope that proponents of these other theories will forgive me for forcing them to do the work of translating these proposals into their own theory of choice.

With this in mind, I will next turn my attention from *want* to *wish* and *regret*, the definitions of which are usually built out of parts scavenged from some definition of *want*. The basic nature of this translation from *want* to *wish* and *regret* can be implemented within any of the above theories, but since I will be adopting the frameworks of von Fintel and Heim, I will only show how these extensions work within their theories in particular.

2.6 From *want* to *wish* and *regret*

So far, we have focused our attention exclusively on two attitudes: *believe* and *want*. We can expand our inventory slightly by looking at *wish* and *regret*, which are similar to *want* in that they also deal with the preferences of some experiencer. But there are obviously some crucial differences as well. In this section, we will briefly explore these similarities and differences, and see how connections can be made in both von Fintel's and Heim's overarching theories.

Let's start with *wish*. Consider the sentence in (79):

(79) Stephanie wishes she had lifted weights this morning.

One of the first things to notice about this sentence is that it carries an inference that Stephanie didn't lift weights this morning. The fact that (80a) and (80b) also carry this inference suggests that the inference is presuppositional, as it appears to project through negation and questions.

(80) a. Stephanie doesn't wish that she had lifted weights this morning.
b. Does Stephanie wish that she had lifted weights this morning?

As it turns out, however, the actual inference isn't necessarily that Stephanie didn't lift weights this morning, but rather that she *believes* that she didn't lift weights this morning. (81), for example, is entirely felicitous because Stephanie believes that she didn't lift weights this morning, even though in actuality she did.

(81) Stephanie lifted weights this morning, but was slipped a memory-altering drug in the afternoon. Now, in her deluded state, she thinks she didn't lift weights this morning, and she wishes she had.

So it seems that α *wishes that p* presupposes that α believes $-p$, i.e., that $\mathrm{Dox}(\alpha, w) \cap p = \varnothing$. (As we will see later, this is an oversimplification, but let's hold on to it for now.) This is a crucial difference between *wish* and *want*, as the latter presupposes that there are p worlds (as well as $-p$ worlds) among the worlds in $\mathrm{Dox}(\alpha, w)$. This also means that whatever the modal domain is for *wish*, it cannot simply be the set of belief worlds. After all, if there are no p belief worlds, then p cannot hold across all ideal belief worlds (von Fintel), and it makes no sense to ask if p belief worlds are better than all of the closest $-p$ worlds (Heim). So what, then, is the modal domain?

Here, roughly, is the idea. When we talk about desires using *want*, we are talking about someone's desires relative to what they perceive to be the current live options. That is, if I currently want to eat some pizza, I am evaluating the preferability of eating pizza relative to the possibilities that I believe to be available to me now. Similarly, if I wanted pizza an hour ago, that desire was relative to the possibilities that I thought were available to me at the time. But if I currently *wish* that I *had* eaten some pizza, I am instead evaluating the preferability of those options that I believe *were* available to me at some (presumably contextually-determined) previous time. So if I currently wish that I had eaten pizza an hour ago, the possibilities at play are those that I thought were available an hour or so ago, or perhaps longer if delivery would have been involved.

A reasonable assumption that often comes up in work on tense and modals (and attitudes) is that the set of possibilities for what the world may come to look like shrinks over time. That is, time has a way of closing doors, but not of opening them. So if it is now possible that my phone will run out of batteries at 1:00pm today, then it was possible ten minutes ago (and ten years ago) that my phone would run out of batteries at 1:00pm today. But if it was possible in 2000 that I might be drafted as a professional football player in 2011, that by no means entails that that is a current possibility—in fact, since 2011 has come and gone without a Pasternak draft pick, this is now a complete *im*possibility. Adopting this perspective for *want* and *wish*, this means that the modal domain for *wish* will be a proper superset of the set of worlds used for *want*, since *want* is about perceived contemporaneous possibilities, while *wish* is about perceived prior possibilities.

As Heim (1992) notes, this difference between *want* and *wish* has an interesting analog in the realm of conditionals. Consider the contrast between (82a) and (82b):

(82) a. If Sara takes Cara to lunch, Rivka will cover the costs.
b. If Sara had taken Cara to lunch, Rivka would have covered the costs.

Adopting the Stalnaker-Lewis analysis of conditionals, (82a) is true iff, among all those worlds circumstantially accessible from the actual world given the way things are now, the most similar ones in which Sara takes Cara to lunch are all worlds in which Rivka covers the costs. (82b), meanwhile, is true iff among all those circumstantially acessible worlds given the way things were *at some previous time*, the closest ones in which Sara took Cara to lunch are all worlds in which Rivka covers the cost. Thus, counterfactuals make use of a wider domain than indicative conditionals in evaluating their modal claims, similar to the relationship between *wish* and *want*. In fact, as Iatridou (2000) notes, this semantic parallel comes with a morphosyntactic parallel as well: cross-linguistically, the verbal morphology used in the complement of *wish* looks like the verbal morphology in the antecedents of counterfactual conditionals. (In English, this can be seen in the "extra layer" of past tense in (79) and (82b).)

With this in mind, in both the von Fintel and Heim frameworks, the transition from *want* to *wish* can be effected in two steps. First, all references to belief worlds are replaced with the extended domain discussed above, which will be a proper superset of the set of belief worlds. I will refer to this set as $\mathrm{Dox}^+(x, w)$ for experiencer x and world of evaluation w. Second, the aforementioned presupposition that the propositional argument is incompatible with the experiencer's beliefs is added. The result in von Fintel's framework can be seen in (83a), and for Heim's framework it will be as in (83b):[18]

[18] Note that I no longer use sets of propositions to generate world-orderings in von Fintel's framework.

(83) a. $\lambda p \lambda x \lambda w : \text{Dox}(x,w) \cap p = \varnothing \wedge \exists w', w'' \in \text{Dox}^+(x,w)[p(w') \wedge \neg p(w'')].$
$\forall w' \in \textsc{Best}(\text{Dox}^+(x,w), \precsim_{x,w})[p(w')]$

b. $\lambda p \lambda x \lambda w : \text{Dox}(x,w) \cap p = \varnothing \wedge \exists w', w'' \in \text{Dox}^+(x,w)[p(w') \wedge \neg p(w'')].$
$\forall w' \in \text{Dox}^+(x,w)[\text{Sim}_{w'}(\text{Dox}^+(x,w) \cap p) \prec_{x,w}$
$\text{Sim}_{w'}(\text{Dox}^+(x,w) \cap -p)]$

Note, by the way, that while *want* and *wish* have different modal domains, the ordering over worlds remains the same, namely, $\precsim_{x,w}$. The reason for this is that for both attitudes, worlds are ordered with respect to their current preferability to the experiencer. The difference is that while *want* focuses on preferences between live options, *wish* includes in its comparisons options that are no longer actually attainable.

Pivoting from *wish* to *regret*, things seem to be mostly the same, but with a switch in polarity. Thus, whereas α *wishes* p presupposes that α believes p to be *false*, α *regrets* p presupposes that α believes p to be *true*. And while α *wishes* p asserts that α has a retrospective preference *in favor* of p, α *regrets* p asserts that α's preference is *against* p. Hence, (84) seems to mean the same thing as (79) with respect to both assertion and presupposition:

(84) Stephanie regrets that she didn't lift weights this morning.

Naturally, then, *regret* can be defined as in (85), independently of the framework used for defining *wish* (and *want*):

(85) $[\![\text{regret}]\!] = \lambda p \lambda x \lambda w.\ [\![\text{wish}]\!](-p)(x)(w)$

That is, to regret p is to wish that $-p$.

2.7 Where's the quantification?

Before tying a bow on this chapter, I'd like to discuss an interesting new direction in the semantics of attitudes that I will not directly address in this dissertation, but that I believe warrants some mention. In all of the theories of attitudes discussed so far, all of the semantic work—and in particular the quantification over possible worlds—is done in the lexical semantics of individual attitude verbs. However, Kratzer (2006) introduces an alternative possibility: namely, for at least some, and perhaps all attitude verbs, it is not the attitude itself that introduces the quantification over possible worlds, but rather a head in the left periphery of the embedded clause.

More specifically, Kratzer adopts a neo-Davidsonian framework, in which the semantics of verbs hovers around an event variable. In conjunction with the view that the external argument is introduced by a separate voice head (Kratzer 1996), the

Outside of the appendix to Chapter 4, this will remain the case through the rest of this dissertation.

result is that the denotation of *believe* is simply a relation between a thing-believed and a belief state:

(86) $[\![\text{believe}]\!] = \lambda x \lambda e.\ \text{believe}(e, x)$

The internal argument (x in (86)) can be saturated by DPs whose denotations have truth-conditional content, such as *my story* or *the president's lies*. Thus, *Maria believes my story* will have the simple interpretation in (87):

(87) $[\![\text{Maria believes my story}]\!] = 1$ iff
$$\exists e[\text{Exp}(e) = \text{maria} \wedge \text{believe}(e, \iota x[\text{story-of-mine}(x)])]$$

But what about something like (88)?

(88) Maria believes that she is a good person.

Naturally, *she is a good person* will denote a proposition. The key work will be done by *that*, which turns this proposition into a predicate. This predicate will be true of an entity x with truth-conditional content iff in all worlds compatible with x's content, Maria is a good person. Such a denotation for *that* can be seen in (89):

(89) $[\![\text{that}]\!] = \lambda p \lambda x.\ \forall w \in \text{content}(x)[p(w)]$

But now we seem to have a type mismatch: *believe* denotes a relation between an eventuality and an entity, while *that she is a good person* denotes a predicate. To fix this, Kratzer uses the compositional rule Restrict (Chung & Ladusaw 2004), which in short states that in this scenario the predicate denoted by the embedded clause (plus complementizer) restricts the first argument of *believe*, leading to something like (90):

(90) $[\![\text{believe that she is a good person}]\!] =$
$$\lambda x \lambda e.\ \text{believe}(e, x) \wedge \forall w \in \text{content}(x)[\text{good-person}(w)]$$

After adding the external argument, as well as existential closure of both the entity and eventuality arguments, we are left with the interpretation in (91):

(91) $[\![\text{Maria believes that she is a good person}]\!] = 1$ iff
$$\exists x \exists e[\text{Exp}(e) = \text{maria} \wedge \text{believe}(e, x) \wedge \forall w \in \text{content}(x)[\text{good-person}(w)]]$$

This approach of injecting world-quantification into the semantics of a nearby head, rather than into the allegedly opaque verb itself, has gained some traction in subsequent work on intensional verbal constructions like attitudes (see, e.g., Deal 2008, Moulton 2009, Bogal-Allbritten 2016, Moltmann 2017). But what's the benefit of such modal relocation? Arguments abound, but I will only go over a couple, as many of the arguments are of a parallel nature: namely, pointing out particular cases where either the presence of world-quantification depends on the existence of an embedded clause, or changing the structure of the embedded clause changes the nature of the quantification over worlds.

As examples of the former sort, Kratzer (2016) notes that there are many verbs that act as *verba dicendi* (verbs of saying) in certain syntactic contexts, but needn't do so in others:

(92) a. Joseph screamed.
b. Joseph screamed that the killer was after him.

(93) a. Janet sighed.
b. Janet sighed that her stocks were underperforming.

(94) a. The patient groaned.
b. The patient groaned that the medication was wearing off.

Notice that in the (a) examples, the scream, sigh, or groan does not have to have truth-conditional content, and can simply be a noise made by the referent of the subject. The (b) sentences, meanwhile, seem to claim that the scream (for example) was also a speech act: Joseph said, in a screaming voice, that the killer was after him. The fact that not all acts of screaming are acts of saying strongly suggests that any quantification over worlds oughtn't be in the lexical semantics of *scream*. Rather, it should be placed either in the *that* clause or in some head that comes along with it, since it is the presence of this clause that adds the requirement that the screaming have truth-conditional content.

As for arguments of the second kind—cases where changing the embedded clause changes the nature of world-quantification—the English *wish* may be a good example. As noted above, the complement of *wish* often has counterfactual morphology in the embedded clause, which in English means an extra layer of past tense, as in (79), repeated below:

(79) Stephanie wishes she had lifted weights this morning.

Perhaps not coincidentally, the semantics of *wish* seems to be counterfactual in nature, in that it uses an expanded domain of worlds that are ordered by bouletic preferability (see the previous section). But as it turns out, getting rid of the counterfactual morphology can in turn mean losing the counterfactual-like interpretation, as in (95), which seems to be a normal claim about the captain's present desires:

(95) The captain wishes to speak to you.

Thus, the choice of modal domain for *wish* seems to depend in part on the syntactic structure of the embedded clause.

Venturing beyond English, Bogal-Allbritten (2016) points out that in Navajo, the verb *nisin* can be used to mean a variety of attitudes, including 'believe', 'want', and 'wish'. (96) is an example where multiple meanings are available, as *nisin* can mean either 'believe' or 'want' (all glosses adapted from Bogal-Allbritten 2016):

(96) Kii nahodoołtį́į́ł nízin.
Kii ArealS.rain.FUT 3S.ATT.IMPF
Ambiguous: Kii **thinks** it will rain, or Kii **wants** it to rain.
(Bogal-Allbritten 2016: 113)

Bogal-Allbritten employs a variety of tests to argue that the availability of these multiple readings is due to a genuine syntactic-semantic ambiguity, and not just a weak meaning of *nisin*. That is, the two interpretations above are genuinely two interpretations, and (96) doesn't just mean something like 'Kii has an attitude toward the proposition that it will rain.' For example, if Ron thinks that Obama will win but wants him to lose, while Kii wants Obama to win but thinks that he will lose, (97) is false:

(97) Ron dóó Kii Obama hodínóołnééł nízin.
Ron and Kii Obama 3S.win.FUT 3S.ATT.IMPF
Intended: Ron and Kii have some feeling about Obama winning.
(Bogal-Allbritten 2016: 147)

Importantly, the inclusion of certain particles in the embedded clause disambiguates between readings: the sentences in (98), unlike (96), are unambiguous, due to the inclusion of the particles *sha'shin* and *laanaa*, respectively.

(98) a. Kii nahałtin sha'shin nízin.
Kii ArealS.rain.IMPF probably 3S.ATT.IMPF
Unambiguous: Kii **thinks** it must be raining.
(Bogal-Allbritten 2016: 80)

b. Níneez laanaa nisin.
2S.tall wishful 1S.ATT.IMPF
Unambiguous: I **wish** you were tall. (Bogal-Allbritten 2016: 64)

Much like the English *wish* case above, these seem to be cases where some element in the embedded clause (in this case, the particles) partially determines the domain of quantification over possible worlds.

That being said, all of the above data are compatible with the view that it is still *nisin* that is responsible for the world-quantification, and that the ambiguity above is due to simple homophony: there are multiple *nisin*s, including (at least) one that means 'believe', one that means 'want', etc. The facts in (96) and (97) are then readily accounted for, since (96) would just be a lexical ambiguity, and (97) would indeed require that either Ron and Kii both want Obama to win, or they both think Obama will win, depending on the choice of *nisin*. As for (98a), Bogal-Allbritten glosses *sha'shin* as 'probably', suggesting that it is something like an epistemic modal. But there are independently attested restrictions on which embedding verbs can take epistemic modals in their complement clause, as discussed by Hacquard (2006) and Anand & Hacquard (2008, 2013):

(99) a. Arin thinks that Danny might be in LA.
b. # Arin commanded that Danny might be in LA.

Importantly, Hacquard and Anand & Hacquard account for the facts in (99) while keeping the quantification over worlds tucked away in the semantics of the verb. So a homophony account can claim that in (98a), each of the forms of *nisin* is in principle available, but only 'think' *nisin* leads to anything other than semantic deviance. It is not implausible that a similar account can be given for *laanaa* in (98b).

But what seems like the clincher in favor of a view in which it is indeed the embedded clause that determines the choice of attitude is cases like (100):

(100) Alice Bill Kinłánígóó 'ííná dóó bich'į deesháał nízin.
Alice Bill Flagstaff.to 3S.move.PERF and 3O.to 1S.go.FUT 3S.ATT.IMPF
Alice thinks Bill moved to Flagstaff and she wants to go see him.
(Bogal-Allbritten 2016: 152)

In (97), there was one *nisin*, one embedded clause, and a conjoined subject, and the result was that there could only be a single attitude, i.e., Ron and Kii have to both be wanting or both be believing. In (100), there is still one *nisin*, but this time it is the embedded clauses that are conjoined. As a result, (100) allows a reading where one proposition is believed, and another is desired. Such an example is not predicted on a homophony account: there is only one *nisin* in (100), so both propositions should be believed or both wanted.[19] But it is fully compatible with embedded clauses being the origin of the quantification over worlds, since the presence of two embedded clauses means, or at least can mean, the presence of two embedded-clause-contained world-quantification operators.[20]

[19] Note that claiming that (100) is an instance of right-node raising, as in (i), does not help the homophony theorist:

(i) Jesse loves $\_\_\__i$, but Katja hates [the photo of the bank on my wall]$_i$.

Regardless of one's account of right-node raising, the lexically ambiguous *bank* in (i) has to be interpreted the same way in the gap as in the pronounced object DP. Hence, if I have a photo of a financial institution and a photo of some land alongside a river, (i) cannot mean that Jesse loves the former photo, while Katja hates the latter.

[20] The fact that (i) below is odd suggests that the cases of English *wish* and Navajo *nisin* are not entirely parallel:

We thus see that a variety of evidence suggests, at the very least, that a significant proportion of the semantics associated with attitudes comes from "downstairs", i.e., in the clausal complement. While this is of great interest in the semantics of attitudes more generally, I will opt for a more traditional view in which it is the semantics of the attitude itself that is responsible for the quantification over possible worlds. I believe that the ideas proposed in this dissertation are compatible with a more downstairs-oriented view of the semantics of attitudes, but attempts at such extensions and revisions will be left for future work.

2.8 Conclusion

In this chapter, I have offered a panoramic overview of (one slice of) the history of the semantics of attitudes, mostly sticking to *believe*/*think* and *want*. Along the way, many questions arose, including how presuppositions project in attitudes, whether the clausal complement of *want* is a (Strawson) upward-entailing environment, how certain entailments (such as those based on believed equivalence) can be constrained, and whether *want* and other desiderative attitudes are semantically focus-sensitive. While I have sometimes offered my own opinions on such matters, in general the topics addressed in the rest of this dissertation will be neutral with respect to these questions. Thus, for the reasons stated previously I will use as my foundation Hintikka's (1969) definition of *believe*, and the definitions for *want* (and *wish* and *regret*) of Heim (1992) and von Fintel (1999).

(i) ?? The captain wishes that the seas had been smoother yesterday, and to speak to you this afternoon.

However, I'm not convinced that this means that the case of English *wish* should be treated as lexical ambiguity along the lines rejected for *nisin*. Rather, the two embedded clauses in (i) may be different kinds of phrases: the first appears to be a CP (hence the complementizer *that*), while the second may be a TP or something smaller. If this is true, then they cannot be conjoined for strictly syntactic reasons, as the two conjuncts must be the same type of phrase.

Chapter 3: Intensity is monotonic

The theories of attitude semantics discussed in the previous chapter were generally non-Davidsonian in character: ⟦want⟧, for example, was treated as a relation between a proposition, an individual, and a world, with no argument for events or states. Starting in this chapter, we will make the Davidsonian turn, adding eventualities into the denotations of verbs (including attitudes). Thus, in the same way that for Bill to run is for there to be an *event* of Bill running, it will also be the case that for Bill to want to leave is for there to be a *state* of Bill wanting to leave.

In this chapter I will put aside the issue of just how the theories of attitude semantics discussed in the previous chapter should be translated into a Davidsonian framework. This is unproblematic because the hypothesis for which I will be arguing is focused very narrowly on the part-whole structure of attitude states themselves, to the exclusion of almost all of the rest of the semantic bells and whistles involved in the denotations of attitudes. So while the particular implementation of my proposal in the next chapter will require a more fleshed-out attitude semantics, for the purposes of this chapter all that we need is the Davidsonian argument, as well as some basic ideas about clausal compositionality and argument structure that will be introduced shortly.

The hypothesis that I will be arguing for in this chapter is that in the natural language ontology of Davidsonian mental states, including attitude states as a particular instance, the intensity of such a state correlates with its part-whole structure in a particular dimension. Put simply, a more intense psychological state is "bigger" along this dimension than another, less intense psychological state. The reasoning underlying this is as follows. As I discuss in Sections 3.1 and 3.2, there is a class of nominal and verbal measure constructions in which the measurement used must track part-whole relations within a particular domain; to use Schwarzschild's (2002, 2006) term, the measurement must be *monotonic*. This class of constructions includes pseudopartitives (*twelve ounces of gold*), the measurement idioms *out/up the wazoo* and *in spades*, adverbial measure phrases (*Chuck ran <u>a lot</u> yesterday*), and nominal and verbal comparatives.[1]

[1] In consulting with other native English speakers, there seems to be variation in whether *out* or *up* is the

As an example, consider the verbal comparative in (1):

(1) Dee ran more than Evan did.

Depending on context, (1) can serve as a comparison of the distance of Dee's and Evan's running, or of temporal duration. However, it cannot serve as a comparison of the *speed* of Dee's and Evan's running. If Dee ran one mile in four minutes, while Evan ran three miles in thirty minutes, (1) is simply false, even though Dee ran faster than Evan did. The reason for this, as observed by Wellwood et al. (2012) and Wellwood (2014, 2015) (building on work by Nakanishi (2007) and Bale & Barner (2009)), is that distance and temporal duration respect the part-whole relations of running events in a way that speed does not: a running event covers more distance and time than any of its proper parts, but it will not have a greater speed than all of its proper parts.

To the extent that this monotonicity requirement for verbal comparatives is robust, it provides an argument for a connection between psychological intensity and the part-whole structure of mental states, as mental state verbs can appear in verbal comparatives in which intensity is measured. This is shown in (2) with transitive mental state verbs like *like* and *hate*, and in (3) with desiderative attitudes (i.e., attitudes based on preferability).

(2)
a. Fiona likes football more than she does baseball.
b. Gavin fears clowns less than he does sharks.
c. Helen hates country music as much as she does rap.
d. Ina respects her teachers more than she does her friends.
e. Jorge admires the CEO less than he does his co-workers.
f. Kwame trusts the poor as much as he does the rich.
g. Marvin loves biology more than he does history.

(3)
a. Jo wants to leave more than Ben wants to stay.
b. Stan wished he'd won more than he wished he'd stayed healthy.
c. Paul regrets buying his car more than Nora regrets selling hers.

After providing a compositional semantics for verbal comparatives like (1) in Section 3.3, in Section 3.4 I will show what it looks like for this compositional semantics to extend to *want* comparatives like (3a) if we assume that they compose just like other verbal comparatives.

But before adopting the view that intensity is a monotonic measure of desire states, there is another, equally plausible hypothesis worth considering: namely, that

preferred preposition in *out/up the wazoo*. I will stick to *out the wazoo* for the rest of this dissertation.

mental state comparatives *don't* compose like other verbal comparatives, and whatever structure imposes the monotonicity requirement for verbal comparatives like (1) is absent in verbal comparatives with mental state verbs. In Section 3.5, I flesh out what such an approach might look like, and discuss evidence from Chinese that indeed illustrates such a distinction between "normal" verbal comparatives and intensity comparatives. However, in Section 3.6 I will argue that this distinction does not suffice as a counterargument to a view in which intensity of mental states is monotonic, for two reasons. First, in English, the intensity of mental states can be measured using not only verbal comparatives, but all five of the normally monotonicity-requiring constructions discussed in Section 3.2, requiring the positing of a wide-ranging structural distinction—with no overt evidence in its favor—across all five constructions. Second, Chinese has at least two other normally monotonicity-requiring constructions that can be used to measure intensity of mental states, and the structural distinction in verbal comparatives that motivated the counterargument to begin with disappears in these constructions. With this in mind, I show at the end of Section 3.6 that a proposal in which intensity correlates with part-whole relations of mental states can readily account for the similarities and differences across languages and constructions, while a view in which intensity is non-monotonic faces an uphill battle. In 3.7 I offer some concluding remarks.

In this chapter, I focus exclusively on the arguments in favor of a mereological approach to the intensity of attitudes and other mental states. Having put the arguments forward in this chapter, in the next chapter I will show how a monotonic natural language metaphysics of psychological intensity can actually be implemented.

3.1 What is monotonicity?

There are many ways one can measure a chunk of gold: by volume, weight, temperature, purity, density, etc. But there is a fundamental difference between weight and purity, for example. If a given chunk of gold weighs twelve ounces, we know for certain that if we chip off a piece of that chunk and weigh it, it will weigh less than twelve ounces. But if the purity of that chunk of gold is eighteen carats, it is not guaranteed that by chipping off a piece, we will end up with a chunk of a lesser purity. It is not impossible, as we might happen to be left with a particularly impure bit of the gold, but importantly, it is not guaranteed.

Similar facts hold, for example, of the volume and temperature of a collection of water. If I start with three liters of water and pour some out, I am certain to be left with less than three liters of water. But if my water is 30° Celsius, then there is no guarantee that after pouring some out, I will be left with water with a lower or higher temperature than 30°C; if anything, the smart money would be on still having water that is 30°C.

Now consider the case of the depth of a collection of snow. There is a sense in which depth is like weight and volume, and a sense in which it is not. Let's say that Baltimore got two feet of snow, with each part of Baltimore having received the same amount of snow. It is not the case that if we remove any bit of snow, we are guaranteed to be left with snow that is less than two feet deep: if we remove all and only the snow in East Baltimore, the remaining snow will still have a depth of two feet. However, if we are only allowed to remove snow in "sheets", removing thin layers of snow that cover the whole area of Baltimore, then it will indeed be the case that by removing some snow, we will be left with snow of a depth less than two feet. An illustration of these two ways of removing snow can be seen in Figure 3.1.

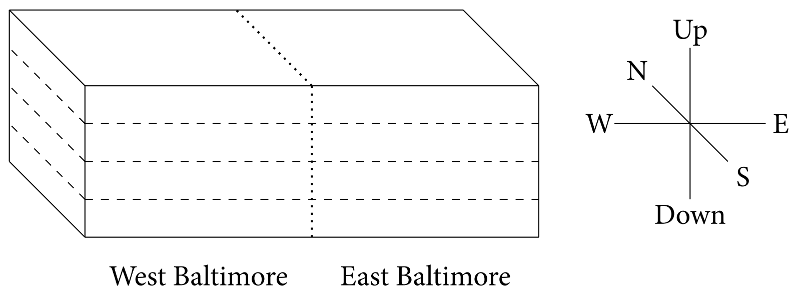

Figure 3.1: Illustration of two ways of removing snow from Baltimore: "chopping off" the snow from East Baltimore (dotted line), and slicing off layers (dashed lines).

Rather than speaking in terms of measuring, removing a portion, and remeasuring, we can instead talk about these measure functions in terms of whether they track certain part-whole relations. Weight tracks part-whole relations of gold, since a bit of gold necessarily weighs more than any of its proper parts; purity, however, does not, since a chunk of gold will not necessarily be purer than a given proper part of it. Similarly, volume tracks part-whole relations of water, while temperature does not. Meanwhile, depth tracks some, but not all, part-whole relations of snow. If the part-whole relation under question is that between the snow in West Baltimore and the snow in all of Baltimore, depth does not track part-whole relations. But depth does track the part-whole relations between layers of snow and their sums, since the sum of two layers of snow is guaranteed to have a greater depth than each of those layers individually.

It will be useful to refer to measure functions like weight, volume, and depth as members of a single class that excludes, e.g., temperature and purity. Several ways of doing this have been proposed in the literature; I will use Schwarzschild's (2002, 2006) notion of a *monotonic* measure function, formally defined in (4):[2]

[2]See also Krifka's (1989) reference to *extensive measure functions*. Champollion (2015b) defines a sim-

(4) Let μ be a measure function, A a domain of entities, and $\sqsubseteq^c$ a contextually salient part-whole relation. μ is **monotonic** on $\sqsubseteq^c$ in A iff for all $x, y \in A$, if $x \sqsubset^c y$, then $\mu(x) < \mu(y)$.[3]

Notice that by the definition in (4), a measure function is not monotonic (or non-monotonic) *simpliciter*, but rather is (non-)monotonic *on* a salient part-whole relation, *in* a domain. So μ_{weight}, which takes an entity and returns the degree that is its weight, is monotonic on pretty much any part-whole relation in $[\![\text{gold}]\!]$, while μ_{purity} is not; *mutatis mutandis* for $\mu_{\text{volume}}/\mu_{\text{temperature}}$ and $[\![\text{water}]\!]$. As for μ_{depth}, whether or not it is monotonic on a part-whole relation in $[\![\text{snow}]\!]$ depends on the part-whole relation. But if $\sqsubseteq^c$ is the part-whole relation between layers of snow and their sums, then μ_{depth} is indeed monotonic on $\sqsubseteq^c$ in $[\![\text{snow}]\!]$. All this being said, in cases where the part-whole relation and domain are clear or irrelevant, I will frequently refer to a measure function as simply being (non-)monotonic.

In the next section, I will discuss five measurement-related constructions that impose monotonicity requirements on the chosen measure function. Before doing so, however, it is worth noting that the requirement they impose is actually slightly stronger than monotonicity, in that the measure function must be *non-trivially* monotonic, as defined in (5):

(5) Let μ be a measure function, A a domain of entities, and $\sqsubseteq^c$ a contextually salient part-whole relation. μ is **non-trivially monotonic** on $\sqsubseteq^c$ in A iff (i) μ is monotonic on $\sqsubseteq^c$ in A, and (ii) there exist $x, y \in A$ such that $x \sqsubset^c y$.

To illustrate the difference between (4) and (5), consider the domain $[\![\text{boy}]\!]$, which contains all and only atomic (i.e., individual) boys. Obviously, no boy is a mereological proper part of any other boy. Thus, the universal quantification in the definition in (4) ends up being vacuous in this case, irrespective of the choice of measure function. (5) safeguards against this vacuity. That being said, in the examples discussed in this chapter, triviality is sidestepped, and the natural language metaphysics for attitude intensity discussed in the next chapter also renders attitude intensity non-trivially monotonic.

3.2 Constructions with monotonicity requirements

3.2.1 Pseudopartitives

One example of the grammatical relevance of monotonicity is pseudopartitives like *twelve ounces of gold* (Krifka 1989; Schwarzschild 2002, 2006; Brasoveanu 2009). As

ilar (but non-identical) concept of *stratified reference*, used to a similar effect.

[3] The more mathematically inclined will note that what Schwarzschild refers to as "monotonicity" is in fact *strict upward monotonicity*. However, I will stick to Schwarzschild's simpler term.

an illustration, consider the sentences in (6):

(6) a. i. Louise bought twelve ounces of gold.
ii. # Louise bought eighteen carats of gold.
b. i. Max poured three liters of water into the tub.
ii. # Max poured 30°C of water into the tub.
c. i. Baltimore got two feet of snow.
ii. # Baltimore got 20°F of snow.

The examples with monotonic measure functions—weight in (6a-i), volume in (6b-i), and depth in (6c-i)—are all acceptable, while those with non-monotonic measure functions—purity in (6a-ii), and temperature in (6b-ii) and (6c-ii)—are out.

In the examples in (6), all of the measure phrases unambiguously denoted a particular degree on a particular scale. But this needn't necessarily be the case, as pseudopartitives with vague measure phrases like *a great deal*, *a lot*, and *a ton* (on a non-literal interpretation) are all permissible:

(7) a. Nevin bought {a great deal/a lot/a ton} of coffee.
b. Baltimore got {a great deal/a lot/a ton} of snow last week.

In these cases, the measure phrases are not only vague, but also capable of denoting degrees on distinct scales: *a lot* can denote a degree of volume in (7a) and a degree of depth in (7b). This flexibility in interpretation can be further illustrated by fixing the measure function (and thus the scale) by means of *in terms of NP*, where *NP* is a type of measurement.

(8) In terms of volume, Owen ate a lot of pudding. But in terms of weight, he didn't eat very much.

(8) essentially means that Owen ate pudding that was not very dense: there was a large volume of it, but it did not weigh very much.

When looking at contextually-determined measure phrases like these, it can be a bit tricky to check for monotonicity requirements, since unlike in (6), the predicted difference is in available readings, rather than acceptability. Of course, one way to check would be by virtue of truth value judgments. For example, if Nevin bought a small volume and weight of coffee, but the coffee was exceptionally dark, (7a) is straightforwardly false, presumably because darkness is non-monotonic on part-whole relations in ⟦coffee⟧. Similarly, if Baltimore only got an inch of snow last week, but the snow was exceptionally cold, (7b) is still false. However, given that the choice of measure function is sensitive to context, it is conceivable that non-monotonic measure functions like temperature are not ruled out by the grammar *per* *se*, but are strongly dispreferred for pragmatic reasons, so that a great deal of contextual setup has to take place in order for such readings to be sufficiently salient. Ideally, then, we would have a test in which the difference is in acceptability, rather than truth conditions, so that we can rule out the possibility of a dispreferred but nonetheless available reading with a non-monotonic measure function.

Fortunately, such a test exists. As mentioned above, *in terms of NP* can be used to fix the choice of measure function. Therefore, if we try to use *in terms of NP* to force the use of a non-monotonic measure function, the result is predicted to be odd. As can be seen in (9), this prediction is in fact borne out:

(9) a. In terms of {volume/??darkness}, Nevin bought a lot of coffee.
b. In terms of {depth/??coldness}, Baltimore got a ton of snow.
c. In terms of {weight/??viscosity}, Owen ate a great deal of pudding.

So in cases with a vague measure phrase, we now have two ways to test whether a particular measure function is available. The first is by means of standard truth value judgments. The second is to see whether the sentence remains felicitous when trying to force a reading with that measure function by means of *in terms of NP*.

Since it will be relevant later, it is worth noting that pseudopartitives can be used to measure not only entities, but eventualities, as can be seen in (10) with the deverbal nominalizations *driving* and *acceleration*:

(10) a. i. Otto did {twenty minutes/ten miles} of driving yesterday.
ii. # Otto did thirty miles per hour of driving yesterday.
b. i. Nell's car only managed {three seconds/five miles per hour} of acceleration before breaking down.
ii. # Nell's car only managed 5°F of acceleration before breaking down.

Once again the measure functions used must be monotonic. A driving event covers more distance and time than its proper parts, but is not necessarily faster, so (10a-i) is acceptable, while (10a-ii) is not. As can be seen in (10b-i), the unacceptability of (10a-ii) is not because measurements involving speed are somehow bad in and of themselves. After all, while speed is not a monotonic measure of driving events, the change in speed of an object *is* a monotonic measure of acceleration events, since a bigger acceleration event will lead to a greater change of speed than any of its proper parts. Hence, a pseudopartitive in which the change of speed (or temporal duration) of an acceleration event is measured is acceptable. Meanwhile, as illustrated in (10b-ii), measuring the change in *temperature* of the object undergoing acceleration is not permissible in a pseudopartitive, even if the acceleration is assumed to be the direct cause of the change in temperature. This is because accelerating does not entail heating up, so it is not the case that an event of acceleration will always involve a greater increase in temperature than any of that event's proper parts; in other words, change of temperature is not monotonic in the domain of ⟦acceleration⟧.

3.2.2 *Out the wazoo* and *in spades*

In addition to pseudopartitives, English has a variety of idioms used to indicate a large amount of something, such as *NP out the wazoo* and *NP in spades*. Naturally-occurring examples of these expressions retrieved from the Internet can be seen in (11) and (12) (emphasis my own):

(11) a. Right now, most of Texas has water out the wazoo.[4]
 b. We have snow out the wazoo and all I have is some Bridgestone all season tires on our vehicles.[5]
 c. Soon, we had milk out the wazoo, and I had to figure out what to do with all of it.[6]

(12) a. Behana Gorge delivers rainforest beauty and…water in spades.[7]
 b. They have snow in spades![8]
 c. I had extra milk in spades, so she mixed it [in] her food.[9]

The same monotonicity requirement seen in pseudopartitives arises here as well. For example, if Texas only has a small amount of water, but that water is very pure or cold, (11a) is false; what is required is that Texas have a very large amount of water, by depth or by volume. In the case of (11b) and (11c), it is necessary that there be a significant amount of snow or milk, rather than a very hot, cold, viscous, tasty, or nutritious portion.

Similar facts hold for *NP in spades* as for *NP out the wazoo*: (12a) requires that Behana Gorge have a large amount of water, and cannot mean that it has particularly hot or cold water. Furthermore, (12b) and (12c) again disallow measurements based on temperature, viscosity, etc. In other words, both *out the wazoo* and *in spades* require the use of measure functions that are monotonic in the domain of the modified NP.

[4] http://www.tribtalk.org/2015/06/08/the-texas-drought-is-over-but-what-about-the-next/

[5] http://bobistheoilguy.com/forums/ubbthreads.php/topics/2114123/Snow_shoes_for_the_Thunder(sno

[6] http://www.ourcoop.com/ourcoop08/headlines/viewNews.aspx?artID=3433

[7] http://www.aussiedrifterz.com.au/
(This example was slightly revised to better illustrate the point at hand.)

[8] http://stuebysoutdoorjournal.blogspot.com/2012/12/head-for-high-country-to-find-snow-plan.html

[9] https://www.reddit.com/r/JUSTNOMIL/comments/4frtoy/yes_mil_i_know_my_boobs_are_bigger/

3.2.3 Adverbial measure phrases

Earlier, we saw that the monotonicity requirement for pseudopartitives extended to cases where the noun denoted a set of events, as illustrated in (10) above. This extension from entities to events is further exemplified by the adverbial use of vague measure phrases like *a lot*, as in (13):

(13) a. Mara swam {a great deal/a lot/a ton} yesterday.
 b. It rained {a great deal/a lot/a ton} in London last week.

If Mara swam for two seconds at breakneck speed, (13a) is false, since she has to have swum a great distance or for a long time in order for (13a) to be true. Similarly, if a small amount of rain fell in London over a small amount of time, but that rain was highly acidic, (13b) is false. Once again, this correlates with the (non-)monotonicity of the chosen measure function, since speed of swimming and acidity of rain are not monotonic measure functions.

The *in terms of NP* test used for pseudopartitives with vague measure phrases provides further evidence that these adverbial measure phrases only allow monotonic measure functions. As can be seen in (14), *in terms of NP* can be used to fix the measure function used:

(14) a. In terms of time, Mara swam a lot yesterday. But in terms of distance, she only swam an average amount.
 b. In terms of time, it rained a great deal in London last week. But in terms of amount, it didn't rain all that much.

(14a) means that Mara swam slower than average, as she covered an average distance in a large amount of time. (14b) would likewise be true if there was a light drizzle in London that lasted for a long time. Importantly, when trying to use *in terms of NP* to force a non-monotonic measure function, the result is once again odd.

(15) a. ?? In terms of speed, Mara swam a lot yesterday.
 b. ?? In terms of acidity, it rained a great deal in London last week.

So it appears that adverbial measure phrases, much like their pseudopartitive counterparts in the nominal domain, impose a monotonicity requirement on the measure function used.

3.2.4 Nominal comparatives

As noted by Schwarzschild (2002, 2006), Wellwood et al. (2012), and Wellwood (2014, 2015), nominal comparatives also exhibit a monotonicity requirement. As an example, consider (16) below:

(16) Baltimore got more snow than Williamstown did.

(16) can be interpreted as comparing depth or overall volume of snow, but not coldness. The fact that both depth and overall volume are available means of comparison can be seen in (17), which makes use of *in terms of NP*.

(17) In terms of depth, Williamstown got more snow than Baltimore did, but in terms of overall volume, Baltimore got more snow than Williamstown did.

As before, the unavailability of coldness as a choice of measurement can be shown in two ways. The first is by truth value judgment: if the depth and overall volume of the snow in Williamstown exceed those of the snow in Baltimore, but the snow in Baltimore is colder than that in Williamstown, then (16) remains simply false. Second, the *in terms of NP* test once again differentiates between depth and volume on the one hand, and temperature on the other:

(18) ?? In terms of coldness, Baltimore got more snow than Williamstown did.

The same sort of reasoning can be applied to (19), which allows for a comparison of weight or volume, but not viscosity, of pudding:

(19) Pauline ate more pudding than Owen did.

As can be seen in (20), both weight and volume can be specified by *in terms of NP*. (21) further shows that trying to use *in terms of NP* to force a reading in which viscosity is compared leads to oddity.

(20) In terms of weight, Owen ate more pudding than Pauline did, but in terms of volume, Pauline ate more pudding than Owen did.

(21) ?? In terms of viscosity, Pauline ate more pudding than Owen did.

Finally, just like with pseudopartitives, where the monotonicity requirement extended to nouns with eventive denotations, nominal comparatives involving such nouns again retain the monotonicity requirement, as seen in (22) and (23):

(22) In terms of time, Otto did more driving yesterday than Rhonda did, but in terms of distance, Rhonda did more driving than Otto did.

(23) ?? In terms of speed, Otto did more driving yesterday than Rhonda did.

3.2.5 Verbal comparatives

We saw earlier that a monotonicity requirement in a particular nominal measurement construction (pseudopartitives) extended to a seemingly structurally parallel verbal measurement construction (adverbial measure phrases). Along similar lines, Wellwood et al. (2012) and Wellwood (2014, 2015) observe that the monotonicity constraint seen in nominal comparatives also arises in the case of verbal comparatives. Consider again (1), repeated below:

(1) Dee ran more than Evan did.

It was previously noted that (1) could serve as a comparison of time or distance of running, but not of speed, based on truth value judgments: if Dee ran for less time and less distance than Evan, but she ran faster, (1) remains false. The same restriction can be illustrated by means of the *in terms of NP* test:

(24) ?? In terms of speed, Dee ran more than Evan did.

The same story plays out with *rain*. As the *in terms of NP* test confirms, both temporal duration and amount of rain are available for *rain* comparatives, while acidity is not:

(25) In terms of amount, it rained more in London than it did in Paris, but in terms of time, it rained more in Paris than it did in London.

(26) ?? In terms of acidity, it rained more in London than it did in Paris.

Yet again, a restriction to monotonic measure functions gets the facts right here: the monotonic measure functions (time and distance in the case of *run*, time and amount in the case of *rain*) are permissible, while the non-monotonic measure functions are not.

3.2.6 Summary

In this section, we have seen that a variety of measurement constructions, including pseudopartitives, the measurement idioms *out the wazoo* and *in spades*, adverbial measure phrases, and nominal and verbal comparatives, have a requirement that the measure function used must be monotonic. In the next section, we will narrow our focus a bit and look specifically at verbal comparatives, providing a compositional semantics that enforces the monotonicity requirement illustrated above.

3.3 A compositional semantics for (verbal) comparatives

In this section, I will provide a compositional semantics for verbal comparatives, based on a version of the traditional analyses of comparatives by von Stechow (1984), Heim (1985, 2000), and Rullmann (1995), in conjunction with additional insights from Wellwood (2014, 2015). My choices at various points are meant for the most part to maximize the analysis's simplicity, as well as its familarity to readers acquainted with the "standard" analysis of comparatives. While this analysis has certain well-known faults, especially with respect to the interpretation of quantifiers, the proposal in this section will suffice for all of the cases at hand.

I will approach the compositional semantics of verbal comparatives in three steps. First, I will show how the von Stechow-Heim-Rullmann (SHR) analysis works in the case of simple adjectival comparatives. I will then demonstrate how this analysis can be easily extended to adverbial comparatives, and finally, we will approach verbal comparatives in an analogous manner.

3.3.1 Adjectival comparatives

In illustrating the SHR semantics for adjectival comparatives, I will use (27) as a toy example (where ~~strikethrough~~ indicates ellipsis):

(27) Lana is taller than Archer is ~~tall~~.

The final denotation assigned to (27) will be as in (28), which is true iff there is a degree d such that Lana's height is at least d, and d exceeds the maximal degree d' such that Archer is at least d'-tall. Naturally, this comes out as equivalent to (29), which just states that Lana's height exceeds Archer's. However, we will see that in more complex cases, such a simple conversion is not always available, and in these cases we will need to start from something like (28), rather than (29).

(28) $\exists d[\text{height}(\text{lana}) \geq d \wedge d > \max(\{d' \mid \text{height}(\text{archer}) \geq d'\})]$

(29) $\text{height}(\text{lana}) > \text{height}(\text{archer})$

So how is (28) derived from (27)? First, let's start with the denotation of *tall* in (30), a relation between a degree and an individual:

(30) $[\![\text{tall}]\!] = \lambda d \lambda x.\, \text{height}(x) \geq d$

For the instance of *tall* in the matrix clause, the degree argument is filled (in a roundabout way to be discussed soon) by the phrase *-er than Archer is ~~tall~~*, while the entity argument is filled by *Lana*. For the elided instance of *tall* in the comparative clause, the entity argument is clearly *Archer*, but what saturates the degree argument? For this, I follow Chomsky (1977) in positing the existence of a wh-element, which I will call WH, merging in the appropriate degree position and subsequently undergoing normal wh-movement to the left periphery. Thus, the comparison clause will look as in (31), with the node λd_1 lambda-abstracting over the free variable over degrees denoted by the trace t_1:

(31) 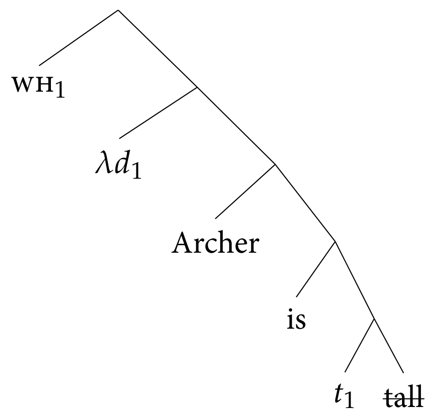

The denotation up to where WH is moved will thus be $\lambda d.\ \text{height}(\text{archer}) \geq d$, the characteristic function of the set of all degrees that do not exceed Archer's height. The denotation of WH will then return the maximal degree meeting this description, as in (32) (cf. Rullmann 1995). (This makes $[\![\text{WH}]\!]$ of type $\langle\langle d,t\rangle,d\rangle$.)

(32) $[\![\text{WH}]\!] = \lambda D.\ \max(\{d \mid D(d)\})$

Notice that assigning this sort of denotation to WH provides two possible analyses for why WH moves to the left periphery. The first is simply syntactic: WH is a wh-phrase, and in English wh-phrases move to spec-CP unless the latter is otherwise occupied. The second is semantic: there is a type mismatch between WH, which is of type $\langle\langle d,t\rangle,d\rangle$, and *tall*, which is of type $\langle d,\langle e,t\rangle\rangle$, so WH must move out, with its trace saturating the degree argument of *tall*.

Treating *than* as semantically vacuous, this means that *than* WH *Archer is ~~tall~~* will have the meaning in (33a), which (conveniently) is equivalent to (33b):

(33) a. $\max(\{d \mid \text{height}(\text{archer}) \geq d\})$
b. $\text{height}(\text{archer})$

What about the rest of the sentence? Let's say that (27) as a whole has the (oversimplified) syntactic structure in (34).

(34)

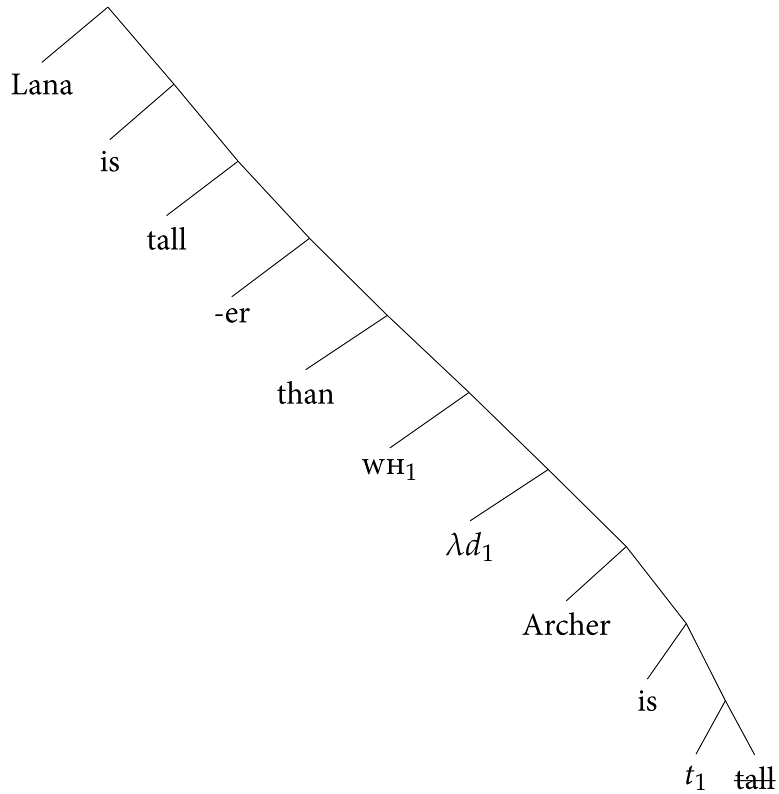

Our next step is thus to assign a denotation to the comparative morpheme *-er*. Our definition will treat *-er* as relating a degree d and (the characteristic function of) a set D of degrees, and returning true iff there is a degree d' such that $D(d')$ and $d' > d$. This can be seen in (35a), with the result of combining this with *than* WH *Archer is* ~~*tall*~~ visible in (35b):

(35) a. $[\![\text{-er}]\!] = \lambda d \lambda D.\ \exists d'[D(d') \wedge d' > d]$

b. [[-er than WH Archer is ~~tall~~]] =

$$\lambda D.\ \exists d'[D(d') \wedge d' > \max(\{d \mid \text{height}(\text{archer}) \geq d\})]$$

The resulting denotation here is akin to a quantificational DP, but quantifying over degrees instead of entities: it is of type $\langle\langle d,t\rangle,t\rangle$, rather than $\langle\langle e,t\rangle,t\rangle$. And in the same way that (operating under certain assumptions) quantificational DPs often have to undergo QR due to mismatches in type, this degree quantifier must undergo QR due to a mismatch in type between it and *tall*. The resulting syntactic representation after this iteration of QR is as in (36):

(36)

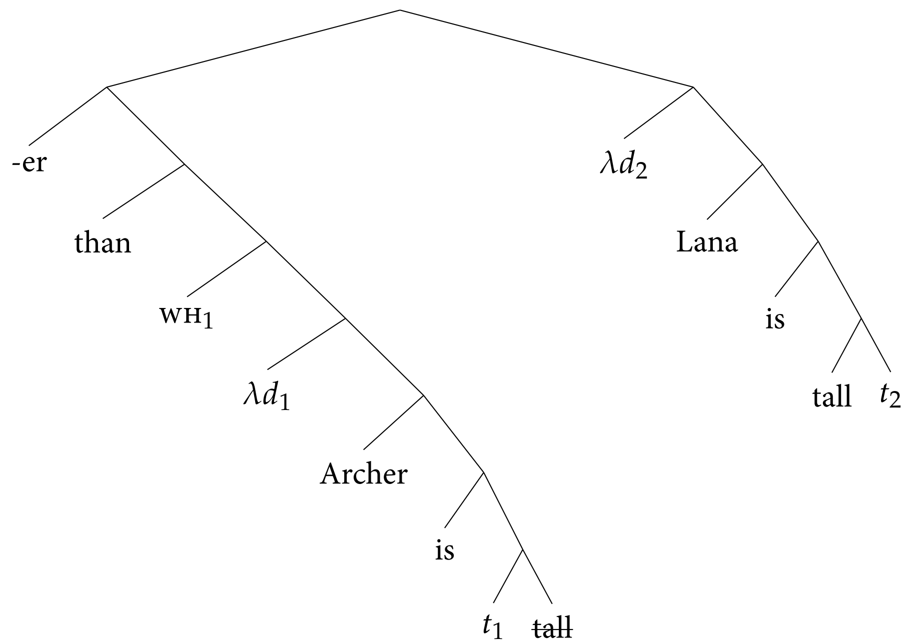

From here, the rest of the denotation falls out naturally. The free variable denoted by t_2 saturates *tall*'s degree argument, *Lana* takes care of the entity argument, and λd_2 abstracts over the free degree variable, giving us $\lambda d.$ height(lana) $\geq d$. This readily composes with (35b), giving us the desired (28).

This, in a nutshell, is the SHR analysis of comparatives. Before moving on to adverbial comparatives, a couple of things are worth noting. First, in simplifying the analysis, I have left out any means of handling differential comparatives like (37a), which in the SHR analysis has an interpretation along the lines of (37b):[10]

(37) a. Lana is two centimeters taller than Archer is ~~tall~~.

b. $\exists d[\text{height}(\text{lana}) \geq d \wedge d = 2\text{cm} + \max(\{d' \mid \text{height}(\text{archer}) \geq d'\})]$

Since differential comparatives are also possible in non-adjectival comparatives—see (38a) and (38b) for differential adverbial and verbal comparatives, respectively—any

[10]Note that according to (37b), (37a) is true iff Lana is *at least* two centimeters taller than Archer is. Any inference that Lana is *exactly* two centimeters taller than Archer is is then predicted to be due to scalar implicature. Based on the default interpretation of (i), this prediction is desirable:

(i) Any spy who is two centimeters taller than Archer (is) can go on this ride.

If the *exactly*-inference is due to the semantics of differential comparatives, then (i) says only that those spies who are *exactly* two centimeters taller than Archer can go on the ride. However, the more salient reading is the (stronger) *at least* reading: any spy who is *at least* two centimeters taller than Archer can go on the ride.

account of differential comparatives will also need to be transported to these cases as well.

(38) a. Dee ran two miles per hour faster than Evan did.
b. Dee ran two miles more than Evan did.

Second, the SHR analysis has well-known problems with quantifiers interpreted within the comparison clause, as in (39):

(39) Lana is taller than every male spy is.

According to the analysis currently on the table, (39) has the interpretation in (40):

(40) $\exists d[\text{height}(\text{lana}) \geq d \wedge d > \max(\{d' \mid \forall x : \text{male-spy}(x)[\text{height}(x) \geq d']\})]$

But these truth conditions are far too weak. According to (40), (39) is true iff Lana's height exceeds the maximal degree d such that every male spy's height is at or exceeds d. This maximal degree is the height of the shortest spy, so (40) is true iff Lana is taller than the shortest male spy. But this, of course, is not the meaning of (39), which requires that Lana be taller than the *tallest* male spy.

In this dissertation, I will not be covering the intricate relationship between quantifiers and comparatives, so for the most part such issues need not concern us here.[11] However, this issue is worth bearing in mind because introducing events means adding a quantificational element to both the matrix and comparison clauses: namely, the existential event-quantifier. Thus, when we next turn our attention to adverbial comparatives, we will have to make sure the inclusion of this event-quantifier doesn't lead to odd results.

3.3.2 Adverbial comparatives

Next up are adverbial comparatives, which we will see are semantically composed in a manner quite similar to adjectival comparatives. For our discussion of adverbial comparatives, I will use the sentence in (41):

(41) Dee climbed Mt. Fuji faster than Evan did ~~climb Mt. Fuji fast~~.

3.3.2.1 Event compositionality

In order to analyze (41), we first need to develop an understanding of how events and event-quantification fit into the semantics of verbal projections and the clausal spine. I will therefore first discuss the syntax and semantics of the simpler sentence

[11] For detailed discussion on the semantics of quantifiers in comparison clauses, see, e.g., von Stechow 1984; Larson 1988; Schwarzschild & Wilkinson 2002; Schwarzschild 2004, 2008; Heim 2006; Gajewski 2008; Krasikova 2008b; van Rooij 2008; Beck 2010, 2011, 2014.

Dee climbed Mt. Fuji. Syntactically, I take this sentence to have the structure in (42) (ignoring head movement, which is irrelevant for our purposes):

(42)

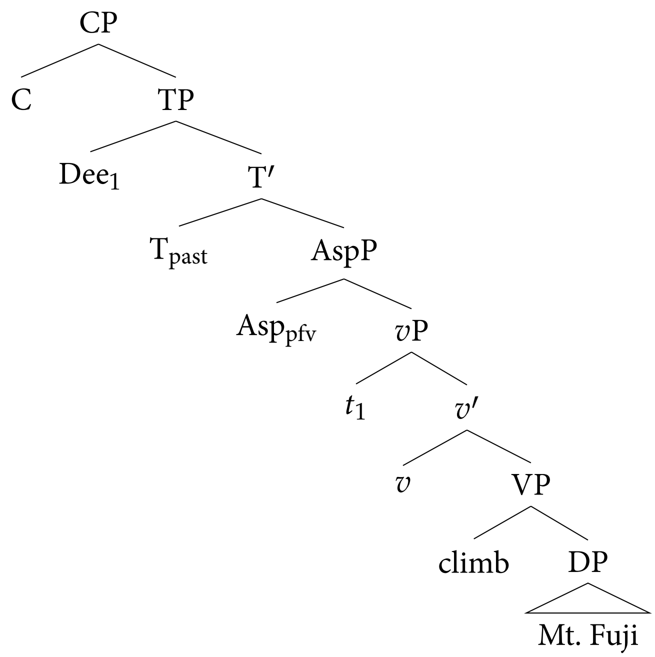

The verb *climb* merges with its internal argument (*Mt. Fuji*), creating the VP *climb Mt. Fuji*. Next, following Kratzer (1996), the voice head v merges with VP and takes in its specifier the external argument (*Dee*), while simultaneously assigning accusative case to *Mt. Fuji*. Next comes the perfective aspectual head $\mathrm{Asp}_{\mathrm{pfv}}$ and the past tense head $\mathrm{T}_{\mathrm{past}}$, with the external argument undergoing EPP movement to the subject position in spec-TP. Finally, C takes TP as its complement, resulting in a full CP.

So how does this sentence compose semantically? The first step in figuring this out is to decide upon an interpretation for *climb*. Since v is responsible for introducing the external argument both syntactically and semantically, the external argument will be absent from the definition of ⟦climb⟧. But while such syntactic and semantic separation of the *external* argument has been more or less agreed upon, how a verb combines with its *internal* arguments is less settled. One possibility, adopted by Kratzer (1996, 2003), is that unlike the external argument, internal arguments of the verb compose straightforwardly, so that *climb* has a meaning like (43):

(43) $\lambda x \lambda e.\, \mathrm{climb}(e, x)$

On the other end of the spectrum, one could instead posit that *all* arguments of the verb, and not just the external argument, are introduced by separate thematic heads, so that *climb* has the simpler interpretation in (44) (see, e.g., Schein 1993, Champollion 2015a):

(44) $\lambda e.\ \text{climb}(e)$

The result of combining *climb* with the lower thematic head and direct object will then have an interpretation like (45) (where $\text{Thm}(e)$ is the theme of e):

(45) $\lambda e.\ \text{climb}(e) \wedge \text{Thm}(e) = \text{mt-fuji}$

A third possibility is that the verb combines directly with its direct object, but predicate decomposition is still present in $[\![\text{climb}]\!]$, as in (46):

(46) $\lambda x \lambda e.\ \text{climb}(e) \wedge \text{Thm}(e) = x$

Given that in (42) *Mt. Fuji* is the complement of *climb*, with no intervening thematic head, I am thereby committing myself to either (43) or (46). Since these two options are identical as far as the compositional semantics is concerned, there is no relevant difference between them; I will opt for (46) for ease of reading.

After *climb* composes with *Mt. Fuji*, v introduces the external argument into the mix. This can be accomplished by defining $[\![v]\!]$ as in (47a). According to this definition, $[\![v]\!]$ takes an event predicate V (denoted by the VP) and an entity x (denoted by the external argument) and returns an event predicate that is true of an event iff it is a V-event with x as its agent.[12] The result of combining this with the VP above can be seen in (47b).

(47) a. $[\![v]\!] = \lambda V \lambda x \lambda e.\ \text{Agt}(e) = x \wedge V(e)$
b. $\lambda x \lambda e.\ \text{Agt}(e) = x \wedge \text{climb}(e) \wedge \text{Thm}(e) = \text{mt-fuji}$

In our syntactic tree in (42), the occupant of spec-vP is not *Dee*, but its trace. This is important in the case of quantificational subjects, since there would otherwise be a type mismatch between a quantificational DP like *every boy* (type $\langle\langle e,t\rangle,t\rangle$) and v' (type $\langle e,\langle v,t\rangle\rangle$). However, if the subject is referential, then interpreting the trace as a variable that is subsequently lambda-abstracted over and saturated by the subject generates the same result as if the subject were simply interpreted in its initial merge position. With this in mind, in cases where the subject is referential I will adopt the simplifying assumption that it is interpreted in its merge position. The denotation of vP will therefore be (48):

(48) $[\![\text{Dee climb Mt. Fuji}]\!] = \lambda e.\ \text{Agt}(e) = \text{dee} \wedge \text{climb}(e) \wedge \text{Thm}(e) = \text{mt-fuji}$

[12] This defintion of $[\![v]\!]$ is distinct from Kratzer's, which is the simpler $\lambda x \lambda e.\ \text{Agt}(e) = x$. The trade-off for Kratzer is that she is forced to introduce a new rule of semantic composition, since otherwise there would be a type mismatch between the VP (type $\langle v,t\rangle$) and v (type $\langle e,\langle v,t\rangle\rangle$). This new rule is what she calls *Event Identification*, and goes as follows: If $[\![\alpha]\!]$ is of type $\langle e,\langle v,t\rangle\rangle$, and $[\![\beta]\!]$ is of type $\langle v,t\rangle$, then the result of composing α and β is $\lambda x \lambda e.\ [\![\beta]\!](e) \wedge [\![\alpha]\!](x)(e)$. The ensuing denotation of vP is the same as my own.

Since I will not be discussing tense or aspect in great deal in this dissertation, I will for the most part be ignoring the semantic contributions of these heads, with one exception: something needs to introduce existential quantification over the event predicate built up by the vP. I will follow Kratzer (1998) and Hacquard (2006) in locating this existential quantification in the denotations of aspectual heads. Since I am otherwise ignoring the semantics of aspect, I will simply treat the denotations of aspectual heads as existential event-quantifiers:

(49) $[\![\text{Asp}]\!] = \lambda V.\, \exists e[V(e)]$

As a result, the interpretation of AspP, and the sentence as a whole, is as in (50):

(50) $\exists e[\text{Agt}(e) = \text{dee} \wedge \text{climb}(e) \wedge \text{Thm}(e) = \text{mt-fuji}]$

3.3.2.2 Back to comparatives

Now that we know how a simple transitive sentence composes after an event variable is introduced, we can successfully analyze (41). First, the syntax. I take (41) to have the syntactic structure in (51), minus any QR; the tree is divided into two subtrees for readability. (It is worth noting that while I take *faster than Evan did ~~climb Mt. Fuji fast~~* to be a vP-adjunct, as far as the semantics is concerned there seems to be no difference between its being a vP- or VP-adjunct.)

(51) a.

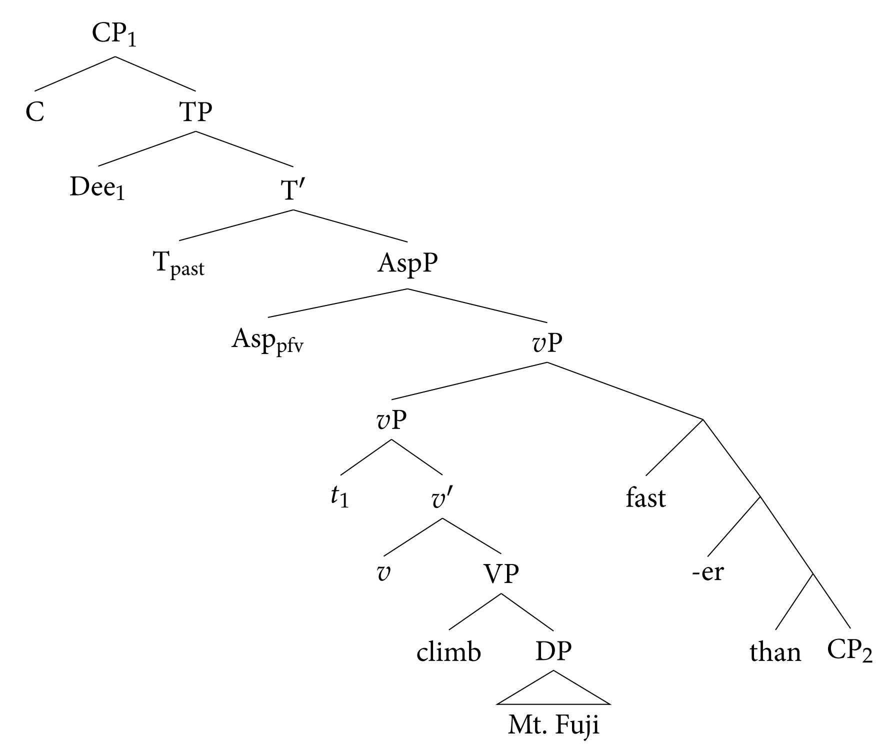

b.

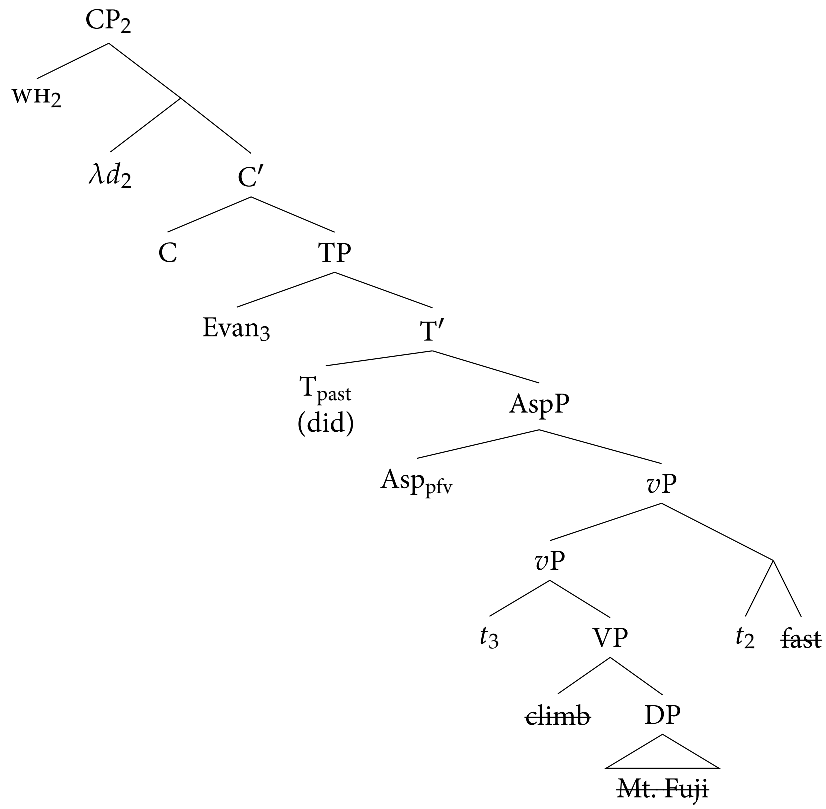

Semantically, combining the SHR theory of comparatives with the compositional event semantics discussed above leads to a relatively straightforward analysis of adverbial comparatives. Looking at the comparison clause (51b), the denotation of the pre-adjunction vP is $\lambda e.\ \mathrm{Agt}(e) = \mathrm{evan} \wedge \mathrm{climb}(e) \wedge \mathrm{Thm}(e) = \text{mt-fuji}$ (again assuming the subject is interpreted in its merge position). The gradable adverb *fast* will have the interpretation $\lambda d \lambda e.\ \mathrm{speed}(e) \geq d$, parallel to our previous definition of *tall*. The trace of the wh-element WH saturates the degree argument with the free variable d_2. Since both the pre-adjunction vP and *fast*+t_2 denote predicates of events, they are composed by means of predicate modification (Heim & Kratzer 1998), which conjoins these predicates, resulting in $\lambda e.\ \mathrm{Agt}(e) = \mathrm{evan} \wedge \mathrm{climb}(e) \wedge \mathrm{Thm}(e) = \text{mt-fuji} \wedge \mathrm{speed}(e) \geq d_2$. Next comes existential closure of the event predicate by Asp$_{pfv}$, followed by lambda abstraction over d_2 and maximalization by WH. Thus, the final denotation of (51b) is as in (52):

(52) $$\max(\{d \mid \exists e[\mathrm{Agt}(e) = \mathrm{evan} \wedge \mathrm{climb}(e) \wedge \mathrm{Thm}(e) = \text{mt-fuji} \wedge \mathrm{speed}(e) \geq d]\})$$

Notice that as discussed previously, in (52) there is an existential quantifier over events within the scope of the maximalizing wh-element WH. Given the troubles the

SHR analysis has with (certain) quantifiers, once we have a full denotation for (41) we will have to make sure that this existential quantifier does not do any harm.

When combined with the comparative morpheme *-er*, which has the same semantics as before, the result is again of type $\langle\langle d, t\rangle, t\rangle$, a quantifier over degrees. A type mismatch between this degree quantifier and *fast* means that the latter undergoes QR, leading to a structure like (53):

(53)

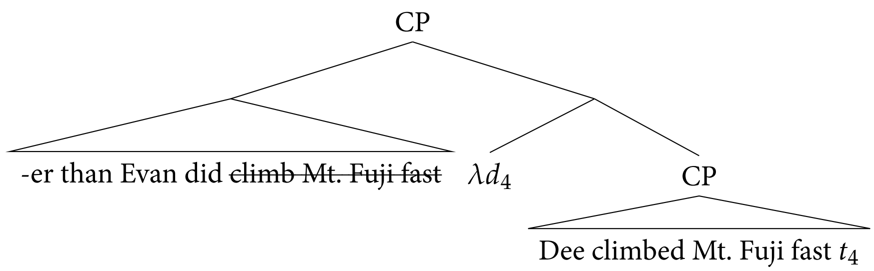

From here on out there are no big surpises. *fast* combines with the trace t_4, with the result undergoing predicate modification with the vP to which it is adjoined. The resulting event predicate is existentially closed by Asp_{pfv}, and then the free degree variable is lambda-abstracted, with the resulting degree predicate quantified over by the degree quantifier. The final denotation for (41) thus comes out to (54):

(54) $\exists d \exists e[\text{Agt}(e) = \text{dee} \wedge \text{climb}(e) \wedge \text{Thm}(e) = \text{mt-fuji} \wedge \text{speed}(e) \geq d \wedge$
$d > \max(\{d' \mid \exists e'[\text{Agt}(e') = \text{evan} \wedge \text{climb}(e') \wedge \text{Thm}(e') = \text{mt-fuji} \wedge$
$\text{speed}(e') \geq d']\})]$

In plain English, (54) is true iff there is an event of Dee climbing Mt. Fuji whose speed exceeds the maximal degree of speed obtained by any event of Evan climbing Mt. Fuji. Notice that unlike the problematic case of universal quantification discussed above, the existential event quantification in the comparison clause generates the right result: Dee's climb must be faster than *all* of Evan's climbs, not just the slowest one.

3.3.3 Verbal comparatives

We can now take the final step from adverbial to verbal comparatives. One important aspect of verbal comparatives that any theory thereof must address is how degrees enter into the semantic derivation. In the case of adjectival and adverbial comparatives, this question receives a pedestrian answer: gradable adjectives and adverbs simply have their own degree arguments. But (many of) the verbs that appear in verbal comparatives do not show any morphosyntactic or semantic signs of being lexically gradable. In fact, the evidence points the opposite way, with perhaps the clearest argument for this coming from the monotonicity requirement itself. If verbal comparatives owe their existence to the lexical gradability of verbs, then presumably the possible dimensions of measurement for each verb are similarly lexically determined. In other words, the fact that *run* comparatives like (1) (repeated below) can compare distance and time, but not speed, should be due to the semantics of *run* and its constraints on its degree argument.

(1) Dee ran more than Evan did.

But this reduces the monotonicity requirement to a mere quirk of the lexicon: verbal comparatives obey the monotonicity requirement only because each verb individually obeys the monotonicity requirement. While this is certainly possible, it is at best exceedingly unlikely. A much more likely explanation is that (many) verbs are *not* lexically gradable, and that whatever head introduces the degree argument needed for verbal comparatives simultaneously imposes a monotonicity requirement. This way, the monotonicity requirement stems from the semantics of a single head, rather than the cumulative semantic representations of all of the verbs in the lexicon.

While my analysis of verbal comparatives is in certain respects heavily influenced by that of Wellwood (2014, 2015), a wedge can be driven between our proposals with respect to where and how degrees are introduced. Doing Wellwood the injustice of a brief summary, her main ideas can be paraphrased as consisting of a few core proposals. First, morphosyntactically, it is always the case that *more* = *much* + *-er*. Second, *much* always appears in comparatives, so that (27) comes out to something like *Lana is -er much tall than Archer is*; when *much* does not appear overtly, it is because some morphosyntactic process deletes it at PF (cf. Bresnan 1973). Third, *much* is always what introduces degrees—even for gradable adjectives like *tall*—and universally imposes a monotonicity requirement. That is, much like how $[\![\text{run}]\!] = \lambda e.\, \text{run}(e)$, for Wellwood, $[\![\text{tall}]\!]$ is just $\lambda e.\, \text{tall}(e)$, with *much* introducing degrees of height—a monotonic measure of states of tallness.

Several objections can be raised against such an account. One clear potential objection is to the expansive ontological commitments that it engenders. After all, for Wellwood, not only must each object have its own neo-Davidsonian state of hardness, hotness, etc., but since each of these states must be monotonic in the appropriate dimension, an object's state of hardness must itself be made up of many smaller states of hardness. Moreover, it seems that one must not only have a state of *tallness*, but a state of *shortness* as well, and these states must stand in a particular metaphysical relationship, since Lana's being taller than Archer entails Archer's being shorter than Lana. But while this may be unpleasant to some, this large ontology may simply be the pill one has to swallow in order get the semantics right. Nothing says that the folk metaphysics underlying natural language interpretation has to be as elegant—or even accurate—as that devised by philosophers. Perhaps, then, ontological parsimony is simply not a relevant desideratum.

More worryingly, Wellwood's analysis seems to predict an ambiguity in simple adjectival comparatives like (55) that does not actually arise.

(55) The mad scientist was taller than her sister was.

Recall that when we were looking at verbal comparatives, temporal duration was generally available as a means of comparison due to its being a monotonic measure of many types of events. But a five-hour state of the mad scientist's being six feet tall, for example, can also be chopped up along the temporal dimension in the same way, meaning that both height and temporal duration are monotonic measures of Wellwoodian tallness states. Since the only constraint that *much* imposes is monotonicity, this means that (55) should have a (non-existent) reading comparing how long the two people's states of tallness lasted.

To this one might respond that this supposedly non-existent reading would be highly odd anyway: if one has a tallness state for as long as one exists (since one always has a height), then the relevant reading of (55) would reduce to the claim that the mad scientist existed for longer than her sister. It therefore might be that this reading is not barred by the grammar, but by any of a broad range of pragmatic principles. But this retort falls flat for cases like (56):

(56) ?? My dog was more pregnant than my cat was.

This sentence is odd because gradable readings of *pregnant* are difficult to get, a fact that most theories would attribute to *pregnant* not being lexically gradable, requiring coercion for a gradable interpretation. But notice that this time, the absent temporal comparison is completely reasonable: my dog's state of pregnancy lasted longer than my cat's did. The fact that such a reading is unavailable, even in the face of an alternative that is clearly odd, suggests that it is indeed the grammar that blocks it.

The presumed next step in countering this argument would then be to say that either (i) temporal duration must not be a monotonic measure of tallness states, or (ii) both height and duration are monotonic, and something else prevents states described by adjectives from being measured temporally. But if one keeps the adjectives and changes the type of comparative from adjectival to verbal, suddenly temporal measurement pops up:

(57) a. (While the height-changing machine was on the fritz,) the mad scientist was tall more than her sister was.
b. (Over the course of the last three years,) my dog was pregnant more than my cat was.

What's more, neither of these sentences has an interpretation identical to the corresponding adjectival comparative: each either compares duration/frequency, or serves as a so-called "metalinguistic comparative" comparing the relative appropriateness of calling certain individuals "tall" or "pregnant".[13] It is difficult to see how this divergence in the availability of temporal interpretations and "plain" degree interpertations can be accounted for within Wellwood's theory, since for her both boil down to monotonic measurement of states.

This conflict can be resolved by adopting an alternative hypothesis: gradable adjectives have *both* a degree argument *and* an eventuality argument, essentially leading to a Davidsonian version of SHR. Assuming continued separation of the external argument, this means that $[\![\text{tall}]\!] = \lambda d \lambda e.\ \text{tall}(e) \geq d$. For adjectival comparatives the adjectives' degree arguments are used to generate the compared degrees, essentially in the manner described above. Since the scales for degree arguments are lexically determined, only the dimension fixed by the denotation of the adjective will be an available means of comparison, leading to the non-ambiguity observed for (55).

Meanwhile, for comparatives like (57a) and (57b), the degree argument is saturated by the covert positive morpheme POS, which in general provides the contextually determined degree of comparison for positive adjectival attributions like *Lana is tall* (cf. Cresswell 1976 and many since). The semantic result of this combination is (the characteristic function of) a set of eventualities, the same semantic type as a VP. We can then treat the temporal comparison readings of (57a) and (57b) as simple verbal comparatives, requiring monotonic measurements of states of being pregnant or having at least such-and-such a height. Since, as discussed above, temporal measurement is presumably a monotonic measurement of these states, the right result is obtained.

On the morphosyntactic end of things, rather than saying that *much* is omnipresent in comparison constructions and is deleted under the right circumstances, we can instead say that *much* is *inserted* when there is nothing that the comparative morpheme *-er* can attach to (cf. Bhatt & Pancheva 2004, Embick 2007). Thus, in the case of *smart* + *-er*, there is no need to insert *much*, and we get *smarter*, while in the case of *intelligent*, the fact that *-er* cannot attach directly to *intelligent* means that our only option is *much* + *-er* (= *more*) *intelligent*. For both nominal and verbal comparatives, since there is no comparative form for nouns or verbs (in English), *much*-insertion appears across the board.

But now we are left with the same question we started with: verbs don't have degree arguments, and the semantics of *-er* trades in degrees. So how are degrees introduced into the semantic computation? For this I will posit the covert preposition FOR—analogous to the *for* in durative adverbials like *run for an hour*—which will be responsible for introducing degrees and imposing the monotonicity requirement.

[13] For discussion of the syntax of metalinguistic comparatives—including why the synthetic comparative form (e.g., *more smart* instead of *smarter*) is required for metalinguistic comparatives—see Embick 2007. For semantic discussion, see Giannakidou & Stavrou 2009, Morzycki 2011, Giannakidou & Yoon 2011, Wellwood 2014.

The (pre-QR) syntactic structure of (1) will thus look like (58), again broken up into two trees for readability:

(58) a.

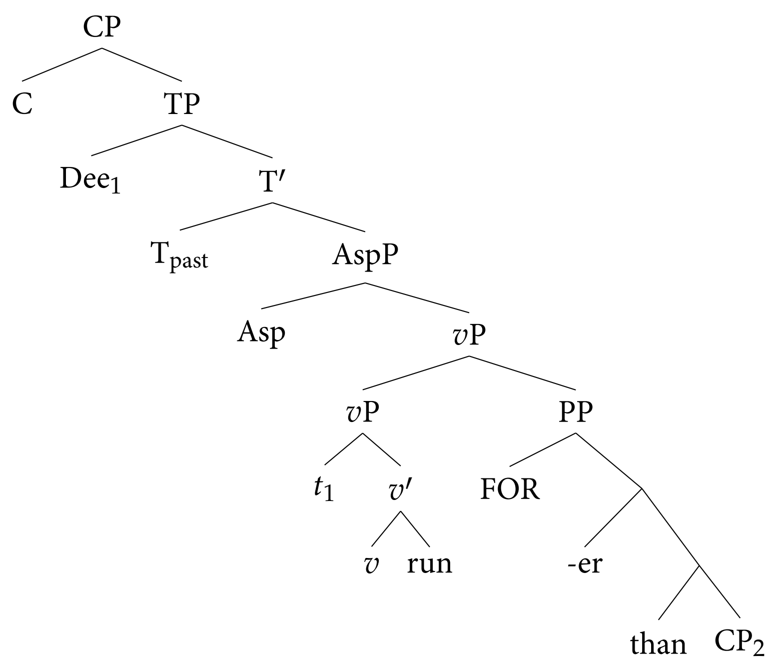

b.

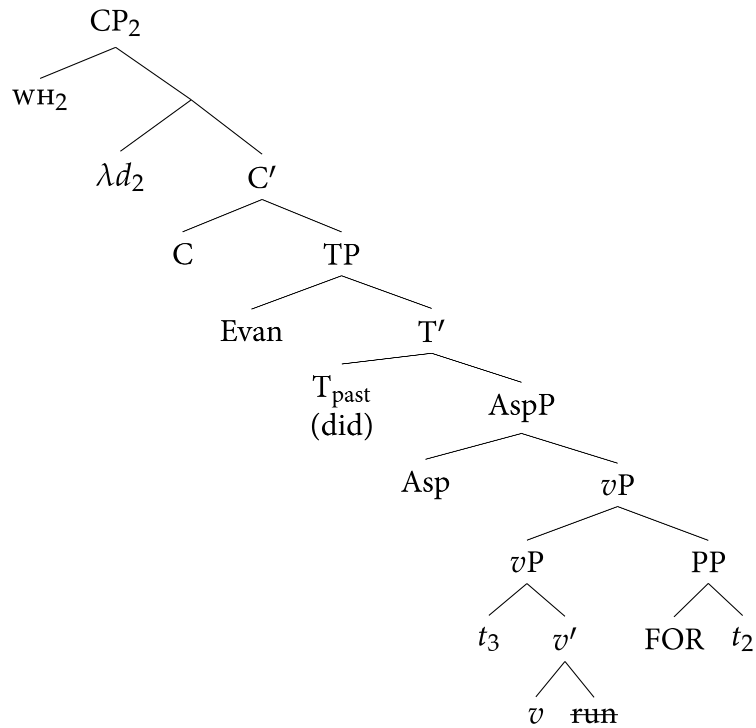

As per the proposed morphosyntax, since there is nothing for *-er* to attach to, *much* is inserted to create *more*, giving us the surface string (1).

On the semantic end of things, FOR takes a degree d and event predicate V, and returns an event predicate true of an event iff V holds of it, and its measurement by the contextually-determined measure function μ^c exceeds d. FOR also imposes the presupposition that μ^c is monotonic on salient part-whole relations in V.

(59) $[\![\text{FOR}]\!]^c = \lambda d \lambda V \lambda e : \mu^c$ is monotonic on $\sqsubseteq^c$ in V. $V(e) \wedge \mu^c(e) \geq d$

Let's see how this all pans out compositionally. Starting in the comparison clause, the first argument of FOR is saturated by the free variable over degrees denoted by t_2. The ensuing combination then takes the denotation of the vP as an argument. Putting aside the monotonicity presupposition, the result is (60):

(60) $\lambda e.\ \text{Agt}(e) = \text{evan} \wedge \text{run}(e) \wedge \mu^c(e) \geq d_2$

The monotonicity requirement is checked relative to the vP event predicate, $\lambda e.\ \text{Agt}(e) = \text{evan} \wedge \text{run}(e)$. On a reading comparing distance, this monotonicity requirement is satisfied, since μ_{distance} is monotonic on $\sqsubseteq^c$ in this event domain.

Everything then continues as normal, with lambda abstraction over d_2 and maximalization by WH. The resulting maximal degree then serves as the first argument of *-er*. Once again, *-er than* WH *Evan did ~~run~~ FOR* undergoes QR, and from there there are no surprises. The end result of this computation is as in (61):

(61) $\exists d \exists e[\text{Agt}(e) = \text{dee} \wedge \text{run}(e) \wedge \mu^c(e) \geq d \wedge d >$
$\max(\{d' \mid \exists e'[\text{Agt}(e') = \text{evan} \wedge \text{run}(e') \wedge \mu^c(e') \geq d']\})]$

As desired, (61) is true iff there is an event of Dee running that exceeds by the contextually determined measure function μ^c any event of Evan running, with the two iterations of FOR imposing monotonicity requirements with respect to the event domains denoted by the two vPs.

As a final note, this account can be readily extended to adverbial measure phrases like *a lot* in *Dee ran a lot*. Assuming *a lot* denotes a contextually determined (high) degree, the simple inclusion of FOR as in (62) generates the semantic result in (63), so long as the monotonicity presupposition of FOR is satisfied.

(62)

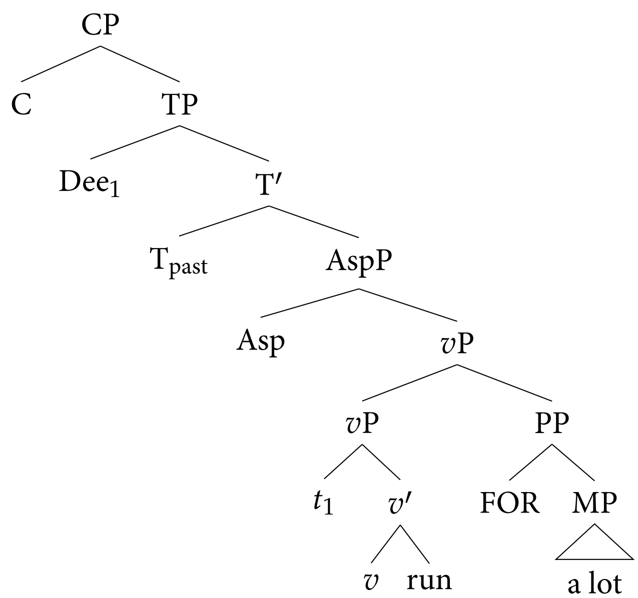

(63) $[\![(62)]\!]^c = 1$ iff $\exists e[\text{Agt}(e) = \text{dee} \wedge \text{run}(e) \wedge \mu^c(e) \geq [\![\text{a lot}]\!]^c]$

Extending this analysis to non-upward-entailing measure phrase adjuncts like *very little* requires a bit more work. A plausible analysis of $[\![\text{very little}]\!]$ is as in (64), where it denotes a downward-entailing quantifier over degrees, rather than referring to an individual degree. (d^c_{low} is a contextually determined low degree.)

(64) $[\![\text{very little}]\!]^c = \lambda D.\ \max(\{d \mid D(d)\}) < d^c_{\text{low}}$

Due to a type mismatch with FOR, *very little* must undergo QR:

(65) $[[\text{very little}]_1\ [\lambda d_1\ [\text{Dee ran FOR } t_1]]]$

And as a result, *Dee ran very little* gets the plausible interpretation in (66), again excluding the monotonicity presupposition:

(66) $[\![\text{Dee ran very little}]\!]^c = 1$ iff

$$\max(\{d \mid \exists e[\text{Agt}(e) = \text{dee} \wedge \text{run}(e) \wedge \mu^c(e) \geq d]\}) < d^c_{\text{low}}$$

It may be that *a lot*, *a great deal*, etc., similarly denote degree quantifiers instead of a contextually-determined high degree. That is, $[\![\text{a lot}]\!]$ may be as in (67), which is essentially (64), but flipping < to > and replacing d^c_{low} with d^c_{high}.

(67) $[\![\text{a lot}]\!]_{\text{quant}} = \lambda D.\ \max(\{d \mid D(d)\}) > d^c_{\text{high}}$

However, for simplicity's sake I will continue with the assumption that they just denote contextually-determined high degrees.

3.4 Back to intensity comparatives

Assume for the time being that intensity comparatives like (2–3), repeated below, compose just like other verbal comparatives.

(2)
a. Fiona <u>likes</u> football more than she does baseball.
b. Gavin <u>fears</u> clowns less than he does sharks.
c. Helen <u>hates</u> country music as much as she does rap.
d. Ina <u>respects</u> her teachers more than she does her friends.
e. Jorge <u>admires</u> the CEO less than he does his co-workers.
f. Kwame <u>trusts</u> the poor as much as he does the rich.
g. Marvin <u>loves</u> biology more than he does history.

(3)
a. Jo <u>wants</u> to leave more than Ben <u>wants</u> to stay.
b. Stan <u>wished</u> he'd won more than he <u>wished</u> he'd stayed healthy.
c. Paul <u>regrets</u> buying his car more than Nora <u>regrets</u> selling hers.

As discussed above, something in the structure of verbal comparatives imposes a monotonicity requirement on the measure function used. Therefore, if intensity comparatives have the same structure as other verbal comparatives, this provides support for a natural language ontology in which intensity is a monotonic measure of mental states in general, and attitude states in particular. I happen to have attributed this monotonicity requirement to the presence of the covert preposition FOR, but note that my overall line of argumentation is fully independent of my specific analysis of verbal comparatives. That is, the monotonicity requirement for (non-intensity) verbal comparatives is an empirical observation, not a theory-internal prediction. So even if my proposed structure and semantics of verbal comparatives is not the right one, the facts about monotonicity and verbal comparatives remain.

That being said, let's see what our particular compositional semantics for verbal comparatives would mean for *want* comparatives. For now I will avoid discussion of the lexical semantics of *want*; the important thing is that due to the nature of the compositional semantics, $[\![\text{want}]\!]$ will be a relation between a proposition (denoted by the clausal complement) and an event argument:

(68) $[\![\text{want}]\!] \approx \lambda p \lambda e.\ e$ is a state of wanting p

I will use (3a) as a sample *want* comparative. First, I will show how the simple non-comparative sentence *Jo wants to leave* is composed, and then we will make the leap to (3a). Following along with our prior assumptions on syntactic structure and event compositionality, I take *Jo wants to leave* to have the syntax in (69), where v_{exp} is the experiencer v head, parallel to the agentive v head discussed previously.

(69)

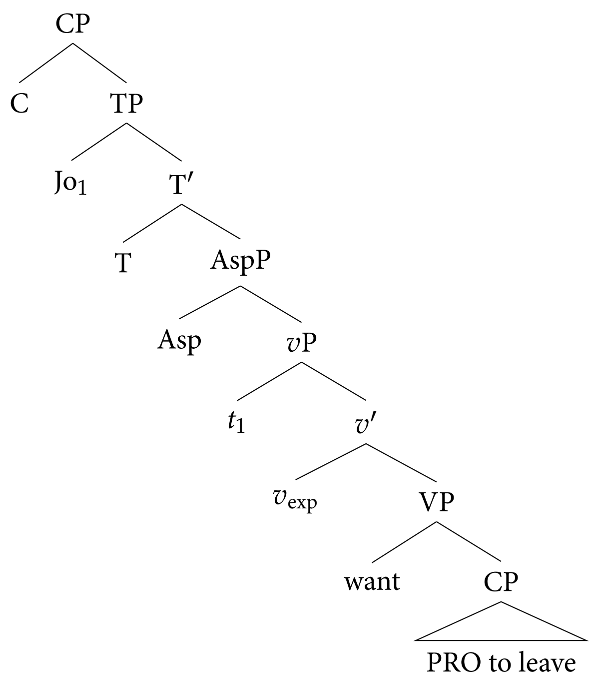

For our purposes we can avoid a lengthy discussion about the semantics of control, and simply assume that the denotation of the embedded CP is the set of possible worlds in which Jo leaves. With this established, the sentence composes in a relatively pedestrian manner. First, *want* composes with its CP argument, resulting in (70):

(70) $[\![\text{VP}]\!] = \lambda e.\ [\![\text{want}]\!](\text{jo-leave})(e)$

Next up is v_{exp}, which I take to have the denotation in (71), parallel to the definition of v in (47a). The result of combining this with (70) is in (72).

(71) $[\![v_{\text{exp}}]\!] = \lambda V \lambda x \lambda e.\ \text{Exp}(e) = x \wedge V(e)$

(72) $[\![v']\!] = \lambda x \lambda e.\ \text{Exp}(e) = x \wedge [\![\text{want}]\!](\text{jo-leave})(e)$

To round the derivation out, *Jo* (or rather, its trace) saturates v_{exp}'s entity argument, with Asp existentially closing off the ensuing event predicate. The result is (73):

(73) $[\![\text{Jo wants to leave}]\!] = 1$ iff $\exists e[\text{Exp}(e) = \text{jo} \wedge [\![\text{want}]\!](\text{jo-leave})(e)]$

Let us move on now to the *want* comparative (3a). Following our previously established narrative for the syntax of verbal comparatives, I take (3a) to have the (pre-QR) syntactic representation in (74). Notice that there is nothing special here about the syntax of (3a) in contrast to other verbal comparatives, and that the covert preposition FOR introduces degree arguments and imposes a monotonicity requirement.

(74) a.

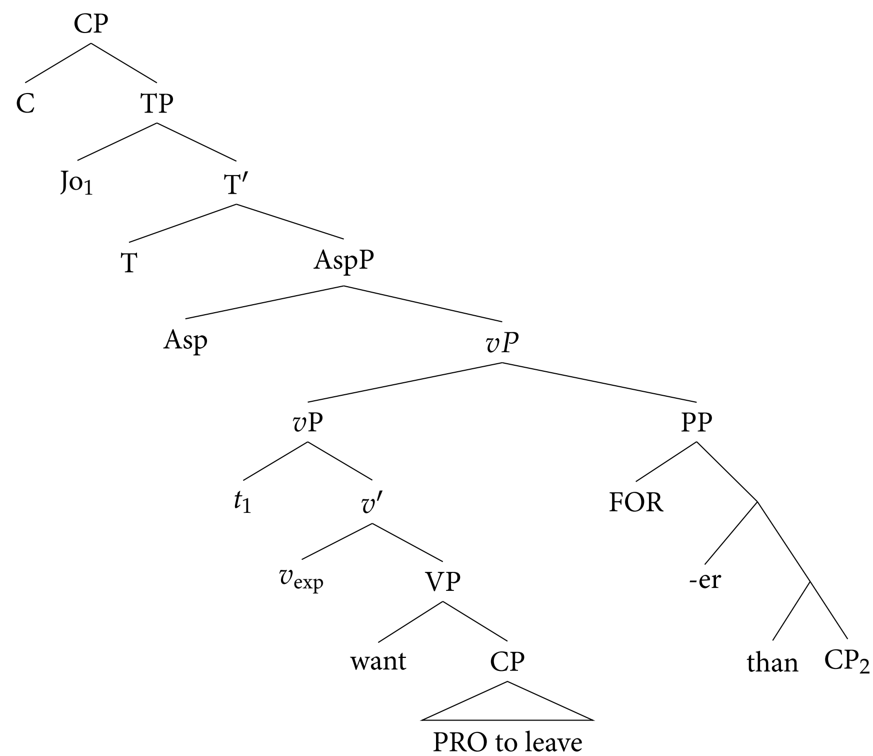

b.

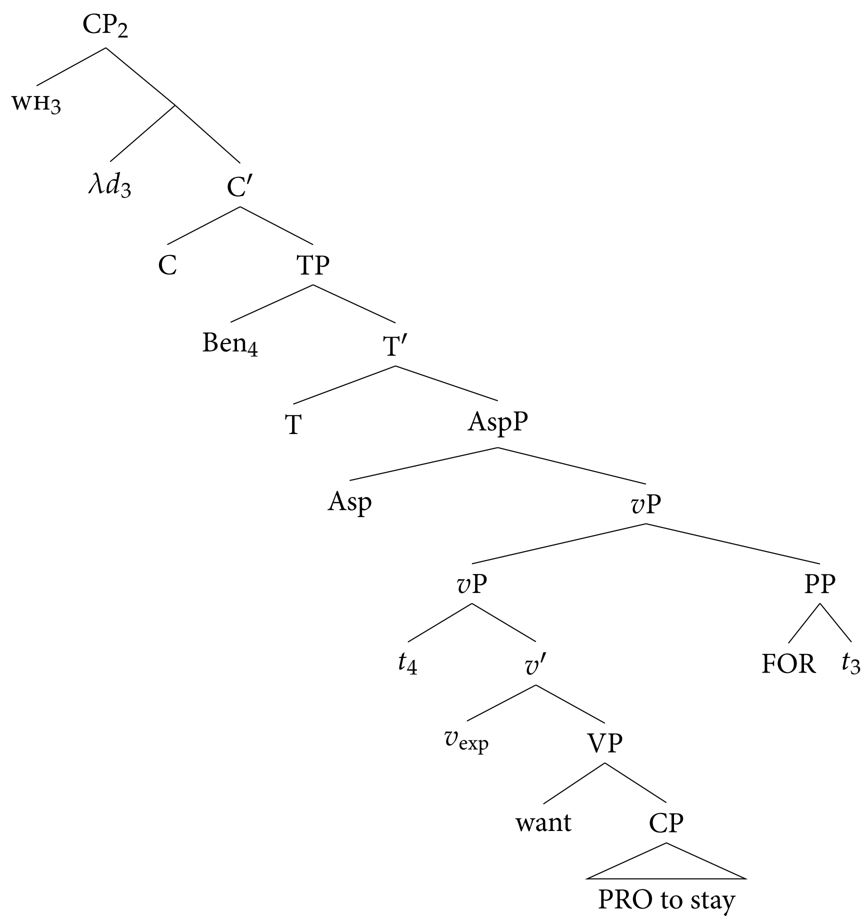

Much like the syntax, the semantics of (3a) brings nothing of substantive novelty with it: everything composes for (3a) just like it did for (1). The end result, minus the monotonicity presuppositions, is as in (75).

(75) $\exists d \exists e[\text{Exp}(e) = \text{jo} \wedge [\![\text{want}]\!](\text{jo-leave})(e) \wedge \mu^c(e) \geq d \wedge d >$
$\max(\{d' \mid \exists e'[\text{Exp}(e') = \text{ben} \wedge [\![\text{want}]\!](\text{ben-stay})(e') \wedge \mu^c(e') \geq d']\})]$

The assertion of (75) is that there is a state of Jo wanting to leave that exceeds by the contextually determined measure function μ^c any state of Ben wanting to stay. Assuming that μ^c is determined to be the intensity measure function, this gives us the intended interpretation of (3a). Moreover, there are two monotonicity presuppositions imposed by the two instances of FOR. Since the FOR PPs are adjoined at the vP level, the monotonicity requirements are checked with respect to the event predicates denoted by the vPs. Thus, the requirement is that μ^c be (non-trivially) monotonic on part-whole relations in the domains of states of Jo wanting to leave, and states of Ben wanting to stay. (If we instead say that FOR PPs adjoin to VPs, as assumed for convenience by Pasternak (in revision), then the event predicate against which the monotonicity requirement is checked is smaller, excluding the ascription of experiencer-hood.)

In this section, we have seen what *want* comparatives will look like if we assume that they compose like other verbal comparatives. Naturally, given that the monotonicity requirement was part of the motivation for the proposed semantics of verbal comparatives in the first place, it should be unsurprising to see that in this instance a monotonicity requirement is imposed as well. Since mental state comparatives allow for readings in which intensity is compared, this entails that intensity must be a monotonic measure of mental states. But this of course presupposes that (2–3) compose like other verbal comparatives to begin with. In the next section, I will offer an alternative proposal, in which the structure of intensity comparatives differs in a way that would suffice to skate around the monotonicity requirement, thereby potentially invalidating the argument in favor of intensity being a monotonic measure of mental states.

3.5 An alternative hypothesis: Lexical gradability

Suppose we want to avoid claiming that intensity is monotonic, but we don't want to let go of our generalizations about other verbal comparatives. How can we do this? Here's one path we can rule out right off the bat: a division of FOR into FOR_1 and FOR_2, with the former imposing a monotonicity requirement (and being used in non-intensity comparatives), and the latter not. After all, the two FORs would presumably be of the same semantic type and syntactic category, so there is no way beyond sheer stipulation by which we could prevent FOR_2 from being used in non-intensity verbal comparatives. And if FOR_2 is always available, we will generate non-monotonic readings where they don't belong.

It seems that the best way to avoid asserting that intensity is monotonic isn't to use a different FOR for intensity comparatives, but rather to avoid including FOR altogether. Since FOR is both what imposes the monotonicity requirement and what introduces the degree argument in non-intensity verbal comparatives, we can get rid of the former by obviating the latter: if mental state verbs are *lexically* gradable, then we can build intensity comparatives without FOR, and thus without imposing a monotonicity requirement.

Let's see what this would look like, again using (3a) as our example. Say that we informally define *want* as in (76):

(76) $[\![\text{want}]\!]_{\text{grad}} \approx \lambda p \lambda d \lambda e.\ e$ is a state of wanting p to at least degree d

Obviously, nothing in this definition states that intensity of desire is monotonic. And if we throw this *want* into the syntactic structure in (77), we have no trouble generating an intensity comparison reading without FOR entering into the derivation.

(77) a.

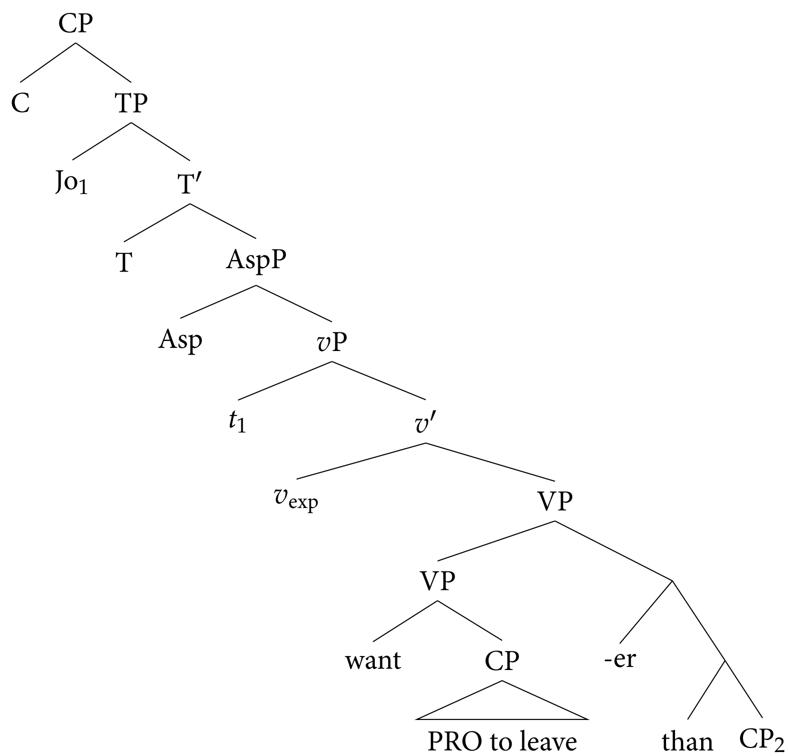

b.

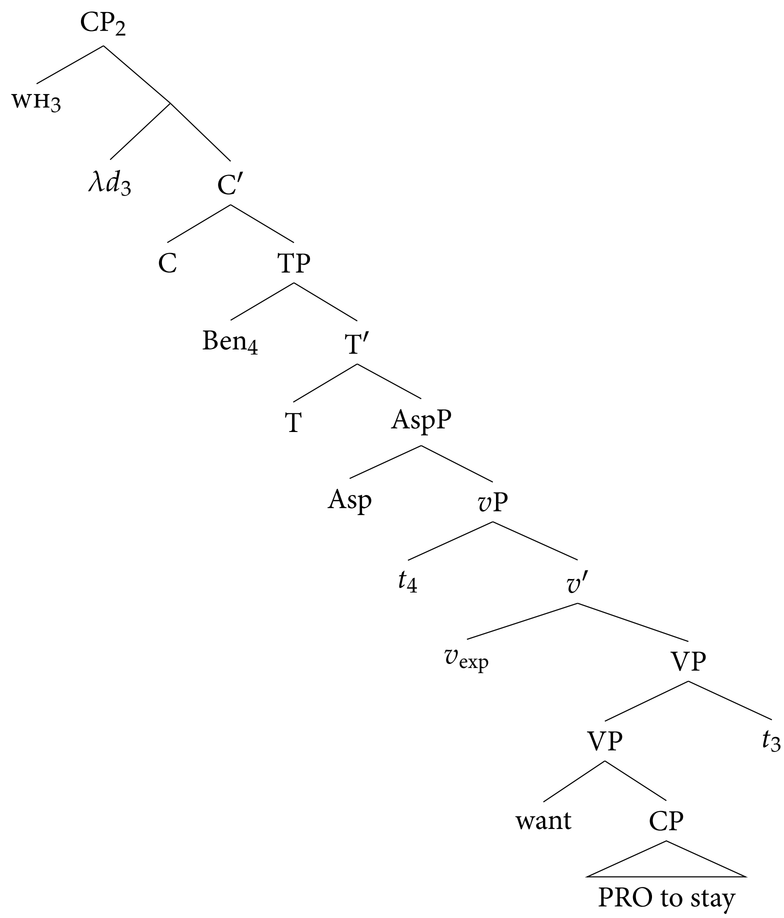

Starting in the comparison clause (77b), *want* first composes with its clausal complement, followed by the trace of the maximalizing wh- operator WH. Given that the trace denotes a free degree variable d_3, the VP as a whole comes out to (78):

(78) $[\![\text{VP}]\!] \approx \lambda e.\ e$ is a state of wanting ben-stay to at least degree d_3

This is of the right semantic type to combine with v_{exp}, so things proceed as normal. After lambda abstraction over d_3 and maximalization by WH, what we get is (79), the maximal degree to which Ben wants to stay:

(79) $[\![\text{CP}_2]\!] \approx \max(\{d \mid \exists e[\text{Exp}(e) = \text{ben} \wedge$ e is a state of wanting ben-stay to at least degree $d]\})$

Since $[\![\text{CP}_2]\!]$ is a degree, it can combine with $[\![\text{-er}]\!]$, denoting a generalized quantifier over degrees. This constituent undergoes the typical QR, leading to the syntax in (80):

(80) [[-er than WH Ben wants to stay] [λd_5 [Jo wants to leave t_5]]]

Just like in the comparison clause, matrix *want* combines with its clausal complement and the degree-variable-denoting trace, and the semantics of the matrix clause is built up from there, up until the lambda-abstraction over d_5. The resulting predicate of degrees—true of a degree d iff there is a state of Jo wanting to leave to at least degree d—can then serve as the argument of *-er than* CP_2, leading to the final result in (81):

(81) $\exists d \exists e[\text{Exp}(e) = \text{jo} \wedge$
e is a state of wanting jo-leave to at least degree $d \wedge d >$
$\max(\{d' \mid \exists e'[\text{Exp}(e') = \text{ben} \wedge$
e' is a state of wanting ben-stay to at least degree $d']\})]$

We thus see that by making intensity verbs gradable, we can get a fully compositional semantics for intensity comparatives that completely excludes FOR, thereby avoiding an imposition of the monotonicity requirement. We also retain the generalization about other verbal comparatives, since non-intensity verbs are not gradable, and thus need FOR to do the work of introducing degrees. Furthermore, on the morphosyntactic level, we still predict *much*-insertion, since the fact that *want* is gradable does not change the fact that it lacks a comparative form.

What about simple positive attributions of desire, as in *Jo wants to leave*? In this case, something will have to saturate the degree argument of *want*. Given the parallel with positive adjectives (e.g., *Lana is tall*), we can simply introduce a verbal version of the aforementioned positive morpheme POS (cf. Piñon 2005, Kennedy & Levin 2008), either as a specifier of VP or as a VP adjunct, as represented in (82):

(82)

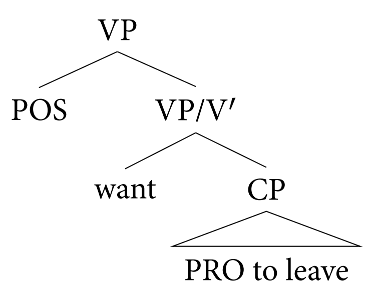

Semantically, POS takes a relation g between degrees and eventualities, and returns an event predicate true of an event e iff it is an event of g-ing to a degree greater than some contextually determined standard for g-ing events.[14] This is stated formally in (83), and the resulting denotation for (82) is as in (84).

(83) $[\![\text{POS}]\!]^c = \lambda g_{\langle d,\langle v,t\rangle\rangle} \lambda e.\, \exists d[g(d)(e) \wedge d > \text{std}^c(g)]$

[14] For discussion of how such standard degrees are determined, and in particular how they relate to the scale structure of gradable elements, see (among others) Kennedy & McNally 1999, 2005; Rotstein & Winter 2004; Kennedy 2007; McNally 2011; Lassiter & Goodman 2013.

(84) $[\![(82)]\!] = \lambda e.\ \exists d[[\![\text{want}]\!]_{\text{grad}}(\text{jo-leave})(d)(e) \wedge d > \text{std}^c([\![\text{want}]\!]_{\text{grad}}(\text{jo-leave}))]$

This is now of the right type to combine with v_{exp}, so the rest of the semantic derivation can go off without a hitch. The interpretation of the sentence as a whole will thus be true iff there is a state of Jo wanting to leave to a degree exceeding some contextually-determined standard.

We now have a more or less fully fleshed-out alternative to the view in which intensity is monotonic: namely, intensity comparatives compose differently from other verbal comparatives due to their lexical gradability. Let's call this proposal the **lexical gradability hypothesis (LGH)**. The question, then, is whether there is any morphosyntactic evidence suggesting that intensity comparatives compose differently from other verbal comparatives, and in particular whether any evidence suggests that mental state verbs are lexically gradable.

As it turns out, there is overt evidence from Chinese suggesting the plausibility of LGH. For verbal comparatives measuring something other than intensity, Chinese requires the inclusion of *duo* ('much'), along with a concomitant particle *de*. This is demonstrated with *pao* ('run') in (85), which has the same range of meanings as the English (1).

(85) Zhangsan bi Lisi pao *(de duo).
Zhangsan than[15] Lisi run *(DE much)
Zhangsan ran more than Lisi.

With adjectival comparisons, on the other hand, *duo* is absent:

(86) Zhangsan bi Lisi gao.
Zhangsan than Lisi tall
Zhangsan is taller than Lisi.

Importantly, intensity comparatives pattern with adjectival comparatives, and not with other verbal comparatives: they lack *duo*, as can be seen in (87) with *xiang* ('want').

(87) Zhangsan bi Lisi xiang likai.
Zhangsan than Lisi want leave
Zhangsan wants to leave more than Lisi does.

[15] While I follow Liu (1996) and Xiang (2003) in glossing *bi* as 'than', the syntactic category (and thus the proper gloss) of *bi* remains unclear; Liu (1996) and Xiang (2003) analyze it as a preposition, Erlewine (2007) argues that it is a functional verbal head, and Erlewine (2017) proposes that it is a semantically asymmetric conjunction.

If we assume some version of LGH, and that *duo* is (or can be) an indicator of the presence of FOR—perhaps there is "*duo* insertion", somewhat like the aforementioned possibility of *much* insertion—then the facts in (85)–(87) fall out nicely. In this account, adjectives and mental state verbs, which carry their own degree argument, do not combine with FOR to form comparatives, while other verbs must combine with FOR, which introduces a degree argument, imposes a monotonicity requirement, and brings about *duo* insertion.

As further evidence, mental state verbs, like gradable adjectives and unlike other verbs, can be directly modified by degree modifiers like *hen* ('very'):

(88) Zhangsan bu hen gao.
Zhangsan NEG very tall
Zhangsan is not very tall.[16]

(89) # Zhangsan hen pao (-le).
Zhangsan very run (-PERF)

(90) Zhangsan hen xiang likai.
Zhangsan very want leave
Zhangsan wants to leave very much.

If *hen* can only combine with something that carries a degree argument, this is again expected under LGH: the degree-carrying gradable adjectives and mental state verbs accept modification by *hen*, while other verbs do not.

However, in the next section I will argue that the similarities and differences between Chinese and English are best accounted for by positing that intensity is, in fact, a monotonic measure of mental states. I will start with English, showing that the use of an apparently monotonicity-requiring construction to measure intensity of psychological states is not restricted to verbal comparatives, and actually extends to all of the constructions discussed earlier. As a result, proponents of using LGH as a counterproposal to a monotonic account of psychological intensity must strengthen their claim, so that a distinction in lexical gradability must be posited across all of these constructions. I then turn back to Chinese, demonstrating that when we look beyond verbal measurement constructions, the contrast between intensity and (other) monotonic measure functions evaporates: where *duo* appears, it appears across the board, even when measuring intensity. I then show that while LGH struggles to account for these facts, a monotonic proposal faces no difficulty in doing so.

[16] I use the negated form of this sentence because *hen* is obligatory in non-negated positive adjectival predications: **Zhangsan gao* is ungrammatical, while *Zhangsan hen gao* can be true if Zhangsan is tall, but not *very* tall. When under negation or modifying a mental state verb, however, *hen* makes its expected semantic contribution. See Krasikova 2008a for further discussion.

3.6 Against the LGH

3.6.1 The evidence from English

Earlier in this chapter, five English constructions were shown to have monotonicity requirements, barring adoption of the LGH: pseudopartitives, the measurement idioms *out the wazoo* and *in spades*, adverbial measure phrases, and nominal and verbal comparatives. Examples (2) and (3) already showed that verbal comparatives allow for measurements of intensity of mental states. Adverbial measure phrases, the other verbal measurement construction, can also be used to measure mental states, as illustrated in (91) with *hate*, *respect*, and *want*:

(91) a. Zelda hates Yoshi a great deal.
b. In that moment, Waldo respected Xavier a ton.
c. At the end of the meeting, Vince wanted the CEO to be fired, and he wanted it a lot.[17]

As can be seen in (92), these adverbial measure phrases are measuring the same thing as what is measured in the case of verbal comparatives, i.e., intensity:

(92) a. Zelda hates Yoshi a great deal, while Claire only hates him a little bit. #But Claire hates him more than Zelda does.
b. Waldo respected Xavier a ton, while Charlotte only respected him a little bit. #But Charlotte respected him more than Waldo did.
c. As for firing the CEO, Vince wanted it a lot, while Tabby only wanted it a little bit. #But Tabby wanted it more than Vince did.

It is worth noting, however, that the LGH also predicts these to be possible. If we continue to assume that vague measure phrases like *a great deal* denote contextually determined degrees, these degrees can simply saturate the degree arguments of mental state verbs directly, without FOR as an intermediary.

What about the other three monotonicity-requiring constructions? For those, we will switch from the verbs *hate*, *respect*, and *want* to the nouns *hatred*, *respect*, and *desire*. First, pseudopartitives:

(93) a. Zelda has a great deal of hatred for Yoshi.
b. Waldo had a ton of respect for Xavier.
c. There was a lot of desire on Vince's part for a change in leadership.

[17] The reason for the somewhat cumbersome wording here is that English has a preference for low adjunct attachment, so the preferred reading of *Vince wanted the CEO to be fired a lot* is one in which *a lot* modifies *be fired*, rather than *want the CEO to be fired*. The inclusion of *at the end of the meeting* is to prevent a reading involving the frequency, rather than intensity, of Vince's desire.

The examples in (93) are all well-formed and mean what one would expect: for instance, that Vince had an intense desire for a change in leadership. The fact that the adverbial measure phrases in (91) and the pseudopartitives in (93) use the same measure function is made clear by the contradictory nature of the sentences in (94):

(94) a. Zelda has a great deal of hatred for Yoshi, #but she doesn't hate him a great deal.
b. Waldo had a ton of respect for Xavier, #but he didn't respect him a ton.
c. There was a lot of desire on Vince's part for a change in leadership, #but he didn't want it a lot.

Turning next to the measurement idioms, (95) and (96) provide cases where *out the wazoo* and *in spades* (respectively) are felicitously used to measure the intensity of states of hatred, respect, and desire:

(95) a. Zelda has hatred out the wazoo for Yoshi.
b. Waldo had respect out the wazoo for Xavier.
c. Vince had desire out the wazoo for a change in leadership.

(96) a. Zelda has hatred in spades for the newly formed government.
b. Waldo had respect in spades for anyone who would risk their own life to save someone else's.
c. I love her phrase, too, "a desire to know more and still more." As a therapist, I've got that desire in spades.[18]

Again, the fact that the examples in (95) involve intensity measurements can be seen in (97). (The same is true of the sentences in (96).)

(97) a. Zelda has hatred out the wazoo for Yoshi, #but she doesn't hate him very much.
b. Waldo had respect out the wazoo for Xavier, #but he didn't respect him very much.
c. Vince had desire out the wazoo for a change in leadership, #but he didn't want it very much.

Last but not least, in (98) we see that nominal comparatives also allow for the measurement of psychological states in terms of intensity:

(98) a. Zelda has more hatred for Yoshi than Claire does. (#But Claire hates him more than Zelda does.)

[18] https://www.psychologytoday.com/blog/headshrinkers-guide-the-galaxy/201208/got-curiosity

b. Waldo had more respect for Xavier than Charlotte did. (#But Charlotte respected him more than Waldo did.)

c. There was more desire on Vince's part than on Tabby's part for a change in leadership. (#But Tabby wanted it more than Vince did.)

In summary, all five of the normally monotonicity-requiring English constructions discussed earlier can be used to measure the intensity of psychological states. Meanwhile, only two of these constructions—verbal comparatives and adverbial measure phrases—can be directly accounted for by adopting LGH. This means that in order for LGH to be viable as a counterproposal to a monotonic account of intensity, nouns like *desire* and *hatred* need to be gradable in the same way that *want* and *hate* allegedly are, and the distinction in the presence or absence of something like FOR needs to cut across all five constructions. As a result, LGH becomes a much stronger hypothesis than it was when looked at solely through the lens of verbal comparatives, especially given that there is no overt evidence in English for such a widespread structural distinction.

3.6.2 The evidence from Chinese

By placing more demands on LGH, we also place more demands on what Chinese has to look like in order to constitute overt evidence in favor of LGH. If by hypothesis FOR (or something like it) adds a degree argument and imposes the monotonicity requirement, and it is the presence of FOR that (somehow) triggers *duo* insertion, then by LGH any normally monotonicity-requiring construction with *duo* should be *duo*-less when used to measure psychological intensity, for the same reason that *duo* was absent in the intensity comparative (87). I will show that this is not the case, based on evidence from a nominal measure construction roughly analogous to the pseudopartitive, as well as from nominal comparatives.

Jiang (2009) observes that in Chinese, pre-nominal measure phrases have different syntactic properties depending on whether the measure function is monotonic or not. When the measure function is monotonic, there is an option to include or exclude the particle *de* between the measure phrase and the noun, as in (99a). When the measure function is not monotonic, as in (99b), *de* is obligatory.

(99) a. si sheng (de) dui
four liter (DE) water
four liters of water

b. si du *(de) shui
four degree *(DE) water
four-degree water

With this in mind, consider (100), in which *duo* is necessary, while *de* is optional:

(100) Zhangsan mai -le hen *(duo) (de) kafei.
Zhangsan buy -PERF very *(much) (DE) coffee
Zhangsan bought a lot of coffee.

hen duo (de) kafei is interpreted like the pseudopartitive *a lot of coffee* in that the degree is vague and the measure function is context-dependent, with a requirement for monotonicity. A natural broad-strokes analysis of (100) that fits with the assumptions underlying LGH is that (something like) FOR introduces a degree argument, imposes a monotonicity requirement, and triggers *duo* insertion. Because FOR brings a degree argument with it, modification by *hen* becomes permissible, and since the result must be monotonic, *de* is optional, as per Jiang's observation.

We now have another Chinese measurement construction that imposes a monotonicity requirement in a fashion that brings *duo* along for the ride. Thus, the prediction of LGH is that if intensity of mental states is measurable in this construction, then it should be measurable without *duo*, as the noun should come with its own degree argument and permit direct modification by *hen*. However, this turns out not to be the case. Consider the examples of love and respect. As can be seen in (101), the verbs *ai* ('love') and *zunjing* ('respect') pattern with *xiang* ('want') in being directly modifiable by *hen* and appearing in verbal comparatives without *duo*:

(101) a. Zhangsan {hen / bi Lisi} ai Chong.
Zhangsan {very / than Lisi} love Chong
Zhangsan loves Chong {very much/more than Lisi does}.

b. Zhangsan {hen / bi Lisi} zunjing jingli.
Zhangsan {very / than Lisi} respect manager
Zhangsan respects the manager {very much/more than Lisi does}.

But when we turn to *hen duo (de)*, what we see is that just like in (100), *de* is optional, while *duo* is required.

(102) a. Zhangsan dui Chong you hen *(duo) (de) ai.
Zhangsan to Chong have very *(much) (DE) love
Zhangsan has a lot of love for Chong.

b. Zhangsan dui jingli you hen *(duo) (de) jingyi.
Zhangsan to manager have very *(much) (DE) respect
Zhangsan has a lot of respect for the manager.

Since the examples in (102) allow for—in fact, prefer—a reading in which what is measured is the intensity of love/respect, this spells trouble for the strengthened

LGH. If FOR (+ *duo*) brings monotonicity with it as LGH predicts, then the contrast between (101) and (102) leads to the awkward prediction that intensity of love and respect both is and is not monotonic.

Similar facts can be gleaned from nominal comparatives, which in Chinese also impose a monotonicity requirement (as in English), and also require *duo*:

(103) Zhangsan bi Lisi mai -le geng *(duo) de kafei.
Zhangsan than Lisi buy -PERF GENG[19] *(much) DE coffee
Zhangsan bought more coffee than Lisi did.

Once again, LGH makes the prediction that when switching from coffee to love and respect, *duo* should disappear. But again, this prediction fails, and *duo* is obligatory:

(104) a. Zhangsan bi Lisi dui Chong you geng *(duo) de ai.
Zhangsan than Lisi to Chong have GENG *(much) DE love
Zhangsan has more love for Chong than Lisi does.

b. Dui jingli Zhangsan bi Lisi you geng *(duo) de jingyi.
to manager Zhangsan than Lisi have GENG *(much) DE respect
Zhangsan has more respect for the manager than Lisi does.

To summarize, both *hen duo (de)* and nominal comparatives generally require that the measure function used be monotonic, and for both constructions, *duo* appears across the board, including when psychological intensity is measured. This does serious damage to the claim that Chinese provides overt evidence for a version of LGH strong enough to oppose a monotonic account of intensity. In order to keep LGH afloat, one would have to abandon the claim that it is the presence of FOR that triggers *duo* insertion; otherwise, (102) and (104) go unaccounted for. But then the whole explanation for the difference between (85) and (87)—the verbal comparatives—goes out the window, and it is back to square one. This, of course, is not to say that an LGH-based account is impossible, as the right combination of covert elements can no doubt bring about the desired result. But the ensuing proposal would be no less stipulative than it would have been for English, and the distribution of *duo* becomes a mystery.

Meanwhile, if intensity *is* taken to be a monotonic measure of mental states, then the facts in this section are readily accounted for. Let us start with English. Intensity comparatives in English are not overtly distinct from other verbal comparatives, so they can be analyzed as composing in the same way: neither type of verb carries its own degree argument, so FOR always introduces the degree argument and adds the monotonicity requirement. Since the intensity measure function is monotonic, an

[19] Much like with *bi*, how *geng* should be glossed is not obvious. I leave it unglossed, but see Krasikova 2008a for arguments that it is an intensifier like English *even* or *still*.

intensity comparative reading can arise. The same holds of nominal measurement constructions: nouns like *desire*, *love*, etc. do not carry their own degree argument, and so they compose just like other nouns in these measure constructions. The result is that on a compositional level, there is no difference in English between measurements of intensity and other monotonic measurements.

A monotonic proposal can also account for the Chinese data, while simultaneously preserving the intuition that FOR adds a degree argument and brings *duo* and monotonicity with it. In the nominal realm, where measurements of intensity do not stand out from other monotonic measurements, Chinese looks just like English: nouns like *ai* ('love') and *jingyi* ('respect') do not have their own degree argument, and as a result they compose like other nouns. This accounts for (102) and (104): both contain *duo* because FOR is needed to add the degree argument, and both have intensity readings because intensity is monotonic. As for mental state verbs, we can take a page from the LGH book and simply say that they come with a built-in degree argument. The rest plays out just like in LGH, with the pre-existing degree argument obviating the need for FOR and *duo* and enabling direct modification by *hen*. So while the English verb *respect* will have the degreeless denotation in (105a), the analogous Chinese verb *zunjing* will look like (105b), where μ_{int} is the intensity measure function:[20]

(105) a. $[\![\text{respect}]\!] = \lambda x \lambda e.\ \text{respect}(e) \wedge \text{Thm}(e) = x$

b. $[\![\text{zunjing}]\!] = \lambda x \lambda d \lambda e.\ \text{respect}(e) \wedge \text{Thm}(e) = x \wedge \mu_{\text{int}}(e) \geq d$

In this proposal we see that a monotonic view of intensity is not inherently at odds with a view in which there is variation across verbs (and across languages) in the presence or absence of a degree argument. After all, a denotation along the lines of (105b) can still conform to a monotonic account if μ_{int} is monotonic. The difference is that under a monotonic account, the predictions are more lax on the compositional level, since the LGH predicts universal presence of a built-in degree argument for mental state verbs and nouns, while a monotonic account is agnostic about its presence or absence for a given lexical item.

Of course, this begs the question of *why* certain verbs should look like (105b), but not others. While I do not have a complete answer to this question, I can at least offer some speculation. Notice that outside of mental state verbs, in all of the ex-

[20] Krasikova (2008a) and Erlewine (2017) argue, following the pioneering work of Beck et al. (2004) on similar phenomena in Japanese, that degrees do not enter the compositional semantics of Chinese comparatives like they do in English. In particular, they argue that while English permits lambda-abstraction over degrees in the manner specified by von Stechow (1984), Heim (1985, 2000), and others, Chinese does not, leading to certain notable differences in interpretation between Chinese and English comparatives. If this is true, then the denotations of mental state verbs and the Chinese analog to FOR may have to be tweaked accordingly.

amples of verbal measurement constructions discussed in this chapter the choice of measurement was either temporal duration or something that aligned with temporal duration. For example, while running events can be measured in terms of duration or distance, the part-whole relations under consideration are still the same: in checking for monotonicity, we compare the measurement of a longer event to its shorter subevents, which have both a shorter duration and a shorter distance traveled. But mental state intensity and temporal duration do not pattern with each other in this manner, and measuring the intensity of a mental state involves looking at wholly different part-whole relations from the ones involved in temporal measurement. Maybe, then, a verb is more likely to be lexically gradable if its measurement is along a dimension whose part-whole relations are not the same as those for temporal measurement. While this claim is of course highly speculative, it is at least falsifiable, though I must leave the task of such falsification or verification for future work.

3.7 Conclusion

This chapter has been devoted to arguing at length for a single, simple hypothesis: the intensity of mental states tracks their part-whole structure. In English, the evidence in favor of this hypothesis is substantial, as the reviewed constructions imposing monotonicity requirements on the domain of measurement—pseudopartitives, nominal comparatives, verbal comparatives, adverbial measure phrases, and the measurement idioms *out/up the wazoo* and *in spades*—can all be used to measure the intensity of mental states. In Chinese, the evidence was more mixed. Evidence from verbal measurement constructions suggests that Chinese mental state verbs are lexically gradable, rendering the lexical gradability hypothesis a viable (and perhaps pretheoretically preferable) alternative to the ontological account that I propose. However, when turning our attention to nominal measurement constructions, the evidence for lexical gradability disappears. The best analysis, I argued, was one in which the ontological condition holds, and which posits that there is cross-linguistic variation in which verbs are or are not lexically gradable. Thus, in English, mental state verbs are not lexically gradable, and hence compose in measurement constructions like other verbs, while in Chinese, they are in fact lexically gradable. In both languages, mental state nominals are not lexically gradable, so both compose in the expected manner in nominal measurement constructions.

In the next chapter, we will see what a natural language metaphysics of mental states (and especially attitudes) exemplifying the ontological requirement argued for in this chapter might look like.

Chapter 4: Two-dimensional attitudes

In the previous chapter, I argued that in an adequate natural language metaphysics, intensity is a monotonic measure of mental states, including attitudes. In this chapter, I will explore in depth what such a natural language metaphysics might look like. I will start in Section 4.1 by looking at the simpler case of transitive mental state verbs like *hate*, going over the basic proposed ontology and its repercussions for the semantics of mental state verbal comparatives. In Section 4.2 I discuss how the proposals of von Fintel (1999) and Heim (1992) can be incorporated into the semantics and ontology adopted in this dissertation, and in particular how the semantics of *want* should interact with the part-whole structure of desire states. Section 4.3 is dedicated to the question of how, given the possibilities for the semantics of *want* discussed in the previous section, cases can be analyzed where a single experiencer wants multiple things with varying intensities. I show in Section 4.4 that the previous results easily extend to *wish* and *regret*, as the differences between these attitudes and *want* have no effect on the analysis at hand. In Section 4.5 I discuss some of the semantic predictions of my account outside of the mereological claims of monotonicity, as well as evidence suggesting that these predictions are correct; I also discuss how some other proposals fail to make some of these predictions. After some concluding remarks in Section 4.6, I show in an appendix how my proposal can be effected in a premise semantics, wherein sets of propositions are used to generate orderings over worlds (Kratzer 1981a, Lewis 1981).

4.1 Non-attitude mental states: The case of *hate*

Before discussing *want*, we will go through the analysis for the simpler case of transitive psychological verbs like *hate*. In Section 4.1.1 I introduce the basics of the proposed ontology, including the structure of psychological states and the relationship between ⟦hate⟧ and the mereological structure of states of hatred. I then show in Section 4.1.2 how the proposed ontology, in conjunction with the analysis of verbal comparatives offered in the previous chapter, generates the right results for *hate* comparatives.

4.1.1 The ontology of intensity: Going vertical

In order to make μ_{int} monotonic, I will treat mental states as extending in two dimensions. The first, "horizontal" dimension is time; the fact that such states exist in time is intuitively obvious, as well as necessary for the interpretation of tense and aspect. The second, "vertical" dimension will be the one along which intensity is measured.

Before talking about two-dimensional states, it will help to clearly establish the ontology and terminology of the more commonly discussed horizontal dimension of time. In most implementations, a timeline is an ordered pair $\langle T, \leq_T \rangle$, where T is a set of *moments* in time, and $\leq_T$ is a dense ordering on T, usually with no minimal element (i.e., no "first moment") or maximal element. Events can then be situated on this timeline. For example, let's say that Dee's running event e occupies the bit of timeline seen in Figure 4.1:

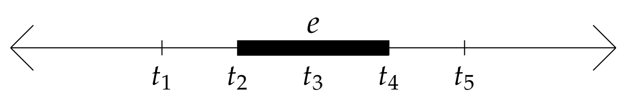

Figure 4.1: Dee's running event, situated in time

The set of moments that e spans is the closed interval $[t_2, t_4]$, which is traditionally referred to as e's *temporal trace* ($\tau(e)$). Notice that τ is *not* a measure function on events. A measure function takes an entity or eventuality and returns a degree; τ, on the other hand, takes an eventuality and returns a set of moments. For example, if t_2 is 2PM, and t_4 is 4PM, then $\tau(e)$ is the set of moments from 2PM to 4PM, inclusive. The temporal measure function μ_{dur}, on the other hand, returns the degree denoted by the measure phrase *two hours*. That being said, there is a clear relationship between τ and μ_{dur}: if $\tau(e')$ is the same as $\tau(e)$, then $\mu_{\text{dur}}(e') = \mu_{\text{dur}}(e)$ (= two hours), and if $\tau(e')$ is the set of moments from 2PM to 3PM (inclusive), then $\tau(e') \subset \tau(e)$, and $\mu_{\text{dur}}(e') < \mu_{\text{dur}}(e)$.

We thus have at our disposal three ways of talking about time: moments (e.g., t_1), intervals ($[t_2, t_4]$), and degrees of duration (two hours). In moving from one- to two-dimensional eventualities, each of these notions will have an analog in the vertical dimension.

In the same way that the (horizontal) timeline was a pair $\langle T, \leq_T \rangle$, the vertical analog to a timeline will be an ordered pair $\langle K, \leq_K \rangle$, where K is a set of *altitudes*, and $\leq_K$ is a dense ordering over K such that $k_a \leq_K k_b$ iff k_b is at least as high an altitude as k_a. However, I will assume two important distinctions between moments (and their ordering) and altitudes (and their ordering). First, whereas $\leq_T$ was taken to have no minimum, I will assume that there is in fact a minimum, "sea level" altitude k_0. Second, whereas eventualities can start and end at arbitrary times, mental states will always start at k_0 and extend upwards. The reason for these stipulations is for the sake of clarity. In the horizontal dimension of time, we have a very clear idea of what it means for two events to start at different times, but have the same duration, such as if one event goes from 1–3PM, and another goes from 2-4PM. In the vertical dimension of intensity, on the other hand, it is less obvious what it would mean for two mental states to start and end at different altitudes, but have the same intensity. I will therefore side-step this issue by stipulating that mental states simply start at k_0, and leave the exploration of alternative possibilities for another time.

Since mental states are two-dimensional objects, they occupy spaces in a two-dimensional coordinate system of moments and altitudes. Hence, the temporal trace function τ can be replaced by the more general function π, which takes a psychological state and returns the set of pairs (t, k) of a moment t and altitude k such that e occupies k at t. τ can then be redefined based on π as in (1a), in which τ takes an eventuality and returns the set of times such that that eventuality occupies some altitude at that time. Similarly, $\kappa(e)$—e's *vertical span*, the vertical analog to its temporal trace—can be defined as in (1b).

(1) a. $\tau(e) = \{t \mid \exists k[(t, k) \in \pi(e)]\}$
  b. $\kappa(e) = \{k \mid \exists t[(t, k) \in \pi(e)]\}$

Note that κ, like τ, is not a measure function, since neither returns a degree. But much like the aforementioned relationship between τ and μ_{dur}, I assume a close-knit relationship between κ and μ_{int}: if two mental states e_1 and e_2 start at k_0, with e_1 extending up to k_1 and e_2 reaching k_2 (where $k_1 <_K k_2$), then $\mu_{\text{int}}(e_1) < \mu_{\text{int}}(e_2)$.

It will help to consider an example. Figure 4.2 illustrates a psychological state e that grows more intense, reaches a peak, and then rapidly dissipates. In this example, $\tau(e) = [t_2, t_4]$, since e occupies every moment from t_2 to t_4. Similarly, $\kappa(e) = [k_0, k_2]$, since for every altitude k in that range, there is some t such that $(t, k) \in \pi(e)$. As stated above, at each moment the state starts at k_0 and extends upward. As for $\mu_{\text{int}}(e)$, what matters is not what we label the degree assigned to it, but rather that μ_{int} and κ are related in a manner parallel to μ_{dur} and τ, as discussed above.

Finally, some remarks are in order about the relationship between mental states and their proper parts. If there is a state of Ann hating Bill, and that state goes from 1PM to 3PM, then clearly the part of this state from 1PM to 2PM is also a state of Ann hating Bill, as are the parts from 1:10PM to 1:11PM and from 1:10:12PM to 1:10:13PM. In other words, mental state ascriptions appear to obey some version of the subinterval property (Bennett & Partee 1972), at least down to a certain granularity. How fine the granularity is is not obvious, but I will assume that at least for statives (of which mental state ascriptions are exemplars) the subinterval property extends down even to individual moments. This should be interpreted more as a simplifying assumption

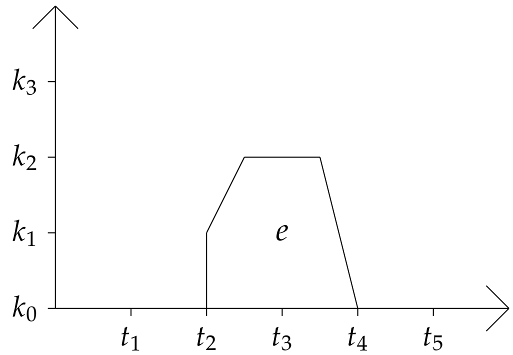

Figure 4.2: A sample mental state, situated horizontally and vertically

than as an independently motivated claim; a coarser-grained version of the subinterval property would do just as well for our purposes.

More importantly, in switching to two-dimensional mental states, I will extend the subinterval property to the vertical dimension as well, so that if e is Ann's state of hating Bill, and $\kappa(e) = [k_0, k_2]$, then the portion of e from k_0 to k_1 (where $k_1 <_K k_2$) will also be a state of Ann hating Bill. The result of combining the horizontal and vertical versions of the subinterval property is that any part of a state of Ann hating Bill will itself be a state of Ann hating Bill. I will refer to this property of mental states as *two-dimensional subdivision*.

Two-dimensional subdivision is a claim about what can be inferred about the parts of a mental state, given certain information about the whole. Similar inferences can be made in the opposite direction as well. If e_1 is a state of Ann hating Bill that goes from 1pm to 2pm, and e_2 is an Ann-hating-Bill state going from 2pm to 3pm, then clearly $e_1 \sqcup e_2$, the sum of e_1 and e_2, is also a state of Ann hating Bill. This is the familiar trait of *cumulativity* (Krifka 1989), which holds of a property iff it is closed under mereological sum. Like two-dimensional subdivision, I will assume that cumulativity is not restricted to the horizontal dimension, but is also true vertically: the sum of two "stacked" Ann-hating-Bill states is also a state of Ann hating Bill.

The conjunction of two-dimensional subdivision and cumulativity leads to a biconditional constraint that I will refer to as *mental state homogeneity*, defined in (2) (where $v\text{P}_{\text{men}}$ is a vP whose verb is a mental state verb).

(2) **Mental State Homogeneity:**
$[\![v\text{P}_{\text{men}}]\!](e) \leftrightarrow \forall e' \sqsubseteq e[[\![v\text{P}_{\text{men}}]\!](e')]$

Thus, mental state homogeneity requires that a state is a state of Ann hating Bill if and only if all of its substates are states of Ann hating Bill.

4.1.2 Deriving *hate* comparatives

With the ontology of intensity now in place, we can see how the semantics of verbal comparatives interacts with the part-whole structure of mental states in order to derive readings in which intensity is compared. The sentence under consideration will be (3):

(3) Ann hates Bill more than Matt hates Jeff.

As per the analysis in the previous chapter, (3) has the pre-QR syntactic representation in (4):

(4) a.

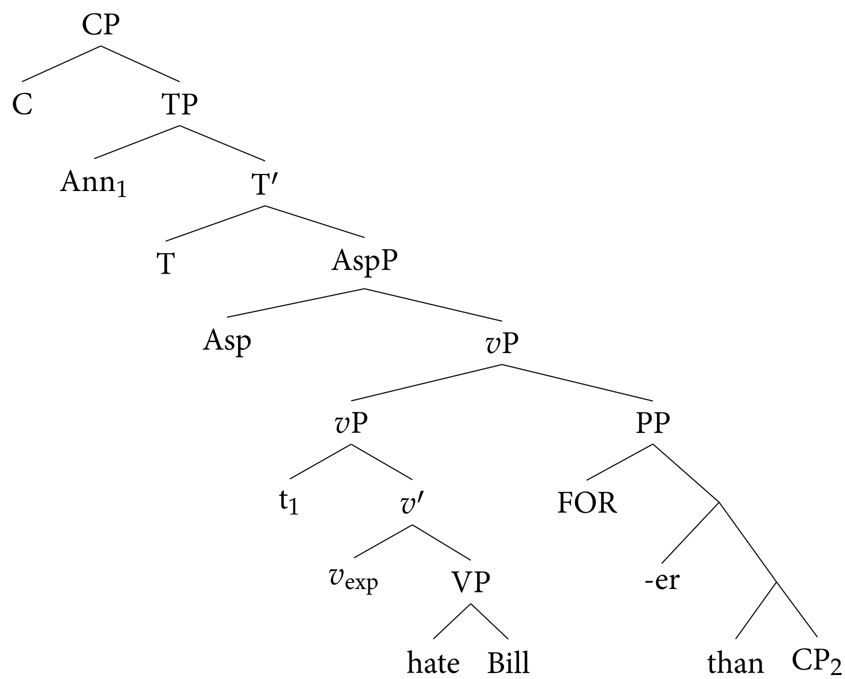

b.

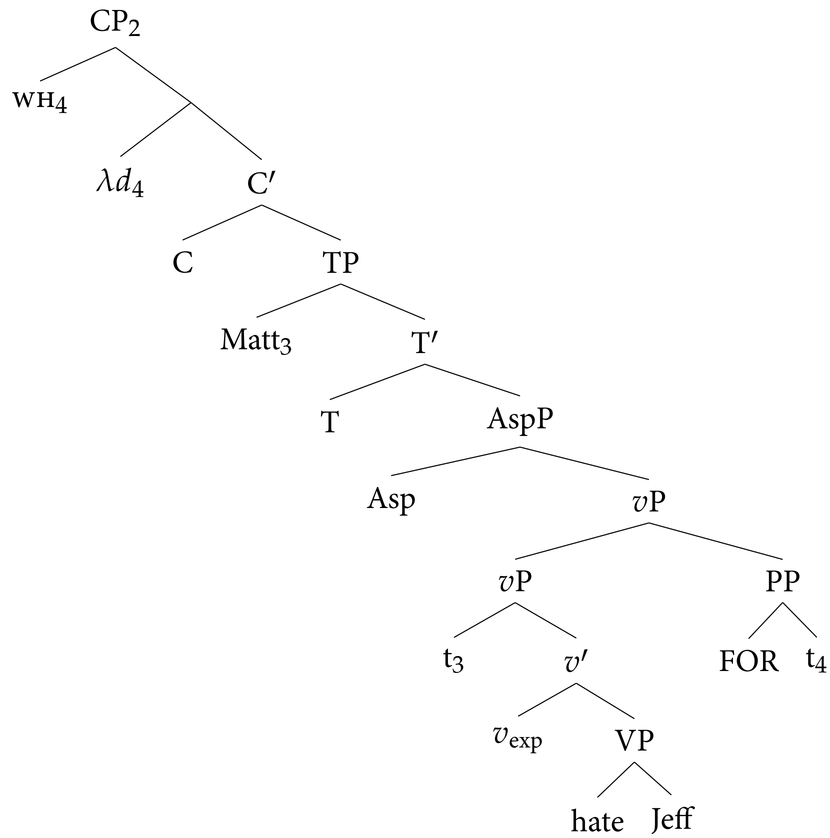

After QR of the degree phrase, we get the structure in (5):

(5) [[-er than WH$_4$ Matt hates Jeff FOR t$_4$]$_5$ [λd_5 [Ann hates Bill FOR t$_5$]]]

As for the compositional semantics, starting in the matrix clause, FOR combines with t$_5$ to get the denotation in (6a) (where d_5 is the free variable denoted by t$_5$), while the denotation of the pre-adjunction matrix vP is (6b) (reviving our simplifying assumption that referential DPs are interpreted in their merge position):

(6) a. $\lambda V \lambda e : \mu^c$ is monotonic on $\sqsubseteq^c$ in V. $V(e) \wedge \mu^c(e) \geq d_5$
b. λe. Exp(e) = ann $\wedge$ hate(e) $\wedge$ Thm(e) = bill

Assuming that the contextually determined μ^c is μ_{int}, the assertive component of the composition of these by function application is (7). There will also be a presupposition, brought in by FOR, that μ_{int} is monotonic on salient part-whole relations in (6b), the set of states of Ann hating Bill.

(7) λe. Exp(e) = ann $\wedge$ hate(e) $\wedge$ Thm(e) = bill $\wedge$ $\mu_{\text{int}}(e) \geq d_5$

After existential closure of the event predicate by Asp, as well as lambda abstraction over d_5, we are left with the denotation in (8):

(8) $\lambda d.\, \exists e[\mathrm{Exp}(e) = \mathrm{ann} \wedge \mathrm{hate}(e) \wedge \mathrm{Thm}(e) = \mathrm{bill} \wedge \mu_{\mathrm{int}}(e) \geq d]$

Moving to the comparison clause, things go similarly up until and including lambda abstraction over d_4, leading to (9), again with μ_{int} serving as the contextually determined measure function. This time, the presupposition will be that μ_{int} is monotonic on $\sqsubseteq^c$ in the domain of states of Matt hating Jeff.

(9) $\lambda d.\, \exists e[\mathrm{Exp}(e) = \mathrm{matt} \wedge \mathrm{hate}(e) \wedge \mathrm{Thm}(e) = \mathrm{jeff} \wedge \mu_{\mathrm{int}}(e) \geq d]$

This (characteristic function of a) set of degrees is fed to WH, which finds the maximum degree:

(10) $\max(\{d \mid \exists e[\mathrm{Exp}(e) = \mathrm{matt} \wedge \mathrm{hate}(e) \wedge \mathrm{Thm}(e) = \mathrm{jeff} \wedge \mu_{\mathrm{int}}(e) \geq d]\})$

This degree is fed to $[\![\text{-er}]\!]$, followed by the set of degrees in (8). The result, which will be the denotation of the entire sentence, is as in (11):

(11) $\exists d \exists e[\mathrm{Exp}(e) = \mathrm{ann} \wedge \mathrm{hate}(e) \wedge \mathrm{Thm}(e) = \mathrm{bill} \wedge \mu_{\mathrm{int}}(e) \geq d \wedge d >$
$\max(\{d' \mid \exists e'[\mathrm{Exp}(e') = \mathrm{matt} \wedge \mathrm{hate}(e') \wedge \mathrm{Thm}(e') = \mathrm{jeff} \wedge \mu_{\mathrm{int}}(e') \geq d']\})]$

(11) is true iff there is a state of Ann hating Bill that exceeds by μ_{int}—that is, in intensity—any state of Matt hating Jeff. Along the way, we generated presuppositions that μ_{int} was monotonic on $\sqsubseteq^c$ in the domains of states of Ann hating Bill, and of Matt hating Jeff.

Now say that e_{a} is Ann's (maximal) state of hating Bill, so that $\mathrm{Exp}(e_{\mathrm{a}}) = \mathrm{ann}$, $\mathrm{hate}(e_{\mathrm{a}})$, and $\mathrm{Thm}(e_{\mathrm{a}}) = \mathrm{bill}$ are all true. Similarly, e_{m} is Matt's state of hating Jeff, with $\mathrm{Exp}(e_{\mathrm{m}}) = \mathrm{matt}$, $\mathrm{hate}(e_{\mathrm{m}})$, and $\mathrm{Thm}(e_{\mathrm{m}}) = \mathrm{jeff}$ all being true. Consider the scenario in which e_{a} and e_{m} are as diagrammed in Figure 4.3. (While I place the states side by side, these states should be thought of as simultaneous.) As can be seen from the diagram, $\kappa(e_{\mathrm{a}}) = [k_0, k_2]$, while $\kappa(e_{\mathrm{m}}) = [k_0, k_1]$, where $k_1 <_K k_2$. Hence, $\kappa(e_{\mathrm{a}}) \supset \kappa(e_{\mathrm{m}})$, so given the relationship between μ_{int} and κ, $\mu_{\mathrm{int}}(e_{\mathrm{a}}) > \mu_{\mathrm{int}}(e_{\mathrm{m}})$. The assertion in (11) is thus true: the highest degree of intensity manifested in a state of Matt hating Jeff is $\mu_{\mathrm{int}}(e_{\mathrm{m}})$, and there is a state of Ann hating Bill that exceeds e_{m} in intensity, namely e_{a}.

In addition, the monotonicity presupposition is satisfied as well, so long as the salient part-whole relation is set in the right way. Consider what happens when we look at horizontal "strips" of e_{a} and e_{m}—indicated in Figure 4.3 by dashed lines—and their sums. Courtesy of mental state homogeneity, each strip of e_{a} will itself be an Ann-hating-Bill state. Furthermore, just like the thin layers used to measure the depth of snow in the previous chapter, the sum of any two strips of e_{a} will have a greater measurement in the vertical dimension (i.e., a greater intensity) than each of its parts. Thus, μ_{int} is indeed monotonic on such a part-whole relation.

Naturally, the same sort of analysis extends to sentences with adverbial measure phrases like *Ann hates Bill a lot*, with the difference being that the degree of comparison is not the degree to which Matt hates Jeff, but the degree denoted by the measure

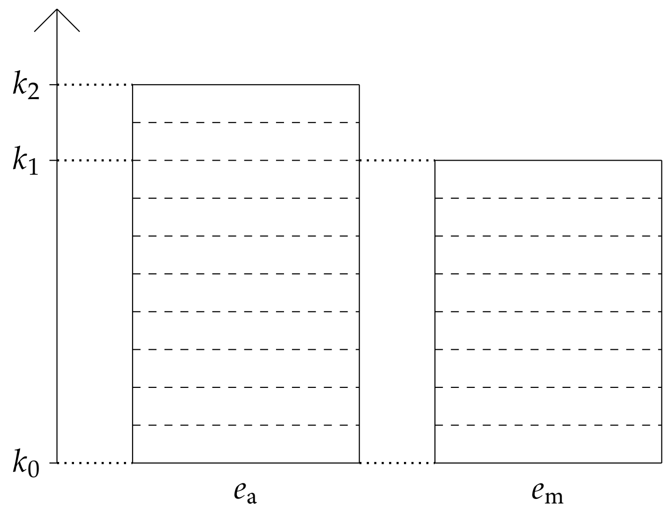

Figure 4.3: Diagram of Ann's state of hating Bill and Matt's state of hating Jeff

phrase *a lot* (though see Section 3.3). The proposed syntax for this sentence can be seen in (12):

(12)

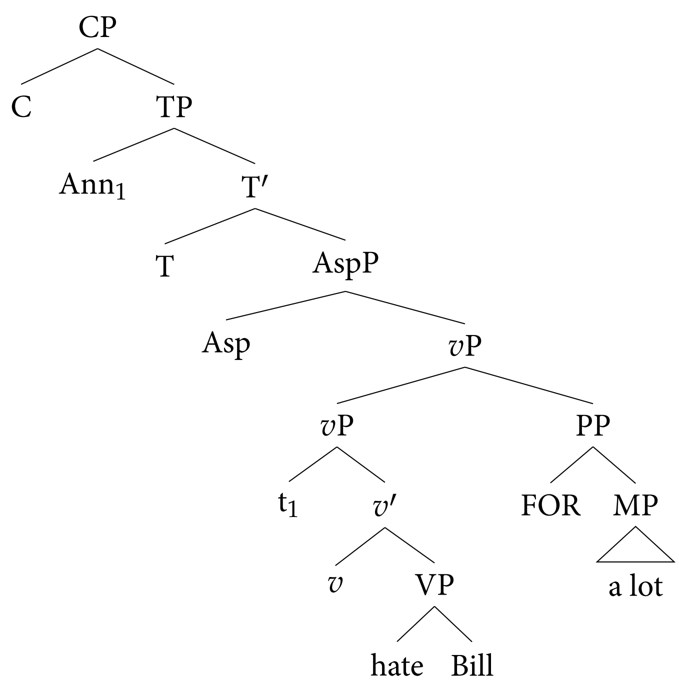

The semantic result on current assumptions will be as in (13) if the contextually determined measure function is μ_{int}.

(13) $\exists e[\text{Exp}(e) = \text{ann} \wedge \text{hate}(e) \wedge \text{Thm}(e) = \text{bill} \wedge \mu_{\text{int}}(e) > [\![\text{a lot}]\!]^c]$

Thanks to FOR, we again generate a monotonicity presupposition, i.e., a presuppositoin that μ_{int} is monotonic on $\sqsubseteq^c$ in the domain of states of Ann hating Bill. As per the previous discussion, this presupposition is indeed satisfied.

4.2 Transitioning to *want*

We have seen that the proposed ontology and semantics get us the right results when looking at cases like *hate*. In this section I extend the analysis to the more complex case of *want*. I start in 4.2.1 by discussing the relationship between $[\![\text{want}]\!]$ and the part-whole structure of desire states. In 4.2.2, I show how the proposals of von Fintel (1999) and Heim (1992) on the semantics of *want* can be adapted to our ontology and semantics. Finally, in 4.2.3 we will see what happens when this semantics for *want* is inserted into verbal comparatives.

4.2.1 Semanticizing homogeneity

In the previous section, mental state homogeneity was posited as a constraint on the model used for interpretation, rather than being explicitly included in the denotations of verbs like *hate*. However, the additional complexities of *want* will require a more intricate relationship between semantics and part-whole structure, so for these, homogeneity will be baked directly into the denotation of the verb: $[\![\text{want}]\!]$ will break up a desire state into very small parts and universally quantify over those parts. How small these parts are depends on how fine-grained one takes two-dimensional subdivision to be; since we are assuming an extremely fine-grained version, these parts will only occupy a single moment and a single altitude. I will refer to these tiny parts of a state e as *point-states* of e, with $\text{PT}(e)$ being the set of such point-states:

(14) $\text{PT}(e) = \{e/(t,k) \mid (t,k) \in \pi(e)\}$,
where $e/(t,k) = \iota e' \sqsubseteq e[\pi(e') = \{(t,k)\}]$[1]

Thus, if WANT is everything in the denotation of *want* other than this quantification over point-states, $[\![\text{want}]\!]$ will be as in (15), where p is the proposition denoted by the clausal complement of *want*:

(15) $[\![\text{want}]\!] = \lambda p \lambda e.\ \forall e' \in \text{PT}(e)[\text{WANT}(p)(e')]$

While the primary motivation for building homogeneity into the denotation of *want* is to provide a direct means of weaving the part-whole structure of desire states into the semantics, it is worth noting that independent evidence in favor of such

[1] I assume that the ι operator ends up being well-defined, i.e., each desire state has exactly one part occupying a given moment-altitude pair. If this is not the case, ι can be replaced with Link's (1983) σ operator, which would return the *maximal* such substate.

breaking up and universal quantification can be found in the temporal relationship between matrix and embedded clauses. It has often been noted that *want* requires its embedded clause to be interpreted in the future relative to the desiring itself; to put it in Condoravdi's (2002) terms, *want* has a *future temporal orientation*. Thus, (16) is fully acceptable with *tomorrow* in the embedded clause, but replacing *tomorrow* with *yesterday* requires either a play on words or a time travel scenario:

(16) Heinrich wants to leave {tomorrow/#yesterday}.

Proposed explanations for this fact vary, but it has generally been taken for granted that the future-shifting of the embedded clause is relative to the temporal trace of the desire state in the matrix clause (though many of these proposals do not cash this intuition out in neo-Davidsonian terms). (16) is thus predicted to be bad with *yesterday* because the embedded clause must be future-shifted with respect to the temporal trace of Heinrich's current desire state, meaning that the embedded clause in (16) requires that Heinrich's potential leaving be both in the future and yesterday, a contradiction.

With this in mind, consider the sentence in (17):

(17) (At 8PM,) Heinrich wanted to leave immediately.

Here is a rough translation for (17): There was a desire state e (at 8PM), with experiencer Heinrich, that was a state of wanting to leave immediately after $\tau(e)$. The future-shifting takes place relative to $\tau(e)$, with *immediately* serving to relate the future-shifted time to $\tau(e)$ by adding a requirement that they be temporally proximate. Now consider (18):

(18) For three hours, Heinrich wanted to leave immediately.

Here is what (18) does *not* mean: There was a three-hour desire state e, with experiencer Heinrich, that was a state of wanting to leave immediately after $\tau(e)$. Such an analysis would predict (18) to be true in a scenario in which Heinrich's desire from 8PM to 11PM was that he leave right after 11PM, but in this scenario (18) is in fact false. Instead, what must be the case is that for each (near-)momentary substate e' of Heinrich's three-hour desire state, Heinrich's desire in e' is to leave immediately after $\tau(e')$, so that at 8 he wants to leave right after 8, at 9:30 he wants to leave right after 9:30, etc.

So the interpretation required for (18) is one in which the actual proposition desired changes over the course of Heinrich's three-hour desire state. A denotation like (15) allows this to happen in a composition-friendly manner, so long as the future-shifting of the embedded clause is relative to e' (the quantified-over substates), rather than e (the larger desire state). This phenomenon of shifting goalposts thus provides direct evidence in favor of quantifying over small substates, at least along the horizontal dimension; I take extending this analysis into the vertical dimension to be a harmless stipulation.[2]

4.2.2 Two options for WANT

Now that the relationship between ⟦want⟧ and the mereology of attitude states has been established, we are left with the task of defining WANT, which is responsible for the bulk of the work in the semantics of *want*. Naturally, one's choice of WANT depends on what one takes the appropriate base semantics for *want* to be. For reasons discussed in Chapter 2, I will explore two possibilities: a WANT along the lines of von Fintel's (1999) Kratzerian theory of *want*, and a WANT based on Heim's (1992) double-ordering definition of *want*.

First, von Fintel. Recall von Fintel's definition of *want* in (19):

(19) $$[\![\text{want}]\!]_{\text{von Fintel}} = \lambda p \lambda x \lambda w : \exists w', w'' \in \text{Dox}(x,w)[p(w') \wedge \neg p(w'')]. \\ \forall w' \in \textsc{Best}(\text{Dox}(x,w), \precsim_{g(x,w)})[p(w')]$$

As discussed in Chapter 2, $\text{Dox}(x,w)$ is the set of possible worlds compatible with x's beliefs in w (x's *belief worlds*). g is a Kratzerian ordering source that is individual-relative, generating a set of propositions used to order worlds based on x's preferences in w. The presupposition of *want* is its diversity condition: the proposition denoted by the clausal complement must be compatible with, but not entailed by, the experiencer's beliefs. The assertion is that among those worlds compatible with x's beliefs, the best worlds as determined by $\precsim_{g(x,w)}$ are all worlds in which p holds.

In order for von Fintel's definition of *want* to serve as our WANT, some furniture needs to be rearranged, as the semantic type of von Fintel's *want* ($\langle\langle s,t\rangle,\langle e,\langle s,t\rangle\rangle\rangle$) is not the same as that of WANT ($\langle\langle s,t\rangle,\langle v,t\rangle\rangle$). The result of this rearrangement can be seen in (20), where $\text{Dox}(e)$ is the set of belief worlds of e, and $\precsim_e$ is e's ordering over worlds based on bouletic preferability. The diversity condition will not be relevant for the rest of this chapter, so it is excluded; its translation from von Fintel's definition is straightforward.

(20) $$\text{WANT}_{\text{vF}} = \lambda p \lambda e.\ \forall w \in \textsc{Best}(\text{Dox}(e), \precsim_e)[p(w)]$$

[2]There is an alternative analysis that, to my knowledge, generates the same result in this regard. Say that ⟦want⟧ has the simpler denotation $\lambda p \lambda e.\ \text{WANT}(p)(e)$, and that it is only true of individual point-states to begin with. Hence, temporal shifting is only relative to point-states, rather than their sums, as desired. There could then be a higher head that contributes Link's (1983) * operator, closing the eventuality predicate under mereological sum, as would seem to be required for a durative adverbial like *for three hours* in (18). I will leave for future work a choice between (15) and this alternative; the proposal in this chapter is compatible with either one.

Note that we are no longer using sets of propositions to generate orderings over worlds, and are simply positing that each desire state comes with a world-ordering. For those who wish to see how the proposals in this chapter, and especially the next section, can be translated into a premise semantics, see the Appendix to this chapter.

If WANT_{vF} is used as our WANT, the resulting denotation for *want* will be as in (21):

(21) $\lambda p \lambda e.\ \forall e' \in \textsc{pt}(e)[\forall w \in \textsc{Best}(\text{Dox}(e'), \precsim_{e'})[p(w)]]$

(21) is true of a state e if and only if at every point-state in e, all of the bouletically ideal belief worlds at that point-state are worlds in which p is true.

Let us next move on to Heim's (1992) definition of *want*, repeated below:

(22) Preliminary definitions
   a. $\text{Sim}_w(p) =_{\text{def}}$ the set of p worlds most similar to w.
   b. $w_1 <_{x,w} w_2$ iff w_1 is more preferable to x in w than w_2 is.
   c. $A_1 <_{x,w} A_2$ iff for all $w_1 \in A_1$ and $w_2 \in A_2$, $w_1 <_{x,w} w_2$.

(23) $[\![\text{want}]\!]_{\text{Heim}} = \lambda p \lambda x \lambda w : \exists w', w'' \in \text{Dox}(x, w)[p(w') \wedge \neg p(w'')].$
$\forall w' \in \text{Dox}(x, w)[\text{Sim}_{w'}(\text{Dox}(x, w) \cap p) <_{x,w} \text{Sim}_{w'}(\text{Dox}(x, w) - p)]$

By Heim's definition, x wants p in w iff for every one of x's belief worlds w', the closest belief worlds to w' in which p is true are all better than the closest belief worlds to w' in which p is false.

Making Heim's $[\![\text{want}]\!]$ the right type to serve as our WANT is more or less straightforward given the translation for von Fintel's $[\![\text{want}]\!]$. This can be seen in (24); note that the diversity presupposition is once again excluded.

(24) $\text{WANT}_{\text{H}} = \lambda p \lambda e.\ \forall w \in \text{Dox}(e)[\text{Sim}_{w'}(\text{Dox}(e) \cap p) <_e \text{Sim}_{w'}(\text{Dox}(e) - p)]$

The ensuing denotation for *want* will then be as in (25):

(25) $\lambda p \lambda e.\ \forall e' \in \textsc{pt}(e)[\forall w \in \text{Dox}(e')[\text{Sim}_w(\text{Dox}(e') \cap p) <_{e'}$
$\text{Sim}_{w'}(\text{Dox}(e') - p)]]$

We thus have two possibilities for the semantics of *want*, both translated into the compositional semantics and ontology adopted in this dissertation. As stated in Chapter 2, I will not try to choose between these two options, as the problems that arise and the proposed solutions that are adopted are identical for both. That being said, I will for the most part use von Fintel's definition of *want* in cases where the choice of theory does not matter, since von Fintel's is the formally simpler of the two theories.

4.2.3 Verbal comparatives with *want*

Next we will see what happens when a definition of *want* along the lines just proposed is inserted into a verbal comparative. To demonstrate, I will use the sentence in (26):

(26) Jo wants to leave more than Ben wants to stay.

A complete bottom-to-top derivation of (26) was provided in Chapter 3. The assertive component of the result was as in (27), assuming μ^c is set to μ_{int}:

(27) $\exists d \exists e[\text{Exp}(e) = \text{jo} \wedge [\![\text{want}]\!](\text{jo-leave})(e) \wedge \mu_{\text{int}}(e) \geq d \wedge d > \max(\{d' \mid \exists e'[\text{Exp}(e') = \text{ben} \wedge [\![\text{want}]\!](\text{ben-stay})(e') \wedge \mu_{\text{int}}(e') \geq d']\})]$

According to (27), (26) is true iff there is a state of Jo wanting to leave that exceeds in intensity any state of Ben wanting to stay. The monotonicity presuppositions, introduced by FOR, are that μ_{int} is monotonic on $\sqsubseteq^c$ in the domain of states of Jo wanting to leave, and that it is monotonic on $\sqsubseteq^c$ in the domain of states of Ben wanting to stay.

We can now fill in the gaps by inserting our definition of *want*. I will use the von Fintel-based definition, but the same logic applies to a Heim-based approach. The result is in (28):

(28) $\exists d \exists e[\text{Exp}(e) = \text{jo} \wedge \forall e' \in \textsc{pt}(e)[\forall w \in \textsc{Best}(\text{Dox}(e'), \precsim_{e'})[\text{jo-leave}(w)]] \wedge \mu_{\text{int}}(e) \geq d \wedge d > \max(\{d' \mid \exists e''[\text{Exp}(e'') = \text{ben} \wedge \forall e''' \in \textsc{pt}(e'')[\forall w' \in \textsc{Best}(\text{Dox}(e'''), \precsim_{e'''})[\text{ben-stay}(w')]] \wedge \mu_{\text{int}}(e'') \geq d']\})]$

(28) is true iff there is a state with experiencer Jo, such that at all point-states of that state all ideal worlds are Jo-leaving worlds, and such that the measurement of that state by intensity exceeds that of any state with experiencer Ben, such that at all point-states of that state all ideal worlds are Ben-staying worlds.

Now suppose that e_{j} is Jo's (maximal) desire state, and likewise for e_{b} and Ben. Furthermore, let's say for simplicity's sake that for every point-state in $\textsc{pt}(e_{\text{j}})$, worlds in which Jo leaves are better than worlds in which she doesn't, and that's that. Likewise, assume that for every point-state in $\textsc{pt}(e_{\text{b}})$, worlds in which Ben stays are better than worlds in which he doesn't. Finally, suppose that Jo and Ben's desire states are as in Figure 4.4.

In this case, the assertion of (26) is predicted to be true: the maximal degree to which Ben wants to stay is $\mu_{\text{int}}(e_{\text{b}})$, and there is a state of Jo wanting to leave whose intensity exceeds this, namely, e_{j}. In addition, the presupposition that μ_{int} is monotonic on $\sqsubseteq^c$ in the domains of states of Jo wanting to leave and states of Ben wanting to stay is satisfied. After all, as per the discussion in 4.2.1, the quantification over point-states in the definition of *want* puts mental state homogeneity into the

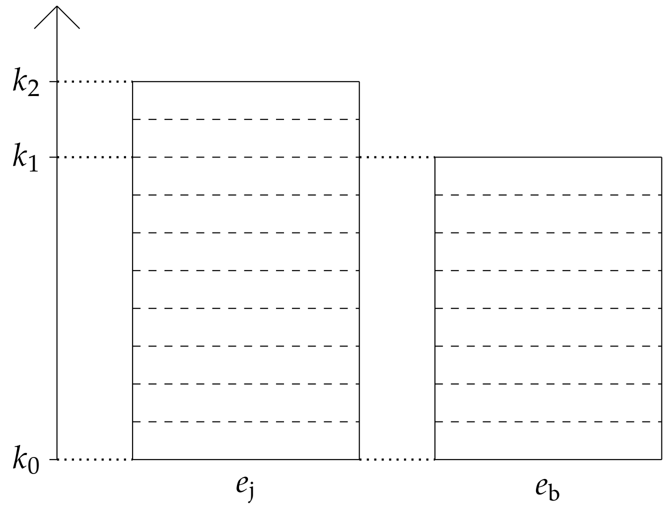

Figure 4.4: Diagram of Jo's and Ben's desire states

semantics. If all of the point-states of a state e have some property, then all of the point-states of a substate e' of e have that property, since every point-state of e' is a point-state of e. This gets us two-dimensional subdivision (the two-dimensional version of the subinterval property), meaning that just like with *hate*, the thin, snow-like layers of e_j are themselves states of Jo wanting to leave, and likewise for e_b and Ben wanting to stay. Thus, we can once again say that $\sqsubseteq^c$ is the relationship between layers of desire states and their sums, and just like with states of hatred (and, by analogy, collections of snow), vertical measurement—that is, intensity—satisfies the monotonicity property of verbal comparatives.

We thus see that, perhaps not shockingly, semanticizing mental state homogeneity garners as a result that intensity is a monotonic measurement of desire states, thereby allowing *want* comparatives to compose in the same manner as other verbal comparatives. But there was a little quirk to our example scenario: for each of e_j and e_b, all of the point-states had the same ordering over possible worlds. Since nothing is changing across point-states within a given experiencer's desire state, our definitions of *want* entail that Jo's desires are all equally strong, and likewise for Ben's desires. That is, there are no propositions p and q such that Jo wants p more than she wants q. This observation raises an obvious follow-up question: what about when this is *not* the case, i.e., situations where a single experiencer wants different things with differing intensities? It is this problem that I address in the next section.

4.3 Desires of varying intensity

In this section, I will discuss how the von Fintel- and Heim-based semantics for *want* can be used to capture a single experiencer's varying intensity of desires in the ontological and semantic framework of this dissertation. As noted previously, I will mostly stick to von Fintel's best-worlds analysis to illustrate the problem and the proposed solution, but will note where relevant how my proposal extends to Heim's double-ordering theory.

4.3.1 The problem

Framing things in terms of WANT_{vF} (the version of WANT based on von Fintel's analysis of *want*), suppose that Ron has three relevant desires that he believes to be mutually compatible: he wants to eat some peanuts (p), he wants to visit Quebec (q), and he wants to learn Russian (r). Naturally, in all of Ron's bouletically ideal belief worlds, all three happen. But this does not entail that Ron wants all three equally. It might be the case that while Ron wants (and believes that he can get) all three, his desire to learn Russian is stronger than his desire to visit Quebec, which is stronger than his desire to eat peanuts. There are thus two pressing questions to answer. First, since WANT_{vF} only cares about the set of bouletically ideal belief worlds, how can one proposition be wanted more than another if both hold in all ideal worlds? And second, how can this be handled in a way that retains the monotonicity of intensity of desire?

A tempting answer to the first question is to follow Kratzer (1981a, 1991, 2012) in positing that Ron's three desires can be differentiated by widening our lens and looking at worlds that are less than ideal. For example, imagine that there are only eight worlds compatible with Ron's beliefs: w_{pqr}, where all three propositions hold, w_{qr}, where only q and r hold, and so on, for each combination of truth and falsehood of p, q, and r. Furthermore, imagine that Ron's bouletic ranking of worlds is as in (29), where (i) all r worlds are better than all $-r$ worlds, (ii) q serves as a tiebreaker for r, and (iii) p is a tiebreaker for q. (Of course, such a rigid ranking of priorities is something of an idealization, but it is a useful one for expository purposes.)

(29)
$$\underbrace{\underbrace{\underbrace{w_{pqr}}_{p} < \underbrace{w_{qr}}_{-p}}_{q} < \underbrace{\underbrace{w_{pr}}_{p} < \underbrace{w_{r}}_{-p}}_{-q}}_{r} < \underbrace{\underbrace{\underbrace{w_{pq}}_{p} < \underbrace{w_{q}}_{-p}}_{q} < \underbrace{\underbrace{w_{p}}_{p} < \underbrace{w_{\varnothing}}_{-p}}_{-q}}_{-r}$$

As observed above, looking only at the singleton set of ideal worlds $\{w_{pqr}\}$ does not provide enough information to tell whether Ron wants p, q, or r more. However, notice that in the ordering in (29), the best world in which Ron learns Russian but does not go to Quebec (w_{pr}) is more ideal than the best world in which Ron goes to Quebec but does not learn Russian (w_{pq}). In this sense, it can be said that Ron wants to learn Russian more than he wants to go to Quebec: while any Russian-but-no-Quebec or Quebec-but-no-Russian world is less than ideal, Ron finds the best worlds of the former sort more tolerable than the best worlds of the latter sort. The same game can be played in comparing learning Russian and eating peanuts, or going to Quebec and eating peanuts: the best $r-p$ world (w_{qr}) outranks the best $p-r$ world (w_{pq}), so Ron wants to learn Russian more than he wants to eat peanuts, and the best $q-p$ world (w_{qr}) outranks the best $p-q$ world (w_{pr}), so he wants to visit Quebec more than he wants to eat peanuts.

Considering things in this way allows us to reframe our above questions. What we now need to find out is this: how can the semantics and/or ontology of desire be sensitive to the relative rankings of worlds that the experiencer considers sub-ideal, while at the same time retaining a definition of *want* in terms of the best-worlds quantification seen in WANT_{vF}? Notice that while I have framed this problem in terms of WANT_{vF}, the same basic issue arises for WANT_{H}. Each definition of WANT incorporates a graded notion of preferability, since both utilize a ranking of worlds in terms of their comparative preferability to the experiencer. Moreover, this graded preferability of worlds intuitively ought to correlate with which propositions are wanted more or less than others. But at the same time, each definition of WANT is itself *non*-graded, and only uses a small amount of this preference information in order to determine which propositions are wanted or not. The question, then, is how we recover this lost information. To shed some light on how to answer this question, it will help to take a brief detour and look at a formally similar problem faced by certain theories of gradable adjectives.

4.3.2 Lessons from extension gap theories of adjectives

Most contemporary theories of gradability (including that adopted in this dissertation) take the basic meaning of a scalar adjective to be comparative in nature. The denotation of *tall*, for example, might take a degree and an individual as arguments and compare the individual's height to the degree:

(30) $[\![\text{tall}_1]\!]^c = \lambda d \lambda x.\ \text{height}(x) \geq d$

A positive use of *tall*, as in *Steph is tall*, then requires a silent morpheme POS to provide a degree of comparison that Steph's height must exceed in order to count as tall. But what if we instead wish to make an adjective's *positive* interpretation its basic meaning, so that $[\![\text{tall}]\!]$ is just an $\langle e,t \rangle$-type predicate true of those individuals that qualify as tall?

(31) $[\![\text{tall}_2]\!]^c = \lambda x.\ x$ qualifies as tall by the standards in c

How can we derive from this the comparative interpretation that arises in *Steph is taller than Mark is*?

Here we start to see problems quite similar to those discussed above. After all, clearly this alternative definition of *tall* somehow has to be sensitive to the relative heights of the objects in its domain, much like how WANT is sensitive to the preference rankings of possible worlds. But at the same time, this definition of *tall* is itself non-graded, only using that information about comparative heights that is required to tell us what counts as tall and what does not. Hence, if Steph and Mark both qualify as tall, then this "yes-or-no" denotation of *tall* does not differentiate between them, just as WANT_{vF} is incapable of differentiating between two propositions that both hold throughout all bouletically ideal worlds.

There are thus lessons to be learned from looking at those theories that endorse just such an $\langle e, t\rangle$-type semantics for gradable adjectives. While this particular cat has been skinned in many ways, I will use a simplified form of Kamp's (1975) proposal to illustrate the gist. In a given conversational context, an adjective like *tall* can be thought of as dividing its domain into three groups: those entities that definitely qualify as tall (the *positive extension*), those that definitely qualify as not tall (the *negative extension*), and those in the middle that qualify neither as tall nor as not tall (what Klein (1980) calls the *extension gap*). If x is in the positive extension, then $[\![\text{tall}]\!]^c(x) = 1$; if x is in the negative extension, then $[\![\text{tall}]\!]^c(x) = 0$; and if x is in the extension gap, then $[\![\text{tall}]\!]^c(x)$ is undefined. In contexts more precise than c, the members of the extension gap can be assigned freely to either the positive or negative extension, with the following caveat: if x is at least as tall as y, then y can only be in the positive extension of *tall* if x also is, and x can only be in the negative extension of *tall* if y is. Thus, the positive extension of *tall* in a given context will contain those individuals above a certain height, the negative extension those individuals below a certain (possibly lower) height, and the extension gap those individuals with heights between the bottom of the positive extension and the top of the negative extension. Given this constraint, we can say that for Steph to be taller than Mark is for there to be some context c such that Steph is in the positive extension of $[\![\text{tall}]\!]^c$, and Mark is in its negative extension. So long as there is a sufficiently large class of contexts, this will be true if and only if Steph is indeed taller than Mark.

We therefore see that gradability of height can be captured while keeping a definition of *tall* that is simply a (partial) function from entities to truth values. The key is to allow ourselves to toy with the standards for what does or does not count as tall. If x is taller than y, then there will be some way of setting the standard for tallness so that x is in the positive extension and y is in the negative extension. I will adopt a formally similar approach in accounting for desires of varying intensities: while WANT will continue to just be a relation between propositions and point-states, the standard for what counts as desired will be lowered or raised to weed out differences between propositions. In order to do so, there are two further complications that need to be addressed. First, this raising and lowering of standards must somehow be integrated with the quantification over possible worlds used in our definitions of WANT. Second, it must also be integrated with the mereology of attitude states adopted in this chapter. I address these two problems in turn.

4.3.3 The nitpicker's guide to graded intensionality

As per the discussion of *tall* above, we need to find a way to appropriately manipulate the standard for the positive extension of WANT. If the standard for WANT is lowered, then every proposition that was previously WANTed needs to still be WANTed, with the possible addition of some newly-WANTed propositions. That is, the positive extension of WANT must be a (possibly proper) superset of what it used to be. Naturally, the opposite must happen when raising the standard: the positive extension must become a (possibly proper) *subset* of what it used to be.

As far as integrating this with quantification over worlds, some recent work on gradability in the Kratzerian tradition of best-worlds quantification has exploited the fact that if $A \subset B$, then universal quantification over A is weaker than universal quantification over B. Put another way, the smaller the domain of world-quantification, the more propositions will hold in all worlds throughout that domain. Hence, if we generate a sequence of progressively shrinking domains of world-quantification, then as we go through the sequence, more and more propositions will hold throughout the modal domain, thereby lowering the standard for the positive extension of the world-quantifying operator. The relative importance of a proposition can then be tied to how early in the sequence that proposition starts to hold throughout the modal domain: propositions whose necessity is established early in the sequence are of greater importance than those that require further shrinking of the modal domain in order for their necessity to be established. Or on a more intuitive level, the more nitpicky we have to be in order to mandate p, the less important p is.

The *locus classicus* for this approach to weakening modal world-quantification is von Fintel & Iatridou's (2008) work on the contrast between so-called "weak" necessity modals like *should* and strong necessity modals like *must*. In short, the idea is that (deontic) *must* p is true iff p holds in all *acceptable* worlds, while *should* p is true iff p holds in all *ideal* worlds. Since all ideal worlds are acceptable worlds (but not vice versa), more propositions will hold throughout all ideal worlds than will hold throughout all acceptable worlds. As a result, we rightly predict *should* p to be weaker than *must* p:

(32) a. You should do your homework, but it's not the case that you must.

b. # You must do your homework, but it's not the case that you should.

von Fintel & Iatridou formalize this idea roughly as follows. Operating in a Kratzerian apparatus, say that there are in fact two ordering sources at play, g_1 and g_2. Where w is the world of evaluation, $\precsim_{g_1(w)}$ serves to distinguish those worlds that are morally acceptable from those that are not: worlds are ranked with respect to weighty matters like whether murder is committed, basic safety standards are met, etc. Meanwhile, $\precsim_{g_2(w)}$ only ranks worlds with respect to those lower-level priorities that distinguish ideal worlds from merely acceptable worlds, like whether waitstaff are properly tipped and people do not cut in line. In this case, *must* can be defined à la Kratzer, but sticking in the high-priority ordering source g_1.

(33) $[\![\text{must}]\!] = \lambda p \lambda w.\ \forall w' \in \textsc{Best}(\bigcap f(w), \precsim_{g_1(w)})[p(w')]$

Since g_1 only cares about important matters, *You must not commit murder* is predicted to be true, while *You must not cut in line* is false.

As for *should*, $\precsim_{g_2(w)}$'s ranking of worlds based on little things like cutting in line and tipping waitstaff occurs in addition to, and not instead of, $\precsim_{g_1(w)}$'s differentiation based on more important criteria. Thus, $[\![\text{should}]\!]$ is not simply the result of replacing $\precsim_{g_1(w)}$ with $\precsim_{g_2(w)}$ in $[\![\text{must}]\!]$. Instead, $[\![\text{should}]\!]$ takes the set of acceptable worlds as determined by $\precsim_{g_1(w)}$, i.e., $\textsc{Best}(\bigcap f(w), \precsim_{g_1(w)})$, and finds the best worlds among these as determined by the ordering $\precsim_{g_2(w)}$, thereby deriving the set of ideal worlds:

(34) $[\![\text{should}]\!] = \lambda p \lambda w.\ \forall w' \in \textsc{Best}(\textsc{Best}(\bigcap f(w), \precsim_{g_1(w)}), \precsim_{g_2(w)})[p(w')]$

Since it is necessarily the case that $\textsc{Best}(A, \precsim) \subseteq A$, we derive the right entailment relations, as the domain for modal quantification for *should* will be a (usually proper) subset of that for *must*.

While I will retain the core intuition behind von Fintel & Iatridou's proposal—that weakening of world-quantification is due to lower-level priorities having their say in determining what counts as a more ideal world—I will formalize this intuition differently. The reason for this is that while von Fintel & Iatridou's formalism works fine for a best-worlds semantics like Kratzer's *must* or WANT_{vF}, it is unclear how the same idea can be extended to WANT_{H}, which does not make reference to sets of ideal worlds. Of course, if von Fintel is right and Heim is wrong, then this is not a problem at all. But as it turns out, the same basic idea can be accomplished without forcing us to choose between WANT_{vF} and WANT_{H}.

More recent proposals in the tradition of weakening-by-domain-shrinkage have explored how much can be accomplished just by manipulating the world-ordering on its own (Katz et al. 2012, Portner & Rubinstein 2016, Pasternak 2016). That is, rather than first finding the $\precsim_{g_1(w)}$-ideal worlds and subsequently feeding those to $\precsim_{g_2(w)}$, we could instead combine g_1 and g_2 into a single ordering source $g_{1,2}$ that simulates g_2 serving as a tiebreaker for g_1. So if $\precsim_{g_1(w)}$ just ranks murder-free worlds as better than murder-ful ones, as semi-formally diagrammed in (35a), and $\precsim_{g_2(w)}$ only ranks no-line-cutting worlds above line-cutting ones as in (35b), then $\precsim_{g_{1,2}(w)}$ will still rank all murder-free worlds as better than all murder-ful ones, but worlds with the same status with respect to murder are further differentiated based on the existence or non-existence of line-cutting, as seen in (35c).

(35) a. $-\text{murder} \prec_{g_1(w)} \text{murder}$

b. $-\text{cut} \prec_{g_2(w)} \text{cut}$

c. $\underbrace{\underbrace{-\text{murder}, -\text{cut}}_{-\text{cut}} \prec_{g_{1,2}(w)} \underbrace{-\text{murder}, \text{cut}}_{\text{cut}}}_{-\text{murder}} \prec_{g_{1,2}(w)} \underbrace{\underbrace{\text{murder}, -\text{cut}}_{-\text{cut}} \prec_{g_{1,2}(w)} \underbrace{\text{murder}, \text{cut}}_{\text{cut}}}_{\text{murder}}$

We can then assign *should* the interpretation in (36):

(36) $[\![\text{should}]\!] = \lambda p \lambda w.\ \forall w' \in \textsc{Best}(\bigcap f(w), \precsim_{g_{1,2}(w)})[p(w')]$

In this case, we successfully weaken our quantification over worlds by manipulating only the world-ordering, since all of the $\precsim_{g_{1,2}(w)}$-ideal worlds will be devoid of both murder and line-cutting, in contrast to $\precsim_{g_1(w)}$.

There is a sense in which $\precsim_{g_{1,2}(w)}$ can be thought of as a "choosier" ordering than $\precsim_{g_1(w)}$, in that it is the result of replacing certain equivalences with strict orderings. In general, we can say that an ordering $\precsim_a$ is *more fine-grained* than $\precsim_b$ if $\precsim_a$ is the result of taking some equivalences in $\precsim_b$ and replacing them either with strict orderings or incomparabilities. This is stated more formally in (37).

(37) a. If $\precsim_a$ and $\precsim_b$ are preorders over the same set of worlds, then $\precsim_a$ is *at least as fine-grained* as $\precsim_b$ iff the following two conditions hold for all worlds u and v in the domains of $\precsim_a$ and $\precsim_b$:

i. If $u \prec_b v$, then $u \prec_a v$.

ii. If $u \parallel_b v$, then $u \parallel_a v$.

b. $\precsim_a$ is *more fine-grained* than $\precsim_b$ iff $\precsim_a$ is at least as fine-grained as $\precsim_b$, but not vice versa.

c. $\precsim_a$ is *at least as coarse* as $\precsim_b$ iff $\precsim_b$ is at least as fine-grained as $\precsim_a$. $\precsim_a$ is *coarser* than $\precsim_b$ iff $\precsim_b$ is more fine-grained than $\precsim_a$.

(35) illustrates how increasing fine-grainedness can shrink the set of ideal worlds, thereby weakening the quantification over worlds seen in a best-worlds semantics. In fact, it turns out that for *both* definitions of WANT, if the fine-grainedness of the ordering over worlds is increased, then it will necessarily be the case that all of the propositions that used to be WANTed will still be WANTed, and often new propositions will join the ranks of the WANTed as well. Note that the definitions of WANT in (20) and (24) are equivalent to those in (38):

(38) a. $\text{WANT}_{\text{vF}} = \lambda p \lambda e.\ \forall w \in \text{Dox}(e)[(\neg\exists w' \in \text{Dox}(e)[w' \prec_e w]) \rightarrow p(w)]$

b. $\text{WANT}_{\text{H}} = \lambda p \lambda e.\ \forall w \in \text{Dox}(e)[\forall w' \in \text{Sim}_w(\text{Dox}(e) \cap p)[$
$\forall w'' \in \text{Sim}_w(\text{Dox}(e) - p)[w' \prec_e w'']]]$

Both of these definitions make use specifically of the strict ordering relation ($\prec_e$), and moreover, this relation between worlds occurs in an upward-entailing environment. In WANT_{H}, it is multiply embedded in the scope of various universal quantifiers, which is upward-entailing. In the case of WANT_{vF}, the direction of entailment is twice reversed: the scope of negation and the antecedents of material implications are both downward-entailing environments, so embedding a negation in an antecedent leads to an upward entailing environment. So no matter which WANT you choose, retaining or expanding the set of strict ordering relations—which, by (37a-i), necessarily occurs when an ordering is made more fine-grained—will never render a previously WANTed proposition unWANTed, in much the same way that retaining or lowering the standard for tallness will never render a previously tall individual non-tall. What's more, increasing strict ordering relations can render previously unWANTed propositions newly WANTed, similar to how lowering the standard for tallness can make some individuals qualify as tall that hadn't previously qualified.

To make things more concrete, let us return to the Ron example and switch back to WANT_{vF}. Ron's three desires can be differentiated by generating a series of increasingly fine-grained orderings, as in (39):

(39) Step 1: $\underbrace{w_{pqr}, w_{qr}, w_{pr}, w_r}_{r} \prec_1 \underbrace{w_{pq}, w_q, w_p, w_\varnothing}_{-r}$

Step 2: $\underbrace{\underbrace{w_{pqr}, w_{qr}}_{q} \prec_2 \underbrace{w_{pr}, w_r}_{-q}}_{r} \prec_2 \underbrace{\underbrace{w_{pq}, w_q}_{q} \prec_2 \underbrace{w_p, w_\varnothing}_{-q}}_{-r}$

Step 3: $\precsim_3$ is as in (29)

Notice that for all three orderings, Ron learns Russian (r) in all ideal worlds. As for visiting Quebec (q), this happens in all ideal worlds with respect to $\precsim_2$ and $\precsim_3$, but not $\precsim_1$, which has as ideal the Quebec-less worlds w_{pr} and w_r. Finally, the only one of these three orderings in which Ron eats peanuts (p) in all ideal worlds is $\precsim_3$. Unsurprisingly, as the standard for WANT_{vF} is lowered by making the ordering more fine-grained, more propositions become WANTed. Since r is wanted earliest—it is true in all ideal worlds with respect to the coarsest ordering—it is wanted most, followed by q, then p.

To summarize, we have seen that for both WANT_{vF} and WANT_{H}, the standard for inclusion in the positive extension of WANT can be altered by manipulating the coarseness of the ordering over worlds. Such manipulations allow us to tease apart desires with differing intensities while keeping WANT as a relation between a proposition and an eventuality, in a similar fashion to how extension gap theories of adjectives like *tall* handle facts about gradability. The final question we must answer is how this raising and lowering of the standard for WANT can be integrated with the mereology of states of desire.

4.3.4 Integration with the mereology

The way in which the manipulation of fine-grainedness of orderings will be integrated into the mereology of attitude states will be by means of a natural language metaphysical principle relating world-orderings at various point-states. More specifically, the principle, which I refer to as *Downward Ordering Generation (DOG)*, states that world-orderings at lower and lower point-states get progressively more fine-grained. The formulation of this principle in (40) comes in two parts. (40a) imposes the requirement about fine-grainedness, while (40b) adds a requirement that the set of belief worlds does not change across altitudes.

(40) **Downward Ordering Generation:**
If $k_a \leq_K k_b$, and if $e_a = e/(t, k_a)$ and $e_b = e/(t, k_b)$ for some desire state e and moment $t \in \tau(e)$, then:
  a. $\precsim_{e_a}$ is at least as fine-grained as $\precsim_{e_b}$, and
  b. $\text{Dox}(e_a) = \text{Dox}(e_b)$.

Note that (40a) and (40b) only impose constraints on simultaneous point-states; the relationship between beliefs and desires at different times is unconstrained. Note in addition that since the at-least-as-fine-grained relation is reflexive (i.e., each ordering is at least as fine-grained as itself), $\precsim_{e_a}$ and $\precsim_{e_b}$ are permitted to be identical.

As an illustration, Figure 4.5 shows a potential structure for Ron's desire state that obeys DOG. At the highest altitudes, point-states have the coarsest world-ordering $\precsim_1$, and at the lowest altitudes, worlds are ordered as in $\precsim_3$, the finest-grained of the three orderings. Because r holds in all ideal worlds with respect to all three world-orderings, the state as a whole is a state of Ron wanting r. Since q holds in all $\precsim_2$- and $\precsim_3$-ideal worlds, but not all $\precsim_1$-ideal worlds, only the bottom two-thirds of this state is a state of Ron wanting q. And finally, just the bottom third of this state will be a state of Ron wanting p.

Assuming a structure like Figure 4.5 for Ron's desire state generates the correct predictions for both positive and comparative desire attributions. As far as positive desire attributions go, all of *Ron wants to learn Russian*, *Ron wants to visit Quebec*, and *Ron wants to eat peanuts* are rightly predicted to be true.

(41) $[\![\text{Ron wants to eat peanuts}]\!]^c = 1$ iff
$$\exists e[\text{Exp}(e, \text{ron}) \wedge \forall e' \in \textsc{pt}(e)[\text{WANT}(p)(e')]]$$

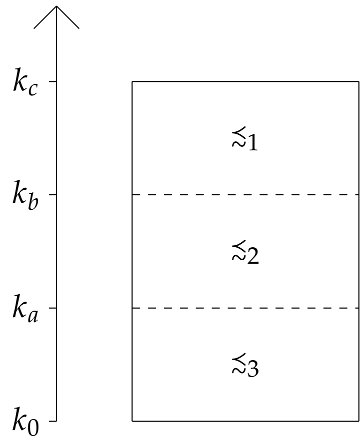

Figure 4.5: Ron's desire state

After all, for each of these three propositions, there is some state of Ron wanting that proposition.

Turning to *want* comparatives, following the discussion in the previous section, (43) is the predicted interpretation of (42) when the contextually determined measure function is set to μ_{int}.

(42) Ron wants to learn Russian more than he wants to visit Quebec.

(43) Assertion:

$$\exists d \exists e[\text{Exp}(e) = \text{ron} \wedge \forall e' \in \textsc{pt}(e)[\text{WANT}(r)(e')] \wedge \mu_{\text{int}}(e) \geq d \wedge d > \max(\{d' \mid \exists e''[\text{Exp}(e'') = \text{ron} \wedge \forall e''' \in \textsc{pt}(e'')[\text{WANT}(q)(e''')] \wedge \mu_{\text{int}}(e'') \geq d']\})]$$

Ron's largest state of wanting r is the whole desire state, and his largest state of wanting q is the lower two-thirds of his desire state. Thus, the assertion in (42) is true: there is a state of Ron wanting to learn Russian that exceeds in intensity (i.e., is larger than) any state of Ron wanting to visit Quebec. Naturally, the previously established monotonicity of intensity means that the presupposition is satisfied as well, so (42) is indeed true.

4.3.5 Summary

In this section, we have seen that cases where a single experiencer has desires with varying intensity can be captured without revising the core semantics and ontology adopted in this dissertation. The way this was done was by means of a natural language metaphysical principle of Downward Ordering Generation (DOG), which simulates the lowering of standards for WANT and integrates this lowering of standards into the part-whole structure of desire states.

4.4 Extension to *wish* and *regret*

So much for *want*. But what about *wish* and *regret*? Starting with *wish*, we can retain our quantification over point-states as before, but replace WANT with WISH, as in (44):

(44) $[\![\text{wish}]\!] = \lambda p \lambda e.\ \forall e' \in \textsc{pt}(e)[\text{WISH}(p)(e')]$

The task now is to define WISH. Recall the von Fintel- and Heim-style definitions for *wish* from Chapter 2 ((45a) and (45b), respectively):

(45) a. $\lambda p \lambda x \lambda w : \text{Dox}(x,w) \cap p = \varnothing \wedge \exists w', w'' \in \text{Dox}^+(x,w)[p(w') \wedge \neg p(w'')].$
$\forall w' \in \textsc{Best}(\text{Dox}^+(x,w), \precsim_{x,w})[p(w')]$

b. $\lambda p \lambda x \lambda w : \text{Dox}(x,w) \cap p = \varnothing \wedge \exists w', w'' \in \text{Dox}^+(x,w)[p(w') \wedge \neg p(w'')].$
$\forall w' \in \text{Dox}^+(x,w)[\text{Sim}_{w'}(\text{Dox}^+(x,w) \cap p) \prec_{x,w}$
$\text{Sim}_{w'}(\text{Dox}^+(x,w) - p)]$

$\text{Dox}^+(x,w)$ is the expanded modal domain used for *wish*: it is the set of worlds compatible with what x believes was previously possible. This is a proper superset of $\text{Dox}(x,w)$, the set of worlds that x believes are *curently* possible. Both definitions presuppose that p is compatible with, but not entailed by, the expanded modal domain (the diversity condition), and both presuppose that $-p$ is entailed by x's beliefs (the counterfactuality presupposition). In the assertive component, both look just like their *want* counterparts, excepting the replacement of Dox with Dox^+. By now it should be clear how WISH_{vF} and WISH_{H} are to be defined. This can be seen in (46); the presuppositions are again excluded, but translating them is again straightforward.

(46) a. $\text{WISH}_{\text{vF}} = \lambda p \lambda e.\ \forall w \in \textsc{Best}(\text{Dox}^+(e), \precsim_e)[p(w)]$

b. $\text{WISH}_{\text{H}} = \lambda p \lambda e.\ \forall w \in \text{Dox}^+(e)[\text{Sim}_w(\text{Dox}^+(e) \cap p) \prec_e$
$\text{Sim}_w(\text{Dox}^+(e) - p)]$

Using WISH_{vF} and translating the scenario from before, imagine that Ron never learns Russian, visits Quebec, or eats peanuts, and he is aware of this fact, so that each of the propositions p, q, and r satisfies the counterfactuality presupposition of *wish*. Furthermore, assume that the expanded set of worlds used for *wish* is as diverse in its range of possible outcomes as the set of belief worlds was in the *want* scenario above, and that the rankings of worlds are just like in Figure 4.5. In this scenario, (47a) and (47b) are predicted to be true for the same reason that (41) and (42) were before.

(47) a. Ron wishes he had {learned Russian/visited Quebec/eaten peanuts}.

b. Ron wishes he had learned Russian more than he wishes he had visited Quebec.

Similar facts hold for *regret*. In Chapter 2 it was posited that with respect to both presupposition and assertion, to regret p is to wish that $-p$. This means we can simply define $[\![\text{regret}]\!]$ as in (48):

(48) $[\![\text{regret}]\!] = \lambda p \lambda e.\ \forall e' \in \text{PT}(e)[\text{WISH}(-p)(e')]$

Of course, this means that in the scenario for *wish* discussed above, (49a) and (49b) come out as true for the same reason that (47a) and (47b) did.

(49) a. Ron regrets that he didn't {learn Russian/visit Quebec/eat peanuts}.
b. Ron regrets that he didn't learn Russian more than he regrets that he didn't visit Quebec.

4.5 Consequences and comparisons

Now that the proposal has been set forth in its entirety, we can explore some of its non-monotonicity-related semantic predictions, as well as how it stacks up when compared to some other treatments of *want* comparatives.

4.5.1 Some consequences

I will discuss five predictions made by my analysis of mental state verbal comparatives like *α wants p more than β wants q*: (I) an entailment to *α wants p*; (II) a *lack* of entailment to *β wants q*; (III) the possibility of comparison across kinds of mental states; (IV) the (limited) availability of non-intensity readings; and (V) the possibility that two experiencers can want exactly the same things, with exactly the same relative intensities, but with differing absolute intensities. I will provide evidence suggesting that the first four of these predictions are correct. The fifth prediction is difficult to test, but is by all appearances intuitively plausible.

4.5.1.1 The positive entailment

Given the semantics of verbal comparatives adopted in this dissertation, the truth of *α VP_1 more than β VP_2* is predicted to entail the truth of *α VP_1*. After all, the comparative essentially just adds another conjunct within the scope of the matrix clause's existential quantifier over events, with the result being that if the comparative is true, the non-comparative must be true as well. I will call this predicted entailment the *positive entailment*.

For non-intensity comparatives, the positive entailment is clear: if Dee ran more than Evan did, she has to have done at least some running. With transitive psychological verbs like *hate*, this seems to hold as well. For example, if Jack merely likes football, while Jill absolutely adores soccer, (50) is false (or odd), even though Jack's views on football are less positive than Jill's views on soccer.

(50) Jack hates football more than Jill hates soccer.

Similar facts can be shown for attitude verbs, at least in some cases. Suppose that Lana and Sterling are both happily alive, and both want to continue living. Furthermore, suppose that Sterling's desire not to die is stronger than Lana's. In spite of the fact that Lana's views on death are less negative than Sterling's, each of the examples in (51) is false (or odd):

(51) a. Lana wants to die today more than Sterling does.
b. Lana wishes she was dead more than Sterling does.
c. Lana regrets being alive more than Sterling does.

However, not all cases are quite so clear. For instance, recall Villalta's (2008) picnic scenario from Chapter 2: Sofía will bring either chocolate cake, apple pie, or ice cream to the picnic, and Victoria prefers chocolate cake to apple pie, which she in turn prefers to ice cream. In this scenario, (52a) is true, and (52b) seems to be acceptable and accurate, if perhaps slightly odd:

(52) a. Victoria doesn't want Sofía to bring apple pie.
b. ? Victoria wants Sofía to bring apple pie more than she wants her to bring ice cream.

The apparent compatibility of (52a) and (52b) suggests that the positive entailment may not actually hold, an urgent problem for the proposal at hand.

However, the positive entailment can be rescued. As Rubinstein (2012) notes, there seems to be more flexibility to the modal domain for *want* than has sometimes been acknowledged. On top of Rubinstein's own arguments for this claim, here is some additional evidence: it appears that the modal domain for *want* can be restricted by an *if* clause, in a manner similar to the restriction of modals in conditionals (cf. the discussion of conditionals in Chapter 2):

(53) If I become a zombie, I want you to shoot me.

(53) is not a claim about what my desires *would* be if I became a zombie—zombie-Pasternak would no doubt prefer not to be shot—but rather about what my *current* desires are when restricting my attention to just those belief worlds in which I become a zombie.

Given this possibility for overt restriction, it seems plausible to replace $\mathrm{Dox}(e)$ in $\mathrm{WANT}_{\mathrm{vF}}$ and $\mathrm{WANT}_{\mathrm{H}}$ with $\mathrm{Dox}(e) \cap \mathrm{Res}^c$, where Res^c is some contextually-determined domain restriction. In *want* conditionals like (53), Res^c is set to the worlds where the antecedent is true (e.g., worlds where I become a zombie). In cases like (52a) and (52b), however, there is no overt antecedent, and this restriction is filled contextually. What we can then say is the following. In (52a), Res^c imposes no impactful restriction on the modal domain. As a result, some worlds where Sofía brings chocolate cake are included in the modal domain, and Victoria thus wants Sofía to bring chocolate cake and does not want her to bring apple pie. In (52b), meanwhile, the restriction rules out those belief worlds of Victoria's in which Sofía brings cake, meaning the only options are apple pie and ice cream. In this case, Victoria wants apple pie and not ice cream, and she thus wants apple pie *more* than she wants ice cream. But notice that in the context for (52b), given the exclusion of cake worlds, Victoria *does* want apple pie. Thus, what we predict is that the apparent lack of positive entailment from (52b) to (52a) is really due to a shift in context, and not a lack of semantic entailment.

This explanation may seem a bit unwieldy, but some favorable evidence does arise when shifting our attention to *wish* and *regret*. Unlike *want*, *wish* and *regret* do not seem to allow restriction by overt *if* clauses:[3]

(54) a. # If I {became/had become} a zombie, I wish you had shot me.
b. # If I {became/had become} a zombie, I regret that you didn't shoot me.

Note that the absent reading of (54a), for example, is not inherently contradictory or absurd. It could be that I never actually became a zombie, but that if I *had*, then my current preferences would dictate that you shot me. The absence of this entirely plausible reading is therefore most likely to be of a grammatical origin. Perhaps *wish* and *regret*, unlike *want*, lack Res^c in their denotations, meaning that there is no contextual restriction of the modal domain $\mathrm{Dox}^+(e)$ (and thus nothing for the *if* clauses to contribute in (54)). But if there is no such contextual restriction, then by the previous argument cases like (52) should not arise for *wish* and *regret* like they do for *want*. Interestingly, this prediction is borne out: if what Sofía brought was in fact chocolate cake, the sentences in (55) are both false or odd:

(55) a. Victoria wishes Sofía had brought apple pie more than she wishes she had brought ice cream.
b. Victoria regrets that Sofía didn't bring apple pie more than she regrets that she didn't bring ice cream.

In summary, if cases like (52) are attributed to covert restriction of the modal domain, then the facts in (54) and (55) receive a unified explanation: *wish* and *regret*

[3] Note that these sentences do allow a distinct, seemingly epistemic reading. Thus, (54a) has an interpretation like (i):

(i) If it really is true that I became a zombie, then given this knowledge I wish you had shot me.

do not allow for such domain restriction, covert or otherwise. As for (51a), the plain falsehood of this sentence would presumably be due to the fact that the contextual restrictions required in order to make it true would be quite unnatural—we would have to devise contexts in which the options are reduced to either dying today or suffering some unspecified fate worse than death. But without significant contextual set up, this is an entirely arbitrary restriction, in contrast to the more salient and plausible restriction in (52b).

4.5.1.2 The positive non-entailment

While the truth of a *run* comparative like *Dee ran more than Evan did* entails that Dee ran at least a little bit (the positive entailment), it does not entail that Evan ran. This is illustrated by the non-contradictory nature of (56):

(56) Dee ran more than Evan did. In fact, Evan didn't run at all.

Our semantics will have to be slightly revised in order to actually derive this non-entailment. To see why, consider the truth conditions we currently predict for this sentence:

(57) $[\![\text{Dee ran more than Evan did}]\!]^c = 1$ iff:
$$\exists d \exists e[\text{Agt}(e) = \text{dee} \wedge \text{run}(e) \wedge \mu^c(e) \geq d \wedge d > \max(\{d' \mid \exists e'[\text{Agt}(e') = \text{evan} \wedge \text{run}(e') \wedge \mu^c(e') \geq d']\})]$$

If there is no event of Evan running, then the set of degrees fed to the max operator is the empty set, and so the output of the max operator is undefined. We can fix this by modifying our definition of the covert preposition FOR from (58a) to (58b), where $D_0(\mu^c)$ is the "zero degree" (i.e., minimum degree) of the scale that is the codomain of μ^c:

(58) a. $[\![\text{FOR}]\!]^c_{\text{original}} = \lambda d \lambda V \lambda e : \mu^c \text{ is monotonic on } \sqsubseteq^c \text{ in } V.\ V(e) \wedge \mu^c(e) \geq d$

b. $[\![\text{FOR}]\!]^c_{\text{revised}} = \lambda d \lambda V \lambda e : \mu^c \text{ is monotonic on } \sqsubseteq^c \text{ in } V.\ (V(e) \wedge \mu^c(e) \geq d) \vee d = D_0(\mu^c)$

When our new FOR is swapped in for the old one in our example, the result is the somewhat cumbersome (59):

(59)
$$\exists d[\exists e[(\text{Agt}(e) = \text{dee} \wedge \text{run}(e) \wedge \mu^c(e) \geq d) \vee d = D_0(\mu^c)] \wedge d > \max(\{d' \mid \exists e'[(\text{Agt}(e') = \text{evan} \wedge \text{run}(e') \wedge \mu^c(e') \geq d') \vee d' = D_0(\mu^c)]\})]$$

Luckily, this can be simplified. First, notice that the subformulas $d = D_0(\mu^c)$ and $d' = D_0(\mu^c)$ are both within the scope of existential event-quantifiers, but due to the lack of bound event variables the result is the same as if they had been outside of their scope. This means that (59) is equivalent to (60):

(60) $\exists d[(d = D_0(\mu^c) \vee \exists e[\text{Agt}(e) = \text{dee} \wedge \text{run}(e) \wedge \mu^c(e) \geq d]) \wedge d > \max(\{d' \mid d' = D_0(\mu^c) \vee \exists e'[\text{Agt}(e') = \text{evan} \wedge \text{run}(e') \wedge \mu^c(e') \geq d']\})]$

Second, observe that in (60), $D_0(\mu^c)$ will always be in the set of degrees fed to the max operator. But the existential quantifier over the variable d is only true if its witness is greater than the greatest degree in this set. This means that $D_0(\mu^c)$ cannot be this witness degree, so the first disjunct ($d = D_0(\mu^c)$) is redundant. In other words, in the same way that (61a) is equivalent to (61b), (60) is equivalent to (62):

(61) a. $\exists d[(d = d_k \vee \phi) \wedge d > d_k]$
b. $\exists d[\phi \wedge d > d_k]$

(62) $\exists d \exists e[\text{Agt}(e) = \text{dee} \wedge \text{run}(e) \wedge \mu^c(e) \geq d \wedge d > \max(\{d' \mid d' = D_0(\mu^c) \vee \exists e'[\text{Agt}(e') = \text{evan} \wedge \text{run}(e') \wedge \mu^c(e') \geq d']\})]$

For concreteness, let's say that μ^c is set to the distance measure function, meaning that *Dee ran more than Evan did* is being interpreted as comparing the relative distances that Dee and Evan ran. In this case, (62) is true iff there is a distance degree d such that there is an event of Dee running at least d distance, and such that d is greater than the greatest degree in the smallest set that contains (I) the degree of zero length, and (II) every degree d such that there is an event of Evan running whose distance is at least d. If Evan did run, then things pan out exactly as before, since any running event will cover a non-zero distance. If Evan didn't run, then the set of degrees fed to the max operator will be the singleton set containing the degree of zero distance. In this case, the sentence will be true if Dee ran a non-zero distance. Thus, in (62) we retain the positive entailment that Dee ran—outside of the set of degrees fed to the max operator, (62) is identical to (57)—but lose the requirement that Evan also ran.

Turning back to intensity comparatives, (63) illustrates that the lack of positive entailment in the comparison clause extends to intensity comparatives as well.

(63) a. Jack hates football more than Jill does. In fact, Jill doesn't hate football at all.
b. Isabella regrets losing the competition more than Miguel does. In fact, Miguel doesn't regret it at all.

Notice that extending the above analysis to psychological intensity requires positing the well-definedness of $D_0(\mu_{\text{int}})$, i.e., the existence of a zero degree on the intensity scale. This claim is independently supported by naturally occuring examples like the following:[4]

[4] While examples like (64) have something of a casual tone to them, this is presumably because there is no way to confirm such a precise measurement of psychological intensity. If we imagine ourselves in a

(64) a. Arsenal fans hate Cole anyway, but they would hate him twice as much if he moved to arch-rival Spurs.[5]

b. Long weekends are both a blessing and a curse. On the one hand, you don't have work on Monday. On the other, you now want to die twice as much Tuesday morning.[6]

As discussed extensively by Lassiter (2011a, 2017), a necessary condition for multiplicative modifiers like *twice as much* is that the scale for measurement have a meaningful zero degree. To illustrate, consider the contrast between (65a) and (65b):

(65) a. My limo is twice as long as yours is.

b. # My frying pan is twice as hot as yours is.

The important difference between length and temperature is that the former, but not the latter, has a meaningful zero value: whether measured in inches, centimeters, feet, or meters, zero length is zero length. Temperature, meanwhile, does not have a meaningful zero value, hence why zero degrees Celsius is far warmer than zero degrees Fahrenheit.[7] Thus, the current proposal's prediction that there is a meaningful zero degree of intensity is corroborated by the acceptability of (64).

4.5.1.3 Comparison of different types of mental states

Another prediction made by the proposal as it currently stands is that successful comparison of the intensity of mental states should not depend on the type of mental states being compared.[8] That is, while I have so far focused on comparing the intensity of one hatred state to that of another hatred state (for example), the fact that intensity constitutes a single dimension means it should be just as permissible to compare the intensity of a state of hatred to that of a state of love or respect. So far as I can tell, this seems to be the case:

(66) a. I will not skip Lola's wedding. I love her far more than I hate her fiancé.

b. Sandy fears her advisor, but she respects her even more.

c. I wanted to leave less than I would have hated myself for doing so.

world where psychological intensity can be measured by highly precise instruments, the informality of (64) goes away.

[5] https://www.theguardian.com/sport/2008/apr/06/football.comment1

[6] https://www.someecards.com/life/memes/13-memes-you-should-send-your-co-workers-after-a-long-weekend/

[7] This presupposes that the actual existence of an absolute zero temperature (zero degrees Kelvin) is somehow irrelevant as far as the natural language ontology is concerned.

[8] I am grateful to an anonymous reviewer of Pasternak (in revision) for raising this point.

However, if such universality of comparability does not hold, there is a relatively easy fix within the confines of the theory at hand: instead of a single dimension for mental state intensity, there may simply be multiple different intensity dimensions, with different types of mental state occupying different dimensions. If two types of state, such as hatred and love, are comparable, that indicates that they extend along the same intensity dimension, while two incomparable kinds of state extend along different intensity dimensions. This would then weaken the prediction of the proposal to mere transitivity of comparability: if, say, hatred and love are comparable, and love and respect are comparable, then hatred and respect must be comparable, as all three types of state must extend along the same dimension.

4.5.1.4 (Limited) availability of temporal comparison

Here is a possible complaint about the present proposal raised by an anonymous reviewer for Pasternak (in revision): in all of the examples of non-psychological verbal measurement constructions that arose in Chapter 3, temporal duration was available as the contextually determined measure function. As discussed there, this is in fact predicted by a monotonicity-based account of vP measurement adjuncts and verbal comparatives. But as the reviewer notes, for sentences like *Ann hated Bill a lot* or *Ann hated Bill more than Matt hated Jeff*, these temporal readings do not seem to arise, even though mental state homogeneity predicts temporal duration to be a monotonic measure of mental states. Thus, the proposal at hand has no means of differentiating between mental state verbal comparatives, which apparently lack a temporal reading, and other verbal comparatives for which it is available.

While I agree with the reviewer that there is a strong preference for mental state verbal comparatives to be interpreted in terms of psychological intensity, I think the claim that temporal duration readings are categorically unavailable is too strong. For instance, imagine that some psychologists are running a study in which subjects are sent into house parties, where they witness various pleasant and unpleasant scripted events. The subjects are given a remote with a single button, and told to press down the button whenever they decide they want to leave, and let go of the button whenever they stop wanting to leave. The remote does not provide any way of indicating how intense their desire to leave is. In this scenario, it is perfectly reasonable for a scientist looking at two of the subjects' recorded results to say the following:

(67) Over the course of the evening, Subject A wanted to leave more than Subject B did.

In this context, the scientist can be right even if it turns out that Subject B's desire to leave was more intense. This suggests that an intensity reading is not the only reading available, and that a temporal reading is indeed possible, albeit dispreferred.[9]

But if temporal readings are available for these comparatives, why are they so hard to get? It seems like the best explanation is simply that for the most part, the intensity of someone's hatred, desire, love, or regret is more practically relevant than its temporal duration. If this explanation is correct, then we might expect that mental state verbs are not the only ones showing such a strong dispreference for temporal readings in the comparative. This is, in fact, the case. For instance, we observed in Chapter 3 that when measuring events of acceleration, both temporal duration and the object's change of speed are monotonic: an event of acceleration will have a greater duration and involve a greater net change of speed than any of its proper parts. But the sentences in (68) have a very strong preference for using change in speed over temporal duration as the contextually determined measure function. Hence, if Nell's car only had a slight change in speed over a very long time, while my car had a large change in speed over a very short time, both (68a) and (68b) are interpreted as clearly false, except with a great deal of contextual work.

(68) a. Nell's car accelerated a lot.
b. Nell's car accelerated more than mine did.

Thus, the dispreference for a temporal duration reading of verbal comparatives with mental state verbs can be chalked up to pragmatic preferences for some grammatically available readings over others, rather than a bona fide grammatical distinction.

4.5.1.5 Same rankings, different strengths

Finally, an interesting side effect of this proposal is that we predict it to be possible for two experiencers to want all of the same things, and to rank their desires in exactly the same way, but for the intensities of their desires to vary. An illustration of such a situation can be seen in Figure 4.6. In this scenario, both Ron and Rhonda want r more than q, which in turn they want more than p. But Ron's desire for p is more intense than Rhonda's is, his desire for q is exactly as intense as hers is, and his desire for r is less intense than hers is.

It is difficult—at least difficult to me—to find compelling linguistic data that support or contradict this prediction. That being said, at least on an intuitive level it seems plausible that two people could be in exact agreement about what makes one world more desirable than another, while at the same time feeling those desires with differing intensities. I must, however, leave a more thorough exploration of this possibility for future work.

[9] See Wellwood 2014: 216–217 for similar discussion of frequency readings with *want* comparatives.

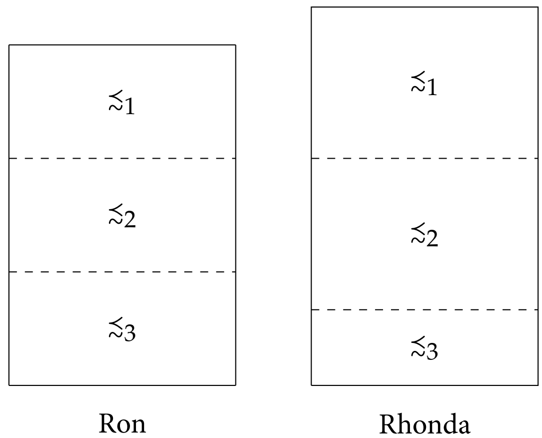

Figure 4.6: Ron and Rhonda's desire states

4.5.2 Some comparisons

We have looked at a few semantic predictions made by the current proposal (besides monotonicity), as well as evidence suggesting that these predictions are indeed correct. Next, we will look at two other prominent theories of *want* comparatives—those of Villalta (2008) on the one hand, and decision theorists like Levinson (2003) and Lassiter (2011a,b) on the other—and see how they stack up. We will see how each of these fares with respect to the facts cited above, and take note of those areas in which they fall short. I will not dive too deeply here into the technical details of each of these theories, as just an informal description of the content of the proposals will suffice.

4.5.2.1 Villalta 2008

As a reminder, according to Villalta (2008), for α to want p more than q is for α to in some sense rank p worlds as better than q worlds. (This is an idea inherited from Kratzer (1981a, 1991, 2012), who uses the same idea in the modal realm to describe notions like *better possibility*; see Katz et al. 2012 for an extension and revision of Kratzer's theory.) For α to want p *simpliciter* is for α to want p more than (s)he wants any of the focus alternatives to p.

A preliminary problem: Comparison across experiencers Before discussing where Villalta makes right and wrong predictions with respect to the facts discussed above, there is an additional, noteworthy problem that her theory faces. Villalta defines what it means for α to want p more than q. But what, on Villalta's analysis, does it mean for α to want p more than β wants q? This is tied to a more general problem, discussed in detail by Pasternak (2016): theories of intensional comparatives based on the direct comparison of idealness of worlds have no clear way to handle cases where comparisons are made across two distinct world-orderings. Pasternak discusses this problem in the context of theories of gradable modal adjectives like *important*, but the point applies just as well to *want* comparatives where there are two experiencers involved.

The only way for Villalta to allow such comparisons would be to have each experiencer's desire rankings mapped onto some independent scale—say, the set of non-negative real numbers—in a way that preserves the relative rankings of preferences for each experiencer, so that if α prefers p to q, then for α, p is mapped to a higher degree than q is. But this won't do, because Villalta's theory provides too little information about the relative desirability of propositions: she tells us what it means for α to want p more than q, but not what it means, say, for α to want *p significantly* more than q, or *slightly* more, or *twice as much*. In other words, she provides rankings, but she provides *only* rankings—adopting a term from measurement theory, she treats desire as an *ordinal scale*. But this means that beyond the preservation of relative rankings, how propositions are mapped to degrees is entirely arbitrary: p and q could equally well be mapped to 2 and 1, or to 100 and .001. This in turn means that the relationship between the mappings of α's desires and β's would be, for all intents and purposes, meaningless, since these mappings are not properly anchored.

Note that this is not a problem for my own proposal, for which the intensities of two different states are comparable in much the same way that their temporal durations are, regardless of whether those states have the same or different experiencers. This is also not a problem for decision-theoretic proposals, for which desires are anchored to the scale of expected utility from the get-go, rather than by means of a *post hoc* mapping. As a result, the relationship between two experiencers' desires on that scale is non-arbitrary and meaningful.

Positive entailments Moving on to those entailments discussed in the previous subsection, Villalta's analysis predicts that *α wants p more than q* entails neither that α wants p, nor that α wants q. After all, even if p and q are both horrible options, it is still possible for p to be a better option than q. Notice that what matters here is Villalta's semantics for *α wants p more than q*, and not so much her semantics for *α wants p*: on any analysis of which I am aware, α ranking p worlds as better than q worlds does not entail that α wants p or q.

As discussed above, the lack of entailment that α wants q is accurate, and also predicted by my own proposal. As for the lack of entailment that α wants p, this seems to work well for her picnic scenario (cf. (52)), but it doesn't account for those cases where we do seem to get a genuine positive entailment, whether that be cases with a less favorable context as in (51), or cases with *wish* and *regret* as in (55).

Comparison of different types of mental states In (66c), we saw an example where the intensity of a state of desire was felicitously compared to the intensity of a state of hatred. While my own theory is fully compatible with this possibility, it is entirely unclear how Villalta's proposal can account for such facts. It seems as though in order for her to get an analysis of (66c) off the ground, *hate myself* must be analyzed as somehow quantifying over an ordered set of possible worlds. But it seems odd to claim that *hate myself* has such an obtuse denotation, and one would have to answer basic questions like what the modal domain is, what the basis for world-ordering is, and how we somehow extract a proposition from the simple object DP *myself*. This semantics would also have to be shown to be superior to the clearly simpler alternative proposal that *hate* is just a transitive verb, along the lines discussed above. Moreover, to the extent that comparison across types of mental state is widespread or universal, one runs the risk of having to do the same exercise for all mental state verbs: what, for example, are the world-orderings and propositional arguments for verbs like *admire* and *respect*? And depending on these results, we are also likely to return again to the problem of comparisons across orderings. In short, then, it seems as though Villalta's proposal would at best require a great deal of stipulation in order to allow for an analysis of sentences like (66c).

Temporal comparison A simple revision to Villalta's theory—namely, one that introduces events into the mix (e.g., Anand & Hacquard 2008)—would allow for temporal comparisons like that seen in (67). However, she would presumably predict that a temporal comparison and an intensity comparison would involve two distinct syntactic structures. The reason for this is that Villalta adopts the lexical gradability hypothesis (LGH) discussed in Chapter 3. Thus, whereas an intensity comparative would make use of the lexical gradability of *want*, thereby having no need of FOR, a temporal comparative would require FOR in order to introduce degrees of temporal duration. While I have argued against LGH in Chapter 3, Villalta's adoption of LGH should not count as a point against her relative to the specific predictions discussed in this section, as it is indeed the case that Villalta's theory allows for temporal comparisons, at least with minimal, reasonable revisions.

4.5.2.2 Decision-theoretic proposals

Next up are the decision theorists. In Chapter 2, I discussed two possibilities for the semantics of α *wants* p: one in which it means that the expected utility of p ($EU_{\alpha,w}(p)$) exceeds some threshold n, and another in which it exceeds the expected utility of $-p$. I will call these the *threshold* and *differential* definitions of *want*, respectively.

As for *α wants p more than β wants q*, there are at least three possible definitions that are worth considering. The first is simply that the expected utility that α assigns to p is greater than the expected utility that β assigns to q, i.e., $\mathrm{EU}_{\alpha,w}(p) > \mathrm{EU}_{\beta,w}(q)$ (cf. Lassiter 2011a). On this definition, the degree to which one wants p is simply the expected utility of p. Another possiblility is that the degree to which one wants p is the difference between the expected value of p and the expected value of $-p$. I will call this the *expected utility differential* (EUD):

(69) $\mathrm{EUD}_{\alpha,w}(p) \equiv \mathrm{EU}_{\alpha,w}(p) - \mathrm{EU}_{\alpha,w}(-p)$

On such an account, instead of comparing expected utilities, a *want* comparative would be comparing expected utility differentials (see, e.g., Levinson 2003). A third possibility is what I will call *non-negative expected utility differential* (EUD^+), which is equivalent to the EUD when the value is positive, but is simply equal to zero when the EUD is zero or less. In other words, EUD^+ returns 0 iff the expected utility of p does not exceed $-p$, and otherwise returns a positive value.

Positive entailments There are currently six combinations of proposals for positive and comparative *want* on the table, or more if one considers potential variation in the threshold n. So far as I can tell, all of these combinations correctly predict a lack of positive entailment in the comparison clause. As for the matrix clause, I will not go through all of them to see if they predict a positive entailment. Instead, I will simply show that at least one plausible combination makes the wrong predictions, and at least one makes the right ones.

One combination that makes the wrong predictions is a threshold semantics for *want*, in conjunction with a semantics for *want* comparatives based on direct comparison of expected utility. Say, for example, that the threshold n is set to zero, so that to want p is to assign it a positive expected utility. Now suppose that α assigns p an expected utility of -5, and q an expected utility of -10. In this case, α is predicted to want p more than q, since p's expected utility exceeds that of q. But α is not predicted to want p *simpliciter*, since p's expected utility is less than zero. Thus, no positive entailment in the matrix clause. This leaves us in a position similar to Villalta (2008): we make plausible predictions for cases like the picnic scenario, but not for those cases where positive entailments do seem to arise.

A combination that makes the *right* predictions is a differential semantics for *want*, conjoined with a semantics of *want* comparatives based on EUD^+. In this case, if *α wants p more than β wants q* is true, then $\mathrm{EUD}^+_{\alpha,w}(p) > \mathrm{EUD}^+_{\beta,w}(q)$. But since zero is the minimum value in the range of EUD^+, this entails that $\mathrm{EUD}^+_{\alpha,w}(p) > 0$. Based on the definition of EUD^+, this entails that $\mathrm{EU}_{\alpha,w}(p) - \mathrm{EU}_{\alpha,w}(-p) > 0$. This in turn entails that $\mathrm{EU}_{\alpha,w}(p) > \mathrm{EU}_{\alpha,w}(-p)$, which is exactly what the differential semantics for positive *want* predicts is equivalent to α wanting p (in w). Hence, we successfuly generate the positive entailment. (Cases with an apparent lack of positive entailment, such as the picnic scenario, can then be handled in a manner similar to my own theory.)

We therefore see that the right combination of positive entailment in the matrix clause and lack of postiive entailment in the comparison clause is predicted by some, but not all, decision-theoretic proposals. Thus, this entailment pattern is not so much a point for or against decision-theoretic proposals as a group, but is instead evidence for or against individual theories within this spectrum.

Comparison of different types of mental states As far as comparison of different types of mental states, decision-theoretic proposals seem to face a conundrum similar to Villalta: expected utilities or (non-negative) expected utility differentials may work for *want*, but what about *hate* or *admire*? How can Ann's hatred of Bill be framed in terms of the expected utility of a proposition, and how can this be done compositionally, given that there is no obvious propositional argument in *Ann hates Bill*?

Furthermore, to the extent that it makes sense to talk about the expected utility involved in hatred, it seems like hating something more involves assigning a *lower* expected utility to some proposition or other, in contrast to loving or wanting something more, which would involve the assignment of a *higher* expected utility value. Thus, if we think of things in terms of expected utility, we would expect *hate* and *love*, for example, to have opposing polarities, similar to the relationship between *short* and *tall*. But then the examples in (66) should be just as bad as other examples of comparison with cross-polar adjectives, such as the sentences in (70), taken from Kennedy (2001: 48) (judgments Kennedy's):

(70) a. ? The ski poles are shorter than the box is wide.
b. ? After she swallowed the potion, Alice discovered that she was two inches shorter than the doorway was high.
c. ? The hole is a foot shallower than the gravedigger is tall.

We therefore see that decision-theoretic proposals, like Villalta's (2008), have a hard time accounting for comparisons of intensity across different kinds of mental states.

Temporal comparison Much like with Villalta's theory, decision-theoretic proposals can easily account for temporal comparisons with the proper insertion of an event argument. But, again like Villalta's theory, this revised theory would still be one that adopts LGH, entailing a predicted structural distinction between temporal and intensity comparisons, since the former use FOR to generate a degree of temporal duration, while degrees of desire simply come from the semantics of $[\![\text{want}]\!]$ itself.

4.5.3 Summary

In this section, we have seen that my own proposal makes a variety of seemingly correct predictions about the semantics of attitude comparatives, independently of the mereological requirement of monotonicity. We also saw that other work on the semantics of want comparatives had more mixed results. Villalta's (2008) semantics struggled with *want* comparatives with two different experiencers, as it provided no mechanism for making comparisons across distinct world-orderings. It also could not account for cases with positive entailments, nor could it tackle comparisons involving the intensity of two different types of mental state. With decision-theoretic proposals, the presence or absence of predicted positive entailments depended on which theory of *want* and *want* comparatives one adopted: we saw one combination that got the wrong results, and one that got the right results. We also saw that the decision-theoretic proposal struggled to account for comparisons across kinds of mental states, since the proposed semantics for *want* could not plausibly generalize to other psychological verbs like *hate* and *admire*. All of this is, of course, in addition to the arguments in favor of monotonicity discussed in the previous chapter, which have not been addressed in any previous work on attitude comparatives.

4.6 Conclusion

Having argued at length in the previous chapter that intensity is a monotonic measure of mental states, I have provided in this chapter a basic natural language metaphysics of intensity that allows for this monotonicity. After showing that this natural language metaphysics works for simpler cases like *hate*, I illustrated a means of integrating ordering and quantification over worlds into the part-whole structure of attitude states, so that attitude comparatives could also enjoy the benefits of this theory. I further showed that my proposal makes certain independent semantic predictions that by all appearances are accurate, and that are not necessarily made by other theories of desire ascriptions and *want* comparatives. In concluding, I will note two areas that warrant further investigation in the short-term, outside of those questions that have already arisen over the course of this chapter.

First, one might reasonably ask whether altitudes deserve an independent existence in the natural language ontology, as asumed here, or whether a metaphysics can be established that generates the same results as the proposal in this chapter without the stipulation of a distinct dimension reserved exclusively for psychological intensity. In the form adopted in this chapter, mental states extend partly into their own little world, existing in a dimension that other kinds of objects seem not to occupy. The spatiotemporal dimensions are not nearly as restricted in this regard, containing a variety of entities and eventualities that cannot clearly be lumped under a single natural kind. Running events, states of happiness, and pieces of paper all have some relationship to time and space, so it is *prima facie* odd that there should be a dimension that only contains mental states. It would therefore be ideal to either independently justify the existence of such a separate dimension within the natural language metaphysics, or integrate psychological intensity into the metaphysics in a way that diminishes or eliminates its unique status.[10]

Second, I have had nothing to say about desires that are, or are believed by the experiencer to be, mutually incompatible. But mutually incompatible desires arise all the time. I can want to spend my summer relaxing on the beach, while also wanting to spend my summer catching up on research. On a certain level, the possibility for incompatible desires to arise is quite easy for a neo-Davidsonian theory like the one in this dissertation to account for. Since desire states are existentially quantified over in desire ascriptions, I can simply have two distinct desire states: one in which I spend my summer at the beach in all ideal worlds, and another in which I work all summer in all ideal worlds. If this is the case, then of course there is a state of me wanting to be at the beach, and there is a state of me wanting to work, so I want both. But such an account brings with it a host of questions, including basic ones about the ontological origins of and relationship between distinct desire states with the same experiencer, as well as the nature of the resolution of conflicting desires in the establishment of one's overall, "all things considered" desires. Furthermore, each of these questions must be paired with the methodologically prior question of whether the given issue is linguistically relevant at all: which folk-psychological beliefs about conflicting desires are reflected in the semantics of natural languages, and which are simply extralinguistic facts about how humans conceive of others' minds, as well as their own? Such questions are empirical in nature, and require further teasing apart of the linguistic tools at speakers' disposal in discussing desire and other mental states.

[10] Francez & Koontz-Garboden (2017) provide an ontology for the domains of nouns like *courage* and *hunger* that is similar to my proposed ontology of mental states, but that makes no use of altitudes. Instead, they propose that these domains simply come with a "size" ordering that respects part-whole relations, much like monotonic measure functions. So far as I can tell, this simpler ontology would work fine for mental state verbs like *hate*. However, it is unclear how a constraint like DOG could then be defined, except by using some alternative means to essentially reconstruct a notion of altitudes. I leave further exploration of this matter for future work.

Appendix: Implementation in premise semantics

Recall from Chapter 2 that in Kratzer's (1981a, 1991, 2012) theory of modality, as well as von Fintel's (1999) Kratzerian theory of $[\![\text{want}]\!]$, world-orderings are generated by an ordering source g. The ordering source generates a set of propositions—$g(w)$ for modals and $g(x,w)$ for *want* (due to the difference in argument structure)—which are themselves used to generate an ordering $\precsim_{g(w)}$ or $\precsim_{g(x,w)}$ over worlds as in (71), where Q is a set of propositions (cf. Lewis 1981):

(71) $w_1 \precsim_Q w_2$ iff $\{p \in Q \mid p(w_1)\} \supseteq \{p \in Q \mid p(w_2)\}$

In my own use of orderings over possible worlds, I have cast aside this use of sets of propositions, instead simply positing that each point-state e has a bouletic world-ordering $\precsim_e$ associated with it.

In this appendix, I will provide an implementation of my proposal that *does* use an ordering source g to generate world-orderings. Of course, switching from talking about world-orderings to talking about premise sets has long been known—at least since the work of Lewis (1981)—not to be problematic. But where I have placed a new burden on my shoulders is in defining DOG in terms of a novel and formally explicit notion of fine-grainedness. With this in mind, the goal of this appendix is not just to restate the denotation of *want* using a premise semantics, but to do so in a way that allows for a simple restatement of DOG in terms of ordering sources. To be clear, I do not offer this reframing because I believe that it is necessary, or even because I believe it to be especially desirable, but because for some, implementability within a premise semantics is an in-and-of-itself desideratum. Thus, let this appendix serve as a reassurance that the proposal in this chapter is not at odds with a premise-semantic denotation for *want*. I will start by articulating the idea behind the analysis, and will then proceed to some relevant formal proofs.

The idea

Here is the idea in a nutshell. Not all premises used to order worlds need carry equal weight: sometimes, some priorities outrank others. Thus, instead of confining ourselves to a flat premise set in which all propositions are weighted equally, we can instead have an *ordered* (i.e., ranked) premise set, in which lower-ranked premises can serve as tiebreakers for higher-ranked ones (cf. von Fintel & Iatridou 2008, Katz et al. 2012, Portner & Rubinstein 2016, Pasternak 2016, Reisinger 2016).[11] Instead of talking about DOG in terms of fine-grainedness, then, we can say that as we get to lower and lower altitudes, more and more premises are added to the bottom of our

[11] As pointed out by Reisinger (2016), this is formally similar to how constraints are used to eliminate potential output forms in Optimality Theory.

ordered premise set, serving to break previously established ties. And if we appropriately define our ordered premise sets, as well as the way these are used to generate orderings over worlds, what we will get is a semantics of *want* and a statement of DOG that generate the exact same range of possibilities as the ones proposed earlier in terms of world-orderings. That is, we are left with a semantics and ontology exactly as powerful as that offered without recourse to premise semantics.

With this in mind, let an **ordered premise set** (**OPS**) be an ordered pair $\langle Q, \leq \rangle$, where Q is a set of propositions, and $\leq$ is a weak (i.e., reflexive, transitive, connected) ordering over Q: two distinct propositions can be equally ranked, but they cannot be incomparably ranked. If $p \leq q$, we will say that q is *at least as high-priority* as p, and if $p < q$, we will say that q is *(strictly) higher-priority* than p. In addition to OPSs, we will need recourse to what I will call **flat upper cuts** and **ordered upper cuts**, as defined in (72).

(72) a. Set Q_2 of propositions is a **flat upper cut** of OPS $\langle Q_1, \leq_1 \rangle$ iff $Q_2 \subseteq Q_1$, and for all $p \in Q_2$ and all $q \in Q_1$ such that $p \leq_1 q$, $q \in Q_2$.

b. OPS $\langle Q_2, \leq_2 \rangle$ is an **ordered upper cut** of $\langle Q_1, \leq_1 \rangle$ iff Q_2 is a flat upper cut of $\langle Q_1, \leq_1 \rangle$, and for all $p, q \in Q_2$, $p \leq_2 q$ iff $p \leq_1 q$.

In short, a flat upper cut of an OPS $\langle Q, \leq \rangle$ is a set of propositions making up the "top portion" of an OPS: if a proposition is in a flat upper cut, any other proposition ranked as high as it is also in that flat upper cut. An ordered upper cut is itself an OPS, and preserves the rankings of propositions from the original OPS.

We can now define how OPSs are used to order worlds, with lower-priority propositions serving as tie-breakers for higher-priority propositions. Given that $\precsim_Q$ is the traditional Lewis-Kratzer method of using sets of propositions to order worlds, we define $\precsim^*_{\langle Q, \leq \rangle}$—the ordering over worlds generated by OPS $\langle Q, \leq \rangle$—as in (73):

(73) $w_1 \precsim^*_{\langle Q, \leq \rangle} w_2$ iff $w_1 \sim_Q w_2$ or there is a flat upper cut Q' of $\langle Q, \leq \rangle$ such that $w_1 \prec_{Q'} w_2$.

If, as is typical, we define $\sim^*_{\langle Q, \leq \rangle}$ and $\prec^*_{\langle Q, \leq \rangle}$ in terms of $\precsim^*_{\langle Q, \leq \rangle}$, it it is provable that the definition in (73) entails the following:

(74) a. $w_1 \prec^*_{\langle Q, \leq \rangle} w_2$ iff there is a flat upper cut Q' of $\langle Q, \leq \rangle$ such that $w_1 \prec_{Q'} w_2$.

b. $w_1 \sim^*_{\langle Q, \leq \rangle} w_2$ iff $w_1 \sim_Q w_2$.

These facts are proved in **Theorem 3** and **Theorem 4** below, respectively.

As an example, let's consider again the case of Ron's ordering of states of affairs, including the same set of worlds as before (w_{pqr}, w_{qr}, etc.). Let $Q_3 = \{p, q, r\}$, and let $\leq_3$ be such that $p <_3 q <_3 r$. (The choice of numbering will become less opaque shortly.) Intuitively, Q_3 is the set of desiderata used to rank worlds in the ordering in (29) (repeated below), while $\leq_3$ is the ranking of these desiderata: p is a tiebreaker for q, which is a tiebreaker for r.

(29) $$\underbrace{\underbrace{\underbrace{w_{pqr}}_{p} < \underbrace{w_{qr}}_{-p}}_{q} < \underbrace{\underbrace{w_{pr}}_{p} < \underbrace{w_{r}}_{-p}}_{-q}}_{r} < \underbrace{\underbrace{\underbrace{w_{pq}}_{p} < \underbrace{w_{q}}_{-p}}_{q} < \underbrace{\underbrace{w_{p}}_{p} < \underbrace{w_{\varnothing}}_{-p}}_{-q}}_{-r}$$

And as it turns out, the world-ordering $\precsim^*_{\langle Q_3,\leq_3\rangle}$ is identical to (29). By (74b), two worlds will be equally ideal iff they have the same status with respect to all three propositions, since they have to be equal with respect to the whole set Q_3. Since the toy example was constructed so that none of the worlds would have the same status with respect to all three propositions, this means that each world is equivalent only to itself. To determine the strict orderings, let's look at the flat upper cuts of $\langle Q_3,\leq_3\rangle$ one by one. The smallest non-empty flat upper cut is simply $\{r\}$, since r is the highest-priority premise. By (74a), any world w_1 better than world w_2 with respect to $\precsim_{\{r\}}$ is better than w_2 with respect to $\precsim^*_{\langle Q_3,\leq_3\rangle}$. Thus, all worlds in which r is true are better than worlds in which r is false. The next smallest flat upper cut is $\{q,r\}$. Given those strict rankings already effected by the first flat upper cut, all this flat upper cut will do is rank a given q world as better than a given $-q$ world *if they have the same status with respect to* r. That is, q serves as a tiebreaker for r. Similarly for the largest flat upper cut, $\{p,q,r\}$, which only breaks ties between worlds that have the same status with respect to both q and r. The result is that we get the exact ordering over worlds seen in (29).

Now consider what happens when we look at $\langle Q_2,\leq_2\rangle$, where $Q_2 = \{q,r\}$ and $q <_2 r$. First, notice that $\langle Q_2,\leq_2\rangle$ is an ordered upper cut of $\langle Q_3,\leq_3\rangle$: Q_2 is a flat upper cut of $\langle Q_3,\leq_3\rangle$, and $\leq_2$ preserves the rankings of $\leq_3$. In other words, $\langle Q_2,\leq_2\rangle$ is the result of removing p from the bottom of $\langle Q_3,\leq_3\rangle$. As for the ordering $\precsim^*_{\langle Q_2,\leq_2\rangle}$, $w_1 \sim^*_{\langle Q_2,\leq_2\rangle} w_2$ iff w_1 and w_2 have the same status with respect to both q and r. Notice that because of the absence of p in Q_2, there will now be non-trivial equivalences: w_{pqr} and w_{qr}, for example, will be equally ideal, since q and r are true in both. Turning to strict rankings, the flat upper cuts will be a proper subset of the ones seen before (namely, $\{r\}$ and $\{q,r\}$). As a result, r worlds will again be ranked as better than $-r$ worlds, with q serving as a tie breaker; the difference with $\langle Q_3,\leq_3\rangle$ will then be in the absence of an additional p tiebreaker.

Finally, we can go one more step by removing q from the bottom of $\langle Q_2,\leq_2\rangle$, leading to $\langle Q_1,\leq_1\rangle$, where Q_1 is simply $\{r\}$. As the reader can confirm, (I) $\langle Q_1,\leq_1\rangle$ is an ordered upper cut of both $\langle Q_2,\leq_2\rangle$ and $\langle Q_3,\leq_3\rangle$, and (II) $\precsim^*_{\langle Q_1,\leq_1\rangle}$ orders worlds only with respect to whether or not r holds in them.

If I have done things right, all of this talk will sound familiar. In Section 4.3.3 I talked about generating the ordering (29) in steps, as in (39), repeated below:

(39) STEP 1: $\underbrace{w_{pqr}, w_{qr}, w_{pr}, w_r}_{r} \prec_1 \underbrace{w_{pq}, w_q, w_p, w_\varnothing}_{-r}$

STEP 2: $\underbrace{\underbrace{w_{pqr}, w_{qr}}_{q} \prec_2 \underbrace{w_{pr}, w_r}_{-q}}_{r} \prec_2 \underbrace{\underbrace{w_{pq}, w_q}_{q} \prec_2 \underbrace{w_p, w_\varnothing}_{-q}}_{-r}$

STEP 3: $\precsim_3$ is as in (29)

But the orderings $\precsim_1$, $\precsim_2$, and $\precsim_3$ in (39) are identical to the orderings $\precsim^*_{\langle Q_1, \leq_1\rangle}$, $\precsim^*_{\langle Q_2, \leq_2\rangle}$, and $\precsim^*_{\langle Q_3, \leq_3\rangle}$. This is perhaps unsurprising: $\precsim_2$ was framed from the get-go as the addition of a q tiebreaker to $\precsim_1$, and the transition from $\langle Q_1, \leq_1\rangle$ to $\langle Q_2, \leq_2\rangle$ is a direct encoding of this idea, with q being added to the bottom of an ordered premise set.

Given the relationship between the ordering sequence (39) and DOG, and given that DOG is defined in terms of fine-grainedness, one might now ask what the relationship is between fine-grainedness of orderings and ordered upper cuts of OPSs. As it turns out, it is a close one. As I prove in **Theorem 7**, if $\langle Q_a, \leq_a\rangle$ is an ordered upper cut of $\langle Q_b, \leq_b\rangle$, then $\precsim^*_{\langle Q_b, \leq_b\rangle}$ is at least as fine-grained as $\precsim^*_{\langle Q_a, \leq_a\rangle}$. Moreover, as proved in **Theorem 8**, for any set of world-orderings linearly orderable by fine-grainedness, there is an OPS such that for every ordering in the set, there is an ordered upper cut of the OPS with an equivalent world-ordering by $\precsim^*$. The conjunction of these two theorems thus entails that talk of orderings with increasing fine-grainedness and talk of OPSs with added propositions on the bottom are completely interchangeable.

With this in mind, say that for each point-state e, $g(e)$ returns an OPS. We can then define WANT_{vF} and WANT_{H} as in (75a) and (75b), respectively:

(75) a. $\text{WANT}_{\text{vF}} = \lambda p \lambda e.\ \forall w \in \textsc{Best}(\text{Dox}(e), \precsim^*_{g(e)})[p(w)]$

b. $\text{WANT}_{\text{H}} =$

$$\lambda p \lambda e.\ \forall w \in \text{Dox}(e)[\text{Sim}_w(\text{Dox}(e) \cap p) \prec^*_{g(e)} \text{Sim}_w(\text{Dox}(e) - p)]$$

DOG can then be redefined as in (76), so that Ron's desire state will have a structure like Figure 4.7:

(76) **DOWNWARD ORDERING GENERATION (PREMISE-SEMANTIC VERSION):** If $k_a \leq_K k_b$, and if $e_a = e/(t, k_a)$ and $e_b = e/(t, k_b)$ for some desire state e and moment $t \in \tau(e)$, then:

a. $g(e_b)$ is an ordered upper cut of $g(e_a)$, and

b. $\text{Dox}(e_a) = \text{Dox}(e_b)$.

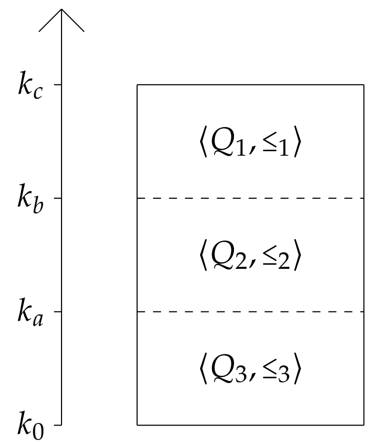

Figure 4.7: Ron's desire state

In summary, because of the tight relationship between OPSs and fine-grainedness of orderings, DOG can be restated in a premise semantics in a more or less intuitive way: at lower and lower point-states, more and more propositions are added to the bottom of the OPS, making the world-ordering more fine-grained. The result, as desired, is exactly as powerful as the theory was when stated in terms of fine-grainedness of orderings. All that's left, then, is to prove the relevant theorems.

Proofs

Theorem 1: If Q_a and Q_b are flat upper cuts of $\langle Q, \leq \rangle$, then $Q_a \subseteq Q_b$ or $Q_b \subseteq Q_a$.

Proof by reductio. Assume that Q_a and Q_b are flat upper cuts of $\langle Q, \leq \rangle$, and that $Q_a \nsubseteq Q_b$ and $Q_b \nsubseteq Q_a$. This requires that there be at least one $p \in Q_a - Q_b$ and at least one $q \in Q_b - Q_a$. Since $p \in Q_a$ and $q \in Q_b$, by the definition of a flat upper cut $p, q \in Q$. By the defintion of an OPS, $p \leq q$ or $q \leq p$. But if $p \leq q$, then there exists a proposition $r \in Q$ such that $p \leq r$ and $r \notin Q_a$ (namely q), so Q_a is not a flat upper cut of $\langle Q, \leq \rangle$. If $q \leq p$, then there exists a proposition $r \in Q$ such that $q \leq r$ and $r \notin Q_b$ (namely p), so Q_b is not a flat upper cut of $\langle Q, \leq \rangle$. So either Q_a or Q_b is not a flat upper cut of $\langle Q, \leq \rangle$, a contradiction. □

. . . .

Theorem 2: For any ordered premise set $\langle Q, \leq \rangle$, $\precsim^{*}_{\langle Q, \leq \rangle}$ is a preorder.

Reflexivity: By the definition of $\precsim^*$, for all $u, v \in W$, if $u \sim_Q v$, then $u \precsim^*_{\langle Q,\leq\rangle} v$. For all worlds u and sets Q of propositions, $u \sim_Q u$, since $\{p \in Q \mid p(u)\} = \{p \in Q \mid p(u)\}$. Therefore, for every world u and OPS $\langle Q, \leq\rangle$, $u \precsim^*_{\langle Q,\leq\rangle} u$.

Transitivity: I will prove that if $u \precsim^*_{\langle Q,\leq\rangle} v$, and $v \precsim^*_{\langle Q,\leq\rangle} w$, then $u \precsim^*_{\langle Q,\leq\rangle} w$. There are four cases to take care of:

- **Case 1:** $u \sim_Q v$ and $v \sim_Q w$.

  If $u \sim_Q v$, then $\{p \in Q \mid p(u)\} = \{p \in Q \mid p(v)\}$. Since $v \sim_Q w$, $\{p \in Q \mid p(v)\} = \{p \in Q \mid p(w)\}$. By the transitivity of set identity, $\{p \in Q \mid p(u)\} = \{p \in Q \mid p(w)\}$, which entails that $u \sim_Q w$, which entails that $u \precsim^*_{\langle Q,\leq\rangle} w$.

- **Case 2:** $u \sim_Q v$ and there exists a flat upper cut Q' of $\langle Q, \leq\rangle$ such that $v \prec_{Q'} w$.

  Since $\{p \in Q \mid p(u)\} = \{p \in Q \mid p(v)\}$, for any subset Q' of Q, $\{p \in Q' \mid p(u)\} = \{p \in Q' \mid p(v)\}$. Therefore, let Q' be a flat upper cut of $\langle Q, \leq\rangle$ such that $v \prec_{Q'} w$. This entails that $\{p \in Q' \mid p(v)\} \supset \{p \in Q' \mid p(w)\}$, which entails that $\{p \in Q' \mid p(u)\} \supset \{p \in Q' \mid p(w)\}$, which entails that there is a flat upper cut Q' of $\langle Q, \leq\rangle$ such that $u \prec_{Q'} w$, which entails that $u \precsim^*_{\langle Q,\leq\rangle} w$.

- **Case 3:** There exists a flat upper cut Q' of $\langle Q, \leq\rangle$ such that $u \prec_{Q'} v$, and $v \sim_Q w$.

  Since $\{p \in Q \mid p(v)\} = \{p \in Q \mid p(w)\}$, for any subset Q' of Q, $\{p \in Q' \mid p(v)\} = \{p \in Q' \mid p(w)\}$. Therefore, let Q' be a flat upper cut of $\langle Q, \leq\rangle$ such that $u \prec_{Q'} v$. This entails that $\{p \in Q' \mid p(u)\} \supset \{p \in Q' \mid p(v)\}$, which entails that $\{p \in Q' \mid p(u)\} \supset \{p \in Q' \mid p(w)\}$, which entails that there is a flat upper cut Q' of $\langle Q, \leq\rangle$ such that $u \prec_{Q'} w$, which entails that $u \precsim^*_{\langle Q,\leq\rangle} w$.

- **Case 4:** There exists a flat upper cut Q' of $\langle Q, \leq\rangle$ such that $u \prec_{Q'} v$, and there exists a flat upper cut Q'' of $\langle Q, \leq\rangle$ such that $v \prec_{Q''} w$.

  By **Theorem 1**, $Q' \subseteq Q''$ or $Q'' \subseteq Q'$. First, assume that $Q' \subseteq Q''$. Since $\{p \in Q'' \mid p(v)\} \supset \{p \in Q'' \mid p(w)\}$, it is also the case that for any $Q''' \subseteq Q''$, $\{p \in Q''' \mid p(v)\} \supseteq \{p \in Q''' \mid p(w)\}$. But since $Q' \subseteq Q''$, this means that $\{p \in Q' \mid p(v)\} \supseteq \{p \in Q' \mid p(w)\}$. Since $\{p \in Q' \mid p(u)\} \supset \{p \in Q' \mid p(v)\}$, this entails that $\{p \in Q' \mid p(u)\} \supset \{p \in Q' \mid p(w)\}$, which entails that $u \prec_{Q'} w$, which entails that $u \precsim^*_{\langle Q,\leq\rangle} w$.

Next, assume that $Q'' \subseteq Q'$. Since $\{p \in Q' \mid p(u)\} \supset \{p \in Q' \mid p(v)\}$, it is also the case that for any $Q''' \subseteq Q'$, $\{p \in Q''' \mid p(u)\} \supseteq \{p \in Q''' \mid p(v)\}$. But since $Q'' \subseteq Q'$, this means that $\{p \in Q'' \mid p(u)\} \supseteq \{p \in Q'' \mid p(v)\}$. Since $\{p \in Q'' \mid p(v)\} \supset \{p \in Q'' \mid p(w)\}$, this entails that $\{p \in Q'' \mid p(u)\} \supset \{p \in Q'' \mid p(w)\}$, which entails that $u \prec_{Q''} w$, which entails that $u \precsim^*_{\langle Q, \leq \rangle} w$. □

....

Theorem 3: $u \prec^*_{\langle Q, \leq \rangle} v$ iff there exists a flat upper cut Q' of $\langle Q, \leq \rangle$ such that $u \prec_{Q'} v$.

Zig: Proof by reductio. Assume that $u \prec^*_{\langle Q, \leq \rangle} v$ and there exists no flat upper cut Q' of $\langle Q, \leq \rangle$ such that $u \prec_{Q'} v$. $u \prec^*_{\langle Q, \leq \rangle} v$ entails $u \precsim^*_{\langle Q, \leq \rangle} v$, so it must be that $u \sim_Q v$. But since $\sim_Q$ is symmetric, $v \sim_Q u$, so $v \precsim^*_{\langle Q, \leq \rangle} u$, so it is not the case that $u \prec^*_{\langle Q, \leq \rangle} v$.

Zag: Assume that there exists a flat upper cut Q' of $\langle Q, \leq \rangle$ such that $u \prec_{Q'} v$. By the definition of $\precsim^*$ this entails that $u \precsim^*_{\langle Q, \leq \rangle} v$, so it now must be shown that it cannot be that $v \precsim^*_{\langle Q, \leq \rangle} u$. First I will show that it is not the case that $v \sim_Q u$, and then I will show that there is no flat upper cut Q'' of $\langle Q, \leq \rangle$ such that $v \prec_{Q''} u$.

For the former, this is simple. Let q be a proposition in Q' such that $q(u)$ and $\neg q(v)$. By the definition of a flat upper cut, $Q' \subseteq Q$, so $q \in Q$. Therefore, there exists a proposition in Q that holds of u and not of v, so $\{p \in Q \mid p(v)\} \neq \{p \in Q \mid p(u)\}$, and therefore it is not the case that $v \sim_Q u$.

For the latter, we will go with a reductio. Assume that there is a flat upper cut Q'' of $\langle Q, \leq \rangle$ such that $v \prec_{Q''} u$. By **Theorem 1**, either $Q' \subseteq Q''$ or $Q'' \subseteq Q'$. For each scenario, I will show that a contradiction results.

First assume that $Q' \subseteq Q''$. Since $u \prec_{Q'} v$, let q be a proposition in Q' such that $q(u)$ and $\neg q(v)$. Since $Q' \subseteq Q''$, $q \in Q''$. But then it cannot be the case that $\{p \in Q'' \mid p(v)\} \supset \{p \in Q'' \mid p(u)\}$, so it cannot be that $v \prec_{Q''} u$, and we have a contradiction.

Next, assume that $Q'' \subseteq Q'$. Since $v \prec_{Q''} u$, let q be a proposition in Q'' such that $q(v)$ and $\neg q(u)$. Since $Q'' \subseteq Q'$, $q \in Q'$. But then it cannot be the case that $\{p \in Q' \mid p(u)\} \supset \{p \in Q' \mid p(v)\}$, so it cannot be that $u \prec_{Q'} v$, and again contradiction ensues. Therefore, there can be no flat upper cut Q'' of $\langle Q, \leq \rangle$ such that $v \prec_{Q''} u$, and so it is not the case that $v \precsim^*_{\langle Q, \leq \rangle} u$. □

....

Theorem 4: $u \sim^*_{\langle Q, \leq \rangle} v$ iff $u \sim_Q v$.

Zig: Assume that $u \sim^*_{\langle Q,\leq\rangle} v$. For any preorder $\preccurlyeq$, $u \sim v$ iff $u \preccurlyeq v$ and $v \preccurlyeq u$, and $u \prec v$ iff $u \preccurlyeq v$ and $v \npreceq u$. Since it must be that $v \precsim u$ or $v \npreceq u$, $u \sim v$ iff $u \precsim v$ and $u \nprec v$. By the definition of $\precsim^*$, $u \precsim^*_{\langle Q,\leq\rangle} v$ iff $u \sim_Q v$ or there exists a flat upper cut Q' of $\langle Q, \leq\rangle$ such that $u \prec_{Q'} v$. But if it's the latter, then by **Theorem 3** $u \prec^*_{\langle Q,\leq\rangle} v$, so if $u \sim^*_{\langle Q,\leq\rangle} v$, then $u \sim_Q v$.

Zag: Assume that $u \sim_Q v$. Therefore, $u \precsim^*_{\langle Q,\leq\rangle} v$. We need to prove that it cannot be the case that $u \prec^*_{\langle Q,\leq\rangle} v$. This proof will be by reductio, so we will assume that $u \prec^*_{\langle Q,\leq\rangle} v$ and derive a contradiction. By **Theorem 3**, $u \prec^*_{\langle Q,\leq\rangle} v$ iff there exists a flat upper cut Q' of $\langle Q, \leq\rangle$ such that $u \prec_{Q'} v$. Let p be a proposition in Q' such that $p(u)$ and $\neg p(v)$. By the definition of a flat upper cut, $Q' \subseteq Q$, so $p \in Q$. but then there is a proposition $q \in Q$ such that $q(u)$ and $\neg q(v)$ (namely p), so it is not the case that $u \sim_Q v$, a contradiction. □

. . . .

Theorem 5: If $\langle Q_1, \leq_1\rangle$ is an ordered upper cut of $\langle Q_2, \leq_2\rangle$, and Q_3 is a flat upper cut of $\langle Q_1, \leq_1\rangle$, then Q_3 is a flat upper cut of $\langle Q_2, \leq_2\rangle$.

By the definitions of upper cuts, $Q_3 \subseteq Q_1$, and $Q_1 \subseteq Q_2$, so by transitivity $Q_3 \subseteq Q_2$. Now we need to show that for all $p, q \in Q_2$, if $p \leq_2 q$ and $p \in Q_3$, then $q \in Q_3$. This proof will be by reductio: I will assume the antecedent of this conditional to be true and the consequent to be false, and derive a contradiction. Therefore, assume that $p \leq_2 q$, $p \in Q_3$, and $q \notin Q_3$. Since $Q_3 \subseteq Q_1$, we know that $p \in Q_1$. Since $\langle Q_1, \leq_1\rangle$ is an ordered upper cut of $\langle Q_2, \leq_2\rangle$, if $p \leq_2 q$, then $q \in Q_1$ and $p \leq_1 q$. But then Q_3 is a flat upper cut of $\langle Q_1, \leq_1\rangle$, $p \leq_1 q$, $p \in Q_3$, and $q \notin Q_3$, a contradiction. □

. . . .

Theorem 6: Let $\langle Q_1, \leq_1\rangle$ be an ordered upper cut of $\langle Q_2, \leq_2\rangle$. If Q is a flat upper cut of $\langle Q_2, \leq_2\rangle$, but not of $\langle Q_1, \leq_1\rangle$, then $Q \supset Q_1$.

Proof by reductio. Assume that $\langle Q_1, \leq_1\rangle$ is an ordered upper cut of $\langle Q_2, \leq_2\rangle$, Q is a flat upper cut of $\langle Q_2, \leq_2\rangle$, Q is not a flat upper cut of $\langle Q_1, \leq_1\rangle$, and $Q \not\supset Q_1$. By the definition of an ordered upper cut, Q_1 is a flat upper cut of $\langle Q_2, \leq_2\rangle$, so by **Theorem 1**, $Q \subseteq Q_1$ or $Q_1 \subseteq Q$. Therefore, since $Q \not\supset Q_1$, $Q \subseteq Q_1$. Since $Q \subseteq Q_1$ and Q is not a flat upper cut of $\langle Q_1, \leq_1\rangle$, there must be propositions $p, q \in Q_1$ such that $p \in Q$, $p \leq_1 q$, and $q \notin Q$. But by the definition of an ordered upper cut, $p, q \in Q_2$, and $p \leq_2 q$, so Q cannot be a flat upper cut of $\langle Q_2, \leq_2\rangle$, leading to a contradiction. □

. . . .

Theorem 7: If $\langle Q_1, \leq_1 \rangle$ is an ordered upper cut of $\langle Q_2, \leq_2 \rangle$, then $\precsim^*_{\langle Q_1, \leq_1 \rangle}$ is at least as coarse as $\precsim^*_{\langle Q_2, \leq_2 \rangle}$. (Equivalently: $\precsim^*_{\langle Q_2, \leq_2 \rangle}$ is at least as fine-grained as $\precsim^*_{\langle Q_1, \leq_1 \rangle}$.)

I will prove that for all u and v, $u \prec^*_{\langle Q_1, \leq_1 \rangle} v$ entails $u \prec^*_{\langle Q_2, \leq_2 \rangle} v$, and $u \parallel^*_{\langle Q_1, \leq_1 \rangle} v$ entails $u \parallel^*_{\langle Q_2, \leq_2 \rangle} v$. By the definition of fine-grainedness this is equivalent to the first statement of the theorem.

First, strict ordering. Assume that $u \prec^*_{\langle Q_1, \leq_1 \rangle} v$. By **Theorem 3**, this entails that there is a flat upper cut Q of $\langle Q_1, \leq_1 \rangle$ such that $u \prec_Q v$. By **Theorem 5**, if Q is a flat upper cut of $\langle Q_1, \leq_1 \rangle$, and $\langle Q_1, \leq_1 \rangle$ is an ordered upper cut of $\langle Q_2, \leq_2 \rangle$, then Q is a flat upper cut of $\langle Q_2, \leq_2 \rangle$. Therefore, Q is a flat upper cut of $\langle Q_2, \leq_2 \rangle$, so there exists a flat upper cut Q of $\langle Q_2, \leq_2 \rangle$ such that $u \prec_Q v$, and so $u \prec^*_{\langle Q_2, \leq_2 \rangle} v$.

Next, incomparability. Assume that $u \parallel^*_{\langle Q_1, \leq_1 \rangle} v$. This is equivalent to the conjunction of the negations of each of $u \sim^*_{\langle Q_1, \leq_1 \rangle} v$, $u \prec^*_{\langle Q_1, \leq_1 \rangle} v$, and $v \prec^*_{\langle Q_1, \leq_1 \rangle} u$. I will now show that it is not the case that $u \sim^*_{\langle Q_2, \leq_2 \rangle} v$, $u \prec^*_{\langle Q_2, \leq_2 \rangle} v$, or $v \prec^*_{\langle Q_2, \leq_2 \rangle} u$.

If it is not the case that $u \sim^*_{\langle Q_1, \leq_1 \rangle} v$, then it is not the case that $u \sim_{Q_1} v$. Furthermore, if it is not the case that $u \prec^*_{\langle Q_1, \leq_1 \rangle} v$ or $v \prec^*_{\langle Q_1, \leq_1 \rangle} u$, then there is no flat upper cut Q of $\langle Q_1, \leq_1 \rangle$ such that $u \prec_Q v$ or $v \prec_Q u$. Therefore, since Q_1 is the maximal flat upper cut of $\langle Q_1, \leq_1 \rangle$, we have that none of the following are true: $u \sim_{Q_1} v$, $u \prec_{Q_1} v$, $v \prec_{Q_1} u$. In order for this to be the case, there must be propositions $p, q \in Q_1$ such that $p(u)$ and $\neg p(v)$, and $q(v)$ and $\neg q(u)$. This entails that for any $Q' \supseteq Q_1$, it will not be the case that $u \sim_{Q'} v$, $u \prec_{Q'} v$, or $v \prec_{Q'} u$. Since $Q_2 \supseteq Q_1$, we know that it is not the case that $u \sim_{Q_2} v$, so it is not the case that $u \sim^*_{\langle Q_2, \leq_2 \rangle} v$. Furthermore, if there is some flat upper cut Q of $\langle Q_2, \leq_2 \rangle$ such that $u \prec_Q v$ or $v \prec_Q u$, then since it is not the case that $u \prec^*_{\langle Q_1, \leq_1 \rangle} v$ or $v \prec^*_{\langle Q_1, \leq_1 \rangle} u$, Q must not be a flat upper cut of $\langle Q_1, \leq_1 \rangle$. By **Theorem 6**, this entails that $Q \supset Q_1$. But as shown above, there can be no $Q \supset Q_1$ such that $u \prec_Q v$ or $v \prec_Q u$, so there can be no flat upper cut Q of $\langle Q_2, \leq_2 \rangle$ such that $u \prec_Q v$ or $v \prec_Q u$, and so it cannot be the case that $u \prec^*_{\langle Q_2, \leq_2 \rangle} v$ or $v \prec^*_{\langle Q_2, \leq_2 \rangle} u$. Therefore, $u \parallel^*_{\langle Q_2, \leq_2 \rangle} v$. □

. . . .

Theorem 8: Let R be a set of world-orderings linearly orderable by fine-grainedness. Then there is an OPS $\langle Q_R, \leq_R \rangle$ such that for all $\precsim_a \in R$, there is an ordered upper cut $\langle Q_a, \leq_a \rangle$ of $\langle Q_R, \leq_R \rangle$ such that $\precsim^*_{\langle Q_a, \leq_a \rangle}$ is equivalent to $\precsim_a$.

Definition: For ordering $\precsim_a$ and world w, let $[w]_{\precsim_a} = \{w' \mid w' \precsim_a w\}$. Let $\mathrm{PS}(\precsim_a) = \{[w]_{\precsim_a} \mid w \in W\}$. Call this $\precsim_a$'s *characteristic premise set*.

Lemma 1: $\precsim_{\mathrm{PS}(\precsim_a)}$ is equivalent to $\precsim_a$.

> **Zig:** Assume that $w_1 \precsim_a w_2$. By the transitivity of $\precsim_a$, for all w', if $w_2 \precsim_a w'$, then $w_1 \precsim_a w'$. Hence, for all w', if $w_2 \in [w']_{\precsim_a}$, then $w_1 \in [w']_{\precsim_a}$. So for all $p \in \mathrm{PS}(\precsim_a)$, if $p(w_2)$, then $p(w_1)$. Therefore, $w_1 \precsim_{\mathrm{PS}(\precsim_a)} w_2$.
>
> **Zag:** Assume that $w_1 \precsim_{\mathrm{PS}(\precsim_a)} w_2$. So for all $p \in \mathrm{PS}(\precsim_a)$, if $p(w_2)$, then $p(w_1)$. By the definition of $\mathrm{PS}(\precsim_a)$, this entails that for all $w' \in W$, if $w_2 \in [w']_{\precsim_a}$, then $w_1 \in [w']_{\precsim_a}$. This in turn entails that for all $w' \in W$, if $w_2 \precsim_a w'$, then $w_1 \precsim_a w'$. By reflexivity of $\precsim_a$, $w_2 \precsim_a w_2$. Substituting w_2 for w', $w_1 \precsim_a w_2$. □

Definition: For $\precsim_a \in R$, let $\mathrm{PS}^+_R(\precsim_a) = \bigcup\{\mathrm{PS}(\precsim_b) \mid \precsim_b \in R$ and $\precsim_b$ is at least as coarse as $\precsim_a\}$.

Definition: Let $Q_R = \bigcup\{\mathrm{PS}(\precsim_a) \mid \precsim_a \in R\}$.

Definition: Let
$$\leq_R = \{(p,q) \in {Q_R}^2 \mid \forall \precsim_a \in R[(p \in \mathrm{PS}^+_R(\precsim_a)) \rightarrow (q \in \mathrm{PS}^+_R(\precsim_a))]\}.$$

Lemma 2: $\langle Q_R, \leq_R \rangle$ is an OPS.

> **Reflexivity of** $\leq_R$**:** For all $p \in Q_R$, it is trivially true that $\forall \precsim_a \in R[(p \in \mathrm{PS}^+_R(\precsim_a)) \rightarrow (p \in \mathrm{PS}^+_R(\precsim_a))]$.
>
> **Transitivity of** $\leq_R$**:** Assume $p, q, r \in Q_R$. Furthermore, assume that $\forall \precsim_a \in R[(p \in \mathrm{PS}^+_R(\precsim_a)) \rightarrow (q \in \mathrm{PS}^+_R(\precsim_a))]$, and that $\forall \precsim_a \in R[(q \in \mathrm{PS}^+_R(\precsim_a)) \rightarrow (r \in \mathrm{PS}^+_R(\precsim_a))]$, and thus that $p \leq_R q \leq_R r$. It is clear that $\forall \precsim_a \in R[(p \in \mathrm{PS}^+_R(\precsim_a)) \rightarrow (r \in \mathrm{PS}^+_R(\precsim_a))]$, and thus that $p \leq_R r$.
>
> **Connectedness of** $\leq_R$**:** Proof by reductio. Assume that $p, q \in Q_R$, and that $p \nleq_R q$ and $q \nleq_R p$. By the definition of $\leq_R$, this is true iff there exists some $\precsim_a \in R$ such that $p \in \mathrm{PS}^+_R(\precsim_a)$ and $q \notin \mathrm{PS}^+_R(\precsim_a)$, and there exists some $\precsim_b \in R$ such that $q \in \mathrm{PS}^+_R(\precsim_b)$ and $p \notin \mathrm{PS}^+_R(\precsim_b)$. Either $\precsim_a$ is at least as fine-grained as $\precsim_b$ or vice versa, by the definition of R. Assume that $\precsim_a$ is at least as fine-grained as $\precsim_b$. Thus, every ordering that is at least as coarse as $\precsim_b$ is at least as coarse as $\precsim_a$. Hence, $\mathrm{PS}^+_R(\precsim_b) \subseteq \mathrm{PS}^+_R(\precsim_a)$. But then if $q \in \mathrm{PS}^+_R(\precsim_b)$, it must be the case that $q \in \mathrm{PS}^+_R(\precsim_a)$, contradicting the assumption above. Assuming instead

that $\precsim_b$ is at least as fine-grained as $\precsim_b$ leads to the same inferences in the opposite direction. □

Definition: For $\precsim_a \in R$, let $Q_a = \mathrm{PS}^+_R(\precsim_a)$.

Definition: For $\precsim_a \in R$, let $\leq_a = \{(p,q) \in \leq_R \mid p,q \in Q_a\}$

Lemma 3: $\langle Q_a, \leq_a \rangle$ is an ordered upper cut of $\langle Q_R, \leq_R \rangle$.

It is trivial that for all $p,q \in Q_a$, $p \leq_a q$ iff $p \leq_R q$, since this is expressly how $\leq_a$ was defined. We thus need only show that Q_a is a flat upper cut of Q_R. That $Q_a \subseteq Q_R$ is trivial, so all that must be shown is that for all $p \in Q_a$ and $q \in Q_R$ such that $p \leq_R q$, $q \in Q_a$.

If $p \in Q_a$, then given the definition of Q_a, there must be some $\precsim_b \in R$ that is at least as coarse as $\precsim_a$, and such that $p \in \mathrm{PS}(\precsim_b)$. Since $\mathrm{PS}(\precsim_b) \subseteq \mathrm{PS}^+_R(\precsim_b)$, we also get that $p \in \mathrm{PS}^+_R(\precsim_b)$. Now suppose that $p \leq_R q$. By the definition of $\leq_R$, this is true iff for every $\precsim_c$ such that $p \in \mathrm{PS}^+_R(\precsim_c)$, $q \in \mathrm{PS}^+_R(\precsim_c)$. Therefore, $q \in \mathrm{PS}^+_R(\precsim_b)$. Since every ordering that is at least as coarse as $\precsim_b$ is at least as coarse as $\precsim_a$, $\mathrm{PS}^+_R(\precsim_b) \subseteq \mathrm{PS}^+_R(\precsim_a)$. Hence, $q \in \mathrm{PS}^+_R(\precsim_a)$, and so $q \in Q_a$. □

Lemma 4: For $\precsim_a \in R$, $\sim_a$ and $\sim^*_{\langle Q_a, \leq_a \rangle}$ are equivalent.

By **Theorem 4**, $\sim^*_{\langle Q_a, \leq_a \rangle}$ is equivalent to $\sim_{Q_a}$, and so I will prove the equivalence of $\sim_a$ and $\sim_{Q_a}$.

Zig: Assume that $w_1 \sim_a w_2$. By the definition of fine-grainedness, for any ordering $\precsim_b$ at least as coarse as $\precsim_a$, $w_1 \sim_b w_2$. Thus, for all $\precsim_b \in R$ at least as coarse as $\precsim_a$, for all $p \in \mathrm{PS}(\precsim_b)$, $p(w_1)$ iff $p(w_2)$. Thus, for all p in $\mathrm{PS}^+_R(\precsim_a)$, $p(w_1)$ iff $p(w_2)$. Therefore, $p \sim_{\mathrm{PS}^+_R(\precsim_a)} q$, and since $Q_a = \mathrm{PS}^+_R(\precsim_a)$, $w_1 \sim_{Q_a} w_2$.

Zag: Assume that $w_1 \sim_{Q_a} w_2$. Therefore, $w_1 \sim_{\mathrm{PS}^+_R(\precsim_a)} w_2$, so for all $p \in \mathrm{PS}^+_R(\precsim_a)$, $p(w_1)$ iff $p(w_2)$. $\mathrm{PS}(\precsim_a) \subseteq \mathrm{PS}^+_R(\precsim_a)$, so for all $p \in \mathrm{PS}(\precsim_a)$, $p(w_1)$ iff $p(w_2)$. Hence, $w_1 \sim_{\mathrm{PS}(\precsim_a)} w_2$, so by **Lemma 1**, $w_1 \sim_a w_2$. □

Lemma 5: For $\precsim_a \in R$, $\prec_a$ and $\prec^*_{\langle Q_a, \leq_a \rangle}$ are equivalent.

As per **Theorem 3**, $w_1 \prec^*_{\langle Q_a, \leq_a \rangle} w_2$ iff there is a flat upper cut Q of $\langle Q_a, \leq_a \rangle$ such that $w_1 \prec_Q w_2$. I will show that Q_a itself is just such a flat upper cut. That Q_a is a flat upper cut of $\langle Q_a, \leq_a \rangle$ is trivial. I will show that $\prec_a$ and $\prec_{Q_a}$ are equivalent.

Zig: Assume that $w_1 \prec_a w_2$. By the definition of fine-grainedness, for any ordering $\precsim_b$ at least as coarse as $\precsim_a$, $w_1 \precsim_b w_2$. So for all $\precsim_b$ at least as coarse as $\precsim_a$, for all propositions $p \in \mathrm{PS}(\precsim_b)$, if $p(w_1)$ then $p(w_2)$. Therefore, for every proposition $p \in \mathrm{PS}^+_R(\precsim_a)$, if $p(w_1)$, then $p(w_2)$. Furthermore, since $w_1 \precsim_a w_1$, but $w_2 \not\precsim_a w_1$, we know that $w_1 \in [w_1]_{\precsim_a}$, but $w_2 \notin [w_1]_{\precsim_a}$. Since $[w_1]_{\precsim_a} \in \mathrm{PS}^+_R(\precsim_a)$, we know that every proposition in $\mathrm{PS}^+_R(\precsim_a)$ that holds of w_2 also holds of w_1, and that at least one proposition in $\mathrm{PS}^+_R(\precsim_a)$ holds of w_1 and not of w_2. Hence, $w_1 \prec_{\mathrm{PS}^+_R(\precsim_a)} w_2$, and so $w_1 \prec_{Q_a} w_2$.

Zag: Assume that $w_1 \prec_{Q_a} w_2$, and thus that $w_1 \prec_{\mathrm{PS}^+_R(\precsim_a)} w_2$. There must be some proposition $p \in \mathrm{PS}^+_R(\precsim_a)$ such that $p(w_1)$ and $\neg p(w_2)$. Furthermore, by the definition of $\mathrm{PS}^+_R(\precsim_a)$, there must be an ordering $\precsim_b \in R$ such that $\precsim_b$ is at least as coarse as $\precsim_a$, and such that $p \in \mathrm{PS}(\precsim_b)$. Furthermore, it must also be the case that for all $q \in \mathrm{PS}^+_R(\precsim_a)$, if $q(w_1)$, then $q(w_2)$. Since $\mathrm{PS}(\precsim_b) \subseteq \mathrm{PS}^+_R(\precsim_a)$, it will also be the case that for all $q \in \mathrm{PS}(\precsim_b)$, if $q(w_1)$, then $q(w_2)$. This in conjunction with the existence of p entails that $w_1 \prec_{\mathrm{PS}(\precsim_b)} w_2$, which by **Lemma 1** entails that $w_1 \prec_b w_2$. Since $\precsim_b$ is at least as fine-grained as $\precsim_a$, by the definition of fine-grainedness, $w_1 \prec_b w_2$ entails that $w_1 \prec_a w_2$. □

From here, the rest of the proof of the proposition at hand is simple. We know that $w_1 \sim_a w_2$ iff $w_1 \sim^*_{\langle Q_a, \leq_a \rangle} w_2$, and $w_1 \prec_a w_2$ iff $w_1 \prec^*_{\langle Q_a, \leq_a \rangle} w_2$. Thus, $w_1 \precsim_a w_2$ iff $w_1 \precsim^*_{\langle Q_a, \leq_a \rangle} w_2$. Since $\langle Q_a, \leq_a \rangle$ is an ordered upper cut of $\langle Q_R, \leq_R \rangle$, we know that for each $\precsim_a \in R$, there is an ordered upper cut $\langle Q_a, \leq_a \rangle$ of $\langle Q_R, \leq_R \rangle$ such that $\precsim_a$ is equivalent to $\precsim^*_{\langle Q_a, \leq_a \rangle}$. □

Chapter 5: Non-distributive ascriptions of belief

5.1 Introduction

The previous two chapters have focused on desiderative attitudes like *want*, *wish*, and *regret*, justifying and exploring an ontology in which intensity is correlated with the part-whole relations of attitude states. In this chapter, which is more preliminary in nature than the previous two, I will shift my attention from desire to belief, simultaneously pivoting from attitude intensity to a more well-trodden area of mereological inquiry: plurality.

It is by this point a truism that properties can sometimes be attributed to pluralities that cannot be attributed to the individuals of which they are constituted. For example, on its most salient reading the sentence in (1) does not entail (2a) or (2b):

(1) Rick and Morty ate the whole pie.

(2) a. Rick ate the whole pie.
b. Morty ate the whole pie.

The lack of entailment from (1) to (2a) and (2b) indicates that the salient interpretation of (1) is *non-distributive*. This is in contrast to the distributive interpretation of *be tall* in (3), which entails (4a) and (4b):

(3) Rick and Morty are tall.

(4) a. Rick is tall.
b. Morty is tall.

While (1) and (3) present cases where the choice of (non-)distributivity is relatively clear, (5) displays an ambiguity between the two readings, as it is not obvious whether each board member received a million dollars (distributive), or whether the board members' combined income was a million dollars (non-distributive):

(5) The board members were paid a million dollars last year.

Normally, and perhaps unsurprisingly, attitude ascriptions are interpreted distributively. Take, for example, the sentences in (6):

(6) a. Rick and Morty believe that Summer left.
b. Rick and Morty want Summer to leave.
c. Rick and Morty wish that Summer had left.
d. Rick and Morty regret that Summer didn't leave.

On a default interpretation of (6a), there is an entailment that Rick believes that Summer left, and that Morty does too. *Mutatis mutandis* for *want*, *wish*, and *regret* in (6b), (6c), and (6d), respectively.

Given the discussion above, and focusing our attention on the semantically simplest case of belief, we are faced with two intriguing questions. First, on a purely conceptual level, does it make sense to talk about the non-distributive attribution of belief? That is, is it possible for Rick and Morty to have a belief that p, without either of them individually believing that p? And second, if the answer to the first question is "yes", then are non-distributive ascriptions of belief available as part of the grammar? Note that a "yes" to the first question does not *a priori* entail a "yes" to the second: it is conceivable that plurally-experienced beliefs exist "in the wild", but that *believe* is lexically distributive.

In this chapter, I argue that the answer to both of these questions is "yes", and provide an analysis of the semantics and natural language ontology of plurally-experienced belief. I start in Section 5.2 by introducing the basic assumptions about the ontology of plural individuals and the semantics of plural DPs that underlie my own analysis. In Section 5.3 I discuss the evidence suggesting that non-distributive ascription of belief is possible, and provide a basic syntax and semantics for distributive and non-distributive interpretations of both non-attitude and attitude constructions. In Section 5.4 I turn to the main issue facing an analysis of non-distributive belief ascription: what is the relationship between the beliefs of a plurality and the beliefs of the individuals it comprises? The solution I propose is that the various compatibilities and incompatibilities between experiencers' beliefs are negotiated in a manner formally identical to the classical premise semantics of Lewis (1981) and Kratzer (1981a). However, I show at the end of Section 5.4 that this analysis has problems differentiating between relevant and irrelevant disagreements between experiencers. In Section 5.5 I propose to fix this by means of a notion of "aboutness": beliefs are about situations, and it is by fixing our choice of "about-situations" that we filter out those beliefs that are irrelevant. Section 5.6 concludes.

But before moving on, two caveats are worth mentioning. Caveat number one: My focus in this chapter is only on the distinction between distributive and non-distributive interpretations of sentences containing plural DPs. I am therefore setting aside the possibility that there may be multiple kinds of non-distributive reading (e.g, cumulative vs. collective vs. group). Determining precisely which type of reading occurs in the non-distributive belief ascriptions in this chapter is beyond the scope of this dissertation. I have therefore opted for the simplest possibility: if ascribing a belief to an individual involves inspecting the content of their belief state, then ascribing a belief to a plurality involves inspecting the content of the mereological sum of its members' belief states. This is, essentially, a cumulative interpretation. But even if it turns out that the examples discussed in Section 5.3 are really some other kind of non-distributive interpetation, the basic features of the analysis in Sections 5.4 and 5.5 can in all probability be tailored to fit into the appropriate semantics and/or ontology.

Caveat number two: After finishing work on this dissertation, I was made aware of work by Viola Schmitt analyzing data bearing a *prima facie* resemblance to those discussed in this chapter (see, e.g., Schmitt 2017). The theory she offers is both very interesting and very different from my own. However, for reasons of space and time I must content myself with sticking to my own analysis, leaving for another occasion an exploration of the similarities and differences between our proposals.

5.2 Basic assumptions about plurals

First, let us establish some basic assumptions about the semantics of plural DPs. One could of course write a whole dissertation just on this topic alone. Fortunately, many theories of the intepretation of plurals are differentiated by predictions subtler than what will be required for the purposes of this chapter. All we need to fulfill the task at hand is a theory of plurality powerful enough to capture the distributive/non-distributive distinction. As a result, the core ideas proposed in the rest of this chapter should fairly readily translate to whatever one's favorite framework is for the semantics of plurals.

The theory I will adopt for the semantics of plurals is in the tradition of Link (1983), who places a significant amount of the work into the ontology by positing the existence of plural individuals of type e. That is, in addition to individuals x and y, there is an individual $x \sqcup y$ that is the mereological sum of x and y. There is also a part-whole relation $\sqsubseteq$ between entities, with $\sqcup$ and $\sqsubseteq$ being interdefinable: $x \sqcup y$ is the smallest (i.e., minimal by $\sqsubseteq$) individual z such that $x \sqsubseteq z$ and $y \sqsubseteq z$. Hopefully this all sounds familiar, as I have been operating under essentially the same assumptions about the domain of events, as per the extensions of Link's theory by Bach (1986a) and Krifka (1989).[1]

In addition to these familar characters, three new concepts need to be introduced. First is the notion of an *atom*, which is a non-plural individual (i.e., an indi-

[1] It is worth noting that accepting a Linkian mereology for the domain of events does not necessarily commit one to a Linkian mereology for the domain of entities. Schein (1993), for example, adopts a Link-style mereology for events, but not for individuals.

vidual with no proper parts):

(7) $\text{Atom}(x)$ iff $\neg\exists y[y \sqsubset x]$

Next up is the $*$ operator, which closes an $\langle e,t\rangle$-type predicate under mereological sum. That is, if $P(x)$ and $P(y)$, then it will also be the case that $^*P(x)$, $^*P(y)$, and $^*P(x \sqcup y)$.

(8) $^*P \equiv \lambda x.\ \exists A[x = \bigsqcup A \wedge \forall y \in A[P(y)]]$

(where $\bigsqcup A$ is the sum of the elements of A)

Finally, there is the σ operator, a maximizing referential operator akin to the Russellian ι. My definition of the σ operator in (9) is slightly different from Link's, but is in the same spirit:[2]

(9) $\sigma x[P(x)] \equiv \iota x[P(x) \wedge \forall y[P(y) \rightarrow y \sqsubseteq x]]$

Let's look at some toy examples: the DPs *the boy*, *the boys*, and *the six boys*. First, *the boy*. I take *the* and *boy* to have the denotations in (10), meaning that *the boy* will have the denotation in (11):

(10) a. $[\![\text{the}]\!] = \lambda P.\ \sigma x[P(x)]$

b. $[\![\text{boy}]\!] = \lambda x.\ \text{boy}(x)$

(11) $[\![\text{the boy}]\!] = \sigma x[\text{boy}(x)]$

By the definition of the σ operator, (11) is equivalent to (12):

(12) $\iota x[\text{boy}(x) \wedge \forall y[\text{boy}(y) \rightarrow y \sqsubseteq x]]$

Importantly, without the $*$ operator, the predicate $\lambda x.\ \text{boy}(x)$ is true only of atomic (i.e., individual) boys, and not of their mereological sums. Thus, (12) will return the single boy x such that for all boys y, y is a part of x. But it is clear that no boy is a mereological part of any boy other than himself. As a result, (12), and thus (11), will only be defined if there is exactly one boy in the domain, meaning that we get the same result as if we had used the Russellian ι.

As for *the boys*, I take the plural morphology on *boys*—or perhaps the head responsible for its presence—to contribute the $*$ operator, meaning that $[\![\text{the boys}]\!]$ will look like (13):

[2] Link's definition can be seen below. Note the replacement of P with $^*P$:

(i) $\sigma x[P(x)] \equiv \iota x[^*P(x) \wedge \forall y[^*P(y) \rightarrow y \sqsubseteq x]]$

On Link's definition, $\sigma x[P(x)]$ returns the sum of all Ps, while on mine it returns the maximal P (if there is one). As will be seen in the body of the text, mine allows an easy definition of *the* as contributing the σ operator. Link's definition cannot do this: if there are three boys, for him $\sigma x[\text{boy}(x)]$ (my denotation for *the boy*) will return the sum of the three boys, rather than being undefined.

(13) $[\![\text{the boys}]\!] = \sigma x[{}^*\text{boy}(x)]$

$\lambda x.\ {}^*\text{boy}(x)$, unlike $\lambda x.\ \text{boy}(x)$, is closed under mereological sum, meaning that it includes plural individuals whose atoms are boys, as well as the individual boys themselves.[3] (13) will thus return the individual x that is a sum of boys, and that contains every sum of boys as a part. That is, (13) will return the mereological sum of all boys.

Finally, there's *the six boys*. I assume $[\![\text{six}]\!]$ to be a predicate of plural individuals, true of an individual x iff the set of atoms of which it is composed has a cardinality of six:

(14) $[\![\text{six}]\!] = \lambda x.\ |x| = 6$ (where $|x| \equiv |\{y \sqsubseteq x \mid \text{Atom}(y)\}|$)

As a result, *the six boys* has the denotation in (15):

(15) $[\![\text{the six boys}]\!] = \sigma x[{}^*\text{boy}(x) \wedge |x| = 6]$

(15) denotes the six-boy collection that includes as a part every other six-boy collection. But as was the case with *the boy*, this is only defined if there is exactly one six-boy collection: if x is a collection of six boys, and y is a collection of six boys, and $x \neq y$, then neither can be a part of the other.[4] We therefore predict *the six boys* to presuppose that there are, in fact, exactly six boys.

5.3 Evidence of non-distributivity and compositional semantics of the two interpretations

With a semantics for plural definite DPs now in place, in this section I will do for non-distributively ascribed beliefs what I did for the mereological basis of attitude intensity in Chapter 3: I will provide evidence of its existence and discuss most of the strictly semantic side of the analysis. Once the semantic analysis is in place, I will set up the ontological problem to be addressed in the next section, where can be found the majority of the work of my theory.

5.3.1 Non-distributive ascriptions of belief

Consider sentence (16) in the context provided:

> *Sam had six clients, none of whom knew of the others' existence. She convinced each of her six clients that she would build a house for him. In reality, she was a con artist and built no houses at all.*

[3] It is up for debate whether $[\![\text{boys}]\!]$, for example, should include individual boys in its domain, or only strict pluralities. For arguments for the latter, see Chierchia 1998; for arguments for the former, see Sauerland 2003. While I side with Sauerland, this choice is not relevant for my own proposal.

[4] Krifka (1989) refers to such predicates as *quantized* predicates.

(16) (In total,) Sam's six clients thought she built six houses for them.

There is an interpretation of (16) that is true here. However, the proposition denoted by the complement of *think* is not believed by each of Sam's six clients, meaning that this interpretation cannot be a distributive one. None of Sam's clients knows of the others' existence, so no client has a belief that Sam built six houses for her six clients. Moreover, the truth of (16) cannot be attributed to a wide scope interpretation of *six houses* in conjunction with some other mereological wizardry at the matrix-clause level, as could conceivably be done for a sentence like (17):

(17) There are six houses that Sam's six clients think she built for them.

After all, as stated in the scenario, there are no houses to begin with. It seems that the most likely analysis for (16), then, is belief accumulation: client one thinks he got one house, client two thinks he got one house, etc., so the total number of houses the clients cumulatively believe themselves to have received is six.

As another example, consider the sentences in (18), again with the context provided:

> *Paul just got married, and his cousins Arnie and Beatrice, who have never met, just caught wind of it. Arnie suspects that Paul's husband is rich, and has no other relevant opinions. Beatrice thinks he's a New Yorker, and has no other relevant opinions.*

(18) a. Paul's cousins think he married a rich New Yorker.
b. Arnie and Beatrice think Paul married a rich New Yorker.

Once again, we have a belief that is attributed to the plurality, but that cannot be attributed to the individuals of which it is composed. Since Arnie is agnostic about where Paul's husband is from, he does not have a belief that Paul married a rich *New Yorker*. Likewise, since Beatrice is agnostic about Paul's husband's wealth, it is not true that she believes that Paul married a *rich* New Yorker. Thus, the truth of (18) can only be attributed to the existence of a non-distributive reading, in which Arnie and Beatrice's beliefs are combined into a single belief that Paul married a rich New Yorker.

So far, we have looked at cases where the relevant beliefs of the individual experiencers do not conflict. In the house-building scenario, each client is agnostic about whether there are other clients, and whether they are having houses built for them. In the scenario of Paul's wedding, Arnie and Beatrice are each agnostic about a particular trait of Paul's husband's. But when we turn to cases where there are relevant disagreements between the experiencers, interesting things happen. For example, in the context below, the sentences in (18) are false, while the sentences in (19) are true:

Arnie thinks that Paul married a rich Marylander, while Beatrice thinks he married a poor New Yorker.

(19) a. Paul's cousins think he married either a rich Marylander or a poor New Yorker.

b. Arnie and Beatrice think Paul married either a rich Marylander or a poor New Yorker.

In the conflict-free scenarios, the beliefs of the plurality were a conjunction of the beliefs of the atomic individuals: six beliefs in individual houses were combined into a single belief in six houses, while the traits that Arnie and Beatrice ascribed to Paul's husband were conjoined in their combined belief state. The examples in (19), meanwhile, seem to suggest that when there is conflict, the beliefs of the plurality are a *dis*junction of the beliefs of the atomic individuals.

But there's a catch: unlike in the previous cases, the sentences in (19) are actually true on a distributive reading as well, as Arnie and Beatrice both believe the weak proposition denoted by the complement of *think*. If we're going to truly show that agreement leads to conjunction and disagreement leads to disjunction, we will have to come up with a scenario in which there is still disjunction of beliefs, but the distributive reading is false. With this in mind, consider (20) with its associated context:

Paul has three cousins, Arnie, Beatrice, and Kate. Arnie still thinks Paul married a rich man, and Beatrice thinks he married a New Yorker. Kate, like Beatrice, is unopinionated about Paul's husband's wealth, but she thinks he's from California, not New York.

(20) Paul's cousins think he married a rich man from either California or New York.

In the context for (20), there is no conflict when it comes to Paul's husband's wealth, but there is conflict when it comes to his place of origin. Importantly, the proposition denoted by the embedded clause in (20) is not believed by Arnie, Beatrice, or Kate individually, since Arnie is still agnostic about place of origin, and Beatrice and Kate are both agnostic about wealth. Thus, the sentence in (20) is a genuine case of a non-distributive ascription of belief, and we indeed get disjunction of beliefs where beliefs conflict, in addition to conjunction of beliefs where there is no conflict.

It is these observations, as well as slight extensions thereto, that I will be accounting for in the rest of this chapter. As was the case in the previous chapters, the majority of the interesting work will be done in the natural language ontology. That being said, a certain semantic foundation is needed in order to proceed with the analysis, and so in the rest of this section I will put forward a basic syntax and compositional semantics for sentences with distributive and non-distributive interpretations of plural DPs.

5.3.2 Compositionality of distributive interpretation

In discussing the syntax and semantics of (non-)distributive interpretations of sentences with plural DPs, I will use two sentences as illustrative examples. The first is (21), which provides a simple case of an ambiguous non-attitude sentence.

(21) The students lifted the table.

For an ambiguous belief ascription, I will use (16).

I will treat the distributive interpretation of DPs as being due to the presence of a distributive operator DIST, which adjoins to the referential DP to generate a distributive interpretation. Thus, the syntactic representation for the distributive interpretation of (21) will be as in (22):

(22)

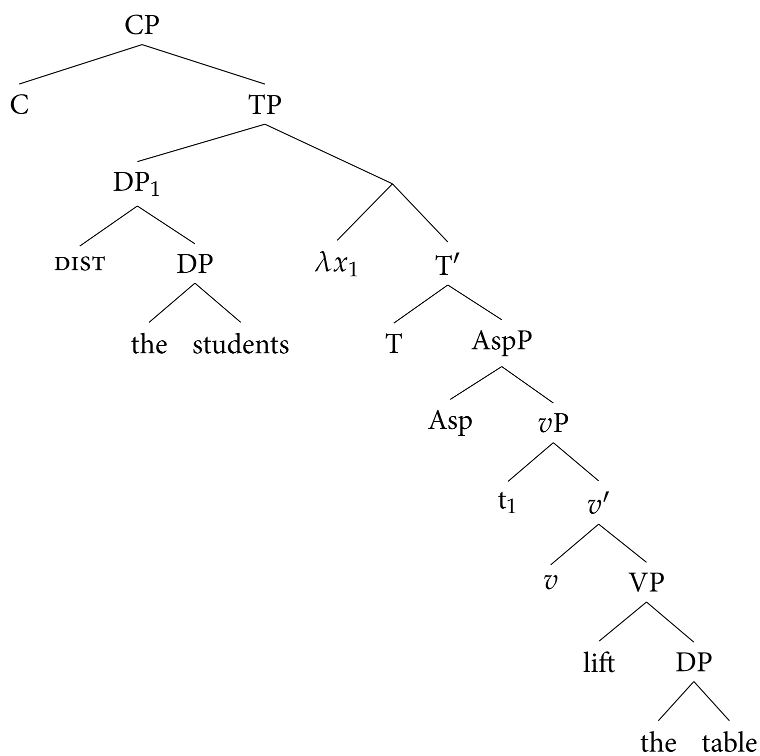

It is worth noting that there are strong arguments, dating back at least to the work of Dowty (1987), that the distributive/non-distributive distinction cannot be due only to a difference in the interpretation of the plural DP. For example, consider (23):

(23) The students lifted the table and then gathered in the weight room.

(23) is compatible with an interpretation according to which each student lifted the table, and then all of the students gathered in the weight room. On this reading,

within the conjoined VP in (23), the first conjunct has a distributive interpretation, since each student is lifting the table. Meanwhile, the second conjunct includes *gather*, which can only be predicated of pluralities and groups, and is thus infelicitous on a distributive interpretation of *the students*. This is further illustrated in (24), which is out due to the presence of the lexically distributive *each*:

(24) # Each of the students gathered in the weight room.

Thus, *the students* appears to receive a distributive interpretation with *lift the table*, and a non-distributive interpretation with *gather in the weight room*. But there is of course only one iteration of *the students*, meaning that the presence or absence of distributivity cannot be due to a difference in interpretation of *the students*.

To this problem, my response is that it does not really matter for my own proposal where the distributivity operator attaches: as a DP-adjunct, as I have it, or as a VP adjunct, or elsewhere.[5] All that is needed for our purposes is the ability to generate both a distributive and a non-distributive reading in the simplest cases. Since I treat belief ascriptions as composing in an unexceptional manner on the way up the (matrix) clausal spine, wherever one wants to introduce distributivity in simpler cases like (21), the same method ought to work for belief cases like (16). I will thus proceed with the DP-adjunct analysis, for the sole reason that it lends itself to a relatively simple compositional semantics.

Speaking of which, the semantic interpretation for DIST can be seen in (25):

(25) $[\![\textsc{dist}]\!] = \lambda x \lambda P.\ \forall y[(\text{Atom}(y) \wedge y \sqsubseteq x) \rightarrow P(y)]$

$[\![\textsc{dist}]\!]$ takes a (non-atomic) individual and returns a quantifier (type $\langle\langle e,t\rangle,t\rangle$) that is true of any predicate that holds of every atom of which the individual is composed. For example, DIST *the students* will have the denotation in (26). This is equivalent to (27), which is the traditional denotation for *every student*:

(26) $[\![\textsc{dist}\ \text{the students}]\!] = \lambda P.\ \forall y[(\text{Atom}(y) \wedge y \sqsubseteq \sigma x[{}^*\text{student}(x)]) \rightarrow P(y)]$

(27) $= \lambda P.\ \forall y[\text{student}(y) \rightarrow P(y)]$

Because $[\![\textsc{dist}\ \text{the students}]\!]$ is of type $\langle\langle e,t\rangle,t\rangle$, rather than type e, we no longer have the luxury of not caring whether the subject of (22) is interpreted in its merge or its post-movement position. There is a type mismatch at the merge position of spec-vP, since $[\![v']\!]$ is of type $\langle e,\langle v,t\rangle\rangle$. QR is therefore obligatory, hence why there is lambda abstraction over x_1 (the variable denoted by the subject's trace) just below the subject's landing spot in spec-TP. Assuming the same basic clausal compositionality

[5] In fact, there needn't even be a distributivity operator at all: the basic ideas in this chapter are equally compatible with theories in which the distributive/non-distributive distinction is attributed strictly to context (e.g., Gillon 1987, 1990; Schwarzschild 1996).

discussed in Chapter 3, this means that the predicate with which ⟦DIST the students⟧ combines is as in (28):

(28) $\lambda x.\ \exists e[\mathrm{Agt}(e) = x \wedge \mathrm{lift}(e) \wedge \mathrm{Thm}(e) = \sigma z[\mathrm{table}(z)]]$

This in turn means that the final interpretation of (22) will be as in (29):

(29) $[\![(22)]\!] = 1$ iff $\forall y[\mathrm{student}(y) \rightarrow \exists e[\mathrm{Agt}(e) = y \wedge \mathrm{lift}(e) \wedge \mathrm{Thm}(e) = \sigma z[\mathrm{table}(z)]]]$

This is the interpretation we want, as the distributive reading of (21) is predicted to be true iff for every student y, there is an event of y lifting the table.

Moving on to belief, the distributive interpretation of (16) will have the syntactic representation in (30):

(30)

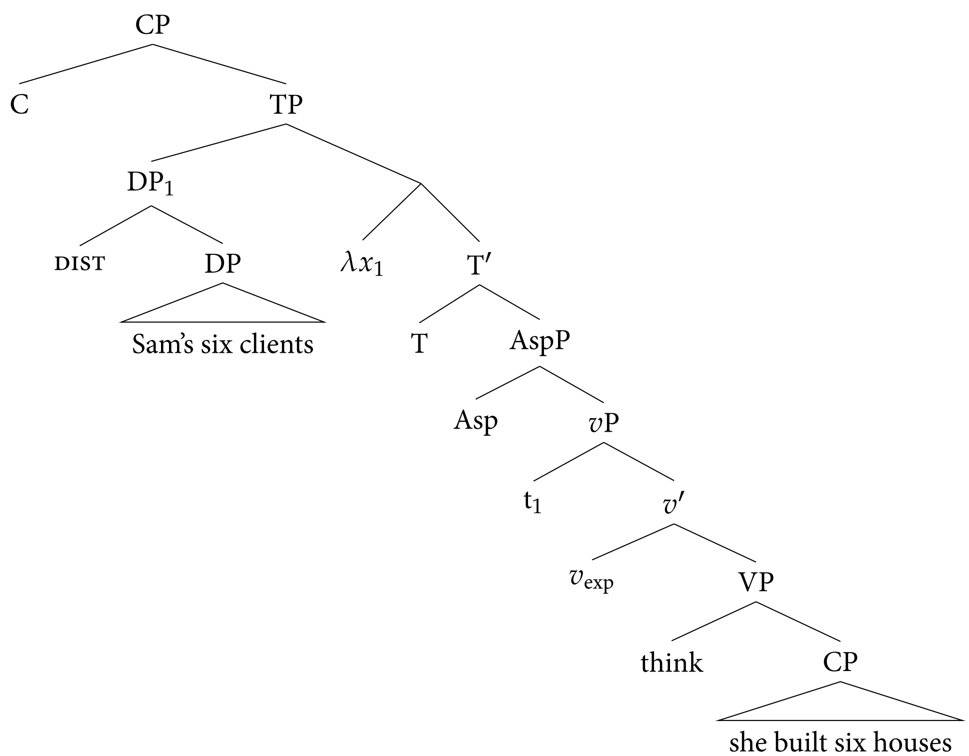

Starting from the bottom of the clausal spine, I take *think* and *believe* to have a neo-Davidsonian version of a Hintikkan world-quantificational semantics, as in (31). (I call this definition "take 1" because there will be further revisions in Section 5.5.)

(31) $[\![\text{believe/think}]\!]_{\text{take 1}} = \lambda p \lambda e.\ \forall w \in \mathrm{Dox}(e)[p(w)]$

Up until and including lambda abstraction over x_1, there are no surprises. The result is as in (32):

(32) $\lambda x.\ \exists e[\mathrm{Exp}(e) = x \wedge \forall w \in \mathrm{Dox}(e)[\text{six\_houses}(w)]]$

Turning to the subject DP, the denotation for *Sam's six clients* (without DIST) is as in (33):

(33) $[\![\text{Sam's six clients}]\!] = \sigma y[{}^{*}\text{client-of-sam}(y) \wedge |y| = 6]$

As per the discussion of *the six boys* above, (33) is well-defined only if Sam has exactly six clients, and denotes the plural individual whose atoms are Sam's clients. When combined with DIST, the result is as in (34), which is equivalent to (35) if (33) is well-defined:

(34) $[\![\text{DIST Sam's six clients}]\!] = \lambda P.\ \forall y[(\mathrm{Atom}(y) \wedge y \sqsubseteq \sigma z[{}^{*}\text{client-of-sam}(z) \wedge |z| = 6]) \rightarrow P(y)]$

(35) $= \lambda P.\ \forall y[\text{client-of-sam}(y) \rightarrow P(y)]$

The final result of applying (35) to (32) is as in (36):

(36) $[\![(30)]\!] = 1 \text{ iff } \forall y[\text{client-of-sam}(y) \rightarrow \exists e[\mathrm{Exp}(e) = y \wedge \forall w \in \mathrm{Dox}(e)[\text{six\_houses}(w)]]]$

(36) is true iff for every client y of Sam's, there is a state of y thinking that Sam built six houses. Of course, this reading is false in the context provided for (16), but this is precisely because the context for (16) was designed to render a distributive interpretation false.

5.3.3 Compositionality of non-distributive interpretation

Conveniently enough, the non-distributive interpretation of (16) and (21) will be derived simply by excluding the DIST operator. Thus, the non-distributive interpretation of (21) will have the syntactic representation in (37):

(37)

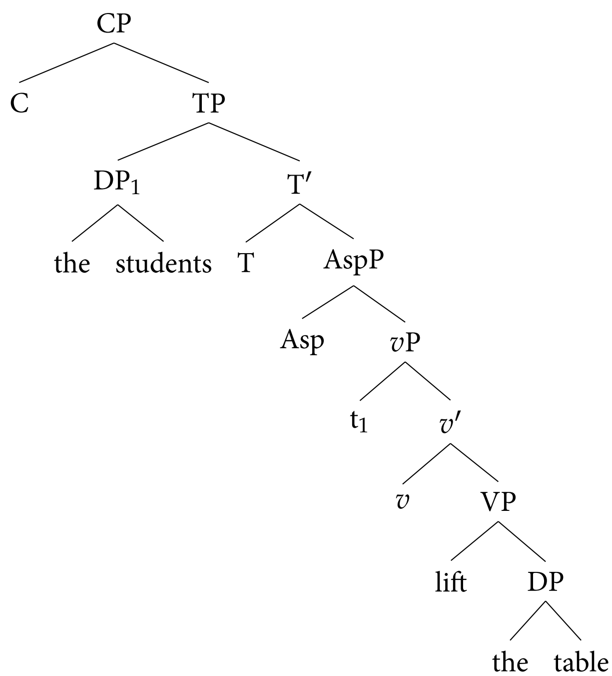

Notice that I no longer include lambda-abstraction over x_1. This is because ⟦the students⟧, unlike ⟦DIST the students⟧, is of type *e*: it is the mereological sum of all students. But this in turn means that we are back in a situation where it does not matter whether the subject is interpreted in its pre- or post-movement position, since there is no type mismatch either way. With this in mind, I will follow my harmless assumption from previous chapters that it is interpreted in its merge position.

Since the DIST-less *the students* is referential, the compositional semantics for (37) has no surprises in store for us, and leads to the result in (38):

(38) ⟦(37)⟧ = 1 iff
$$\exists e[\text{Agt}(e) = \sigma x[{}^{*}\text{student}(x)] \wedge \text{lift}(e) \wedge \text{Thm}(e) = \sigma y[\text{table}(y)]]$$

(38) predicts (21) to be true on its non-distributive reading if there is an event of lifting the table whose agent is the plural individual denoted by *the students*. This will be true if, for example, the students all cooperated in lifting the table together.

As for the non-distributive reading of (16), there are again no surprises on the compositional front. The syntactic representation will be as in (39):

(39)

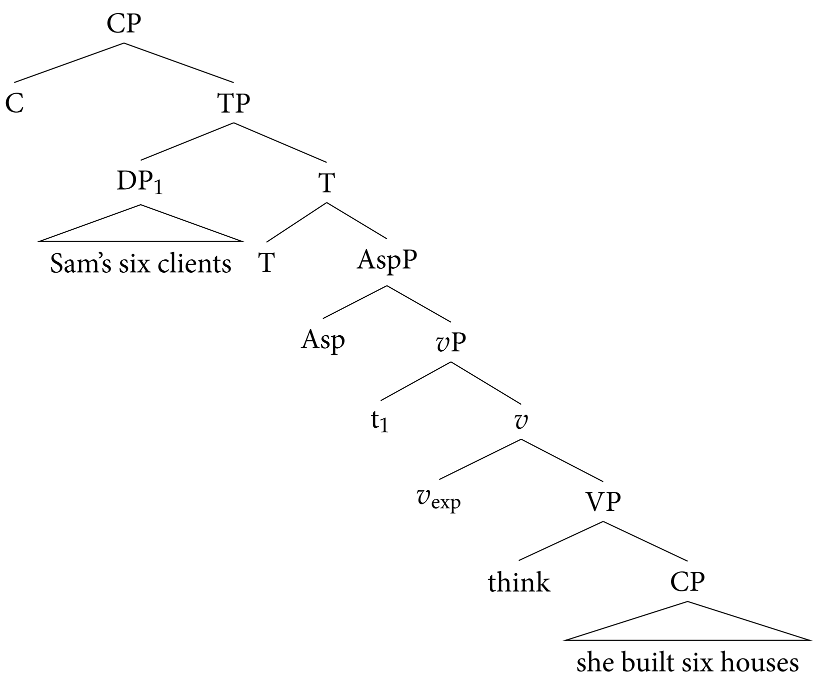

And the resulting semantic denotation given the tree in (39) is as in (40):

(40) $[\![(39)]\!] = 1$ iff $\exists e[\text{Exp}(e) = \sigma x[^*\text{client-of-sam}(x) \wedge |x| = 6] \wedge$
$\forall w \in \text{Dox}(e)[\text{six\_houses}(w)]]$

(40) states that on a non-distributive reading, (16) is true iff there is a belief state whose experiencer is the mereological sum of Sam's six clients, and is such that in all of that state's doxastically accessible worlds, Sam built six houses for her six clients.

5.3.4 Defining the problem

We now have enough of the semantics at our fingertips to define the problem at hand. Consider again the sentence in (18b):

(18b) Arnie and Beatrice think Paul married a rich New Yorker.

According to our semantics, the non-distributive reading of (18b) will be as in (41), assuming that $[\![\text{Arnie and Beatrice}]\!] = a \sqcup b$:

(41) $\exists e[\text{Exp}(e) = a \sqcup b \wedge \forall w \in \text{Dox}(e)[\text{rich\_NYer}(w)]]$

Say that e_a is Arnie's belief state, and e_b is Beatrice's belief state. A reasonable assumption would then be that the experiencer of $e_a \sqcup e_b$ is $a \sqcup b$. That is, the mereological sum of Arnie and Beatrice's belief states has as an experiencer the mereological sum of Arnie and Beatrice.[6] Thus, if (41) is true, it is presumably because $e_a \sqcup e_b$ serves as the witness to the existential quantification over eventualities, meaning that in all worlds in $\mathrm{Dox}(e_a \sqcup e_b)$, Paul married a rich New Yorker.

Now recall that the initial context for (18b) was articulated in terms of the individual beliefs of Arnie and Beatrice. The relevant portion of the provided context, along with the features of e_a and e_b that they speak to, can be seen in (42–43):

(42) a. *Arnie suspects that Paul's husband is rich…*
$\forall w \in \mathrm{Dox}(e_a)[\mathrm{rich}(w)]$

b. *…and has no other relevant opinions.*
$\exists w, w' \in \mathrm{Dox}(e_a)[\mathrm{NYer}(w) \wedge \neg\mathrm{NYer}(w')]$

(43) a. *Beatrice thinks he's a New Yorker…*
$\forall w \in \mathrm{Dox}(e_b)[\mathrm{NYer}(w)]$

b. *…and has no other relevant opinions.*
$\exists w, w' \in \mathrm{Dox}(e_b)[\mathrm{rich}(w) \wedge \neg\mathrm{rich}(w')]$

The fact that we find (18) to be true indicates that the information about $\mathrm{Dox}(e_a)$ and $\mathrm{Dox}(e_b)$ conveyed in (42–43) is somehow enough for us to make the crucial inference about $\mathrm{Dox}(e_a \sqcup e_b)$. This is precisely the problem that I will be addressing in the rest of this chapter: what is the nature of the relationship between $\mathrm{Dox}(e_a)$, $\mathrm{Dox}(e_b)$, and $\mathrm{Dox}(e_a \sqcup e_b)$?

5.4 Premise semantics and belief-summing

In this section, I will offer a first stab at the puzzle introduced in the previous section. While the proposal has some success, I show at the end of this section that some revisions are in order due to the theory's inability to differentiate between relevant and irrelevant disagreements between experiencers.

5.4.1 A Kratzerian analogy

Recall the general observation from the previous section: when the beliefs of the individual experiencers are not in conflict, their combined belief is the *conjunction* of the individuals' beliefs. Meanwhile, if their beliefs *are* in conflict, then their combined belief is the *disjunction* of the individuals' beliefs. Or, stated in terms of the discussion above, if $\mathrm{Dox}(e_1) \cap \mathrm{Dox}(e_2) \neq \varnothing$ (beliefs do not conflict), then $\mathrm{Dox}(e_1 \sqcup e_2) = \mathrm{Dox}(e_1) \cap \mathrm{Dox}(e_2)$. But if $\mathrm{Dox}(e_1) \cap \mathrm{Dox}(e_2) = \varnothing$ (beliefs do conflict), then $\mathrm{Dox}(e_1 \sqcup e_2) = \mathrm{Dox}(e_1) \cup \mathrm{Dox}(e_2)$.

[6] See Kratzer 2003 for a similar point with respect to the agent argument.

For those familar with Kratzer's (1977, 1981a, 1991, 2012) work on modals, these facts may have a certain air of familiarity to them. Consider the following scenario from Kratzer (1991) (originally discussed in Kratzer 1977):

> Let us imagine a country where the only source of law is the judgments which are handed down. There are no hierarchies of judges, and all judgments have equal weight. There are no majorities to be considered. There is one judgment which provides that murder is a crime. Never in the whole history of the country has anyone dared to attack this judgment. Sometimes, judges have not agreed. Here is an example of such a disagreement: Judge A decided that owners of goats are liable for damage their animals inflict on flowers and vegetables. Judge B handed down a judgment providing that owners of goats are not liable for damage caused by their animals. Owners of gardens have to construct adequate fencing. This means that the set of propositions corresponding to the judgments handed down in the country we are considering is an inconsistent set of propositions.
>
> (Kratzer 1991: 642)

As Kratzer notes, in this scenario, the following sentences are true:

(44) (In view of what the judgments provide,) Murder is necessarily a crime.

(45) (In view of what the judgments provide,)

a. Owners of goats are possibly liable for damage caused by their animals.

b. Owners of goats are possibly not liable for damage caused by their animals.

In terms of quantification over worlds, the truth of the sentences in (44–45) suggests that in all accessible worlds, murder is a crime, but that the set of accessible worlds is a mixture of worlds where goat owners are and are not liable. Thus, the set of accessible worlds is decided with respect to murder, where there is no conflict, but in the case of the conflicting views on goat owner liability, we have disjunction: there are some worlds where owners are liable, and some where they are not.

I'd like to revise Kratzer's scenario slightly, for two reasons. The first is that her example is in a sense doubly modalized: on the one hand, there are the adverbial epistemic modals *necessarily* and *possibly*, and on the other there are the inherently deontic concepts of crime and liability. Thus, to simplify the example a bit, I will ditch the talk of crime. The second reason for revising the scenario is that the disagreement about goat owner liability exhausts all of the logically available options: goat owners simply either are or are not responsible for their animals. In order to show that disjunction arises in cases of conflict, it will help to have a scenario where the conflicting views do not collective exhaust the logical space.

With this in mind, say that laws are not determined by judges, but are instead codified in a national constitution. There has just been an election, where citizens voted both for president and for local representatives. In interpreting election law, each judge has a say about the presidential election (with equal weight, no majorities, etc.), but for local elections each judge only makes judgments about their own district. Due to certain facts on the ground, in conjunction with vagueness in the constitution, it is not entirely settled who has won the presidency: some judges believe that Paulson won, while others believe it's Quincy. None of the judges have decided in favor of any of the other presidential candidates. But in the election for the capital district's local representative, things are cut and dry: all the local judges agree that Rainier has won.

Following Kratzer's example, in this scenario the sentences in (46–47) are all true:

(46) (In view of what the judgments provide,) Rainier certainly won the capital district election.

(47) (In view of what the judgments provide,)

  a. Paulson may have won the presidency.

  b. Quincy may have won the presidency.

Moreover, the following sentence is interpreted as true in this scenario:

(48) (In view of what the judgments provide,) It is certain that Rainier won the capital district and Paulson or Quincy won the presidency.

There is a striking similarity between this case and (20), repeated below:

> *Paul has three cousins, Arnie, Beatrice, and Kate. Arnie still thinks Paul married a rich man, and Beatrice thinks he married a New Yorker. Kate, like Beatrice, is unopinionated about Paul's husband's wealth, but she thinks he's from California, not New York.*

(20) Paul's cousins think he married a rich man from either California or New York.

In both scenarios, one judgment goes unquestioned—Rainier's victory in the judge scenario, Paul's husband's wealth in the wedding scenario—and one judgment involves a split decision between two options. Moreover, in both cases, the result is the same: the uncontroversial belief/judgment is conjoined with the disjunction of the two opposing views.

It therefore makes sense to explore whether Kratzer's solution to the judge puzzle will work for our non-distributively ascribed beliefs. Recall that for Kratzer, worlds are ordered by means of a set of propositions generated by the ordering source, as follows (see Lewis 1981, Kratzer 1981a):

(49) $w_1 \precsim_Q w_2$ iff $\{p \in Q \mid p(w_1)\} \supseteq \{p \in Q \mid p(w_2)\}$

Necessity modals (e.g., *must*) universally quantify over ideal worlds with respect to this ordering in a given domain, while possibility modals existentially quantify over this same set of worlds. Now say that p is the set of worlds in which Paulson wins the presidential race, q is the set of worlds in which Quincy wins the presidential race, and r is the set of worlds in which Rainier wins the capital district's local election. Assume that it is not possible for both Paulson and Quincy to win the presidential election. What will $\precsim_Q$ look like if $Q = \{p, q, r\}$? Figure 5.1 provides an illustration, where better worlds are toward the top, and the set of ideal worlds is surrounded in a blue ellipse. Equivalence classes of worlds are indicated by a list of those propositions in Q true in those worlds.

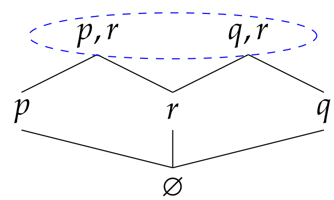

Figure 5.1: The ranking of worlds as provided by the judges.

Notice that all ideal worlds are ones in which Rainier wins, some are worlds in which Paulson wins, some are worlds in which Quincy wins, and none are worlds in which anyone else wins. Since Rainier wins in all ideal worlds, (46) is true, and since Paulson and Quincy each have some ideal worlds in which they win, (47) are true. Furthermore, (48) is also true, since all ideal worlds are such that Rainier and either Paulson or Quincy win. Thus, when it comes to the judge scenario, the Lewis-Kratzer premise-based world-ordering is a success.

So far, so good. Next, let us try using the same formal principle with non-distributive belief ascriptions.

5.4.2 Extension to belief-summing

Given that the set of doxastically accessible worlds is itself a proposition, I propose that when going from individuals' beliefs to summed beliefs, the individuals' belief worlds relate to each other in the same way that the propositions used to order worlds did above. Thus, I adopt the principle in (50), where E is a set of belief states:

(50) $\mathrm{Dox}(\bigsqcup E) = \textsc{Best}_1(\precsim_{\{\mathrm{Dox}(e) \mid e \in E\}})$ (where $\textsc{Best}_1(\precsim) = \{w \mid \neg\exists w'[w' \prec w]\}$)

Let us go through the examples from Section 5.2 one by one in order to see how this works. First, (16), repeated below:

Sam had six clients, none of whom knew of the others' existence. She convinced each of her six clients that she would build a house for him. In reality, she was a con artist and built no houses at all.

(16) (In total,) Sam's six clients thought she built six houses for them.

Let k_1, k_2, etc. be Sam's clients. Furthermore, let e_1, e_2, etc. be the belief states of k_1, k_2, etc. In all the worlds in $\mathrm{Dox}(e_1)$, k_1 gets a house, but whether or not anyone else gets a house is undetermined: in some worlds, other people get houses, and in others, not. *Mutatis mutandis* for the other clients and belief states. Now let e_{1-6} be the sum of the clients' belief states. What do we predict $\mathrm{Dox}(e_{1-6})$ to be? We predict it to be the set of ideal worlds as determined by the premise set $\{\mathrm{Dox}(e_1), \mathrm{Dox}(e_2), \mathrm{Dox}(e_3), \mathrm{Dox}(e_4), \mathrm{Dox}(e_5), \mathrm{Dox}(e_6)\}$. Since these six propositions are mutually compatible, the set of ideal worlds will simply be their intersection. Thus, all worlds in $\mathrm{Dox}(e_{1-6})$ will be such that all six clients have a house built by Sam. Since the experiencer of e_{1-6} is the sum of the experiencers of e_1 through e_6, it is indeed true that there is a belief state whose experiencer is the sum of Sam's clients, and whose ideal worlds are all such that Sam builds six houses for her six clients. We therefore rightly predict (16) to be true.

Similar facts hold for the analogous case of (18), repeated below:

Paul just got married, and his cousins Arnie and Beatrice, who have never met, just caught wind of it. Arnie suspects that Paul's husband is rich, and has no other relevant opinions. Beatrice thinks he's a New Yorker, and has no other relevant opinions.

(18) a. Paul's cousins think he married a rich New Yorker.
b. Arnie and Beatrice think Paul married a rich New Yorker.

Suppose once again that e_a and e_b are Arnie and Beatrice's belief states. $\mathrm{Dox}(e_a \sqcup e_b)$ is predicted to be the set of ideal worlds with respect to the premise set $\{\mathrm{Dox}(e_a), \mathrm{Dox}(e_b)\}$. Since Arnie and Beatrice's beliefs are once again mutually compatible, this means that $\mathrm{Dox}(e_a \sqcup e_b) = \mathrm{Dox}(e_a) \cap \mathrm{Dox}(e_b)$. Because all worlds in this set are worlds in which Paul marries a rich New Yorker, there is in fact a state whose experiencer is $a \sqcup b$, and that is a state of believing that Paul married a rich New Yorker. The sentences in (18) are thus true.

Next up is our first case of disagreement, (19):

Arnie thinks that Paul married a rich Marylander, while Beatrice thinks he married a poor New Yorker.

(19) a. Paul's cousins think he married either a rich Marylander or a poor New Yorker.

b. Arnie and Beatrice think Paul married either a rich Marylander or a poor New Yorker.

Now $\text{Dox}(e_a)$ and $\text{Dox}(e_b)$ are disjoint. Thus, $\precsim_{\{\text{Dox}(e_a),\text{Dox}(e_b)\}}$ will order worlds as in Figure 5.2. We therefore predict $\text{Dox}(e_a \sqcup e_b)$ to be the union of $\text{Dox}(e_a)$ and $\text{Dox}(e_b)$, and not their (empty) intersection. Since all of the worlds in $\text{Dox}(e_a)$ are worlds in which Paul marries a rich Marylander, while those in $\text{Dox}(e_b)$ are worlds in which he marries a poor New Yorker, the worlds in $\text{Dox}(e_a \sqcup e_b)$ will all be worlds in which Paul marries either a rich Marylander or a poor New Yorker, thereby verifying (19). Furthermore, in this context, none of the worlds in $\text{Dox}(e_a \sqcup e_b)$ is a rich New Yorker world, meaning that we also rightly predict (18) to be false here.

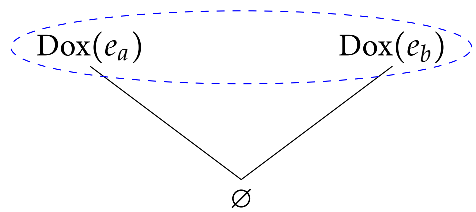

Figure 5.2: World-ordering for Arnie and Beatrice disagreement scenario.

Finally, there's the interesting case of (20). Here, Arnie's beliefs are compatible with each of Beatrice and Kate's, but the latter two have beliefs that are incompatible with each other. If k is Kate and e_k is her belief state, then in determining $\text{Dox}(e_a \sqcup e_b \sqcup e_k)$, we need to find the set of ideal worlds with respect to the premise set $\{\text{Dox}(e_a), \text{Dox}(e_b), \text{Dox}(e_k)\}$. Much like in the judgment scenario above, this will look like Figure 5.3.

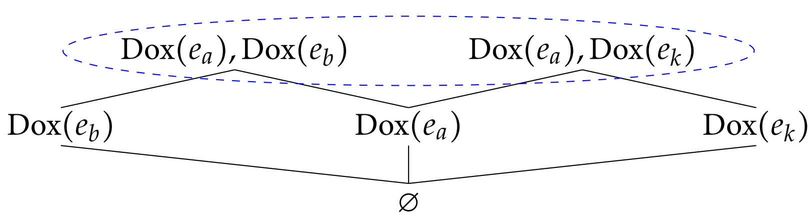

Figure 5.3: The ranking of worlds in the Arnie, Beatrice, and Kate scenario.

What we see in Figure 5.3 is that every ideal world is either in $\text{Dox}(e_a) \cap \text{Dox}(e_b)$ or $\text{Dox}(e_a) \cap \text{Dox}(e_k)$. The result is that the set of ideal worlds—and thus $\text{Dox}(e_a \sqcup e_b \sqcup e_k)$—is the set $\text{Dox}(e_a) \cap (\text{Dox}(e_b) \cup \text{Dox}(e_k))$. Since all of the worlds in $\text{Dox}(e_b)$ are worlds in which Paul's husband is from New York, and the worlds in $\text{Dox}(e_k)$ are worlds in which he's from California, all of the worlds in $\text{Dox}(e_b) \cup \text{Dox}(e_k)$ are such that Paul's husband is from New York or California. Since all worlds in $\mathrm{Dox}(e_a)$ are worlds in which the husband is rich, what we get is that all of the worlds in $\mathrm{Dox}(e_a) \cap (\mathrm{Dox}(e_b) \cup \mathrm{Dox}(e_k))$ are such that Paul's husband is a rich man from California or New York. Thus, there is a belief state whose experiencer is $a \sqcup b \sqcup k$, and is a state of believing that Paul married a rich man from California or New York: namely, $\mathrm{Dox}(e_a \sqcup e_b \sqcup e_k)$. (20) is true.

We thus see that using a Lewis-Kratzer premise semantics as a means of negotiating conflicts (and non-conflicts) between belief states garners us considerable success in accounting for the cases of non-distributively ascribed belief discussed in the previous section. Unfortunately, however, our work is not done: there is a further issue that needs to be addressed.

5.5 (Ir)relevance and aboutness

5.5.1 The problem of irrelevant disagreement

Our analysis from above crucially relied on the apparent observation that if $\mathrm{Dox}(e_1) \cap \mathrm{Dox}(e_2) = \varnothing$, then $\mathrm{Dox}(e_1 \sqcup e_2) = \mathrm{Dox}(e_1) \cup \mathrm{Dox}(e_2)$. That is, if two experiencers disagreed—if there was some proposition p such that e_1 was a state of believing p, and e_2 was a state of believing $-p$—then their beliefs would be disjoined, and not conjoined. But notice that by the principle in (50), this is true *no matter what p is*. In the examples from Section 5.2, disagreements were relevant. But in the new context for (18) provided below, we see a case where there is no relevant conflict between Arnie and Beatrice's beliefs, but there is a completely irrelevant one:

> *Paul just got married, and his cousins Arnie and Beatrice, who have never met, just caught wind of it. Arnie suspects that Paul's husband is rich, and has no other relevant opinions. Beatrice thinks he's a New Yorker, and has no other relevant opinions. In addition, Arnie mistakenly believes that Mozart was born in 1755, while Beatrice correctly believes him to have been born in 1756.*

The sentences in (18) are no less true in this context than they were in the first context in which they were presented. However, we currently predict them to be false. Every world in $\mathrm{Dox}(e_a)$ is such that Mozart was born in 1755. Every world in $\mathrm{Dox}(e_b)$ is such that he was born in 1756. Thus, $\mathrm{Dox}(e_a)$ and $\mathrm{Dox}(e_b)$ are disjoint, and $\mathrm{Dox}(e_a \sqcup e_b)$ should be $\mathrm{Dox}(e_a) \cup \mathrm{Dox}(e_b)$. But then (18) should be false. After all, since Arnie was agnostic about Paul's husband's hometown, there will be worlds compatible with Arnie's beliefs in which he is from, say, California. But since we are taking the union of $\mathrm{Dox}(e_a)$ and $\mathrm{Dox}(e_b)$, these worlds will also be in $\mathrm{Dox}(e_a \sqcup e_b)$, meaning that not all belief worlds will be worlds in which Paul's husband is a New Yorker. The same story can be told for Paul's husband's wealth. So we predict that (18) are false, and that the strongest claim that we can make is the weaker (51):

(51) Paul's cousins think he married a rich man or a New Yorker.

The problem that faces us, then, is the following: how can we retain our analysis's successes when it comes to the negotiation of relevant disagreements, while at the same time filtering out those disagreements that are irrelevant?

5.5.2 Beliefs are about situations

The way in which we will filter out irrelevant disagreements is by recourse to a notion of "aboutness". More specifically, we will say that beliefs are about *situations*, which I take to be partial worlds (cf. Barwise & Perry 1983; Kratzer 1989, 2002, 2017), with their own join $\sqcup_s$ and part-whole relation $\sqsubseteq_s$. (Worlds are maximal situations.) A belief state about some situation s will have among its doxastic alternatives all worlds compatible with what the experiencer believes *specifically about* s. For example, Dee's beliefs about the situation containing only Paddy's Pub and its current inhabitants might entail that Charlie is working the bar, but would not entail that her apartment has recently been broken into. But her beliefs about a situation containing her broken window and cracked safe would entail that her apartment has been broken into, but not that Charlie is working the bar.

In order to incorporate this aboutness into the semantics, the definitions of $[\![\text{believe}]\!]$ and $[\![\text{think}]\!]$ will be revised as follows, where s^c is a situation determined by context:

(52) $[\![\text{believe/think}]\!]^c_{\text{take 2}} = \lambda p \lambda e.\ \text{about}(e) = s^c \wedge \forall w \in \text{Dox}(e)[p(w)]$

Thus, when the ensuing predicate of belief states is existentially quantified over, this existential quantification will be restricted to those states that are about some particular situation as determined by the context.

Adding a notion of aboutness means adding a new dimension to belief-summing. We have previously discussed how $\text{Exp}(e_1 \sqcup e_2)$ relates to $\text{Exp}(e_1)$ and $\text{Exp}(e_2)$ (it is their sum), as well as how $\text{Dox}(e_1 \sqcup e_2)$ relates to $\text{Dox}(e_1)$ and $\text{Dox}(e_2)$ (Lewis-Kratzer premise negotiation). So how does $\text{about}(e_1 \sqcup e_2)$ relate to $\text{about}(e_1)$ and $\text{about}(e_2)$? So far as I can tell, it doesn't really matter, so long as the choice is deterministic. But at least on an intuitive level, it makes sense to again make recourse to summing. Thus, I will operate under the assumption that $\text{about}(e_1 \sqcup e_2) = \text{about}(e_1) \sqcup_s \text{about}(e_2)$.

5.5.3 Fixing the irrelevant disagreement problem

We can now move on to fixing the problem of irrelevant disagreement. As per the new scenario for (18), Arnie believes that Mozart was born in 1755. That is, the sentence in (53) is true if the context is fixed correctly:

(53) $[\![\text{Arnie believes that Mozart was born in 1755}]\!]^c = 1$ iff
$$\exists e[\text{Exp}(e) = a \wedge \text{about}(e) = s^c \wedge \forall w \in \text{Dox}(e)[1755(w)]]$$

What is a—or the—situation about which Arnie believes that Mozart was born in 1755? For our purposes it doesn't really matter. If we cannot target a specific situation, it may be that it is simply a belief about the possible world that Arnie inhabits: w_e, the world containing (the existentially quantified-over) e. All that matters is that there is some situation about which Arnie has this belief, and that this is the situation picked out by the context for s^c. We then get truth: there is a state of Arnie believing about the contextually-determined situation that Mozart was born in 1755. Similar results naturally obtain for Beatrice believing Mozart was born in 1756.

Next are the sentences in (18), whose new denotations will be as in (54):

(54) $[\![(18)]\!]^c = 1$ iff $\exists e[\text{Exp}(e) = a \sqcup b \wedge \text{about}(e) = s^c \wedge$
$$\forall w \in \text{Dox}(e)[\text{rich\_NYer}(w)]]$$

Suppose that s_a is the situation about which Arnie believes that Paul married a rich man. It might contain, say, information about Paul's exorbitant spending habits, his social embedding in wealthy circles, etc. Importantly, while Arnie believes that Mozart was born in 1755, his beliefs *about s_a in particular* do not entail that Mozart was born in 1755, as Mozart's birth date is irrelevant to Arnie's beliefs about s_a. Thus, among the worlds compatible with Arnie's beliefs about s_a, there are some in which Mozart was born in 1755, some in which he was born in 1756, and others in which he was born at other times. Similarly, say that s_b is the situation about which Beatrice believes that Paul married a New Yorker. (Perhaps it includes information about where Paul currently lives, the cultural milieu in which he was raised, etc.) Once again, Beatrice's beliefs about Mozart's birth date are not reflected in her beliefs about s_b, so that the worlds compatible with her beliefs about s_b include some 1756 worlds, but also some 1755 worlds, and some worlds where Mozart was born in other years.

Now let e_a be Arnie's belief state about s_a, and likewise for e_b, Beatrice, and s_b. Because of the filtering out of irrelevant disagreements, $\text{Dox}(e_a)$ and $\text{Dox}(e_b)$ are mutually compatible, and have a non-empty intersection. Thus, the set of ideal worlds as determined by the premise set $\{\text{Dox}(e_a), \text{Dox}(e_b)\}$ is $\text{Dox}(e_a) \cap \text{Dox}(e_b)$, which in turn by (50) means that $\text{Dox}(e_a \sqcup e_b) = \text{Dox}(e_a) \cap \text{Dox}(e_b)$. The set of worlds compatible with $\text{Dox}(e_a \sqcup e_b)$ will thus entail that Paul marries a rich New Yorker, in spite of Arnie and Beatrice's irrelevant disagreement about Mozart's birth date. Since $\text{about}(e_a \sqcup e_b) = \text{about}(e_a) \sqcup_s \text{about}(e_b)$, i.e., $s_a \sqcup_s s_b$, the sentences in (18) are therefore predicted to be true if s^c is set to $s_a \sqcup_s s_b$.

So much for irrelevant disagreements. But what about *relevant* ones? Turning to the rich Marylander/poor New Yorker scenario, say that s'_a is the situation about which Arnie believes that Paul married a rich Marylander, and s'_b the situation about which Beatrice believes that Paul married a poor New Yorker. In this case, $\mathrm{Dox}(e'_a)$ and $\mathrm{Dox}(e'_b)$, the worlds compatible with their respective beliefs about their respective situations, will still be disjoint, so we seemingly retain the results from before: disjunction, and not conjunction.

However, if s'_a is the situation about which Arnie believes that Paul married a rich Marylander, shouldn't there be some subsituation of s'_a that only contains the information that led Arnie to believe that Paul married a rich man? That is, shouldn't there be a situation like s_a from before, and likewise for Beatrice believing that Paul married a not-necessarily-poor New Yorker? In this case, we might expect (18) to come out as true, since Arnie and Beatrice's beliefs about these latter situations are indeed non-disjoint, and thus could be conjoined as before.

There are, I think, two plausible responses to this objection. The first is simply to deny the premise: maybe there is no such situation s_a, and every situation about which Arnie believes that Paul married a rich man is also a situation about which Arnie believes that he married a Marylander. That is, maybe Arnie's beliefs in wealth and Marylanderhood are, in fact, ontologically inseparable. An alternative response would be to emphasize that by our definition, the choice of about-situation is sensitive to context. By framing the context as we have, we are clearly making salient those situations about which Arnie believes that Paul married a rich Marylander (s'_a) and Beatrice believes that Paul married a poor New Yorker (s'_b), and not those about which Arnie only believes that Paul married a rich man (s_a) or about which Beatrice only believes that Paul married a New Yorker (s_b). Thus, the context is heavily—perhaps even indefeasibly—biased towards setting s^c to $s'_a \sqcup_s s'_b$, and not $s_a \sqcup_s s_b$, and we get falsehood for (18).

This second explanation raises an interesting possibility. Given that the work of determining what beliefs are relevant or not is left up to a situation variable whose value is fixed by context, one might think that there would be some cases in which there is a bit more flexibility in determining what does or does not count as relevant. I will next discuss just such a case, as well as how the proposal at hand can account for this flexibility.

5.5.4 Context-dependency

Consider the following scenario, an extension of the original house-building scenario provided for (16):

> *Each of Sam's six clients signed an exclusive contract with her, stating that she would build a house for him and him alone. Sam built no houses.*

It is clear that in this context, (55) is true:

(55) Each of Sam's six clients thought she built a house only for him.

However, suppose that all of Sam's former clients file a joint lawsuit against her for her fraudulent practices. The attorney for the clients, in arguing for a certain amount in total damages, can reasonably and truthfully say (16), repeated below:

(16) (In total,) Sam's six clients thought she built six houses for them.

This raises an apparent conundrum. In order for (55) to be true, Sam's clients must have incompatible beliefs: each client believes that he, and only he, got a house. But as seen previously, getting (16) to be true requires that the beliefs *not* be incompatible: each client has to be agnostic about whether the others receive a house in order for the beliefs to be conjoinable. So how do we reconcile these facts?

Say that Sam's six clients are k_1 through k_6. Allowing ourselves a somewhat simplistic scenario, say that each client k_i has a contract with Sam that has two clauses: (I) Sam will build a house for k_i, and (II) Sam won't build a house for anyone but k_i. For each client k_i, let s_i be the situation that just contains the part of their contract including the first clause, while s_i' contains both clauses, as in Figure 5.4.

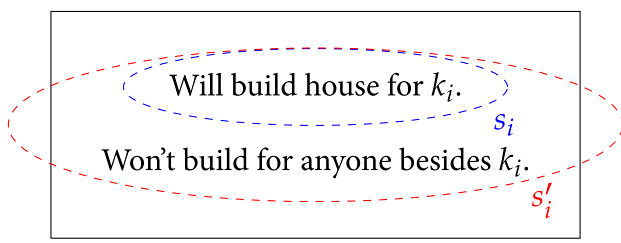

Figure 5.4: About-situations for house-building, with and without exclusivity.

First, let's tackle (55). Each client k_i's beliefs about the larger situation s_i' entail that k_i and only k_i is getting a house from Sam. But in order for this fact to have an appropriate impact on the semantics, we'll need to slightly revise our definitions of $[\![\text{believe}]\!]$ and $[\![\text{think}]\!]$. To see why, consider (56), which is what we currently predict for $[\![(55)]\!]$. (I adopt the assumption that $[\![\text{each of DP}]\!]$ is equivalent to $[\![\textsc{dist}\ \text{DP}]\!]$. only_house$(x, w)$ is true iff Sam built a house for x and only x in w.)

(56) $$\forall x[(\text{Atom}(x) \wedge x \sqsubseteq \sigma y[{}^*\text{client-of-sam}(y) \wedge |y| = 6]) \rightarrow \exists e[\text{Exp}(e) = x \wedge \text{about}(e) = s^c \wedge \forall w \in \text{Dox}(e)[\text{only\_house}(x, w)]]]$$

This is equivalent to (57) if the σ operator is well-defined:

(57) $$\forall x[\text{client-of-sam}(x) \rightarrow \exists e[\text{Exp}(e) = x \wedge \text{about}(e) = s^c \wedge \forall w \in \text{Dox}(e)[\text{only\_house}(x, w)]]]$$

(56–57) require that all of the clients' beliefs be about the same situation (s^c). But in our scenario, k_1's belief is about s_1', k_2's is about s_2', etc. To allow for this possibility, I will revise the definition of *believe*/*think* as in (58), where S^c is a contextually determined *set* of possible about-situations.

(58) $[\![\text{believe/think}]\!]^c_{\text{take 3}} = \lambda p \lambda e.\ \text{about}(e) \in S^c \wedge \forall w \in \text{Dox}(e)[p(w)]$

Note that the cases discussed earlier in this section can still be accounted for by replacing each relevant about-situation s with the singleton set $\{s\}$ of potential about-situations.

As a result of this change, the new denotation for (55) will be as in (59):

(59) $[\![(55)]\!]^c = 1$ iff $\forall x[\text{client-of-sam}(x) \rightarrow \exists e[\text{Exp}(e) = x \wedge \text{about}(e) \in S^c \wedge \forall w \in \text{Dox}(e)[\text{only\_house}(x, w)]]]$

Now suppose that $S^c = \{s'_1, s'_2, s'_3, s'_4, s'_5, s'_6\}$. In this case, it is indeed true that for each client k_i, there is some situation in S^c—namely s'_i—such that k_i's beliefs about that situation entail that k_i and only k_i got a house. We thus rightly predict (55) to be true.

Let us move on to (16). We now predict the truth conditions for the non-distributive reading of (16) to be as in (60):

(60) $\exists e[\text{Exp}(e) = \sigma y[{}^*\text{client-of-sam}(y) \wedge |y| = 6] \wedge \text{about}(e) \in S^c \wedge \forall w \in \text{Dox}(e)[\text{six\_house}(w)]]$

In the case of (55), the exclusivity clause and the clients' belief in its sincerity were pragmatically relevant. But in the context of tabulating damages, exclusivity doesn't matter: we just care about the number of broken house-building promises. The situations that are relevant are thus not s'_1, s'_2, etc., but instead s_1, s_2, etc. With this in mind, let e_1 be k_1's belief state about s_1, e_2 k_2's belief state about s_2, etc. Because s_1 does not contain the exclusivity clause, $\text{Dox}(e_1)$ will entail that k_1 gets a house, but not that *only* k_1 gets a house. We thus return to the original formulation of the housing scenario, in that $\text{Dox}(e_1)$, $\text{Dox}(e_2)$, etc. are all mutually compatible, with their intersection entailing that all six clients get a house. Thus, if $e_{1-6} = e_1 \sqcup e_2 \sqcup \ldots \sqcup e_6$, then $\text{Dox}(e_{1-6})$ will indeed entail that Sam built houses for her six clients. Since $\text{about}(e_{1-6}) = s_1 \sqcup_s s_2 \sqcup_s \ldots \sqcup_s s_6$, setting S^c to the singleton set containing this situation will render (60), and thus (16), true: e_{1-6} is a state that (I) has the sum of Sam's clients as an experiencer, (II) is about the sole situation in S^c, and (III) has as its doxastic alternatives a set of worlds that entails that all six of Sam's clients have houses built for them.

It thus seems that the resort to contextually-determined (sets of) about-situations was well-founded. On an intuitive level, the facts discussed above suggest that what is considered relevant or irrelevant is at least partially determined by the context, and that determinations about what counts as relevant can sometimes affect whether individuals' beliefs qualify as compatible or conflicting.

5.6 Conclusion

In this chapter, I have provided evidence suggesting that non-distributive ascriptions of belief are available as part of the grammar. I have also argued that the best way to account for the relationship between the desires of individuals and the desires of the pluralities in which they are contained is by means of the same formal mechanisms used in a Lewis-Kratzer premise semantics for modals and conditionals. Revisions were then made in order to account for the distinction between relevant and irrelevant conflicts in belief between individuals whose beliefs are summed. In addition to accounting for the problem of irrelevant disagreement, this proposal also made the seemingly correct prediction that what counts as relevant disagreement is a matter sensitive to conversational context.

As for future inquiry, perhaps the most obvious place to look is at other attitudes and clause-embedding verbs. For example, non-distributive readings seem to also arise with so-called *verba dicendi*, or verbs of saying: as illustrated in (61–62), examples can be constructed with *say* that are similar to the examples constructed in this chapter with belief:

Sam has six clients, who do not know each other. Each said that Sam was building a house for him.

(61) Sam's six clients said she was building six houses for them.

Paul has just gotten married. Arnie claims that Paul married a rich man, but has made no other claims about Paul's husband. Beatrice claims that Paul married a New Yorker, and has made no other relevant claims.

(62) Paul's cousins said he married a rich New Yorker.

Beyond serving as an intriguing extension of the analysis in this chapter, non-distributive speech ascriptions could serve as an interesting lens through which to look at the distinction between *saying* and *saying that*. That is, what does the mereological structure of utterance content have to tell us about the relationship between the content of a speech act and the physical event of speech itself?

In addition, as can be seen in (63–64), similar examples can be concocted with desiderative attitudes like *want*:

Sam has six clients, who do not know each other. Each has asked Sam to build a house for him.

(63) Sam's six clients want her to build six houses for them.

Paul has decided to get married, but isn't sure who he'll marry yet. Arnie wants Paul to marry a rich man, and doesn't care about anything else. Beatrice wants him to marry a New Yorker, and doesn't care about anything else.

(64) Paul's cousins want him to marry a rich New Yorker.

Extending our analysis to *want* adds an additional layer of complication. After all, as discussed in Chapter 2, most theories of the semantics of *want* propose that belief is somehow semantically involved, but that there is also an extra layer of comparative deisirability of possible states of affairs that enters into the mix. Thus, figuring out how the desires of individuals relate to the desires of pluralities requires an understanding not just of how their beliefs are negotiated, but how their comparative assessments of idealness are, as well as whether and how the two relate to each other.

Despite all of the gaps that remain, as well as the tentative nature of the proposal at hand, one hopes that the work in this chapter might serve as a foundation for future research in this previously unexplored area.

Chapter 6: Conclusion

In this dissertation I have argued that in the model used for semantic interpretation, attitude states—states of belief, desire, etc.—have non-trivial part-whole relations beyond mere temporal extension, and that these part-whole relations in turn have semantic repercussions. Since this proposal takes us into some fairly new territory, I have had to leave a lot to future work. With this in mind, in tying a bow on this dissertation, I would like to offer a few possible directions for further exploration, beyond those discussed in previous chapters.

Perhaps the most obvious path to revising and extending the analyses adopted here is to expand the set of attitudes accounted for, and perhaps to broaden our focus to include non-attitudes with certain attitude-like semantic components to them. I have focused on a very narrow set of attitudes: namely, *believe/think*, *want*, *wish*, and *regret*. The main reason for this is that these verbs, in addition to a few others like *hope* and *doubt*, have occupied the bulk of the attention of semantic work on attitudes. But what about verbs like *suspect*, *suppose*, *consider*, *prefer*, and *require*? Or verbs of saying, such as *say*, *claim*, and *plead*? Or adjectives with optional clausal modifications like *glad* (e.g., *glad that he left*) and *sad*? Or perhaps even modal auxiliaries like *must*, *should*, and *can*? Looking at a broader range of lexical items with propositional arguments increases the likelihood of finding interesting generalizations or counterexamples, and encourages a greater understanding of additional potential influences from things like grammatical category and semantic type.

Another possible avenue for future research is determining whether there is any relationship between the non-distributively-ascribed attitudes discussed in Chapter 5 and the notion of common ground (or context set) as it appears in the pragmatics literature. After all, the common ground is often discussed in terms of the collective commitments of the conversational participants, a notion with intuitive parallels to beliefs ascribed to non-atomic individuals. That being said, there are well-known cases where interlocutors accept a proposition into the common ground without necessarily committing themselves to actually believing that proposition. But even so, there is a certain relation to belief: while accepting something into the common ground does not commit an interlocutor to believing it, it does seem to commit them to behaving *as if* they believe it. If this is so, then we can reasonably ask whether the feigned beliefs of the individual discourse participants have a similar relationship to the context set as the beliefs of individual epistemic agents have to the beliefs of their pluralities.

Finally, one might wish to explore the extent to which the natural language metaphysics for attitude states proposed in this dissertation aligns with language-external folk metaphysics or folk psychology. In general, when a semantic theory adopts some ontological commitment or other, a reasonable follow-up question to ask is where that ontological commitment comes from. One possibility is simply that there is a complete or near-complete match between one's natural language metaphysics and one's language-external folk metaphysics. But whether this is actually the case is an empirical question. The possibility that there is some structure in the natural language metaphysics that does not exist in one's folk metaphysics seems a bit unlikely, but not wholly implausible. What seems more likely is that there is some structure in the folk metaphysics that does not make its way to the natural language metaphysics, so that the model used for semantic interpretation lacks some of the richness of structure seen in general cognition. If this is the case, then the natural next step is to determine which conceptual structures successfully make the journey from folk ontology to natural language ontology, which ones get left behind, and why.

* * *

Often the value of a line of research lies as much in the questions it raises as in the theories it births. In the case of the work in this dissertation, this may prove a somewhat self-serving perspective. In the likely scenario that many of the central ideas in this dissertation will soon need to be altered or discarded, hopefully the novel questions raised and new paths cleared will have a slightly more permanent value. Either way, there is clearly much work to be done.

Bibliography

Abusch, Dorit. 2010. Presupposition triggering from alternatives. *Journal of Semantics* 27(1). 37–80.

Alonso-Ovalle, Luis. 2006. *Disjunction in alternative semantics*. Amherst, MA: University of Massachusetts Amherst PhD dissertation.

Anand, Pranav & Valentine Hacquard. 2008. Epistemics with attitude. In Tova Friedman & Satoshi Ito (eds.), *Semantics and Linguistic Theory (SALT)* 18, 37–54.

Anand, Pranav & Valentine Hacquard. 2013. Epistemics and attitudes. *Semantics and Pragmatics* 6(8). 1–59.

Asher, Nicholas. 1987. A typology for attitude verbs and their anaphoric properties. *Linguistics and Philosophy* 10(2). 125–197.

Bach, Emmon. 1986a. The algebra of events. *Linguistics and Philosophy* 9. 5–16.

Bach, Emmon. 1986b. Natural language metaphysics. In R. Barcan Marcus, G.J.W. Dorn & P. Weingartner (eds.), *Logic, Methodology, and Philosophy of Science VII*, 573–595. Amsterdam: Elsevier.

Bale, Alan & David Barner. 2009. The interpretation of functional heads: Using comparatives to explore the mass/count distinction. *Journal of Semantics* 26(3). 217–252.

Barwise, Jon & John Perry. 1983. *Situations and Attitudes*. Cambridge, MA: MIT Press.

Beck, Sigrid. 2010. Quantifiers in *than*-clauses. *Semantics & Pragmatics* 3(1). 1–72.

Beck, Sigrid. 2011. Comparison constructions. In Klaus von Heusinger, Claudia Maienborn & Paul Portner (eds.), *Semantics: An International Handbook of Natural Language Meaning* Vol. 2, 1341–1390. Berlin: de Gruyter.

Beck, Sigrid. 2014. Plural predication and quantified 'than'-clauses. In Luka Crnič & Uli Sauerland (eds.), *The Art and Craft of Semantics: A Festschrift for Irene Heim*, 91–115. Cambridge, MA: MIT Working Papers in Linguistics.

Beck, Sigrid, Toshiko Oda & Koji Sugisaki. 2004. Parametric variation in the semantics of comparison: Japanese vs. English. *Journal of East Asian Linguistics* 13(4). 289–344.

Bennett, Michael & Barbara Partee. 1972. Toward the logic of tense and aspect in English. Santa Monica, CA: System Development Corporation.

Bhatt, Rajesh. 1998. Obligation and possession. In Heidi Harley (ed.), *Papers from the UPenn/MIT roundtable on argument structure and aspect* MITWPL 32, 21–40. Cambridge, MA: MIT Press.

Bhatt, Rajesh & Roumyana Pancheva. 2004. Late merger of degree clauses. *Linguistic Inquiry* 35(1). 1–45.

Bogal-Allbritten, Elizabeth. 2016. *Building meaning in Navajo*. Amherst, MA: University of Massachusetts Amherst PhD dissertation.

Brasoveanu, Adrian. 2009. Measure noun polysemy and monotonicity: Evidence from Romanian pseudopartitives. In Anisa Schardl, Martin Walkow & Muhammad Abdurrahman (eds.), *Proceedings of the Thirty-Eighth Conference of the North East Linguistics Society (NELS 38)*, Amherst, MA: GLSA (Graduate Linguistic Student Association, University of Massachusetts).

Brennan, Virginia. 1993. *Root and epistemic modal auxiliary verbs*. Amherst, MA: University of Massachusetts Amherst PhD dissertation.

Bresnan, Joan. 1973. Syntax of the comparative clause construction in English. *Linguistic Inquiry* 4(3). 275–343.

Cariani, Fabrizio, Magdalena Kaufmann & Stefan Kaufmann. 2013. Deliberative modality under epistemic uncertainty. *Linguistics and Philosophy* 36(3). 225–259.

Carlson, Greg. 1984. Thematic roles and their role in semantic interpretation. *Linguistics* 22(3). 259–279.

Carlson, Greg. 1998. Thematic roles and the individuation of events. In Susan Rothstein (ed.), *Events and Grammar*, 35–51. Dordrecht: Springer.

Castañeda, Hector-Neri. 1967. Comments on D. Davidson's "The Logical Form of Action Sentences". In N. Rescher (ed.), *The Logic of Decision and Action*, 104–112. Pittsburgh: Pittsburgh University Press.

Champollion, Lucas. 2015a. The interaction of compositional semantics and event semantics. *Linguistics and Philosophy* 38(1). 31–66.

Champollion, Lucas. 2015b. Stratified reference: the common core of distributivity, aspect and measurement. *Theoretical Linguistics* 41(3–4). 109–149.

Chierchia, Gennaro. 1989. Anaphora and attitudes *de se*. In Renate Bartsch, Johan van Benthem & Peter van Emde Boas (eds.), *Semantics and Contextual Expression*, 1–31. Dordrecht: Foris Publications.

Chierchia, Gennaro. 1998. Plurality of mass nouns and the notion of "semantic parameter". In Susan Rothstein (ed.), *Events and Grammar*, 53–103. Dordrecht: Kluwer.

Chierchia, Gennaro, Danny Fox & Benjamin Spector. 2012. Scalar implicature as a grammatical phenomenon. In Klaus von Heusinger, Claudia Maienborn & Paul Portner (eds.), *Semantics: An International Handbook of Natural Language Meaning, Vol. 3*, 2297–2331. Berlin: de Gruyter.

Chomsky, Noam. 1977. On wh-movement. In P. Culicover, T. Wasow & A. Akmajian

(eds.), *Formal Syntax*, 71–132. New York: Academic Press.
Chung, Sandra & William Ladusaw. 2004. *Restriction and Saturation*. Cambridge, MA: MIT Press.
Condoravdi, Cleo. 2002. Temporal interpretation of modals: Modals for the present and for the past. In David Beaver, Luis Casillas Martinez, Brady Clark & Stefan Kaufmann (eds.), *The Construction of Meaning*, 59–88. CSLI Publications.
Condoravdi, Cleo & Sven Lauer. 2016. Anankastic conditionals are just conditionals. *Semantics & Pragmatics* 9(8). 1–69.
Cresswell, Maxwell J. 1976. The semantics of degree. In Barbara Partee (ed.), *Montague Grammar*, 261–292. New York: Academic Press.
Cresswell, Maxwell J. & Arnim von Stechow. 1982. *De re* belief generalized. *Linguistics and Philosophy* 5(4). 503–535.
Crnič, Luka. 2011. *Getting* even. Cambridge, MA: MIT PhD dissertation.
Davidson, Donald. 1967. The logical form of action sentences. In N. Rescher (ed.), *The Logic of Decision and Action*, 81–95. Pittsburgh: Pittsburgh University Press.
Davidson, Donald. 1969. The individuation of events. In Nicholas Rescher (ed.), *Essays in Honor of Carl G. Hempel: A Tribute on the Occation of his Sixty-Fifth Birthday*, 216–234. Springer Netherlands.
Deal, Amy Rose. 2008. Property-type objects and modal embedding. In Atle Grønn (ed.), *Proceedings of Sinn und Bedeutung 12*, Oslo: ILOS 2008.
Dowty, David. 1987. Collective predicates, distributive predicates, and *all*. In Fred Marshall, Ann Miller & Zheng-sheng Zhang (eds.), *Proceedings of the 1986 Eastern States Conference on Linguistics (ESCOL)*, 97–115. Columbus, OH: Ohio State University.
Embick, David. 2007. Blocking effects and analytic/synthetic alternations. *Natural Language & Linguistic Theory* 25(1). 1–37.
Erlewine, Michael Yoshitaka. 2007. A new syntax-semantics for the Mandarin *bi* comparative. Master's thesis, University of Chicago.
Erlewine, Michael Yoshitaka. 2017. Clausal comparison without degree abstraction in Mandarin Chinese. *Natural Language & Linguistic Theory* 1–38. Advance online publication.
von Fintel, Kai. 1999. NPI licensing, Strawson entailment, and context dependency. *Journal of Semantics* 16(2). 97–148.
von Fintel, Kai & Sabine Iatridou. 2008. How to say *ought* in foreign: The composition of weak necessity modals. In Jacqueline Guéron & Jacqueline Lecarme (eds.), *Time and Modality*, 115–141. Springer.
Fox, Danny. 2007. Free choice and the theory of scalar implicatures. In Uli Sauerland & Penka Stateva (eds.), *Presupposition and Implicature in Compositional Semantics*, 71–120. Hampshire, UK: Palgrave Macmillan.
Francez, Itamar & Andrew Koontz-Garboden. 2017. *Semantics and Morphosyntac-*

tic Variation: Qualities and the Grammar of Property Concepts. Oxford: Oxford University Press.

Gajewski, Jon. 2008. More on quantifiers in comparative clauses. In Tova Friedman & Satoshi Ito (eds.), *Semantics and Linguistic Theory (SALT)* 18, 340–357.

Geurts, Bart. 1998. Presuppositions and anaphors in attitude contexts. *Linguistics and Philosophy* 21(6). 545–601.

Giannakidou, Anastasia & Melita Stavrou. 2009. Metalinguistic comparatives and negation in Greek. In Claire Halpert, Jeremy Hartman & David Hill (eds.), *Proceedings of the 2007 workshop on Greek syntax and semantics (MIT working papers in linguistics 57)*, 57–74. Cambridge: MIT Press.

Giannakidou, Anastasia & Suwon Yoon. 2011. The subjective mode of comparison: Metalinguistic comparatives in Greek and Korean. *Natural Language & Linguistic Theory* 29(3). 621–655.

Gillon, Brendan. 1987. The readings of plural noun phrases in English. *Linguistics and Philosophy* 10(2). 199–219.

Gillon, Brendan. 1990. Plural noun phrases and their readings: A reply to Lasersohn. *Linguistics and Philosophy* 13(4). 477–485.

Goodman, Nelson. 1947. The problem of counterfactual conditionals. *The Journal of Philosophy* 44(5). 113–128.

Groenendijk, Jeroen & Martin Stokhof. 1984. *Studies on the semantics of questions and the pragmatics of answers*. Amsterdam, The Netherlands: University of Amsterdam PhD dissertation.

Hackl, Martin. 1998. On the semantics of ability ascriptions. MIT, Ms.

Hacquard, Valentine. 2006. *Aspects of modality*. Cambridge, MA: MIT PhD dissertation.

Hacquard, Valentine. 2010. On the event relativity of modal auxiliaries. *Natural Language Semantics* 18(1). 79–114.

Heim, Irene. 1983. On the projection problem for presuppositions. In D. Flickinger et al. (ed.), *Proceedings of the Second West Coast Conference on Formal Linguistics*, 114–125. Stanford, CA: Stanford University Press.

Heim, Irene. 1985. Notes on comparatives and related matters. University of Texas at Austin, Ms.

Heim, Irene. 1992. Presupposition projection and the semantics of attitude verbs. *Journal of Semantics* 9(3). 183–221.

Heim, Irene. 2000. Degree operators and scope. In B. Jackson & T. Matthews (eds.), *Semantics and Linguistic Theory (SALT)* 10, 40–64.

Heim, Irene. 2006. Remarks on comparative clauses as generalized quantifiers. MIT, Ms.

Heim, Irene & Angelika Kratzer. 1998. *Semantics in Generative Grammar*. Oxford: Blackwell.

Hintikka, Jaakko. 1969. Semantics for propositional attitudes. In *Models for Modalities*, 87–111. Dordrecht: Reidel.
Hintikka, Jaakko. 1975. Impossible possible worlds vindicated. *Journal of Philosophical Logic* 4(4). 475–484.
Iatridou, Sabine. 2000. The grammatical ingredients of counterfactuality. *Linguistic Inquiry* 31(2). 231–270.
Jackendoff, Ray. 1972. *Semantic Interpretation in Generative Grammar*. Cambridge, MA: MIT Press.
Jiang, Li Julie. 2009. Monotonicity and measure phrases in Chinese. In *Proceedings from the Annual Meeting of the Chicago Linguistics Society* 45(1), 303–317.
Kamp, Hans. 1975. Two theories about adjectives. In Edward Keenan (ed.), *Formal Semantics of Natural Language*, Cambridge: Cambridge University Press.
Karttunen, Lauri. 1974. Presupposition and linguistic context. *Theoretical Linguistics* 1(1–3). 181–194.
Katz, Graham, Paul Portner & Aynat Rubinstein. 2012. Ordering combination for modal comparison. In Anca Chereches (ed.), *Semantics and Linguistic Theory (SALT)* 22, 488–507.
Kennedy, Christopher. 2001. Polar opposition and the ontology of 'degrees'. *Linguistics and Philosophy* 24(1). 33–70.
Kennedy, Christopher. 2007. Vagueness and grammar: The semantics of relative and absolute gradable adjectives. *Linguistics and Philosophy* 30(1). 1–45.
Kennedy, Christopher & Beth Levin. 2008. Measure of change: The adjectival core of degree achievements. In Louise McNally & Christopher Kennedy (eds.), *Adjectives and Adverbs: Syntax, Semantics, and Discourse*, 156–182. Oxford: Oxford University Press.
Kennedy, Christopher & Louise McNally. 1999. From event structure to scale structure: Degree modification in deverbal adjectives. In T. Matthews & D. Strolovitch (eds.), *Semantics and Linguistic Theory (SALT)* 9, 163–180.
Kennedy, Christopher & Louise McNally. 2005. Scale structure, degree modification, and the semantics of gradable predicates. *Language* 81(2). 345–381.
Klein, Ewan. 1980. A semantics for positive and comparative adjectives. *Linguistics and Philosophy* 4(1). 1–45.
Krasikova, Sveta. 2008a. Comparison in Chinese. In Olivier Bonami & Patricia Cabredo Hofherr (eds.), *Empirical Issues in Syntax and Semantics 7*, 263–281.
Krasikova, Sveta. 2008b. Quantifiers in comparatives. In Atle Grønn (ed.), *Proceedings of Sinn Und Bedeutung* 12, 337–352.
Kratzer, Angelika. 1977. What 'must' and 'can' must and can mean. *Linguistics and Philosophy* 1(3). 337–355.
Kratzer, Angelika. 1981a. The notional category of modality. In H.J. Eikmeyer & H. Rieser (eds.), *Words, Worlds, and Contexts: New Approaches in Word Semantics*,

38–74. Berlin: de Gruyter.
Kratzer, Angelika. 1981b. Partition and revision: The semantics of counterfactuals. *Journal of Philosophical Logic* 10(2). 201–216.
Kratzer, Angelika. 1989. An investigation of the lumps of thought. *Linguistics and Philosophy* 12(5). 607–653.
Kratzer, Angelika. 1991. Modality. In Arnim von Stechow & Dieter Wunderlich (eds.), *Semantik/Semantics: An International Handbook of Contemporary Research*, 639–650. Berlin: de Gruyter.
Kratzer, Angelika. 1996. Severing the external argument from its verb. In Johan Rooryck & Laurie Zaring (eds.), *Phrase Structure and the Lexicon*, 109–137. Springer Netherlands.
Kratzer, Angelika. 1998. More structural analogies between pronouns and tenses. In Devon Strolovitch & Aaron Lawson (eds.), *Semantics and Linguistic Theory (SALT)* 8, 92–110.
Kratzer, Angelika. 2002. Facts: Particulars or information units? *Linguistics and Philosophy* 25(5/6). 655–670.
Kratzer, Angelika. 2003. *The Event Argument and the Semantics of Verbs*. Book manuscript, University of Massachusetts Amherst.
Kratzer, Angelika. 2006. Decomposing attitude verbs. Handout of a talk given at The Hebrew University of Jerusalem in honor of Anita Mittwoch.
Kratzer, Angelika. 2012. *Modals and Conditionals: New and Revised Perspectives*. Oxford: Oxford University Press.
Kratzer, Angelika. 2016. Evidential moods in attitude & speech reports. Slides from a talk given at the University of Connecticut.
Kratzer, Angelika. 2017. Situations in natural language semantics. In Edward N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy* (Winter 2017), Metaphysics Research Lab, Stanford University. https://plato.stanford.edu/archives/win2017/entries/situations-semantics/.
Krifka, Manfred. 1989. Nominal reference, temporal constitution and quantification in event semantics. In Renate Bartsch, Johan van Benthem & Peter van Emde Boas (eds.), *Semantics and Contextual Expression*, 75–115. Dordrecht: Foris.
Larson, Richard. 1988. Scope and comparatives. *Linguistics and Philosophy* 11(1). 1–26.
Larson, Richard & Peter Ludlow. 1993. Interpreted logical forms. *Synthese* 95(3). 305–355.
Lassiter, Daniel. 2011a. *Measurement and modality: The scalar basis of modal semantics*. New York, NY: New York University PhD dissertation.
Lassiter, Daniel. 2011b. Nouwen's puzzle and a scalar semantics for obligation, needs, and desires. In Neil Ashton, Anca Chereches & David Lutz (eds.), *Semantics and Linguistic Theory (SALT)* 21, 694–711.

Lassiter, Daniel. 2017. *Graded Modality: Qualitative and Quantitative Perspectives*. Oxford: Oxford University Press.

Lassiter, Daniel & Noah Goodman. 2013. Context, scale structure, and statistics in the interpretation of positive-form adjectives. In Todd Snider (ed.), *Semantics and Linguistic Theory (SALT)* 23, 587–610.

Levinson, Dmitry. 2003. Probabilistic model-theoretic semantics for *want*. In Robert Young & Yuping Zhou (eds.), *Semantics and Linguistic Theory (SALT)* 13, 222–239.

Lewis, David. 1973. *Counterfactuals*. Malden, MA: Blackwell.

Lewis, David. 1979a. Attitudes *de dicto* and *de se*. *The Philosophical Review* 88(4). 513–543.

Lewis, David. 1979b. Scorekeeping in a language game. *Journal of Philosophical Logic* 8(1). 339–359.

Lewis, David. 1981. Ordering semantics and premise semantics for counterfactuals. *Journal of Philosophical Logic* 10(2). 217–234.

Link, Godehard. 1983. The logical analysis of plurals and mass terms. In R. Bäuerle, C. Shwarze & A. von Stechow (eds.), *Meaning, Use, and Interpretation of Language*, 302–323. Walter de Gruyter & Co.

Liu, Chen-Sheng Luther. 1996. A note on Chinese comparatives. *Studies in the Linguistic Sciences* 26(1–2). 217–235.

Maier, Emar. 2015. Parasitic attitudes. *Linguistics and Philosophy* 38(3). 205–236.

McNally, Louise. 2011. The relative role of property type and scale structure in explaining the behavior of gradable adjectives. In Rick Nouwen, Robert van Rooij, Uli Sauerland & Hans-Christian Schmitz (eds.), *Vagueness in Communication*, 151–168. Berlin: Springer.

Moltmann, Friederike. 2017. Cognitive products and the semantics of attitude verbs and deontic modals. In Friederike Moltmann & Mark Textor (eds.), *Act-Based Conceptions of Propositional Content*, 254–289. Oxford: Oxford University Press.

Morzycki, Marcin. 2011. Metalinguistic comparison in an alternative semantics for imprecision. *Natural Language Semantics* 19(1). 39–86.

Moulton, Keir. 2009. *Natural selection and the syntax of clausal complementation*. Amherst, MA: University of Massachusetts Amherst PhD dissertation.

Nakanishi, Kimiko. 2007. Measurement in the nominal and verbal domains. *Linguistics and Philosophy* 30(2). 235–276.

Pasternak, Robert. 2016. A conservative theory of gradable modality. In Mary Moroney, Carol-Rose Little, Jacob Collard & Dan Burgdorf (eds.), *Semantics and Linguistic Theory (SALT)* 26, 371–390.

Pasternak, Robert. in prep. Thinking alone and thinking together. Manuscript in preparation for *Semantics and Linguistic Theory (SALT)* 28.

Pasternak, Robert. in revision. A lot of hatred and a ton of desire: Intensity in the mereology of mental states. Accepted with minor revisions, *Linguistics and Phi-*

losophy.

Phillips-Brown, Milo. 2016. I want to, but…. Pre-print manuscript of a paper to appear in *Proceedings of Sinn und Bedeutung 21*.

Piñon, Christopher. 2005. Adverbs of completion in an event semantics. In Henk Verkuyl, Henriëtte de Swart & Angeliek van Hout (eds.), *Perspectives on Aspect*, 149–166. Dordrecht: Springer.

Portner, Paul. 2009. *Modality*. Oxford: Oxford University Press.

Portner, Paul & Aynat Rubinstein. 2016. Extreme and non-extreme deontic modals. In Nate Charlow & Matthew Chrisman (eds.), *Deontic Modality*, 256–282. Oxford: Oxford University Press.

Reisinger, Drew. 2016. Ranked ordering sources and embedded modality. In Patrick Farrell (ed.), *Proceedings of the Linguistic Society of America* 1, 36:1–12.

Roberts, Craige. 1996. Information structure in discourse: towards an integrated formal theory of pragmatics. In Jae-Hak Yoon & Andreas Kathol (eds.), *Papers in Semantics* (Working Papers in Linguistics 49), Columbus, OH: The Ohio State University.

Ross, John Robert. 1969. Auxiliaries as main verbs. In W. Todd (ed.), *Studies in Philosophical Linguistics* Series 1, 77–102. Evanston, IL: Great Expectations Press.

Rotstein, Carmen & Yoad Winter. 2004. Total adjectives vs. partial adjectives: Scale structure and higher-order modifiers. *Natural Language Semantics* 12(3). 259–288.

Rubinstein, Aynat. 2012. *Roots of modality*. Amherst, MA: University of Massachusetts Amherst PhD dissertation.

Rubinstein, Aynat. 2017. Straddling the line between attitude verbs and necessity modals. In Ana Arregui, María Luisa Rivero & Andrés Salanova (eds.), *Modality Across Syntactic Categories*, 610–633. Oxford: Oxford University Press.

Rullmann, Hotze. 1995. *Maximality in the semantics of wh-constructions*. Amherst, MA: University of Massachusetts Amherst PhD dissertation.

Sauerland, Uli. 2003. A new semantics for number. In Robert Young & Yuping Zhou (eds.), *Semantics and Linguistic Theory (SALT)* 13, 258–275.

Sauerland, Uli. 2004. Scalar implicatures in complex sentences. *Linguistics and Philosophy* 27(3). 367–391.

Schein, Barry. 1993. *Plurals and Events*. Cambridge, MA: MIT Pess.

Schmitt, Viola. 2017. Cross-categorial plurality and plural composition. University of Vienna, Ms.

Schwarzschild, Roger. 1996. *Pluralities*. Dordrecht: Kluwer.

Schwarzschild, Roger. 2002. The grammar of measurement. In Brendan Jackson (ed.), *Semantics and Linguistic Theory (SALT)* 12, 225–245.

Schwarzschild, Roger. 2004. Scope splitting in the comparative. Handout from a colloquium given at MIT.

Schwarzschild, Roger. 2006. The role of dimensions in the syntax of noun phrases. *Syntax* 9(1). 67–110.

Schwarzschild, Roger. 2008. The semantics of comparatives and other degree constructions. *Language and Linguistics Compass* 2(2). 308–331.

Schwarzschild, Roger & Karina Wilkinson. 2002. Quantifiers in comparatives: A semantics of degree based on intervals. *Natural Language Semantics* 10(1). 1–41.

Stalnaker, Robert. 1968. A theory of conditionals. In Nicholas Rescher (ed.), *Studies in Logical Theory* 2 American Philosophical Quarterly Monograph Series, 98–112. Oxford: Blackwell Publishing.

Stalnaker, Robert. 1970. Pragmatics. *Synthese* 22(1–2). 272–289.

Stalnaker, Robert. 1973. Presuppositions. *Journal of Philosophical Logic* 2(4). 447–457.

Stalnaker, Robert. 1975. Indicative conditionals. *Philosophia* 5. 269–286.

Stalnaker, Robert. 1977. Pragmatic presuppositions. In Andy Rogers, Bob Wall & John P. Murphy (eds.), *Proceedings of the Texas Conference on Performatives, Presuppositions, and Implicatures*, 135–147. Arlington, VA: Center for Applied Linguistics.

Stalnaker, Robert. 1978. Assertion. In Peter Cole (ed.), *Syntax and Semantics 9: Pragmatics*, 315–322. New York: Academic Press.

Stalnaker, Robert. 1984. *Inquiry*. Cambridge, MA: MIT Press.

von Stechow, Arnim. 1984. Comparing semantic theories of comparison. *Journal of Semantics* 3(1). 1–77.

van der Sandt, Rob. 1992. Presupposition projection as anaphora resolution. *Journal of Semantics* 9(4). 333–377.

van Rooij, Robert. 2008. Comparatives and quantifiers. In Olivier Bonami & Patricia Cabredo Hofherr (eds.), *Empirical Issues in Syntax and Semantics 7*, 423–444.

Villalta, Elisabeth. 2008. Mood and gradability: an investigation of the subjunctive mood in Spanish. *Linguistics and Philosophy* 31(4). 467–522.

Wellwood, Alexis. 2014. *Measuring predicates*. College Park, MD: University of Maryland PhD dissertation.

Wellwood, Alexis. 2015. On the semantics of comparison across categories. *Linguistics and Philosophy* 38(1). 67–101.

Wellwood, Alexis, Valentine Hacquard & Roumyana Pancheva. 2012. Measuring and comparing individuals and events. *Journal of Semantics* 29(2). 207–228.

Wurmbrand, Susi. 1999. Modal verbs must be raising verbs. In S. Bird, A. Carnie, J. Haugen & P. Norquest (eds.), *Proceedings of WCCFL 18*, 599–612. Somerville, MA: Cascadilla Press.

Xiang, Ming. 2003. A phrasal analysis of Chinese comparatives. In *Proceedings from the Annual Meeting of the Chicago Linguistic Society* 39(1), 739–754.

Yalcin, Seth. 2007. Epistemic modals. *Mind* 116(464). 983–1026.

Yalcin, Seth. 2010. Probability operators. *Philosophy Compass* 5(11). 916–937.
Yalcin, Seth. 2016. Belief as question-sensitive. *Philosophy and Phenomenological Research* 93(2). 1–25.