ABSTRACT

Title of dissertation: LOCAL INFORMATION IN DISCOURSE

Jonathan Caleb Kendrick
Doctor of Philosophy, 2024

Directed by: Professor Fabrizio Cariani
Department of Philosophy

Associate Professor Alexander Williams
Department of Linguistics

This dissertation argues that the interpretation of modals, expressions like *might*, *should*, and *must*, are constrained by their local context. For epistemic modals, local contexts *bound* the admissible domains of modal quantification. In Chapter 2, we use this fact to explain why epistemic *must* is weaker than the $\Box$ operator from epistemic modal logic. For root (i.e., non-deontic) modals, local contexts *restrict* the domain of quantification. In Chapter 3, we show this yields a solution to the Samaritan Paradox concerning why deontic modals do not inherit presuppositions under entailment. In Chapter 4, we propose a solution to the "if p, ought p" problem based on default logic. According to this solution, *ought*'s ordering source consists of default rules and the domain consists of the conclusion of the defaults triggered in the local context.

Local Information in Discourse

by

Jonathan Caleb Kendrick

Dissertation submitted to the Faculty of the Graduate School of the
University of Maryland, College Park in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
2024

Advisory committee:

Associate Professor Alexander Williams, Co-Chair & Advisor
Professor Fabrizio Cariani, Co-Chair
Associate Professor Paolo Santorio
Associate Professor Eric Pacuit
Professor Valentine Hacquard
Assistant Professor Aron Hirsch
Associate Professor Jordan Boyd-Graber-Ying, Dean's Representative

Acknowledgements

I entered grad school three years removed from high school, barely old enough to drink and having never ventured outside the United States. Now six years later I'm leaving with a wife, a passport book that's falling apart, and graying hair. It takes a village (or maybe a small city?) to turn a boy into something resembling an academic and, as a result, these acknowledgements are necessarily non-exhaustive. If I missed you, I apologize.

Alexander played the biggest role in molding that goofy kid from Oklahoma, always ready to entertain my increasingly eccentric proposals and patiently supply devastating counterexamples. His story about a young Alexander attending the release of Metallica's *Ride the Lightning* is the reason why I listened to that album more than a few times writing this dissertation (after all, what's more metal than semantics?) The arrival of Paolo and Fabrizio cemented my desire to stay in philosophy. In the summer of 2019, I had serious doubts that philosophy was right for me, and was considering leaving. But, when Paolo and Fabrizio arrived, they brought an infectious zeal for research that never left. I'm very excited to see what the future holds for the department under their direction with Fabrizio as chair. Valentine set me on the path of writing this dissertation by making me give a presentation on von Fintel and Gillies's "Must … Stay … Strong" in Semantics II, and, like Alexander, has always been quick with a devastating counterexample. Eric taught me so much of what I know about logic, decision theory, game theory, social choice theory, etc. In a nearby possible world where I choose not to study semantics, I believe Eric would have chaired my dissertation. While I have known Aron the shortest of all the members of this committee, I know that he is going to be a valuable part of the meaning community at Maryland with his sunny disposition and remarkable work ethic. Finally, I would like to thank Jordan Boyd-Graber for serving as the Dean's Representative, and serving as MAQT's faculty sponsor.

I would like to thank Jéssica Mendes for being a great friend and co-author through the last several years. I would have never imagined that, when Valentine put us together for a group project in Semantics II when we were baby semanticists, we would be writing papers together and giving talks in Nice and Berkeley, a few years later. Hopefully, this is just the beginning of a long and fruitful friendship. I would like to thank members of the UMD Meaning Meeting–Omar Agha, Yunhui Bai, Sarah Boukendour, Jingyi Chen, Fedya Golosov, Mike McCourt, Mal Shah, Utku Turk, Yu'an Yang, Raven Zhang–and the members of the UMD Logic Seminar–Ilaria Canavotto, Casey Enos, Steve Kuhn, Aleks Knoks, and Masayuki Tashiro. I would especially like to single out Aleks for always being good for a drink and pleasant company no matter where we happened to cross paths–Berlin, Galway, or San Francisco. Finally, I would like to thank the many excellent teachers I've had who are not represented on this committee–Naomi Feldman, Dan Goodhue, Jeff Horty, Chris Morris, and Georges Rey.

I would like to thank the many audiences for their feedback on earlier stages of the current project. Earlier versions of the material in Chapter 2 was presented at symposiums at the 2023 Eastern APA and 2023 Central APA. The material in Chapter 4 was presented at the Institut Jean Nicod and at a colloquium at the 2024 Central APA. For their feedback on various portions of this dissertation, I would like to thank Andy Egan, David Boylan, WooJin Chung, Melissa Fusco, Thony Gillies, Arc Kocurek, Matt Mandelkern, Paul Portner, Rachel Rudolph, Henry Ian Schiller, and two anonymous reviewers for *Semantics & Pragmatics*.

I would like to thank the Maryland Academic Quiz Team (MAQT) for providing a welcome distraction from the tedium of everyday life. If it were not for Justin Hawkins and my other MAQT friends, I likely would have never come to Maryland, and embarked on this chapter of life. To Justin, Graham, and Weijia, we may have fallen short of winning ACF Nationals, but our run at the title will remain one of my fondest memories of UMD.

To Ria, Jack, and Connor, thanks for all the good times.

My partner Leah has been instrumental in making it possible for me to complete my PhD. Thank you for tolerating all the late nights and my pittance of an income! If I had done my PhD anywhere else, I would have never met you and for that I'm forever grateful. Finally, I would like to thank my mother, Elizabeth Kendrick, without whose many sacrifices, none of this would be possible. Mom, I love you to the moon and back.

Table of Contents

Acknowledgements ii

1 Introducing Locality 1

1.1 Introduction 1

1.1.1 A crash course in presupposition and local contexts 2

1.2 Local constraints on epistemics 8

1.3 Outline 11

2 Two Puzzles about Epistemic Must 15

2.1 Introduction 15

2.1.1 Strong theories 17

2.2 Two new arguments against strong *must* 22

2.2.1 The argument from indicative conditionals 22

2.2.2 The argument from anaphora 24

2.3 A weak bounded semantics of *must* 29

2.3.1 Extant weak theories 29

2.3.1.1 A simple fix that doesn't quite work 33

2.3.2 A bounded semantics 35

2.3.2.1 Deriving the anaphora data 38

2.3.2.2 Deriving indirectness as an implicature 40

2.4 Conclusion 42

3 A Local Solution to the Samaritan Paradox 45

3.1 Introduction 45

3.1.1 The real culprit: local satisfaction 46

3.1.2 Overview 48

3.2 The received solution: a diversity presupposition 49

3.2.1 Objection 1: counterfactual and factual *oughts* 50

3.2.2 Objection 2: deontic actuality entailments. 51

3.3 Indirect restriction via local contexts 53

3.3.1 A solution to the Samaritan Paradox. 55

3.3.2 Restricted readings in disjunctions and conjunctions. 55

3.3.3 The Samaritan Inference with epistemics. 57

3.3.4 A solution to Åqvist's Paradox. 59

3.3.5 The Paradox of Gentle Murder 60

3.3.6 The Local Zvolenszky Problem 63

3.4 Extending the solution to bouletic attitude verbs 65

3.5 Conclusion 68

4 Zvolenszky's Problem and Default Premise Semantics 71

4.1 Introduction 71

4.1.1 The problem in Kratzer semantics 72

4.1.1.1 The diversity solution 74

4.1.1.2 The double modal solution 76

4.2 Default premise semantics 78

4.2.1 A first-pass implementation 80

4.2.2 Conflicting defaults and moral dilemmas 81

4.2.2.1 Self-frustrating oughts 84

4.2.2.2 The problem of impermissible antecedents 85

4.2.3 Exploring the logic of default premise semantics 86

4.2.3.1 Failure of Conjunctive Sufficiency 87

4.2.3.2 Failure of Loop. 88

4.2.3.3 Failure of Agglomeration. 90

4.3 Conclusion 91

5 Conclusion 93

A Bounded semantics 95

A.1 The basic system 95

A.2 Definites and indefinites 96

B Indirect restriction semantics 99

B.1 The basic system 99

B.2 Local contexts of conditionals 100

C LFs with overt local contexts 103

C.1 The need for continuations 103

C.2 The formal system 104

C.2.1 Associating local contexts with entailments 106

D Default premise semantics 111

D.1 The basic semantics without conflict 111

D.2 Semantics with tractable conflict 112

D.3 Semantics with intractable conflict 112

Chapter 1: Introducing Locality

1.1 Introduction

This dissertation explores how local information in discourse places various constraints on the interpretation of modals. The core thesis that unifies these chapters is that modals are constrained in various ways by an utterance's *local context*. In particular, I will argue that, while local contexts *bound* the admissible domain of modal quantification for epistemic modals, they *restrict* the domain of modal quantification for root (i.e., non-epistemic) modals. The notion of a local context was introduced in the work of Stalnaker (1974) and Karttunen (1974) to explain presupposition projection. The basic idea is roughly this: if α and β are constituents of Γ such that α precedes β in discourse, then the presuppositions of β must be satisfied in its *local context* $c + \alpha$, the global context c updated with α.[1] This introductory chapter motivates the core thesis that local contexts constrain modal interpretation by discussing Mandelkern (2019a, 2023)'s treatment of Yalcin's epistemic contradictions.

[1]This is essentially the exact definition found in Karttunen, "In compound sentences, the initial context is incremented in a left-to-right fashion giving for each constituent sentence a local context that must satisfy its presuppositions" (1974, p. 187). However, I say 'roughly' for two reasons. First, this definition is asymmetric–local contexts are computed from left-to-right. However, the projection behavior of anaphoric particles like *too* have recently motivated symmetric theories of local context. Second, the definition is based on linear precedence. However, one might consider more sophisticated ways of computing local contexts based on hierarchical relations like, say, c-command (Romoli and Mandelkern, 2018).

1.1.1 A crash course in presupposition and local contexts

Speakers often mark the information they're taking for granted during conversations with presupposition triggers:

(1) a. The German linguist came to the talk. definite description

$\rightsquigarrow$ *presupposes: there's a German linguist.*

b. It was John who stole the bag. *it*-cleft

$\rightsquigarrow$ *presupposes: someone stole the bag.*

c. John knows that it's raining. factive verb

$\rightsquigarrow$ *presupposes: it's raining.*

Presuppositions are distinguished from both from the content of what is asserted and implicatures by the fact that their associated inferences survive or 'project' when their trigger is embedded in entailment-canceling environments:

(2) a. The German linguist didn't come to the talk. negation

$\rightsquigarrow$ *presupposes: there's a German linguist.*

b. Did the German linguist come to the talk? question

$\rightsquigarrow$ *presupposes: there's a German linguist.*

c. The German linguist might have come to the talk. modal

$\rightsquigarrow$ *presupposes: there's a German linguist.*

According to the semantic view of presupposition often attributed to both Strawson (1950) and Frege (1892), presupposition is a relation between sentences.[2] A sentence S presupposes a sentence S' iff whenever sentence S has a truth-value, sentence S' is true. A major stumbling block for semantic theories of presupposition is the *projection problem*–explaining how complex sentences inherit the presuppositions of their constituents. For

[2]While this view is often called the Frege-Strawson view, it's questionable that one should read Frege's discussion of "hinting" in "On Sense and Reference" as endorsing this view of presupposition (Karttunen, 2016). Likewise, while early statements of the semantic view, like van Fraassen (1968), were explicitly inspired by Strawson's discussion of definite descriptions in "On Referring," it's unclear that Strawson endorsed the view.

$\neg$	
T	F
F	T
$\star$	$\star$

$\wedge$	T	$\star$	F
T	T	$\star$	F
$\star$	$\star$	$\star$	F
F	F	F	F

$\vee$	T	$\star$	F
T	T	T	T
$\star$	T	$\star$	$\star$
F	T	$\star$	F

$\Rightarrow$	T	$\star$	F
T	T	$\star$	f
$\star$	T	$\star$	$\star$
F	T	T	T

Table 1.1: Strong Kleene

example, suppose that connectives inherit their presuppositions according to the Strong Kleene semantics given in Table 1.1. Then we would expect, for example, that presupposition projection from conjunctions would be symmetric when it isn't as (3) illustrates.

(3) a. A Maryland linguist entered the room and he insulted the Whorfian.

$\not\rightsquigarrow$ *presupposes: there is a linguist*

b. The linguist insulted the Whorfian and a Maryland linguist entered the room.

$\rightsquigarrow$ *presupposes: there is a linguist*

Given the apparent shortcomings of the semantic view of presupposition, Stalnaker (1974) attempted to provide a pragmatic theory of presupposition. He argued that presuppositions depend on the *common ground*, the propositions commonly accepted by all interlocutors.[3] According to Stalnaker, a speaker in a context presupposes that A if and only if she accepts that A is common ground in the context. This pragmatic view of presupposition is attractive, because it naturally subsumes semantic theories of presupposition given the following bridge principle:

> **Stalnaker's Bridge.** A three-valued sentence A can only be used to update the context c if every world w in c is such that the $[\![A]\!]^w \in \{0, 1\}$.[4]

This bridge principle says that a speaker cannot successfully add an utterance containing a lexically-encoded semantic presupposition to the common ground unless every world

[3] In his early work, Stalnaker (1974, 1978) characterizes this common attitude as common belief or common knowledge; however, in subsequent work Stalnaker (2002, 2014) weakens the relevant condition to common acceptance where acceptance is a doxastic state weaker than belief.

[4] While implicit in Stalnaker's work, this statement which assumes that sentences containing semantic presuppositions are three-valued comes from von Fintel (2008).

in the *context set*–the set of worlds compatible with all the propositions in the common ground–assigns the utterance a classical truth-value–that is, true or false.

However, Stalnaker (1974) offers few details about how to formulate a theory of projection based on his pragmatic theory of presupposition. In his 1974 paper, he only offers a concrete proposal about conjunction, writing:

> "Even if A is not presupposed initially, one may still assert ⌜A and B⌝ since by the time one gets to saying that B, the context has shifted, and it is by then presupposed that A." (Stalnaker, 1974, p. 60)

Stalnaker's proposal for conjunction already illustrates the power of his pragmatic theory. Since conversational update is asymmetric, occurring from left-to-right, Stalnaker's theory correctly predicts that presupposition projection in conjunction is asymmetric.

However, since Stalnaker's analysis proved difficult to extend to contexts more complicated than conjunction, Karttunen (1974) abandoned Stalnaker's original project of giving a purely pragmatic analysis of context change; instead, he specified a set of semantic rules associating certain environments with local contexts. He argued that we should "not ask what the presuppositions of a complex sentence are," but instead we should "ask what it takes to satisfy them" (p. 190). For example, in disjunctions, it often seems that what projects is the presuppositions of the second disjunct conditional on the negation of the first disjunct.[5]

(4) a. Either the king doesn't have a son or the king's son is away from home.

b. Either A or B. $\rightsquigarrow \neg A \Rightarrow$ PRESUP(B)

While in conditionals, the inference that projects is the presuppositions of the antecedent, and the presuppositions of the consequent conditional on the truth of the antecedent.

[5]There may also be an additional inference that projections–the presuppositions of the first disjunct conditional on the negation of the second disjunct.

(5) a. If the king has a son, the king's son is bald.

b. If A, then C. $\rightsquigarrow$ $\textsc{presup}(A) \wedge (A \Rightarrow \textsc{presup}(C))$

Karttunen thought that, instead of attempting to derive these local contexts from first principles à la Stalnaker, they should simply be stipulated.

In response to Karttunen's work, Heim (1982, 1983) attempted to avoid stipulating satisfaction conditions by making semantic the process of context change, originally analyzed as pragmatic in Stalnaker's theory. So, in Heim's semantics, the meaning of an expression was given by its context change potential (CCP) rather than its truth-conditions.

> **Definition 1.1.1.** (Context Change Potential). A *context change potential* (CCP) $\cdot[\cdot] : \mathcal{C} \times \mathcal{L} \rightarrow \mathcal{C}$ is a function from context-expression pairs to contexts.

By redefining meaning in terms of CCPs, Heim could then give a recursively-defined semantics which captured the local contexts posited by Karttunen:

- $c[\neg A] = c \setminus c[A]$
- $c[A \wedge B] = c[A][B]$
- $c[A \vee B] = c[A] \cup c[\neg A][B]$
- $c[A \rightarrow B] = c \setminus (c[A] \setminus c[A][B])$

The resulting account gives a theory of presupposition projection which comes close to descriptive adequacy.[6] However, in moving from a dynamic pragmatics á la Stalnaker to a dynamic semantics, the satisfaction theory loses much of its initial motivation. Soames (1989), in particular, argued that the analysis was no longer explanatory. As he points out,

[6]There are still a few residual problems. The most famous is the Proviso Problem (Geurts, 1996). For example, according to the satisfaction theory, the strongest presupposition of (a) is the material conditional ⌜the speaker picks you up at the airport $\supset$ the speaker has a wife⌝; however, this is weaker than what (a) actually presupposes.

i a Either I won't pick you up at the airport, or I'll bring my wife.

b $\rightsquigarrow$ *presupposes: the speaker has a wife.*

there's no reason we couldn't define a new connective $\wedge^*$ which is truth-conditionally equivalent to conjunction, but has different local contexts:

(6) $c[\mathrm{A} \wedge^* \mathrm{B}] = c[\mathrm{B}][\mathrm{A}]$

This rule would be a disaster. For example, it would incorrectly predict that (7a) is felicitous, while also incorrectly predicting that (7b) is infelicitous.

(7) a. # He$_i$ entered the room and [a Maryland linguist]$_i$ insulted the Whorfian.

b. [A Maryland linguist]$_i$ entered the room and he$_i$ insulted the Whorfian.

Stalnaker's dynamic pragmatics offers us a plausible story about why natural language conjunction doesn't have the local contexts of $\wedge^*$: if agents update their beliefs in real time during a conversation, then presumably they update their beliefs on the basis of the first conjunct before the second conjunct is even uttered. However, since Heim conceives of context change as a semantic process rather than a pragmatic process, it's unclear that this type of explanation is available to her.

Since Heim, the holy grail of presupposition research has been to find an algorithm, which given a sentence's LF and truth-conditions, outputs its constituents' local contexts.

> **Explanatory Challenge (Schlenker, 2008).** Find an algorithm that predicts how any operator transmits presuppositions once its syntax and its classical semantics have been specified.

Schlenker (2007, 2009, 2010a,b) proposed such an algorithm based on the notion of *transparency*. Schlenker's basic intuition is that, in a given syntactic environment, the local context should consist of the worlds in the context set which will remain in the context set, no matter how the utterance is continued.[7] So, an occurrence of an expression p in a

[7] In Schlenker's algorithm, the local context is calculated based on linear order, but we can also consider similar algorithms where the local context is calculated based on hierarchical relations, like c-command (Romoli and Mandelkern, 2018) and, in the appendix, I explore approaches where the local context is calculated based on the evaluation order (Barker, 2022).

syntactic environment a_c is *transparent* if, for any expression b' of the same type as p and any well-formed completion c', $a(p \wedge b')c' \Leftrightarrow ab'c'$. The local context, then, can be computed as the smallest transparent proposition p.[8]

Definition 1.1.2. (Local context á la Schlenker (2009, 2010a,b)). The *local context* of a proposition b occurring in a syntactic environment a_c is the smallest proposition p such that $a(p \wedge b')c' \Leftrightarrow ab'c'$ for all well-formed completions b' and c'.

For example, to calculate the local context of a right disjunct, we need to find the smallest proposition p such that $A \vee (p \wedge B') \Leftrightarrow A \vee B'$ for all B'. It's easy to show that $p = \neg\mathrm{A}$. If A is true, then, for any continuation, the disjunction will be true; so, including any p whatsoever in the right disjunct cannot make the disjunction false. On the other, if A is false, then $\ulcorner \mathrm{A} \vee \mathrm{B} \urcorner$ is true only if B is true, and since $\neg$A is true by assumption, $\ulcorner \mathrm{A} \vee \mathrm{B} \urcorner$ is true just in case $\ulcorner \neg\mathrm{A} \wedge \mathrm{B} \urcorner$ is true. So, including p in the right disjunct will not make the disjunction false.[9] It's straightforward to check that Schlenker's algorithm yields the rules for local contexts in 1.1.3.

Definition 1.1.3. (Local contexts). Let κ be an arbitrary local context. We can then recursively define the local contexts as follows:

- $[\![\neg\mathrm{A}]\!]^{\kappa}$ is defined $\Leftrightarrow$ $[\![\mathrm{A}]\!]^{\kappa}$ is defined
- $[\![\mathrm{A} \wedge \mathrm{B}]\!]^{\kappa}$ is defined $\Leftrightarrow$ $[\![\mathrm{A}]\!]^{\kappa}$ is defined and $[\![\mathrm{B}]\!]^{\kappa+\mathrm{A}}$ is defined
- $[\![\mathrm{A} \vee \mathrm{B}]\!]^{\kappa}$ is defined $\Leftrightarrow$ $[\![\mathrm{A}]\!]^{\kappa}$ is defined and $[\![\mathrm{B}]\!]^{\kappa+\neg\mathrm{A}}$ is defined
- $[\![\mathrm{A} \rightarrow \mathrm{B}]\!]^{\kappa}$ is defined $\Leftrightarrow$ $[\![\mathrm{A}]\!]^{\kappa}$ is defined and $[\![\mathrm{B}]\!]^{\kappa+\mathrm{A}}$ is defined

[8] By 'smallest,' I mean the proposition that's true at the fewest number of worlds–i.e., the logically strongest proposition. So, in this sense of 'small,' $\top$ is the largest proposition and $\bot$ is the smallest proposition.

[9] We still need to show that this is the smallest such proposition. However, for the sake of brevity, I will omit this argument. The reader can consult Schlenker (2009) for a proof.

1.2 Local constraints on epistemics

Let's turn now to the question of how local contexts interact with modals. Yalcin (2007) observed that sentences of the form $\ulcorner A \land \Diamond\neg A \urcorner$, so-called *epistemic contradictions*, like (8a), appear to be incoherent in a way that instances of Moore's Paradox, like (8b), are not.

(8) a. #It's raining and it might not be raining.

b. #It's raining and I don't know it's raining.

Epistemic contradictions appear to be genuine semantic contradictions because they remain infelicitous when embedded in counterfactual attitude verbs and conditional antecedents:

(9) a. #Suppose it is raining and it might not be.

b. Suppose it's raining and I don't know it's raining.

(10) a. #John imagines that it is raining and it might not be raining.

b. John imagines that it's raining and I don't know it's raining.

(11) a. #If it isn't raining and it might be raining, we should stay home.

b. If it isn't raining and I don't know it's raining, we should stay home.

Yalcin (2007) explains this data by adopting a *domain semantics* where *might* quantifies over an information state $s \in \mathcal{P}(W)$ which is a world-independent parameter of the index of evaluation.

(12) $[\![\text{might}]\!]^{w,s} = \lambda p_{\langle s,t \rangle}.\forall w' \in s : p(w') = 1$

This alone does does not explain the infelicity of epistemic contradictions, because epistemic contradictions are still consistent. But, following Stalnaker (1978), we can think of an assertion of A as proposal to make an information state support A.

> **Definition 1.2.1.** (Support).
>
> An information state s *supports* A ($s \models$ A) iff $\forall w' \in s : [\![\mathrm{A}]\!]^{w'} = 1$.

This gives us a natural criterion for pragmatic coherence: an utterance is pragmatically coherent if it is supported by some non-empty information state.

> **Definition 1.2.2.** (Coherence).
>
> An utterance A is *coherent* iff it's supported by a non-empty information set $s \neq \emptyset$.

So, according to Yalcin, epistemic contradictions, while semantically consistent, are not pragmatically coherent, because they're only supported by the empty information state.

Mandelkern (2019a, 2023) observes that epistemic contradictions embedded in disjunctions are incoherent:

(13) #Either I might win, but I won't, or I might lose, but I won't.

However, when coupled with Boolean disjunction, the domain semantics gets this wrong. For example, consider an information state $s = \{w_1, w_2\}$ such that ¬A is true at w_1 and A is true at w_2. It follows that $[\![\text{might A} \wedge \neg\text{A}]\!]^{w_1,s} = 1$ and $[\![\text{might} \neg\text{A} \wedge \text{A}]\!]^{w_2,s} = 1$. Hence, (13) is accepted in s. So, at the very least, the domain semantics is incompatible with Boolean disjunction.

There are even simpler cases where the domain semantics falters:

(14) #It might be raining or snowing, but it's not raining. (Roelofsen, 2016, p. 22)

This has all the hallmarks of an epistemic contradiction–it remains infelicitous when embedded under counterfactual attitude verbs and when placed in the antecedent of conditionals. However, like (13), the domain semantics incorrectly predicts that there are contexts where (14) should be acceptable.[10] For example, suppose that $s = \{w_1, w_2\}$ such

[10]This data is complicated by the fact that *might* scoping over a disjunction licenses Free Choice. Given

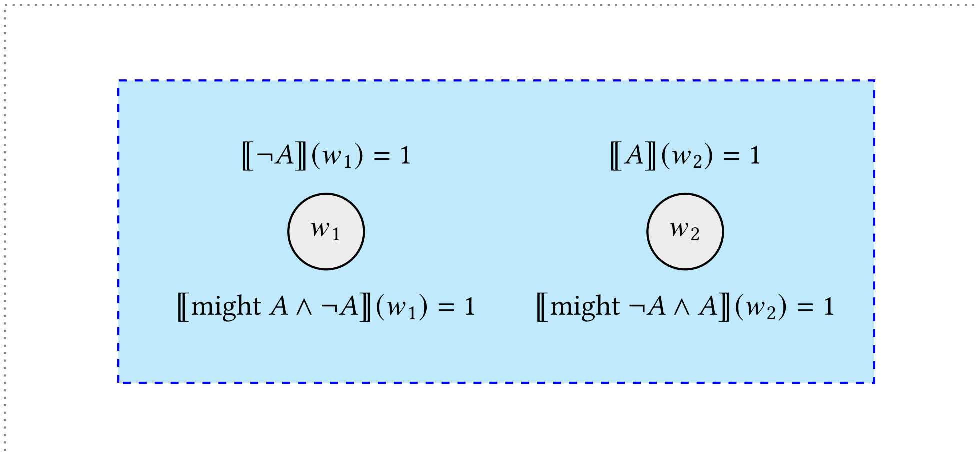

Figure 1.1: Domain semantics model for 'Wittgenstein disjunctions.'

that $\mathbf{A} = \emptyset$, $\mathbf{B} = \{w_1, w_2\}$. Then $s \models \Diamond(\mathrm{A} \vee \mathrm{B}) \wedge \neg\mathrm{A}$.

To solve problems of this sort, Mandelkern (2019a, 2023) proposes that local contexts bound the admissible domains of quantification for epistemic modals. We can recursively define a notion of an expression being assertible at a local context, what Mandelkern calls being "satt," allowing us to derive how these bounds project.

Definition 1.2.3. (Bound projection). Let κ be an arbitrary local context. We can then recursively define satt conditions as follows:

- $[\![\neg\mathrm{A}]\!]^{\kappa}$ is satt $\Leftrightarrow$ $[\![\mathrm{A}]\!]^{\kappa}$ is satt
- $[\![\mathrm{A} \wedge \mathrm{B}]\!]^{\kappa}$ is satt $\Leftrightarrow$ $[\![\mathrm{A}]\!]^{\kappa+\mathrm{B}}$ is satt and $[\![\mathrm{B}]\!]^{\kappa+\mathrm{A}}$ is satt
- $[\![\mathrm{A} \vee \mathrm{B}]\!]^{\kappa}$ is satt $\Leftrightarrow$ $[\![\mathrm{A}]\!]^{\kappa+\neg\mathrm{B}}$ is satt and $[\![\mathrm{B}]\!]^{\kappa+\neg\mathrm{A}}$ is satt
- $[\![Att_S(\mathrm{A})]\!]^{\kappa}$ is satt $\Leftrightarrow$ $\forall w \in \kappa : [\![\mathrm{A}]\!]^{Acc(S,w)}$ is satt

Mandelkern crucially assumes that local contexts are symmetric. Schlenker defines the *symmetric local context* of a proposition b occurring in a syntactic environment a_c as

how the proponent of the domain semantics explains Free Choice, it may be possible for her to correctly predict the infelicity of (14).

the smallest proposition p such that $a(p \wedge b')c \Leftrightarrow ab'c$. The difference between standard local contexts and symmetric local contexts is that the former quantify over all possible syntactic completions, while the latter is only sensitive to the *actual* completion. The need for the symmetric algorithm can be independently motivated by the projection behavior of additive particles like *too*.

(15) If Anne's brother doesn't makes a reasonable decision too, she will not study abroad.

Chemla and Schlenker (2012)

In (15), *too* is sensitive to the consequent proposition and (roughly) presupposes that Anne's brother faces the same choice between study abroad locations as Anne.

Definition 1.2.4. (Bounded *might* à la Mandelkern (2019a, 2023)).

Let $f : W \longrightarrow \mathcal{P}(W)$ be an epistemic modal base.

(16) a. $[\![\text{might}]\!]$ is satt at κ iff $\forall w : f(w) \subseteq \kappa \wedge \kappa \neq \emptyset$.

b. If satt, $[\![\text{might}]\!]^{w,f,\kappa} = \lambda p_{\langle s,t\rangle}.\exists w' f(w) : p(w') = 1$

Here's a derivation of the infelicity of (8a).

(17) $\neg$**raining** $\wedge$ $\Diamond$**raining** is satt at κ iff $\neg$**raining** is satt at $\kappa^{\Diamond\textbf{raining}}$ and $\Diamond$**raining** is satt at $\kappa^{\neg\textbf{raining}}$

Given (17), it follows that, when satt, the domain of $\Diamond$**raining** only contains $\neg$**raining**-worlds from (16a); so, (8a) is false when satt.

1.3 Outline

This dissertation explores the ways in which local information constrains modal interpretation. Our main contribution is that local information interacts with modals in ways that extend beyond simply bounding the admissible domain of modal quantification.

Chapter 2. In Chapter 2, I employ the bounded modality theory of Mandelkern (2019a,

2023) to offer a novel solution to Karttunen's Problem. Karttunen (1972) observed that asserting ⌜must A⌝ signals that a speaker is in a weaker epistemic position than had she simply asserted A. But, according to the semantics for epistemic modals inherited from epistemic logic, *must* is veridical, meaning ⌜must A⌝ entails A. I argue that extant non-veridical solutions to Karttunen's Problem fail to capture that ⌜A and must ¬A⌝ is a contradiction, and propose a new non-veridical semantics based on Mandelkern's bounded theory of modality.

Chapter 3. In Chapter 3, I examine how local contexts interact with root modals and provide a novel solution to the Samaritan Paradox in the process. It's standard to think that ⌜ought A⌝ means (roughly) that, in all the deontically ideal worlds, A is true. Given this semantics, it immediately follows that deontic *ought* is upward monotonic:

> **Upward Monotonicity (UM).** If A $\models$ B, then ⌜ought A⌝ $\models$ ⌜ought B⌝

Prior (1958), however, observed that, given UM, the naïve semantics incorrectly predicts that the following inference is valid:

(18) a. John ought to help the person who was robbed.

b. $\Rightarrow$Someone ought to be mugged.

This is the Samaritan Paradox. Going back to Åqvist (1967), the Samaritan Paradox has often been blamed on the fact that an upward monotonic semantics validates the following inference:

> **Conjunctive Inheritance.** ⌜ought (A $\wedge$ B)⌝ $\models$ ⌜ought A⌝

This diagnosis is incorrect. The Samaritan Paradox is about presupposition, not Conjunctive Inheritance. A speaker who utters (18a) clearly does not assert that a person ought to *both* be mugged and helped. Instead, a speaker who utters (18a) *presupposes* that someone was mugged and *asserts* that they ought to be helped. Once we understand the root cause of the Samaritan Paradox, it becomes apparent that simply rejecting UM will not fully resolve the paradox. To solve the Samaritan Paradox, I argue that local contexts interact with the domain of root modals in way that is unlike how they interact with epistemic modals. In particular, whereas local contexts *bound* the admissible domains of epistemic modals, they *restrict* the domains of root modals.

Chapter 4. In Chapter 4, I defend a novel solution to the "if *p*, ought *p*" problem" also known as Zvolenszky's Problem. Theories of the conditional-modal interaction, like the Lewis-Kratzer restrictor theory, predict that ⌜if A, MODAL A⌝ is trivially true, a prediction borne out for epistemic modals:

(19) If Britney Spears drank a Pepsi, she must have drunk a Pepsi. epistemic

However, for priority modals, like deontic *ought* and bouletic *want*, this prediction is clearly wrong:[11]

(20) a. If Britney Spears drinks a Pepsi, she ought to drink a Pepsi. deontic

b. If Britney Spears drinks a Pepsi, she wants to drink a Pepsi. bouletic

I propose a simple solution to the puzzle: deontic and epistemic conditionals are sensitive to different kinds of reasons. The fact that A is true is not always a reason for doing A; however, the fact that A is true is always a reason for believing A. I make this intuition precise with Cariani's (2023) default premise semantics which integrates a Kratzer-style (1981b; 2012) premise semantics for conditionals with default logic (Reiter, 1980; Horty, 2007b, 2012).

Since the indirect restriction solution to the Samaritan Paradox predicts that restriction by local context generates covert conditionals everywhere, the indirect restriction view faces numerous "local" versions of the Zvolenszky problem. However, default premise semantics allows us to avoid these local Zvolenszky problems, because *ought*'s domain

[11] It's not important whether these conditionals are *actually* false. The important point is that they *can* be false.

will only be calculated with respect to the conclusions of the default rules triggered in the local context.

While many of the themes of each chapter recur in the others, they are written so as to be readable out of order. This has resulted in a certain amount of inevitable overlap.

Chapter 2: Two Puzzles about Epistemic Must

> There is a striking difference between the logical necessity operator and words like *must*.
>
> Karttunen (1972)

2.1 Introduction

Many semantic theories of epistemic modals start from the assumption that they function more or less like the $\Diamond$ and $\Box$ operators from epistemic modal logic. In such logics, the $\Box$ operator is usually assumed to be *veridical*, meaning that it satisfies the modal T Axiom which states that $\Box$A entails A.[1] Karttunen (1972), however, observed that, in natural language, assertions of epistemic necessity modals feel weaker than a bare assertion of their prejacent.[2]

(21) a. It must be raining.

b. It is raining.

[1] Montague (1969) coined the term 'veridical' to describe an operator O such that $O\phi$ entails ϕ. While philosophers and logicians usually call this property 'factivity,' we avoid this terminology, because linguists often reserve the term 'factive' to describe operators which *presuppose* the truth of their prejacent (Kiparsky and Kiparsky, 1971; Karttunen, 1971). While many veridical operators are also factive in this sense, there's good reason to think that these categories do not exactly coincide. For example, *be right* and *demonstrate* are often taken to be veridical, but non-factive.

[2] We use the term 'prejacent' to refer to a modal's argument following von Fintel and Gillies (2008) who, in turn, adapt the term from medieval logic.

While (21a) can be uttered by a speaker with relatively little evidence, an assertion of (21b) typically requires the speaker to be in a stronger epistemic position. While (21a) can be asserted on the basis of, say, having seen a colleague's wet umbrella, (21b) often requires stronger evidence, if not outright knowledge, that it's raining. This is our first puzzle: how should we account for the felt weakness of epistemic *must*?

Our second puzzle concerns the fact that expressions of the form $\ulcorner$A $\wedge$ must $\neg$A$\urcorner$ and $\ulcorner$must A $\wedge$ $\neg$A$\urcorner$ are unacceptable.

(22) a. #It's raining, but it must not be.

b. #It must be raining, but it isn't.

Moreover, unlike instances of Moore's paradox, these expressions appear to be genuine semantic contradictions, which cannot be explained pragmatically, because, as Yalcin (2007) observed, they behave very differently in embedded contexts.

(23) a. #Suppose it's raining, but it must not be.

b. Suppose it's raining, but I don't know it.

(24) a. #If it's raining, but it must not be, ...

b. If it's raining, but I don't know it, ...

It may seem unremarkable that $\ulcorner$A $\wedge$ must $\neg$A$\urcorner$ and $\ulcorner$must A $\wedge$ $\neg$A$\urcorner$ are contradictions, but that's only because we standardly assume that epistemic *must* is veridical. If we drop veridicality in response to Karttunen's observation, it no longer follows immediately that $\ulcorner$A $\wedge$ must $\neg$A$\urcorner$ and $\ulcorner$must A $\wedge$ $\neg$A$\urcorner$ are contradictions.

Taken together, an explanation of these two puzzles places severe constraints on a semantics for epistemic *must*. Karttunen's observation suggests that epistemic *must* is non-veridical, while our second observation suggests that $\ulcorner$must A $\wedge$ $\neg$A$\urcorner$ is a contradiction.

Non-Veridicality. $\ulcorner$must A$\urcorner$ $\not\models$ A

Epistemic Necessity Contradiction. $\ulcorner \text{must } A \wedge \neg A \urcorner \models \bot$

Unfortunately, given a Boolean semantics for negation and conjunction together with a consequence relation that preserves truth-at-a-world, it's impossible to satisfy both constraints because $\phi \wedge \neg\psi \models \bot$ if and only if $\phi \models \psi$.

Extant theories of epistemic necessity bifurcate into two camps. Weak theories secure non-veridicality by restricting the domain of modal quantification to a subset of ideal worlds (Kratzer, 1991; Swanson, 2006; Lassiter, 2016; Goodhue, 2017; Del Pinal, 2021). However, this comes at a cost, because these theories then fail to predict that epistemic necessity contradictions are contradictions. In contrast, strong theories, which treat *must* as veridical, yield a straightforward explanation of why epistemic necessity contradictions are contradictions. They then offer a pragmatic explanation of *must*'s felt weakness (von Fintel and Gillies, 2010, 2021; Mandelkern, 2019b). Since they appear to offer better empirical coverage, let's start by considering strong theories of *must*.

2.1.1 Strong theories

Strong theories maintain veridicality, but explain Karttunen's intuition by positing that epistemic *must* carries an indirect evidential signal. Karttunen himself seems to have considered a theory like this, conjecturing that "the role of *must* …is to indicate that the complement proposition is inferred but not yet known independently" (Karttunen, 1972, p. 13).

von Fintel and Gillies (2010, 2021) defend a semantics along these lines, drawing attention to the following contrast:

(25) **Context:** Seeing a friend come in wearing rain boots with a wet umbrella.

a. It's raining.

b. It must be raining.

(26) **Context:** Looking out the window at rain drops falling.

a. It's raining.

b. #It must be raining.

Unlike a bare assertion of its prejacent, an epistemic necessity claim is incompatible with direct evidence. Strong theories use this datum to argue that epistemic *must* presupposes that the prejacent is known indirectly. So, according to the strong theorist, the weakness intuition arises, because there is "some general conversational principle by which indirect knowledge–that is, knowledge based on logical inferences–is valued less highly than 'direct' knowledge that involves no reasoning" (Karttunen, 1972, p. 13).

Formally, von Fintel and Gillies's account relies on the notion of a kernel borrowed from belief revision theory. A *kernel* for a modal base $f(w)$ is a set of propositions $\mathcal{K}$, which the speaker takes to be direct information in a context, such that $f(w) = \bigcap_{K \in \mathcal{K}(w)} K$. A kernel directly settles a prejacent A if and only if there is some piece of evidence $K \in \mathcal{K}$ such that $K \subseteq [\![A]\!]$ or $K \cap [\![A]\!] = \emptyset$.

Definition 2.1.1. (Strong *must* plus indirect evidentiality).

i $[\![\text{must}]\!]^{w,f,\mathcal{K}} = \lambda p_{\langle s,t \rangle} : \mathcal{K}$ doesn't directly settle $p \,.\, \forall w' \in f(w) : p(w') = 1$

ii $[\![\text{might}]\!]^{w,f,\mathcal{K}} = \lambda p_{\langle s,t \rangle} : \mathcal{K}$ doesn't directly settle $p \,.\, \exists w' \in f(w) : p(w') = 1$

von Fintel and Gillies' theory satisfies our initial desiderata. In particular, it straightforwardly predicts that epistemic necessity contradictions are contradictions, while explaining the felt weakness of *must*.

One objection to strong theories in the literature concerns whether we can even char-

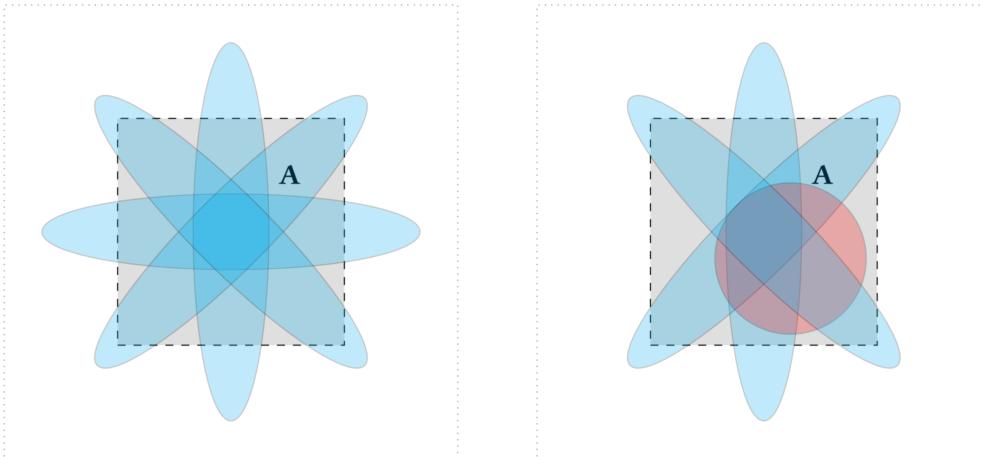

Figure 2.1: A kernel which indirectly settles A versus a kernel which directly settles A

acterize "direct information" independently of speaker knowledge. For example, von Fintel and Gillies acknowledge that, in their original rain context, a *must*-claim uttered by a skeptical epistemologist who doubts her own sensory experience is acceptable (von Fintel and Gillies, 2010, p. 370). Moreover, as Giannakidou and Mari (2016) and Goodhue (2017) emphasize, this phenomenon is quite general: there are often cases where two agents appear to have the same direct information, but one can assert a *must*-claim that the other cannot. Goodhue, for example, imagines a slight variant of von Fintel and Gillies's rain case:

(27) **Fake Rain.** Hilary works next door to a film set. She knows that the crew is shooting a rain scene today and that, if it doesn't rain, the crew will produce fake rain. She knows filming is set to begin at 10:15. At 10:00 she looks outside and sees falling water droplets. She says:

  a. It's raining.

  b. It must be raining.

It's not clear why Hilary's evidence is any less "direct" than the speaker in von Fintel and Gillies's original context. It's tempting to claim that what evidence counts as "direct" depends in some way on Hilary's knowledge. However, von Fintel and Gillies are using the concept of direct evidence to characterize epistemic modal bases; so, if we then attempt to characterize evidence in terms of the speaker's knowledge, there's a genuine worry that the whole analysis will become circular.

Even if these hurdles are surmountable, von Fintel and Gillies' theory appears not to track any pre-theoretic notion of directness. For example, there is nothing obviously "direct" about testimony–indeed, in languages with grammatical evidentials, reportative evidentials always signal indirect evidence (Aikhenvald, 2004; Willett, 1988). However, since it's generally unacceptable to assert ⌜must A⌝ when a trustworthy source has informed the speaker that A is true, von Fintel and Gillies' theory is committed to classing (at least some) testimony as direct evidence (Matthewson, 2015).[3] For example, if you read in the newspaper that the Lakers defeated the Clippers, you cannot then assert, "The Lakers must have won," unless you suspect (for whatever reason) that the newspaper is an unreliable source of information.

In response to this criticism, von Fintel and Gillies bite the bullet, arguing that any attempt to give a first-principles characterization of "direct enough information" would be "ill-suited to do the job that is carved out for it" (von Fintel and Gillies, 2021, p. 105). I find this response unsatisfying, because many languages have devices which signal possession of direct evidence (or lack thereof). Hence, from a functional perspective, it would be strange for epistemic necessity modals alone to make use of an oddly circumscribed notion of "direct enough information."

The second main objection raised in the literature concerns the fact that epistemic

[3]Many theories of testimony in epistemology claim that the hearer inherits whatever evidence originally justified the speaker's belief (Burge, 1993, 1997). These 'transmission' views of testimony, however, remain controversial. Lackey (1999, 2008), for example, notes that a creationist teacher can nonetheless transmit knowledge of evolution to her students. But, even if this type of response is successful, I still believe that the cross-linguistic data gives us good reason to doubt a von Fintel and Gillies-style account of epistemic necessity.

must-claims often appear compatible with a certain amount of speaker ignorance. For example, a corpus study by Lassiter (2016, p. 123) found the following naturally occurring examples:

(28) I have an injected TB42 turbo and don't like the current setup. There is an extra injected located in the piping from the throttle body . . . Must be an old DTS diesel setup but I'm not certain. Why would they have added this extra injector?

(29) This is a very early, very correct Mustang that has been in a private collection for a long time . . . The speedo[meter] shows 38,000 miles and it must be 138,000, but I don't know for sure.

If *must* were veridical, we would expect (28) and (29) to be semantic contradictions, but they're not. von Fintel and Gillies argue that Lassiter, "has found some examples in the wild where epistemic modals undergo shifts in the possibilities deemed relevant, the modal horizon" (von Fintel and Gillies, 2021, p. 99).[4] They present several arguments purporting to show that, if we control for this contextual instability, Lassiter's examples pose no problems for their account.

They argue, for example, that there are contexts where we can force the speaker to choose between modal horizons:

(30) A: That must be an old DTS diesel setup but I'm not certain. Why would they have added this extra injector?

B: So, given that you're not certain, do you still think that it must be an old DTS diesel setup?

A: ?*Like I said: it must be and I'm not certain.

[4] This idea has its roots in Lewis (1979) who claims, "The boundary between the relevant possibilities and the ignored ones (formally, the accessibility relation) is a component of conversational score, which enters into the truth conditions of sentences with "can" or "must" or other modal verbs. It may change in the course of conversation. [. . .] This boundary may also shift in accordance with a rule of accommodation" (p. 354). Gillies (2007) employs this strategy to defend a strict conditional analysis of counterfactuals.

von Fintel and Gillies report that A cannot stick to both conjuncts in (30): she must either retract the initial *must*-claim, or return to the initial modal horizon. I find this data uncompelling. von Fintel and Gillies note that the reviewers of their (2021) found the dialogue in (30) paper perfectly felicitous. If anything, I believe B's response is odd. A just told you that it must be an old DTS diesel setup, but that they're not certain; so, why is B asking if A thinks it must be an old DTS diesel setup?

2.2 Two new arguments against strong *must*

I worry that von Fintel and Gillies' refusal to give an independent characterization of direct evidence coupled with their shifting "modal horizon" response to Lassiter makes their theory unpredictive. Every potential counterexample which assumes two agents have the same direct evidence can simply be recast in an unproblematic light; similarly, every counterexample involving two modals can be said to involve an implicit domain shift. This leaves us at something of an impasse. However, in this section, I present two new arguments against strong *must* which, I believe, are resistant to these moves: one involves interactions with indicative conditionals, the other involves interactions with anaphora.

2.2.1 The argument from indicative conditionals

It's widely accepted that indicative conditionals presuppose that their antecedents are compatible with the context set–the worlds compatible with all the propositions that are common ground (Stalnaker, 1975; von Fintel, 1998, a.o.).

> **Antecedent Compatibility.** An indicative conditional ⌜if A, then C⌝ presupposes that $[\![\text{A}]\!] \cap c \neq \emptyset$.

This constraint explains why the following conjunctions are unacceptable.

(31) a. #Smith isn't the killer, but if he is the killer ...

b. #I failed my logic final, but if I didn't ...

Given the antecedent compatibility constraint, if epistemic *must* were veridical, we would expect a *must*-claim to be incompatible with an indicative conditional whose antecedent is the negation of its prejacent. However, this isn't the case:

(32) a. It must be raining, but if it isn't, bring an umbrella.

b. #It's raining, but if it isn't, bring an umbrella.

Similarly, while it's fine to follow a *must*-claim with an indicative conditional whose antecedent expresses uncertainty about the prejacent's truth, a bare assertion of the prejacent sounds redundant in the very same context.

(33) a. It must be raining, and if I'm right, I'll get soaked.

b. #It's raining, and if I'm right, I'll get soaked.

This suggests that assertions of ⌜must A⌝ have discourse effects that are very different from a bare assertion of A. A bare assertion of the prejacent appears to be eliminative–if A becomes common ground, then all the non-A-worlds are removed from the context set; in contrast, an assertion of ⌜must A⌝ appears to be non-eliminative. Strong theories, therefore, face a major explanatory hurdle: if ⌜must A⌝ entails A, then why doesn't ⌜must A⌝ also eliminate all the non-A-worlds from the context set?

The proponent of veridical *must* might question our deployment of the antecedent compatibility constraint. After all, the antecedent compatibility constraint admits well-known exceptions in special contexts, what Holguín (2020) calls "conspiracy conditionals." Take Adams's (1975) famous example "If Oswald didn't shoot Kennedy, then someone else did." Here the speaker is not necessarily committed to it being possible that Oswald didn't shoot Kennedy; so, a speaker can felicitously assert a conjunction like the following:

(34) Oswald killed Kennedy, but if he didn't, someone else did.

In general, "conspiracy conditionals" require contexts where the speaker has special reason to consider possibilities she takes to be non-actual. For example, in contexts where the stakes for being wrong are high. Consider two thieves mid-heist:[5]

(35) A: What if they check the security camera?

B: Don't worry. They don't check the camera, but if they do, we're wearing masks.

However, I believe conspiracy contexts do little to undermine the plausibility of the antecedent compatibility constraint. In conspiracy contexts, I believe, a speaker attenuates their commitment to the previously asserted proposition by exploiting the very fact that a conditional presupposes its antecedent is compatible with the context set. Regardless of whether this diagnosis is correct, it's clear that appealing to conspiracy contexts will not allow the strong theorist to escape the conclusion of our argument, because our argument merely relies on the fact that there are minimal pairs where an expression of the form $\ulcorner \mathrm{A} \wedge (\neg \mathrm{A} > \mathrm{C}) \urcorner$ is unassertible, but an expression of the form $\ulcorner \Box \mathrm{A} \wedge (\neg \mathrm{A} > \mathrm{C}) \urcorner$ is assertible. However, according to a veridical semantics, the latter should entail the former. So, this is unexpected.

2.2.2 The argument from anaphora

A similar point can be made about the interaction between epistemic *must* and pronominal anaphora. As Kibble (1994, §6) observes, an indefinite in the scope of epistemic *must* may fail to license subsequent pronominal anaphora outside its scope.

(36) a. John must have [a guitar]$_i$.

b. # [It]$_i$'s a Fender Stratocaster. Kibble (1994)

(37) Toby must have written [a letter]$_i$. #[It]$_i$'s very long.

[5] This example is inspired by similar examples in (Ciardelli, 2020, p. 535).

If we accept a strong theory of *must*, this data is puzzling. In general, veridical operators like *know* and *it's clear that* license pronominal anaphora outside their scope:

(38) a. Andy knows Toby wrote [a letter]$_i$. [It]$_i$'s very long.

b. It's clear that Toby wrote [a letter]$_i$. [It]$_i$'s very long.

A link between an operator being veridical and it licensing subsequent anaphora is intuitive. If we adopt the classic theory of Cooper (1983) according to which pronouns presuppose the existence of their referents, then it's reasonable to expect that a veridical operator, which entails the existence of a discourse referent introduced within its scope, will license subsequent pronominal anaphora, while a non-veridical operator, which does not entail the existence of a discourse referent in its scope, will not license subsequent pronominal anaphora. In contrast, whether or not a non-veridical embedding can license pronominal anaphora depends on additional features of the conversational context, especially the "modal subordination" licensed by modals and conditionals (Karttunen, 1976; Roberts, 1989).

(39) a. Toby might have written [a letter]$_i$. #[It]$_i$'s very long.

b. Toby probably wrote [a letter]$_i$. #[It]$_i$'s very long.

(40) a. Toby might have written [a letter]$_i$. If he did, [it]$_i$'s very long.

b. Toby probably wrote [a letter]$_i$. [It]$_i$ had to be very long.

A similar point can be made with additive particles (*too*, *also*, *again*) which are commonly argued to be anaphoric.[6] As with pronominal anaphora, epistemic *must*, like other weak modals, may fail to license subsequent *too*-anaphora[7]:

[6] See Kripke (2009) for arguments that additive particles are anaphoric. Karttunen and Peters (1979) previously argued that the additive particle *too* simply triggers an existential presupposition. See Ruys (2015) for a recent defense of the Karttunen-Peters view.

[7] Likewise, if you replace *too* with *again* or *also*, which are also anaphoric upon earlier discourse, we get the same contrast in felicity.

(41) a. Tim flunked the midterm. Now he's flunked the final too.

b. Tim might have flunked the midterm. #Now he's flunked the final too.

c. John must have flunked the midterm. #Now he's flunked the final too.

If *must* were veridical, we would expect the *must*-discourse to always be felicitous, because it would be common ground that Tim flunked an exam last week. However, in (41b) as in (41c), the use of *too* seems to presuppose a proposition that's not yet common ground. This, again, suggests that asserting ⌜must A⌝ does not have the same effect on the common ground as simply asserting A.

I should emphasize that I'm not claiming non-veridicality is *sufficient* to block anaphoric reference to an indefinite in the scope of an operator; rather, I'm claiming non-veridicality is *necessary* to block anaphoric reference. Since the work of Karttunen (1976), we've known that negation typically blocks anaphoric reference to an indefinite in its scope.

(42) #John doesn't own [a car]$_i$. Paul owns [it]$_i$.

However, this pattern is not exceptionless. In addition to negative subordination (Kibble, 1994), the generalization about negation also admits exceptions when the indefinite refers to an object that's presupposed to have existed in the past (43).[8]

(43) John doesn't have [a car]$_i$ #(anymore). Paul has [it]$_i$. Chierchia (1995)

Chierchia's example doesn't show that the generalization about negation blocking anaphoric reference to an indefinite in its scope is false; rather, (43) shows that the use of *anymore*, which presupposes the existence of a car in the past, makes a discourse referent available as *anymore*'s presupposition projects over negation.

We can create parallel cases with epistemic *must* by modifying the prejacent with the

[8]The phenomenon of *negative subordination*, which Kibble credits to Paul Dekker, is the negation analogue of modal subordination.

i John doesn't own [a guitar]$_i$; so, he doesn't tune [it]$_i$.

Presumably, whatever theory one gives of these cases will resemble closely the theory one gives for modal subordination.

adverbial *still* which presupposes that John had a car.[9]

(44) a. # John must have [a car]$_i$. [It]$_i$ is green.

b. John must still have [a car]$_i$. [It]$_i$ is green.

In (44a), the *must*-claim is naturally interpreted as expressing uncertainty about whether a referent for the indefinite *a car* exists. In contrast, in (44b), the *must*-claim does not express uncertainty about the existence of a referent for the indefinite; rather, it expresses uncertainty about a property of the referent–namely, its ownership by John. The use of *still* in (44b) triggers a presupposition that John had a car; hence, the existence presupposition of *it* will be satisfied.

Similarly, focus-marking can allow an indefinite in the scope of negation to license later anaphora:[10]

(45) John doesn't own [a SUSHI restaurant]$_i$. [It]$_i$'s an Italian restaurant.

Given a standard semantics for focus (Rooth, 1985, 1992; Geurts and Sandt, 2004), focus-marking *sushi* in (45) triggers the presupposition that John owns a restaurant, but not necessarily one that's a sushi restaurant. Even though negation prevents the indefinite from introducing a discourse referent, the presence of focus triggers an existential presupposition that projects over negation ensuring that the subsequent pronoun's existence presupposition is satisfied.

However, perhaps the Kibble discourses are odd, not because the pronoun's existence presupposition isn't satisfied, but because asserting an expression containing a pronoun clashes with a presupposition of the *must*-claim. von Fintel and Gillies's account suggests one such explanation. Uttering (36a) implies that the speaker only has indirect evidence that John has a guitar. However, in asserting (36b), the speaker then implies that she knows

[9] A very similar contrast involving *hope* is discussed by Chierchia (1995) which Pranav and Hacquard (2013) argue has an epistemic modal base.

[10] This example is based on example in Lewis (2021) who uses it to make a slightly different point.

or has good evidence that John has a Stratocaster. This alone causes no problems, because the speaker's indirect evidence for John having a guitar might also support John having a Stratocaster. However, perhaps the norms of cooperative discourse are such that a speaker must explicitly signal that her evidence for (36a) is no better than her evidence for (36b) (e.g., by using a modal). Failure to signal this is interpreted as pragmatic strengthening of (36b)–that is, its use implies the speaker's evidence is more direct. This principle would explain, for example, why (46) is infelicitous.[11]

(46) a. John bought some kind of pizza or other.

b. #It was Hawaiian.

Just as it's hard to imagine how a speaker could have more direct evidence for John having bought a Hawaiian pizza than for John having bought some kind of pizza or other, it's hard to imagine how a speaker could have more direct evidence for John having a Stratocaster than for John having a guitar.

This alternative explanation is hard to test because pronouns almost always admit deictic uses. For example, (47a) sounds felicitous, but as (47b) shows there's no coherent reading of (47a) where the pronoun *it* is interpreted as co-referential with the antecedent indefinite *a new car*; rather, the only available reading is one where *it* is used demonstratively to refer to the salient car.[12]

(47) **Context:** A and B are looking at a new red sports car. Given his evidence, A believes that it's John's new car. A says to B:

a. John must have bought a new car. Do you like its color?

b. #John must have bought [a new car]$_i$. Do you like the color of [that car]$_i$?[13]

[11]Some have argued that constructions like *some … or other* contains implicit epistemic modality (Alonso-Ovalle and Menéndez-Benito, 2010). If this is the case, then one might conclude that (46) is bad for the same reason I've conjectured that the Kibble discourses are bad.

[12]English definites can either be referential or anaphoric; given this, these judgments are quite subtle in English. However, as Alexander Williams (p.c.) suggests, languages, like German, which have 'weak' and 'strong' definites might provide a better contrast case, since only the latter license anaphora (Schwarz, 2009).

[13]In this dialogue, imagine that the demonstrative *that* is used without any kind of gesture that could aid

c. John must have bought a new car. (*points at car*) Do you like the color of that car?

However, insofar as a non-deictic use of the pronoun seems degraded, I take this as good reason to think that the alternative explanation will likely prove unsuccessful.

2.3 A weak bounded semantics of *must*

In this section, I sketch our positive view of epistemic necessity claims. I start by surveying the classic weak theory of Kratzer (1991) and recent measure-theoretic alternatives (Swanson, 2006; Lassiter, 2016; Del Pinal, 2021). I argue that these theories suffer from a common defect–they predict that epistemic necessity contradictions are consistent. I first consider a simple patch inspired by the indicative conditionals literature: *must* presupposes its domain is a subset of the context set. By combining this with a Stalnaker-style dynamic pragmatics, we come close to solving the initial logical puzzle, but ultimately we cannot account for epistemic necessity contradictions with left-embedded *must*-claims nor can we derive embedded epistemic necessity contradictions. To circumvent these problems, I combine Kratzer's weak semantics with the *bounded modality* view of Mandelkern (2019a) giving us a view on which *must* quantifies over the best worlds in its *local context*.

2.3.1 Extant weak theories

In addition to von Fintel and Gillies's strong theory, the other prominent strategy for accommodating Karttunen's observation treats *must* as a weak necessity modal, quantifying over an ideal subset of the modal base (Kratzer, 1991; Swanson, 2006; Lassiter, 2016; Del Pinal, 2021).[14] For example, Kratzer introduces an ordering source function

in fixing the referent.

[14]While domain restriction is the most popular path to non-veridicality, Veltman's (1985) data semantics and descendent approaches like those based on inquisitive semantics (Roelofsen, 2016) secure non-veridicality in a slightly different manner. Since these approaches require a more radical break with standard possible

$g : W \to \mathcal{P}(\mathcal{P}(W))$ which generates a preorder over worlds in the modal base:

$$v \gtrsim_{g(w)} u \Leftrightarrow \{p \in g(w) : u \in p\} \supseteq \{p \in g(w) : v \in p\}$$

We can think of this preorder as ranking worlds by normality. So, instead of quantifying over the entire modal base, Kratzer gives a semantics according to which *must* quantifies over the most normal worlds in the modal base.[15]

$$\textsc{dom}(w, f, g) = \{v \in f(w) : \text{there is no}\, u \in f(w)\, \text{such that}\, u \neq v\, \text{and}\, u \succ_{g(w)} v\}$$

Definition 2.3.1. (Ordering source semantics).

i $[\![\text{must}]\!]^{w,f,g} = \lambda p_{\langle s,t \rangle}.\forall w' \in \textsc{dom}(w, f, g) : p(w') = 1$

ii $[\![\text{might}]\!]^{w,f,g} = \lambda p_{\langle s,t \rangle}.\exists w' \in \textsc{dom}(w, f, g) : p(w') = 1$

Since the actual world will often fail to be among the most normal worlds, epistemic *must* does not entail its prejacent.

Alternatively, one might treat *must* as a scalar adjective expressing a high subjective probability in the prejacent (Swanson, 2006; Lassiter, 2016; Del Pinal, 2021). We can state this view in a measure-theoretic semantics where a *probability measure* $\mu : \mathcal{P}(W) \to [0, 1]$ is a function from an algebra over W to the unit interval satisfying the following axioms:[16]

- $\mu(W) = 1$
- $\mu(\emptyset) = 0$
- $\mu(\cup_i E_i) = \Sigma_i \mu(E_i)$ for a finite partition $\{E_i\}$

worlds semantics, I won't discuss them in detail here.

[15] While Kratzer does not make this assumption herself, I will make the simplifying assumption that this pre-order is well-founded. That is, for any non-empty modal base $f(w) \neq \emptyset$, $\textsc{dom}(w, f_e, g_e) \neq \emptyset$. This assumption, equivalent to Lewis's "Limit Assumption" (1973, §3.4), rules out the possibility of an infinite descending chain of worlds. There is an extensive debate about this assumption–see Kaufmann (2017) for a detailed discussion, but these questions are orthogonal to the current project.

[16] An *algebra* of subsets of W is a collection of sets that contains the empty set $\emptyset$ and is closed under complementation and finite unions.

Definition 2.3.2. (Threshold semantics). Let threshold θ_c be $\theta_c > 0.5$:

i $[\![\text{must}]\!]^{w,\mu} = \lambda p_{\langle s,t\rangle}.\mu(\{w' : p(w') = 1\}) > \theta_c$

ii $[\![\text{might}]\!]^{w,\mu} = \lambda p_{\langle s,t\rangle}.\mu(\{w' : p(w') = 1\}) > 1 - \theta_c$

Since the actual world may not be among the probable worlds, we again have that ⌜must A⌝ does not entail A.

Given the axioms above, the threshold semantics predicts that epistemic *must* fails to validate Agglomeration:[17]

Agglomeration. $\Box$A, $\Box$B $\models$ $\Box$(A $\wedge$ B)

However, this prediction seems wrong. While (48b) is perfectly intelligible, (48a) is incoherent, which we wouldn't expect if Agglomeration failed.

(48) a. # It must rain, and it must snow, but it's not the case that it must both rain and snow.

b. It might rain, and it might snow, but it's not the case that it might both rain and snow.

There may be patches that solve these problems. For example, Del Pinal (2021) defends a modified version of the threshold semantics using a subjective probability measure that is conditionalized on the worlds in the ordering source. In other words, $[\![\text{must A}]\!]^{w,\mu} = 1$ iff $\mu(\{w' : [\![\text{A}]\!]^{w',c} = 1\}| \cap g(w)) > \theta_c$. This allows *must* to have quasi-veridical characteristics and may provide a way around some of these logical problems. However, insofar

[17] Moreover, if we assume that *might* and *must* are duals, then, if $\models$ contraposes, *must* failing to agglomerate is equivalent to *might* failing to distribute over disjunction.

Distribution Over Disjunction. $\Diamond$(A $\vee$ B) $\models$ $\Diamond$A $\vee$ $\Diamond$B

Failure of distribution over disjunction for *might* is even less plausible and has many negative downstream consequences.

as the threshold semantics introduces new theoretical difficulties without increasing explanatory power, I think it's less attractive than traditional Kratzerian approaches.

More generally, I believe extant weak theories are unsatisfactory, because, if correct, they imply that the prejacent of epistemic *must* can be consistent with the negation of its prejacent. Modals like *should* and *ought* on their quasi-epistemic readings, what Yalcin (2016) calls "modals of normality," seem to behave this way, but epistemic necessity modals, in contrast, appear very sensitive to facts about the actual context. So, whereas (49a) sounds contradictory, (49b) does not.[18]

(49) a. # It must be raining, but it isn't.

b. It should be raining, but it isn't.

As I argued earlier, there's compelling evidence that sentences like (49a) are semantic contradictions. First, there's Yalcin's observation that, unlike Moorean sentences, epistemic necessity contradictions remain unacceptable when embedded in counterfactual attitude verbs and conditional antecedents.

(50) a. #Suppose it's raining, but it must not be.

b. Suppose it's raining, but I don't know it.

(51) a. #If it's raining, but it must not be, …

b. If it's raining, but I don't know it, …

Moreover, epistemic necessity contradictions remain infelicitous when embedded in certain disjunctions, what Mandelkern (2019a) calls 'Wittgenstein disjunctions':

[18]There are further differences between epistemic necessity modals and modals of normality. Copley (2003, 2006) observes that, unlike modals of normality, *must*-claims are unacceptable when the speaker has little to no evidence supporting the prejacent.

i a #Max must be there, but I have absolutely no idea whether he is.

b Max should be there, but I have absolutely no idea whether he is.

Copley uses this data to argue that the quasi-epistemic reading of *should* isn't genuinely epistemic.

(52) a. #Either I must have won, but didn't, or I must've lost, but I did't.

b. Either I believe I won, but I didn't or I believe I lost, but I did't.

In the next section, I will consider a simple patch to the standard Kratzerian semantics that doesn't quite work, which involves taking the context set as the modal base. However, its failures will motivate our preferred solution which based on Mandelkern's theory of *bounded modality*.

2.3.1.1 A simple fix that doesn't quite work

It looks as though Kratzer's semantics doesn't derive epistemic necessity contraditions, because epistemic *must*'s modal domain is allowed to include worlds that fall outside of the context set. For modals of normality like *should* and *ought*, this seems correct, but, as we saw, it makes for a problematic analysis of epistemic *must*.

Here's a simple fix which draws on the indicative conditionals literature. Following Stalnaker (1975), it's commonly assumed that indicative conditionals presuppose that their domain lies in the context set (von Fintel, 1998; Iatridou, 2000, a.o.) In contrast, it's commonly assumed that counterfactual morphology signals that we "may reach outside of the context set" (Stalnaker, 1975, p. 276).

(53) If John isn't home, he's working at a café.

$\rightsquigarrow$ *presupposes: there are worlds in the context set where John isn't home*

Since *must* is indicative, we might postulate that it's subject to an analogous constraint that its prejacent has to be evaluated at worlds in the context set:[19]

[19] In English, there is often no clear morphosyntactic difference between the indicative and subjunctive moods. However, English appears to have dedicated morphology signally which modals can be interpreted counterfactually, what von Fintel and Iatridou (2023) call X-marking. In the English modal system, the modals which can be interpreted counterfactually are marked with (usually uninterpreted) past-tense morphology; for example, *would* is the past-marked form of *will*, *should* the past-marked form of *shall*, and *could* the past-marked form of *can*. While *will*, *shall*, and *can* are obligatorily indicative, their past-marked counterparts are not. Since *must* has no past-tense morphology, I will assume it too is obligatory indicative.

Indicative Constraint. An assertion of *must* evaluated with respect to an epistemic modal base f in c presupposes $f(w) \subseteq c$.

Indeed, if we combine this view with the dynamic theory of pragmatics developed by Stalnaker (1978, 2002, 2014), the vanilla ordering source semantics plus the indicative constraint handles most of the problematic data. Following Stalnaker (1978), let's adopt the view that an assertion of A is a proposal to make the context set accept A.

Definition 2.3.3. (Acceptance). A context set c *accepts* A iff $c \subseteq [\![\mathrm{A}]\!]$.

Hence, a successful assertion of A in c updates the context to $c_{\mathbf{A}} = c \cap [\![\mathrm{A}]\!]$. Crucially, acceptance gives us a notion of pragmatic incoherence–an utterance of a sentence A is pragmatically incoherent if A is only accepted by the empty context set.

If we make the common assumption that asserting a conjunction sequentially updates the context set, first with the left conjunct, and then the right conjunct (Stalnaker, 1974; Heim, 1983), we also can derive a partial explanation of Yalcin's puzzle.

$$c[\mathrm{A} \wedge \mathrm{B}] = c[\mathrm{A}][\mathrm{B}]$$

For example, we predict that $\ulcorner\neg\mathrm{A} \wedge \text{must}\,\mathrm{A}\urcorner$ is a semantic contradiction. A successful assertion of $\neg$A eliminates all the A-worlds from the context set; however, by the indicative constraint, this means that the $\ulcorner$must A$\urcorner$ is trivially false because its modal base consists of $\neg$A-worlds. By parity of reasoning, we can also derive that $\ulcorner\mathrm{A} \wedge \text{must}\,\neg\mathrm{A}\urcorner$ is a semantic contradiction. However, since the sequential update rule for conjunction is asymmetric, we fail to predict that $\ulcorner\text{must}\,\mathrm{A} \wedge \neg\mathrm{A}\urcorner$ and $\ulcorner\text{must}\,\neg\mathrm{A} \wedge \mathrm{A}\urcorner$ are semantic contradictions, because assertions of *must* do not eliminate worlds from the context set–they only test whether the most plausible worlds in the context set make their prejacent true.

In addition to failing to derive instances of Yalcin's puzzle involving left-embedded *must*-claims, the vanilla Kratzerian semantics plus indicative constraint offers no clear explanation of why epistemic necessity contradictions are unacceptable in embedded contexts.

2.3.2 A bounded semantics

To solve this problem, I draw on Mandelkern (2019a, 2023) who proposes that local contexts bound the admissible domains of quantification for epistemic modals. Following Schlenker (2009, 2010a,b), I define the *local context* of a proposition b occurring in a syntactic environment a_c as the smallest proposition p such that $a(p \wedge b')c' \Leftrightarrow ab'c'$ for all well-formed completions b' and c'. Mandelkern crucially departs from this definition, assuming instead that local contexts are symmetric. Schlenker defines the *symmetric local context* of a proposition b occurring in a syntactic environment a_c as the smallest proposition p such that $a(p \wedge b')c \Leftrightarrow ab'c$. The difference between standard local contexts and symmetric local contexts is that the former quantify over all possible syntactic completions, while the latter is only sensitive to the *actual* completion. While we saw in Chapter 1 that additive particles like *too* provide some reason for thinking that local contexts are sometimes symmetric, it's clear that they're not *always* symmetric. To avoid having to postulate that local contexts are always symmetric, Mandelkern introduces a new dimension of meaning–conditions specifying when an utterance is assertible at a local context, what Mandelkern calls being "satt"–which only modals and conditionals are sensitive to. Given Schlenker's symmetric algorithm, we can then recursively state satt condition which allow us to derive how bounds project.

Definition 2.3.4. (Bound projection). Let κ be an arbitrary local context. We can then recursively define satt conditions as follows:

- $[\![\neg \mathrm{A}]\!]^{\kappa}$ is satt $\Leftrightarrow$ $[\![\mathrm{A}]\!]^{\kappa}$ is satt
- $[\![\mathrm{A} \wedge \mathrm{B}]\!]^{\kappa}$ is satt $\Leftrightarrow$ $[\![\mathrm{A}]\!]^{\kappa+\mathrm{B}}$ is satt and $[\![\mathrm{B}]\!]^{\kappa+\mathrm{A}}$ is satt
- $[\![\mathrm{A} \vee \mathrm{B}]\!]^{\kappa}$ is satt $\Leftrightarrow$ $[\![\mathrm{A}]\!]^{\kappa+\neg\mathrm{B}}$ is satt and $[\![\mathrm{B}]\!]^{\kappa+\neg\mathrm{A}}$ is satt
- $[\![Att_S(\mathrm{A})]\!]^{\kappa}$ is satt $\Leftrightarrow$ $\forall w \in \kappa : [\![\mathrm{A}]\!]^{Acc(S,w)}$ is satt

Instead of assuming that *must* presupposes that its modal base is a subset of the context set as we did in the previous section, we now require that, for every world, *must*'s modal base be a subset of κ.

> **Strong Locality Constraint.** An assertion of *must* in κ relative to a modal base f presupposes $\forall w : f(w) \subseteq \kappa$ and $\kappa \neq \emptyset$.

This strong locality constraint proposed by Mandelkern (2019a) is, in fact, inconsistent with a veridical semantics for *must*, because, given strong locality, worlds that lie outside the local context will be unable to see themselves since they can only see worlds that lie in the local context. Hence, if we combine strong locality with a standard relational semantics for *must*, we obtain a new version of the weak theory.

Definition 2.3.5. (Bounded weak *must* à la Mandelkern).

(54) a. $[\![\text{must}]\!]$ is satt at κ iff $\forall w : f(w) \subseteq \kappa \wedge \kappa \neq \emptyset$.
b. if satt, $[\![\text{must}]\!]^{w,f,g,\kappa} = \lambda p_{\langle s,t \rangle}.\forall w' \in \textsc{dom}(w, f, g) : p(w') = 1$

Since Mandelkern believes *must* is veridical, this leads him to abandon the strong locality constraint in favor of a weaker version which only quantifies over worlds in the local context rather than *all* the worlds.[20]

[20]These two constraints lead to slightly different empirical predictions. Mandelkern (2019a), crediting Cian

Weak Locality Constraint. An assertion of *must* in κ relative to a modal base f presupposes $\forall w \in \kappa : f(w) \subseteq \kappa$ and $\kappa \neq \emptyset$.

Given that the local contexts of conjunction are symmetric in Mandelkern's bounded theory, our system now correctly derives that instances of Yalcin's puzzle involving left-embedded *must*-claims are contradictions.

(55) $[\![\text{must A} \wedge \neg\text{A}]\!]^{w,f,g,\kappa}$ is satt iff $[\![\text{must A}]\!]^{w,f,g,\kappa+\neg\text{A}}$ is satt and $[\![\neg\text{A}]\!]^{w,f,g,\kappa+\text{must A}}$ is satt

If $\ulcorner$must A $\wedge$ ¬A$\urcorner$ is satt, then $\ulcorner$must A$\urcorner$ has to be a evaluated with respect to a local context which is restricted to ¬A-worlds; hence, $\ulcorner$must A $\wedge$ ¬A$\urcorner$ is predicted to be a semantic contradiction. Furthermore, it's easy to check that the clauses in 2.3.5 plus the locality constraint in (141a) derive the correct predictions about epistemic necessity contradictions in the complement of attitude verbs. For example, $\ulcorner Bel_S(\text{A} \wedge \text{must}\neg\text{A})\urcorner$ is satt iff $\kappa \neq \emptyset$ and, for all $w \in \kappa$, $\ulcorner$A $\wedge$ must ¬A$\urcorner$ is satt at $Dox_{S,w}$. However, since $\ulcorner$A $\wedge$ must ¬A$\urcorner$ isn't true at any κ where it's satt, it's not true at $Dox_{S,w}$.

In addition to solving our initial puzzle, the bounded theory also neatly derives the invalidity of counterexamples to disjunctive syllogism involving embedded *must*-claims. For example, Klinedinst and Rothschild (2012) observe that the following inference is intuitively unacceptable:

Dorr, observes that, since that f is only locally reflexive given strong locality, the following inference fails:

i John must be upstairs, or Susie must be upstairs. $\Box A \vee \Box B$

ii $\Rightarrow$ It's not the case that: both John and Susie are downstairs. $\neg(\neg A \wedge \neg B)$

While this is a problem, these inferences can be accommodated in the current system by appealing to a more pragmatic notion of consequence such as Stalnaker's "reasonable inference." Beyond this, there are other subtle differences between the two constraints. For example, while both predict that there's no context c such that an epistemic contradiction is both true and satt at c if we consider truth-at-a-context (also known as *diagonal truth*), they make different predictions if we just consider truth simpliciter. As Mandelkern (2023) observes, if we consider an epistemic contradiction evaluated at c and therefore a local context κ_c but a world w not in κ_c, then weak locality predicts that epistemic contradictions can be true and satt at $\langle \kappa_c, w\rangle$, while strong locality avoids this prediction. However, since it's standard to assume that assertibility is sensitive to truth-at-a-context rather than truth simpliciter, it's unclear whether this difference is material.

(56) Either the dog is in the backyard or it must be in the kitchen. $A \vee \Box B$

It's not the case that it must be in the kitchen. $\neg\Box B$

Therefore, it is in the backyard. A

Given the clauses in 3.3.1, it follows that $[\![A \vee \Box B]\!]^{w,f,g,\kappa}$ is satt iff $[\![\Box B]\!]^{w,f,g,\kappa+\neg\mathbf{A}}$ is satt. So, given the locality constraint, the right-embedded *must*-claim in the disjunction in (56) is only quantifying over ¬A-worlds while the negated *must*-claim is quantifying over A and ¬A-worlds; so, since the modals are quantifying over different domains, the argument is invalid.

2.3.2.1 Deriving the anaphora data

I can now derive the anaphora data from §2.2 by combining our current theory of *must* with the bounded theory of anaphora proposed by Mandelkern (2022, 2023). Just as the admissible domains for epistemic modals are bounded by their local context, Mandelkern proposes that the domains of definites and indefinites are similarly constrained by their local context. Mandelkern's idea is that indefinites carry a *witness bound*–a definedness condition requiring that if an indefinite scope is true relative to *any* assignment, it must be true and satt relative to the *starting* assignment–and otherwise act like normal existential quantifiers.[21]

> **Witness Bound.** $\ulcorner$an $x(\mathrm{A}, \mathrm{C})\urcorner$ is satt at $\langle w, g, \kappa\rangle$ only if $\exists a \in D_e : [\![\mathrm{A} \wedge \mathrm{C}]\!]^{g[x\to a],w,\kappa} = 1$, then $\mathrm{A} \wedge \mathrm{C}$ is true and satt at $\langle w, g, \kappa\rangle$.

So, for example, the witness bound says that *an elephant ate the banana* is satt only if, if *anything* in w is an elephant who ate the banana, then $g(x)$ is.

Conversely, following Heim (1982) among others, Mandelkern posits that definites satisfy the following *familiarity bound*:

[21]N.B.: Here g refers to an assignment function rather than an ordering source.

Familiarity Bound. $\ulcorner$the x(A, C)$\urcorner$ is satt at $\langle w, g, \kappa\rangle$ only if $\forall\langle w', g'\rangle \in \kappa$: A is true and satt at $\langle w', g', \kappa\rangle$.

Hence, pronouns are satt only if they are true and satt throughout κ and definite descriptions presuppose that their restrictor is true and satt throughout κ.

So far, we don't have a clear story about how indefinites project the definedness conditions of their prejacents. Here Mandelkern makes the following proposal:

(57) $\ulcorner$an x(A, C)$\urcorner$ is satt at $\langle w, g, \kappa\rangle$ only if $\exists g'$: A $\wedge$ C is satt at $\langle w, g', \kappa\rangle$

Note that don't want to make A$\wedge$C satt in $\langle w, g, \kappa\rangle$ because then updating with $\ulcorner\neg$an x(A, C)$\urcorner$ would make A familiar, but, as we saw, negation generally does not license subsequent anaphora, and we also don't want indefinites to have trivial projection conditions, because then a negated indefinite could have novel definite in its scope. Mandelkern (2023) proposes this projection rule as "middle ground" between these two options.

Given Mandelkern's theory of anaphora, it's relatively easy to explain the contrasts noted by Kibble.

(58) John must have [a guitar]$_i$. # [It]$_i$'s a Fender Stratocaster.

In (58), the indefinite is satt only if the witness bound is satisfied at all the world's in *must*'s domain–the best world–assignment pairs in the local context κ. In contrast, the pronominal anaphor is satt only if its familiarity bound is satisfied–that is, only if it's true relative to *all* world-assignment pairs in the local context. Hence, the *must*-claim with the embedded indefinite is insufficient to license the later anaphor, because it does not satisfy its familiarity bound.

2.3.2.2 Deriving indirectness as an implicature

von Fintel and Gillies hard-coded epistemic *must*'s indirect evidence requirement into its lexical meaning as a presupposition. As they are the first to concede, this explanation of the evidential asymmetry is unsatisfying:

> We have not found a language whose expression of epistemic necessity fails to carry an evidential signal of indirect inference. That is, the paradigm illustrated for English in [(25)/(26)] can be replicated in language after language. This should raise the suspicion that what we are dealing with should not be a stipulated, arbitrary part of the lexical meaning of epistemic necessity modals, and so it shouldn't be a lexically specified presupposition or conventional implicature. (2010, p. 367)

The evidence supporting a presupposition analysis of indirectness is, moreover, scant. von Fintel and Gillies claim that the prespposition projects through negation, because, they claim, we can reproduce same evidential asymmetry with epistemic *can't*:

(59) **Context:** Seeing a sweaty friend wearing sunscreen.

a. It's not raining.

b. It can't be raining.

(60) **Context:** Looking out the window at sunshine.

a. It's not raining.

b. #It can't be raining.

von Fintel and Gillies argue "in order to account for it, [the weak theorists] need to either argue that epistemic possibility modals like *can* aren't merely existential modals or they need to break duality between *must* and *can*" (2010, p. 369). However, our new data suggests denying the duality of *must* and *can* might be better motivated than it initially appears. Arc Kocurek (p.c.), for example, observes that, unlike epistemic *must*, epistemic *can't* is unacceptable with an indicative conditional whose antecedent presupposes the prejacent is possible:

(61) a. #It can't be raining, but if it is, bring an umbrella.

b. #It isn't raining, but if it is, bring an umbrella.

Moreover, *can't* appears to asymmetrically entail epistemic *must*, because an assertion $\ulcorner$must $\neg$A$\urcorner$ appears redundant when preceded by an assertion of $\ulcorner$can't A$\urcorner$, while the converse does not hold.

(62) a. #It can't be raining; in fact, it must not be.

b. It must not be raining; in fact, it can't be.

Unfortunately, more compelling arguments are hard to come by. Since epistemic *must* is generally unacceptable in many environments that are presupposition filters, the presupposition analysis is hard to falsify.[22]

A major upshot of my weak theory is that I can derive *must*'s indirect evidential signal using very general pragmatic mechanisms. For example, I can give a standard Gricean derivation of indirectness as a quantity implicature. I assume that $\ulcorner$must A$\urcorner$ is in pragmatic competition with an assertion of A, the latter being informationally stronger, because *must* is non-veridical. By the Maxim of Quantity, a cooperative speaker should be as informative as possible. So, given that the speaker is cooperative, it must be the case that she isn't in a position to assert A, because she doesn't know A.[23]

von Fintel and Gillies (2021) argue that the problem with deriving indirectness as a quantity implicature is that, when information is indirect, "an assertion of the bare preja-

[22]A corpus study by Hacquard and Wellwood (2012) showed that epistemic *must* seldom appears in matrix questions or in conditional antecedents, both environments often used to test for presupposition.

[23]I'm assuming here a knowledge norm of assertion (Williamson, 1996, 2000). However, any theory which assumes that assertion is governed by a strong-ish epistemic norm would suffice.

cent is also OK" (p. 109).[24] I want to make a brief comment about this challenge. While an assertion of the bare prejacent is often equally acceptable when the speaker has indirect evidence, it's easy to construct counterexamples where this isn't the case. Consider the following example from Ninan (2014):

(63) **Proposal**. Suppose A and B are friends with a couple, Laurie and David, who have been dating for a long time and are very like to get married. Suppose B has not heard any news about Laurie and David. A tells B, 'David proposed to Laurie!' B responds:

a. Finally! She must have said 'yes.'

b. #Finally! She said 'yes.'

Here B can assert a *must*-claim, but cannot assert its bare prejacent. I think the explanation for the asymmetry between Ninan's proposal case, and von Fintel and Gillies's original rain case is obvious. In **Proposal**, B clearly does not know that David proposed to Laurie; however, in (25), the speaker *knows* albeit indirectly that it is raining.

2.4 Conclusion

I pointed out an apparent tension between non-veridicality as a solution to Karttunen's problem and the fact that sentences of the $\ulcorner$A $\wedge$ must $\neg$A$\urcorner$ and $\ulcorner$must A $\wedge$ $\neg$A$\urcorner$ are semantic contradictions. While this tension initially appears to favor a strong analysis of must à la von Fintel and Gillies (2010, 2021), I adduce two novel pieces of data for the claim that *must* is weak. First, I showed that epistemic *must* claims are compatible with indicative conditionals whose antecedents contradict their prejacent. Second, I showed that

[24]A different problem facing quantity implicature derivations of *must*'s indirect evidential signal concerns "why interlocutors would systematically compute implicatures which result in enriched readings that are inconsistent with the common ground, in cases when the literal, non-enriched readings would not clash with the common ground" (Del Pinal, 2021, p. 15). For example, why would interlocutors compute the implicature when staring at falling rain? Here I am sympathetic to recent proposals of Magri (2009) and Del Pinal (2021) that the computation of some implicatures is mandatory and blind to information that's common ground.

epistemic *must*-claims do not introduce discourse referents that license later anaphora outside their scope. This data is not easily explained on a strong theory of epistemic *must*. However, the extant weak theories implausibly predict that $\ulcorner$A $\wedge$ must $\neg$A$\urcorner$ and $\ulcorner$must A $\wedge$ $\neg$A$\urcorner$ are consistent. After exploring a simple patch to standard weak theories which involves taking the context set as the modal base, I suggest an alternative based on Mandelkern (2019a, 2022)'s bounded theory of modality and anaphora. This new theory offers a tidy explanation of both the initial puzzle and the anaphora data.

However, I should flag that the account inherits a major theoretical problem from Mandelkern's theory. It's clear that much of discourse appears to be interpreted asymmetrically. So, how do speakers decide when to interpret discourse symmetrically? Mandelkern avoids answering this question by postulating a separate dimension of meaning which only modals and certain quantificational expressions are sensitive to–his "satt" conditions. However, if possible we would like one single story that explains both the bounding effects of local contexts discussed by Mandelkern, and patterns of presupposition projection. I leave this question for future work.

Chapter 3: A Local Solution to the Samaritan Paradox

3.1 Introduction

It's standard to think that ⌜ought A⌝ means (roughly) that, in all the deontically ideal worlds, A is true. Following Kratzer (1981a, 1991, 2012), we can formally implement this idea in terms of a circumstantial modal base $f(w)$, a set of worlds compatible with the contextually relevant facts, and a pre-order $\gtrsim_{g(w)}$ on worlds in that modal base generated by a deontic ordering source $g(w)$, a set of propositions representing laws, moral prescriptions. and the like.[1] Given this definition, we can define the domain of *ought* as the set of undominated worlds in the modal base:

$$\textsc{best}(w, f, g) = \{v \in f(w) : \text{there is no}\, u \in f_e(w) \,\text{such that}\, u \neq v \,\text{and}\, u \succ_{g(w)} v\}$$

We can then analyze *ought* as a universal quantifier over this set of undominated worlds:

(64) $[\![\text{ought}]\!]^{w,f,g} = \lambda p_{\langle s,t \rangle}.\forall w' \in \textsc{best}(w, f, g) : p(w') = 1$

Given this semantics, it immediately follows that deontic *ought* is upward monotonic:

Upward Monotonicity (UM). If A $\models$ B, then ⌜ought A⌝ $\models$ ⌜ought B⌝

[1] We can define the pre-order induced by ordering source as follows:

$$v \gtrsim_{g(w)} u \Leftrightarrow \{p \in g(w) : u \in p\} \supseteq \{p \in g(w) : v \in p\}$$

As in the previous chapter, we will make the simplifying assumption that this pre-order is well-founded. That is, for any non-empty modal base $f(w) \neq \emptyset$, $\textsc{best}(w, f, g) \neq \emptyset$.

Prior (1958), however, observed that, given UM, the naïve semantics coupled with a Russellian analysis of definite descriptions incorrectly predicts that the following inference is valid:

(65) a. John ought to help the person who was robbed.

b. $\not\Rightarrow$ Someone ought to be robbed.

This is the Samaritan Paradox. Going back to Åqvist (1967), the Samaritan Paradox has often been blamed on the fact that an upward monotonic semantics validates the following inference:

Conjunctive Inheritance. $\ulcorner$ought $(A \wedge B)\urcorner \models \ulcorner$ought A$\urcorner$

This diagnosis is incorrect. The Samaritan Paradox is about presupposition, not Conjunctive Inheritance. A speaker who utters (65a) clearly does not assert that a person ought to *both* be mugged and helped. Instead, a speaker who utters (65a) *presupposes* that someone was mugged and *asserts* that they ought to be helped.[2] The misconception that the Samaritan Paradox is a counterexample to UM comes from assuming a Russellian rather than a Strawsonian analysis of definite descriptions.

Samaritan Inference. If A presupposes B, then $\ulcorner$ought A$\urcorner \models \ulcorner$ought B$\urcorner$

Once we understand the root cause of the Samaritan Paradox, it becomes apparent that simply rejecting UM will not fully resolve the paradox.

3.1.1 The real culprit: local satisfaction

According to satisfaction theories of presupposition (Karttunen, 1974; Stalnaker, 1974; Heim, 1982, 1983), presuppositions have to be satisfied in their *local context.* That is,

[2]I am hardly the first to make this argument. It's become the standard gloss of the puzzle in philosophy of language and linguistics (Portner, 2009; Cariani, 2013; Booth, 2022).

presuppositions must be satisfied in the context obtained by updating the global conversational context, the worlds compatible with the propositions that are common ground, Stalnaker's context set, with the expressions that occur earlier in the sentence. For example, given a conjunction ⌜A and B⌝ uttered in a context c, the presuppositions of B have to be satisfied in the local context $c_{\text{A}} = c \cap [\![\text{A}]\!]$, the global context c updated with the proposition expressed by the first conjunct A. For the prejacent of *ought*, many satisfaction theories predict that satisfaction in the local context amounts to the following:[3]

> **Local Satisfaction.** An assertion of ⌜ought A⌝ in context c presupposes that the local context of *ought* satisfies A's presuppositions for all $w \in c$.

There's no consensus about *ought*'s local context; as we will see, some take *ought*'s local context to comprise its modal base, while others treat *ought*'s local context as its entire modal domain. Schlenker (2009, 2010a,b), for example, defines the local context of a proposition b occurring in a syntactic environment a_c as the smallest proposition p such that $a(p \wedge b')c' \Leftrightarrow ab'c'$ for all well-formed completions b' and c'. So, the local context κ is the smallest proposition p such that ⌜$\text{ought}(p \wedge \text{A})$⌝ $\Leftrightarrow$ ⌜$\text{ought}(\text{A})$⌝. Assuming the naïve semantics in (64), this means that that κ is $p = \bigcup\{\textsc{best}(w, f, g) : w \in c\}$.

Given Local Satisfaction, (65a) presupposes that either its modal base or modal domain (depending on our theory of its local context) only contains worlds where someone has been mugged. Given this, the naïve semantics in (64) immediately licenses the Samaritan Inference. So, even if denying UM allows us to block some instances of the Samaritan Paradox, like the inference from (65a) to (65b), it will not give us a solution to the paradox in full if we still have Local Satisfaction. For example, since presuppositions project over

[3] Local Satisfaction, as stated, is contentious. Mandelkern (2016), for example, argues that Local Satisfaction generates a version of the Proviso Problem (Geurts, 1996) for certain attitude verbs. For example, a sentence like "John wants the president to succeed" licenses two inferences both of which behave like presuppositions: John believes there is a president and there is, in fact, a president (Sudo, 2014). However, as Mandelkern points out, Local Satisfaction, only predicts the former weaker inference.

negation, if (66c)'s modal base satisifies the presuppositions of its prejacent, then (66c) will entail (66d).

(66) a. John doesn't know that he insulted Mary.

b. $\rightsquigarrow$ John insulted Mary.

c. John shouldn't know that he insulted Mary.

d. $\not\Rightarrow$ John should insult Mary.

However, simply denying UM will not explain why (66c) does not entail (66d), because what is asserted in (66a) does not entail (66b).[4] So, rather than viewing Samaritan Paradox as a paradox about UM, the Samaritan Paradox is better understood as a paradox concerning Local Satisfaction.[5]

3.1.2 Overview

Here's the plan for the chapter. In §3.2, we examine the received solution to the Samaritan Paradox in semantics which posits that ⌜ought A⌝ carries a diversity presupposition requiring that the modal domain contain both A-worlds and ¬A-worlds. We argue that this solution is unappealing, because it's both *ad hoc* and has clear empirical shortcomings. It's inconsistent with both counterfactual and factual uses of *ought* and the fact that deontics give rise to actuality entailments. In §3.3, we develop a novel *indirect restriction* solution to the Samaritan Paradox which rejects Local Satisfaction and instead posits that deontics quantify over the best worlds compatible with their prejacent's local context. In §3.4, we extend this analysis to bouletic attitude verbs.

---

[4] I am indebted to Fabrizio Cariani (p.c.) for this argument.

[5] Rejecting UM may be required to solve other paradoxes like Ross's Paradox (Ross, 1944), or the Prof. Procrastinate-style failures of Conjunctive Inheritance (Jackson and Pargetter, 1986). However, we believe that rejecting UM is neither a necessary nor sufficient response to the Samaritan Paradox.

3.2 The received solution: a diversity presupposition

The received solution to the Samaritan Paradox, proposed by Heim (1992) and von Fintel (1999) and elevated to textbook status by von Fintel and Heim (2011), combines two assumptions: first, the assumption is that the modal base is the local context of *ought*'s prejacent and second, the assumption that deontics presuppose that the worlds in their modal base are heterogeneous with respect to the truth of their prejacent.[6]

> **Modal Base Satisfaction.** An assertion of ⌜ought A⌝ in context c presupposes that the modal base $f(w)$ satisfies A's presuppositions for all $w \in c$.
>
> **Diversity Condition (DC).** An assertion of ⌜ought A⌝ in a context c presupposes that $f(w) \cap [\![A]\!] \neq \emptyset$ and $f(w) \cap [\![\neg A]\!] \neq \emptyset$ for all $w \in c$.

Given the DC and Modal Base Satisfaction, we can easily derive the failure of the Samaritan Inference. In (65a), the definite in *ought*'s prejacent presupposes that someone was mugged. Since presuppositions hold throughout the modal base by Modal Base Satisfaction, the Samaritan Inference fails, because the inferred *ought* will violate the DC.

However, I believe that the DC is unsatisfying both on both theoretical and empirical grounds. A DC-based solution is unappealing theoretically for two reasons. First, for the argument to succeed, we must stipulate that the local context of a deontic modal is its modal base. The decision to pick the modal base rather than the modal domain as the local context is entirely arbitrary, and given Schlenker's algorithm, the closest thing we have to an explanatory algorithm for presupposition projection, local context is predicted to be the modal domain (i.e., the best worlds in the modal base) rather than the entire modal base. Second, the DC itself is *ad hoc*, lacking any independent motivation.[7] Empirically, a DC-

[6]What I'm calling the "received" solution is rarely spelled out in detail by its proponents, the two clearest statements of the view can be found in Carr (2014) and Booth (2022) who are both critical of the view.

[7]Some might quibble with my charge that the DC is *ad hoc*. The DC was named by Condoravdi (2002) who

based solution generates a number of implausible downstream predictions. In particular, as I will show in the next two section, the DC incorrectly predicts, first, that *ought* should be incompatible with counterfactual and factual readings and, second, that deontic *ought* does not give rise to actuality entailments.

3.2.1 Objection 1: counterfactual and factual *oughts*

First, the DC is at odds with the fact that speakers use evaluative deontic modals counterfactually:

(67) a. There ought to be world peace, but I know it's impossible.

b. It's not possible, but there ought to be world peace.

Here it matters whether the *ought* is evaluative or deliberative. An evaluative *ought* describes the way things ought-to-be, while a deliberative *ought* describes what an agent ought-to-do.[8]

(68) a. There ought to be world peace. evaluative

b. John ought to bring about world peace. deliberative

Deliberative *ought* is usually taken to imply *can*; hence, deliberative *ought* is infelicitous when used counterfactually:[9]

(69) **Context:** It's common knowledge that John is dead. Hence, it's impossible for him to study for the exam. I say:

used the same presupposition to explain the differences in temporal orientation between epistemic modals like *must*, which cannot have future orientation, and root modals like deontic *ought*, which can only have non-past temporal orientation (see also Werner (2006)). Isn't this independent evidence for the DC? In my opinion, it isn't, because Condoravdi's own invocation of the DC is, itself, *ad hoc*. Multiple *ad hoc* applications of the same principle does not provide us with evidence that the principle is true.

[8]We're convinced by the arguments in Schroeder (2011) that the operative difference between evaluative and deliberative *oughts* is syntactic. Whereas deliberative *oughts* are control predicates, evaluative *oughts* are raising predicates. See Bhatt (1998) for counter-arguments that deliberative *oughts* are raising. Since control predicates can't take expletive subjects, we use an expletive subject to force an evaluative reading in (67).

[9]The exact semantic relationship between deliberative *ought* and *can* is unclear. Sinnott-Armstrong (1984) argued that ⌜S ought to A⌝ conversationally implicates that ⌜S can A⌝.

#John ought to study for the exam.

Given the DC, counterfactual evaluative *ought*s should pattern like counterfactual deliberative *ought*s, they should never be assertible when their prejacent is impossible, but this isn't the case. In light of this data, the proponent of the DC might attempt to restrict its scope to deliberative *ought*s. However, this will not succeed for two reasons.

First, the Samaritan Paradox arises with evaluative as well as deliberative *ought*. A speaker can coherently accept (70a), while rejecting (70b).

(70) a. There ought to be a stop to the war in Gaza.

b. $\not\Rightarrow$ There ought to be (an ongoing) war in Gaza.

Second, given the DC, *factual* uses of deliberative *ought* should always be infelicitous. However, as the examples in (71) show, this is not the case.

(71) a. It's certain John will go to college, and he should.

b. John should go to college, and it's certain he will.

In the case of factual uses, there are no special properties of deliberative *ought* which allow us to circumvent the data, because the fact that ⌜*S* ought to A⌝ implies ⌜*S* can A⌝ is fully compatible with a context in which the modal base only contains A-worlds.

3.2.2 Objection 2: deontic actuality entailments.

Bhatt (1999) showed that with perfective aspect ability modals give rise to *actuality entailments* (AEs). With the past perfective, ⌜*S* is able to A⌝ behaves like the implicative verb ⌜*S* managed to A⌝ and, accordingly, entails ⌜*S* did A⌝.

(72) a. John is able to lift the table.

$\not\Rightarrow$ John lifted the table.

b. John was able to lift the table.

$\Rightarrow$ John lifted the table.

In subsequent work, Hacquard (2006, 2009) showed that this phenomenon is not unique to ability modals–nearly all root (i.e., non-epistemic) modals give rise to AEs, including deontics as illustrated by (73).

(73) Sorry, I had to leave early.

$\Rightarrow$ I had to leave early and did.

Since English lacks overt aspectual morphology, the presence of an AE is less obvious in English than, say, a Romance language like Brazilian Portuguese, (74).

(74)	#Desculpa,	eu	tive	que	sair	cedo,	mas	eu	não	saí.
	sorry,	I	have.PST.PFTV	that	leave	early,	but	I	NEG	leave.PST.PFTV

#'Sorry, I had to leave early, but I didn't.'

The fact that deontics generate AEs is hard to reconcile with the DC. It seems that any theory of AEs will need to say that, in the presence of the perfective, deontics have a modal base which is homogeneous with respect to the prejacent to explain why they become veridical; however, when coupled with the DC, any theory of this kind will predict that *should* with perfective aspect leads to presupposition failure due to a violation of the DC.

Borgonovo and Cummins (2007), for example, propose that the perfective makes the modal base *totally realistic*; in other words, in the presence of the perfective the modal base simply becomes the singleton set containing the actual world $f(w) = \{w\}$. However, if we attempted to derive deontic AEs this way, the DC would be violated, leading to presupposition failure. Homer (2011, 2021) alternatively proposes that root modals like *ought* are stative predicates and argues, on this basis, that AEs arise from what he calls "actualistic" aspectual coercion driven by a covert syntactic operator ACT:[10]

(75) $[\![\text{ACT}[\text{ought VP}]]\!] = \lambda e.[\![\text{VP}]\!](e) \wedge \exists e'[\tau(e) \sqsubseteq \tau(e') \wedge \text{ought}([\![\text{VP}]\!](e')]$

Homer's actualistic coercion account predicts that *ought*'s domain will only contain worlds

[10]Here τ is function corresponding to the "temporal trace." τ takes an event e as an argument and returns an interval which can be thought of as the runtime of e.

compatible with the event described by its VP argument; so, again, the DC is violated. Hence, we're skeptical that there's any theory of AEs consistent with the DC.

3.3 Indirect restriction via local contexts

We propose instead that the presuppositions of the prejacent restrict *ought*'s domain of modal quantification.[11]

> **Indirect Restriction.** An assertion of ⌜ought A⌝ is evaluated relative to a modal base f restricted by a local context satisfying A's presuppositions.

Here's a simple toy semantics that illustrates the idea:

(76) $[\![\text{ought}]\!]^{w,f,g,\kappa} = \lambda p_{\langle s,t\rangle}.\forall w' \in \textsc{best}(w, f^{+\kappa}, g) : p(w') = 1$

where $f^{+\kappa}(w) = f(w) \cap \kappa$. In words, given a sentence of the form ⌜ought A⌝, if A presupposes A′, then κ is a local context such that $\kappa \subseteq \mathbf{A'}$. That is, the domain of *ought* is restricted to a local context which entails the prejacent's presuppositions. In contrast, if A presupposes nothing, then the the local context is simply the trivial local context, i.e., $\kappa \equiv \top$. Since a modal with a prejacent which presupposes nothing will quantify over a larger domain than a modal with a prejacent which presupposes something, the Samaritan Inference will not, in general, be valid. For example, while (65a) quantifies over the best worlds where someone has been mugged, (65b) quantifies over *all* the best worlds; hence, (65a) can be true even when (65b) is false.

Our solution has antecedents in the deontic logic literature. Tomberlin (1981), for example, thought that (65a) is a covert conditional which he translated using a conditional deontic operator $O(\cdot \mid \cdot)$ as (77).

[11] A very similar idea has been defended recently by Blumberg and Holguín (2019) to explain readings of attitude reports embedded in conditionals in disjunctions.

(77) $O(S \text{ is helped} \mid S \text{ is mugged})$

While similar, our solutions have a major difference. Tomberlin's solution leaves the Samaritan Paradox a mystery. Given his account, it's unclear why (65a), a sentence without any overt conditional syntax or morphology, is interpreted conditionally.[12] We, in contrast, posit a plausible mechanism for how such a conditional interpretation arises–the local context of the prejacent restricts *ought*'s domain.

For simplicity, we will treat the local context as a parameter of the index of evaluation.[13] Following Schlenker (2009, 2010a,b), we define the local context of a proposition b occurring in a syntactic environment a_c as the smallest proposition p such that $a(p \wedge b')c' \Leftrightarrow ab'c'$ for all well-formed completions b' and c'. One can show that Schlenker's algorithm yields the rules for local contexts in 3.3.1.

Definition 3.3.1. (Local contexts). Let κ be an arbitrary local context. We can then recursively define the local contexts as follows:

- $[\![\neg \mathrm{A}]\!]^{\kappa}$ is defined $\Leftrightarrow$ $[\![\mathrm{A}]\!]^{\kappa}$ is defined
- $[\![\mathrm{A} \wedge \mathrm{B}]\!]^{\kappa}$ is defined $\Leftrightarrow$ $[\![\mathrm{A}]\!]^{\kappa}$ is defined and $[\![\mathrm{B}]\!]^{\kappa+\mathrm{A}}$ is defined
- $[\![\mathrm{A} \vee \mathrm{B}]\!]^{\kappa}$ is defined $\Leftrightarrow$ $[\![\mathrm{A}]\!]^{\kappa}$ is defined and $[\![\mathrm{B}]\!]^{\kappa+\neg\mathrm{A}}$ is defined

For lexical triggers like the change-of-state verb *stop*, we can simply hardwire the precondition it imposes on the local context into the denotation:

(78) $$[\![\text{stop}]\!]^{w,\kappa} = \lambda p_{\langle s,t\rangle} \begin{cases} \text{defined if } \kappa \subseteq \textsc{past}(p)(w) \\ 1 \text{ iff defined and } \mathbf{stop}(p)(w) = 1 \end{cases}$$

[12] Many authors have noted that theories which treat conditional *oughts* as involving a special deontic conditional operator $O(\cdot|\cdot)$ are linguistically implausible. Since they do not attempt to derive the meaning of deontic conditionals from the interaction of *ought* and *if*, we're essentially treating conditional *oughts* as idioms (Thomason, 1981; Kolodny and MacFarlane, 2010).

[13] In an appendix, we show how the current view can be implemented in a continuation-based semantics à la Barker (2022) which, instead of treating the local context as a parameter of the index of evaluation, posits overt local context variables at LF.

3.3.1 A solution to the Samaritan Paradox.

If we combine our entry for *stop* with our new entry for *ought*, it follows that (79a) has the truth-conditions in (79b). So, (79a) is true just in case, in the deontically ideal worlds where John smoked in the past, John stops smoking.

(79) a. John ought to stop smoking.

b. $[\![(79a)]\!]^{w,f,g,\kappa} = \begin{cases} \text{defined if } \kappa \subseteq \textsc{past}(\textbf{smoking}(j))) \\ 1 \text{ iff } \forall w' \in \textsc{best}(w, f^{+\kappa}, g) : \neg(\textbf{smoking}(j))(w') = 1 \end{cases}$

In contrast, (80a) has the truth-conditions in (80b).

(80) a. John ought to have smoked.

b. $[\![80a]\!]^{w,f,g,\kappa} = \begin{cases} \text{defined if } \kappa \subseteq \top \\ 1 \text{ iff } \forall w' \in \textsc{best}(w, f^{+\kappa}, g) : \textsc{past}((\textbf{smoking}(j))(w')) = 1 \end{cases}$

Since (80a) quantifies over a larger domain than (79a), it's clear that (80a) can be false even when (79a) is true.

3.3.2 Restricted readings in disjunctions and conjunctions.

Given our indirect restriction semantics and Schlenker's theory of local context, my solution to the Samaritan Paradox predicts that *ought*-ascriptions in conjunctions and disjunctions generate restricted readings. These predictions are borne out by empirical evidence. Right-embedded *ought* statements in conjunctions are restricted by the information in their left conjunct.

(81) John murdered someone and he ought to go to prison.

$\Rightarrow$ Given that he killed someone, John ought to go prison.

$\not\Rightarrow$ John ought to go prison (whether or not he killed someone).

Likewise, right-embedded *ought* statements in disjunctions are restricted by the complement of the information in their left disjunct.

(82) Either John is innocent, or he ought to go to prison.

$\Rightarrow$ If he isn't innocent, John ought to go to prison.

$\not\Rightarrow$ John ought to go prison (whether or not he killed someone).

Our theory easily accounts for these facts. In ⌜A or ought B⌝, *ought* is evaluated with respect to a domain restricted by the first conjunct, i.e., $\textsc{best}(w, f^{+\mathrm{A}}, g)$. Likewise, in ⌜A or ought B⌝, *ought* is evaluated with respect to domain restricted by the negation of the first disjunct, i.e., $\textsc{best}(w, f^{+\neg\mathrm{A}}, g)$.

Here we can ask an important question: are deontic modals sensitive to symmetric or asymmetric local contexts? Schlenker defines the *symmetric local context* of a proposition b occurring in a syntactic environment a_c as the smallest proposition p such that $a(p \wedge b')c \Leftrightarrow ab'c$. The difference between standard local contexts and symmetric local contexts is that the former quantify over all possible syntactic completions, while the latter is only sensitive to the *actual* completion. This algorithm generates the recursive rules in Definition 3.3.2.

Definition 3.3.2. (Symmetric local contexts). Let κ be an arbitrary local context. We can then recursively define the local contexts as follows:

- $[\![\neg\mathrm{A}]\!]^{\kappa}$ is defined $\Leftrightarrow$ $[\![\mathrm{A}]\!]^{\kappa}$ is defined
- $[\![\mathrm{A} \wedge \mathrm{B}]\!]^{\kappa}$ is defined $\Leftrightarrow$ $[\![\mathrm{A}]\!]^{\kappa+\mathrm{B}}$ is defined and $[\![\mathrm{B}]\!]^{\kappa+\mathrm{A}}$ is defined
- $[\![\mathrm{A} \vee \mathrm{B}]\!]^{\kappa}$ is defined $\Leftrightarrow$ $[\![\mathrm{A}]\!]^{\kappa+\neg\mathrm{B}}$ is defined and $[\![\mathrm{B}]\!]^{\kappa+\neg\mathrm{A}}$ is defined

If the local contexts that restricted the domain of *ought* were symmetric, then we would expect that left-embedded *oughts* in conjunction could generate restricted readings too, but such a reading appears unattested.

(83) John ought to go to prison and he killed someone.

$\Rightarrow$ John ought to go prison (whether or not he killed someone).

For example, in (83), it's not possible to access a reading where John going to prison is conditional on whether or not he murdered someone.

3.3.3 The Samaritan Inference with epistemics.

Booth (2022) points out that epistemic modals, unlike deontics, license the Samaritan Inference, as shown in (84).

(84) a. John might have helped the person who was mugged.

$\Rightarrow$ There might have been a mugging.

b. Susie must regret punching Timmie.

$\Rightarrow$ Susie must have punched Timmie.

To explain various embedded versions of Yalcin's (2007) epistemic contradictions, Mandelkern (2019a, 2023) has recently proposed that local contexts bound the admissible domains of quantification for epistemic modals.

(85) $$[\![\text{might}]\!]^{w,f,g,k} = \lambda p_{\langle s,t\rangle}.\begin{cases}\text{defined if } f(w) \subseteq \kappa \\ 1 \text{ iff } \exists w' \in f(w) : p(w') = 1\end{cases}$$

If we combine Mandelkern's bounded semantics for epistemics with the analysis of factive attitude verbs from the previous section, (84b) has the following truth-conditions:

(86) $$[\![(84b)]\!]^{w,f,g,\kappa} = \begin{cases}\text{defined if } \kappa \subseteq \mathbf{punched}(s,t) \wedge f(w) \subseteq \kappa \\ 1 \text{ iff } \exists w' \in f(w) : \mathbf{regret}(\mathbf{punching}(s,t)) = 1\end{cases}$$

These truth-conditions entail that Susie might have punched Timmie, because the bound presupposition forces the modal base to contain worlds where Susie punched Timmie; so, epistemics license the Samaritan inference. This difference in how epistemic and root modals interact with their local contexts potentially sheds light on the debate regarding whether or not *ought* can have epistemic flavor. Copley (2003, 2006) and Yalcin (2016) have both argued that quasi-epistemic uses of *ought* like (87) are not epistemic. On this reading, ⌜ought A⌝ means, roughly, that A is expected.

(87) It's three. John should be in his office.

Unlike modals taken to be paradigmatically epistemic, like certain uses of *might* and *must*, the quasi-epistemic use of *ought* does not license the Samaritan Inference.[14]

(88) a. John ought to have survived the fall.

$\not\Rightarrow$ John ought to have fallen.

b. John must have survived the fall.

$\Rightarrow$ John must have fallen.

The inference in (88a), for example, is not valid, because John's survival may have been expected even if the fall itself was unexpected. This supports Copley and Yalcin's thesis that the quasi-epistemic reading of *ought* is not a genuine epistemic modal.

However, integrating Mandelkern's bounded semantics with our indirect restriction theory poses one non-trivial problem. Whereas we assume that local contexts are asymmetric, Mandelkern assumes that they're symmetric. This allows Mandelkern to, for example, correctly derive the infelicity of epistemic contradictions with left-embedded epistemic modals. However, since our data seems to support an asymmetric analysis of local contexts, we would like a compositional system that allows us to have both; in an appendix, we sketch such a system drawing on Barker (2022). This still leaves us with a major open question: when can we calculate local contexts symmetrically rather than asymmetrically? Schlenker (2009) posits that local contexts are asymmetric by default, but symmetric local contexts can be accessed at a processing cost. It's not clear that this explanation is sufficient. Given Schlenker's theory, we would expect that speakers could access

[14] Thanks to Fabrizio Cariani (p.c.) for this observation.

symmetric interpretations would *always* be available if the speaker was willing to exert the effort. However, in most contexts, speakers appear only able to access a symmetric or an assymmetric interpretation, not both. For example, in (81), a symmetric interpretation isn't simply dispreferred, it's impossible.

3.3.4 A solution to Åqvist's Paradox.

Åqvist (1967) observed that a monotonic semantics for *ought* coupled with the usual assumption that knowledge is *veridical*–that is, that ⌜*S* knows that A⌝ entails A, generates a paradox, because ⌜ought *S* knows that A⌝ can be true even when ⌜ought A⌝ is false.

(89) a. It ought to be that the night-guard knew about the robbery.

b. ⇏It ought be that there was a robbery.

Åqvist viewed this puzzle as arising from the veridicality of *know* plus UM. In contrast, we believe it arises from the fact that *know* is *factive* plus Local Satisfaction. While philosophers and logicians often use the terms 'veridical' and 'factive' interchangeably to describe an operator O such that $O\phi$ entails ϕ, linguists typically reserve the term 'factive' for operators which *presuppose* the truth of their prejacent (Kiparsky and Kiparsky, 1971; Karttunen, 1971). We can replicate Åqvist's puzzle with virtually any factive attitude verb. For example, both *regret* and *realize* give rise to instances of the puzzle:

(90) a. Susie should regret punching Timmie.

⇏ Susie should have punched Timmie.

b. Carol should realize that she offended Tom.

⇏ Carol should have offended Tom.

However, Åqvist's puzzle does not necessarily arise with any veridical attitude. For example, *demonstrate* is a canonical example of veridical non-factive attitude and it licenses the Samaritan Inference.[15]

[15] Degen and Tonhauser (2022) provide experimental results that are consistent with *demonstrate* being a

(91) John should demonstrate that he feels apologetic.

$\Rightarrow$ John should feel apologetic.

The other canonical veridical non-factive, *be right*, provides us less conclusive evidence, since it can only appear in the prejacent of *ought*, when *ought* has quasi-epistemic flavor.[16]

We can give factive attitude verbs entries where they presuppose that their local context entails their complement.

(92) $$[\![\text{know}]\!]^{w,f,g,\kappa} = \lambda p_{\langle s,t\rangle}. \begin{cases} \text{defined if } \kappa \subseteq p \\ 1 \text{ iff } \textbf{know}(p) = 1 \end{cases}$$

If we combine this analysis of factive attitude verbs with our new entry for *ought*, we obtain the following truth-conditions for (89a):

(93) $$[\![(89a)]\!]^{w,f,g,\kappa} = \begin{cases} \text{defined if } \kappa \subseteq [\![\text{there was a robbery}]\!] \\ 1 \text{ iff } \forall w' \in \textsc{best}(w, f^{+\kappa}, g) : \textbf{know}(j)(w') = 1 \end{cases}$$

However, (89b), in contrast, is restricted by a trivial local context.

(94) $$[\![(89a)]\!]^{w,f,g,\kappa} = \begin{cases} \text{defined if } \kappa \subseteq \top \\ 1 \text{ iff } \forall w' \in \textsc{best}(w, f^{+\kappa}, g) : \textbf{was-a-robbery}(w') = 1 \end{cases}$$

Since (89a) quantifies over a smaller domain than (89b), (89a) does not entail (89b).

3.3.5 The Paradox of Gentle Murder

The Samaritan Paradox also crops up in conditional contexts where it generates the Paradox of Gentle Murder (Forrester, 1984). In most contexts, (95) and (96) are both true and, moreover, seemingly non-contradictory.

non-veridical, non-factive, though this is unclear.

[16] Of the 10 occurrences I found of *be right* in *ought*'s prejacent in the Corpus of Contemporary American English (COCA), *ought* has quasi-epistemic flavor in each.

(95) Smith ought not murder Jones.

(96) If Smith murders Jones, Smith ought to murder Jones gently.

However, suppose that, contrary to the law, Smith does murder Jones.

(97) Smith murders Jones.

Then (98) follows from (96) and (97) by *modus ponens*.

(98) Smith ought to murder Jones gently.

However, given the validity of the Samaritan Inference, we can derive (99), which directly contradicts (95), from (98).

(99) Smith ought to murder Jones.

The restrictor analysis of Kratzer (1991) derives a solution to the Paradox of Gentle Murder based on the failure of Factual Detachment.

Factual Detachment (FD). $\ulcorner$A$\urcorner$, $\ulcorner$if A, ought C$\urcorner$ $\models$ $\ulcorner$ought C$\urcorner$

Kratzer's solution requires that one both adopt a variably strict semantics for the conditional and further assume that the *if*-clause serves as a restrictor of the *ought* in the conditional's consequent (Kratzer, 1986). In contrast, given our indirect restriction theory, FD fails regardless of whether we assume a material, strict, or variably strict semantics for the conditional, and regardless of whether we adopt the restrictor analysis.

Given a material conditional analysis, it follows that $p = \mathbf{A}$ is the smallest proposition such that $A \rightarrow (p \wedge B') \Leftrightarrow A \rightarrow B'$ for all well-formed completions B'. So, assuming a material analysis of the conditional, Schlenker's theory of local context, therefore, predicts that the local context κ of a conditional consequent is simply its antecedent.

(100) $$[\![(96)]\!]^{w,f,g,\kappa} = \begin{cases} \text{defined if } \kappa \subseteq \textbf{murders}(s,j) \\ 1 \text{ iff } \textbf{murders}(s,j) \Rightarrow \forall w' \in \textsc{best}(w, f^{+\kappa}, g) : \\ \qquad \textbf{gently}(\textbf{murders}(s,j))(w') = 1 \end{cases}$$

(101) $$[\![(99)]\!]^{w,f,g,\kappa} = \begin{cases} \text{defined if } \kappa \subseteq \top \\ 1 \text{ iff } \forall w' \in \textsc{best}(w, f^{+\kappa}, g) : \textbf{murders}(s,j)(w') = 1 \end{cases}$$

Since (96) and (97) can be true even when (99) is false, then reasoning from (96) and (97) to (98) by *modus ponens* fails. We can derive a similar failure of Deontic Detachment assuming a strict conditional theory.

(102) $$[\![\text{if}]\!]^{w,f,g} = \lambda p_{\langle s,t\rangle}\lambda q_{\langle s,t\rangle}.\forall w' \in f^{+p}(w) : q(w') = 1$$

Given Schlenker's theory of local context, the local context of a strict conditional's consequent is the smallest proposition p such that $\Box(\mathrm{A} \rightarrow (p \wedge B')) \Leftrightarrow \Box(\mathrm{A} \rightarrow B')$ for all well-formed completions B'. It's straightforward to show that $p = \bigcup\{f^{+\mathrm{A}}(w) : w \in c\}$. This means that, in (96), given a strict conditional analysis, the modal domain is restricted to the circumstantially possible worlds where Smith murders Jones that are accessible from the context set, resulting in the following truth-conditions:

(103) $$[\![(96)]\!]^{w,f,g,\kappa} = \begin{cases} \text{defined if } \kappa \subseteq \bigcup\{f^{+\textbf{murders}(s,j)} : w \in c\} \\ 1 \text{ iff } \forall w' \in f^{+\textbf{murders}(s,j)} : \forall w'' \in \textsc{best}(w', f^{+\kappa}, g) : \\ \qquad \textbf{gently}(\textbf{murders}(s,j))(w') = 1 \end{cases}$$

Likewise, we can derive a similar failure of Deontic Detachment assuming a strict conditional theory where MIN is the set of circumstantially possible worlds that are the closest given the ordering induced by g:

(104) $$[\![\text{if}]\!]^{w,f,g} = \lambda p_{\langle s,t\rangle}\lambda q_{\langle s,t\rangle}.\forall w' \in \textsc{min}(w, f^{+p}, g) : q(w'') = 1$$

Similar reasoning shows that the local context of a variably strict conditional's consequent is $p = \bigcup\{\textsc{min}(w, f^{+\mathrm{A}}, g) : w \in c\}$.[17] Hence, given these assumptions, we arrive at the following truth conditions:[18]

(105) $$[\![(96)]\!]^{w,f,g,\kappa} = \begin{cases} \text{defined if } \kappa = \bigcup\{\textsc{best}(w', f^{+\mathbf{murders}(s,j)}, g) : w' \in c\} \\ 1 \text{ iff } \forall w' \in \textsc{min}(w, f^{+\mathbf{murders}(s,j)}, g) : \forall w'' \in \textsc{best}(w', f^{+\kappa}, g) : \\ \hfill \mathbf{gently}(\mathbf{murders}(s,j))(w'') = 1 \end{cases}$$

This means that the modal domain only quantifies over the best circumstantially possible worlds where Smith murders Jones that are accessible from the context set. So, unlike the solution proposed by Kratzer (1991), the indirect restriction approach can solve the Paradox of Gentle Murder while remaining agnostic about the semantics of the conditional and without adopting the restrictor analysis.

3.3.6 The Local Zvolenszky Problem

When paired with the naïve semantics for *ought*, the Lewis-Kratzer restrictor analysis of conditionals generates a well-known bug: it renders ⌜if A, ought A⌝ trivially true (Zvolenszky, 2002; Carr, 2014). However, unlike, say, epistemic conditionals of the form ⌜if A, must have A⌝, these conditionals do not seem trivially true:

(106) a. If Britney Spears drinks Pepsi, she ought to drink Pepsi. deontic

b. If Britney Spears drank Pepsi, she must have drunk Pepsi. epistemic

Given our analysis of the Samaritan Paradox, an analogous problem arises in our system, call it the *local Zvolenszky problem*. For example, in (107), the local context of the right conjunct consists of the worlds in the context set where someone was mugged; hence, our theory predicts that the second conjunct of (107) should be trivially true.

[17] As Mackay (2019) shows, this follows only given the Limit Assumption. If we drop the Limit Assumption, then the local context of a variably strict conditional's consequent is not defined.

[18] We can also derive a similar resulting assuming the restrictor analysis. In that case, the modal base would be restricted by $\mathbf{murders}(s, j)$; however, since κ would be a subset of the best worlds where Smith murders Jones, this restriction does not effect the sentence's truth-value.

(107) John was mugged and someone ought to be mugged.

This may seem like a serious defect of the current theory. However, given that deontic modals embedded in a right conjunct can generate conditional interpretations as we saw in §3.2, it's not surprising that the Zvolenszky problem reappears in this context. Moreover, the local Zvolenszky problem is not only a problem for our theory. It's also a problem for nearly all dynamic analyses of deontic modality. Since our analysis is built on top of the static translation of dynamic semantics, this is unsurprising. In Willer's (2014; 2016) dynamic semantics, a *deontic context* d determines for every context $c \in C$ a deontically ideal state c_d by providing an ordering source $o \subseteq C^C$, which is a set of context change potentials (CCPs). We can then give deontic necessity modals a test semantics à la Veltman (1996).

(108) $c[\Box_d \mathbf{A}] = \{w \in c : c_d[\mathbf{A}] = c_d\}$

However, if we adopt the standard dynamic entry for conjunction, then Willer's semantics will generate instances of the local Zvolenszky problem. Willer blocks this prediction by adopting a version of the Diversity Condition, imposing the requirement that input contexts "be 'non-trivial' in the sense that its prejacent as well as its negation must be open possibilities" (p. 345). The DC is commonly invoked to solve the Zvolenszky problem. Frank (1997), for example, proposed that ⌜if A, ought A⌝ is infelicitous because, when restricted to A-worlds, the domain of *ought* fails to satisfy the DC.

However, given the DC, Zvolenszky (2002) points out that we face the so-called *flipside problem*: ⌜if A, ought A⌝ is predicted to *always* be defective. However, this prediction is incorrect. Despite being instances of ⌜if A, ought A⌝, the following conditionals are felicitous:

(109) a. The Dalai Lama never gets mad easily, so if the Dalai Lama is mad, then he should be mad. (Zvolenszky, 2002, p. 15)

b. John knows the route well, so if John turns on Exit 49, then he should turn on Exit 49. (Carr, 2014, p. 560)

These conditionals are coherent and seemingly true: (109a) is true because, if the Dalai Lama is mad about something, then there's a good reason for him to be mad about the thing in question, and (109b) is true because, if John takes Exit 49, then there's a good reason for him to take Exit 49. Given that the DC predicts that these conditionals will always be unassertible, the DC seemingly isn't the right way to solve Zvolenszky's problem.

In sum, while the local Zvolenszky problem is a problem, it is a tractable problem and simply an instance of a problem any theory of deontic modality must solve anyways. A number of extant solutions to Zvolenszky's problem can be integrated with our solution to the Samaritan Paradox. For example, in Chapter 4, I propose a solution to the ⌜if A, ought A⌝ problem that involves enriching the ordering source with default rules drawing on resources from non-monotonic logic.[19]

3.4 Extending the solution to bouletic attitude verbs

We can also extend this analysis to bouletic attitude verbs. If we give bouletic attitude verbs a semantics based on the relational semantics for attitude verbs proposed by Hintikka (1969), then ⌜S wants A⌝ is true just in case A is true in every world compatible with S's desires.

(110) $[\![\text{want}]\!]^w = \lambda x_e.\lambda p_{\langle s,t\rangle}.\forall w' \in \textsc{Des}(x,w) : p(w') = 1$

where $\textsc{Des}(x,w) = \{w' : \text{all of } x\text{'s desires in } w \text{ are satisfied in } w'\}$

[19] A similar reason-based analysis has also been suggested by Chung (2020) who instead employs causal models. However, Chung's solution relies on positing that conditionals sometimes introduce a second covert epistemic modal (Geurts, 2004); so, it's unclear if it can be generalized to cover cases of the Zvolenszky problem when it arises in conjunction and disjunction.

This naïve Hintikkan semantics is upward monotonic. So, it generates bouletic analogues of the Samaritan Paradox.

(111) John wants a free flight on the Concorde. Asher (1987)

$\not\Rightarrow$ John wants a flight on the Concorde.

(112) Alice wants to die peacefully. Stalnaker (1984)

$\not\Rightarrow$ Alice wants to die.

Here again simply denying UM won't fully block the paradox, because Local Satisfaction alone suffices to generate the paradox. For example, according to Schlenker's algorithm, the local context κ of an attitude verb like *want* is the strongest proposition such that (113a) and (113b) are equivalent.

(113) a. S wants $(p \wedge \text{A})$

b. S wants A

It's easy to check that, assuming a Hintikkan semantics for *want*, κ is $p = \bigcup\{\textsc{Des}(S, w) : w \in c\}$ according to Schlenker's algorithm.[20]

von Fintel (1999) proposed a solution to the Samaritan Paradox, which abandoned the Hintikka semantics in favor of a Kratzer-style semantics where *want* was evaluated with respect to a doxastic modal base f_{DOX} and a preference ordering $\gtrsim_{g(w)}$ on worlds.[21]

(114) $[\![\text{want}]\!]^{w,f,g} = \lambda x_e.\lambda p_{\langle s,t\rangle}.\forall w' \in \textsc{best}(x, w, f_{\text{DOX}}, g) : p(w') = 1$

Given the entry in (129), *want* is monotonic. To circumvent this prediction, von Fintel assumes that bouletic attitude verbs carry a diversity presupposition. von Fintel then explains the failure of bouletic versions of the Samaritan inference using the standard Diversity Condition-based reasoning we discussed in §3.2.

[20] See Blumberg and Goldstein (2023) for an in depth examination of how Schlenker's theory interfaces with various semantics for *want*.

[21] The choice of a doxastic modal base is motivated by the fact that desire ascriptions give rise to *belief projection*. When A presupposes A', then ⌜S wants A⌝ presupposes ⌜S believes A$'$⌝.

In the bouletic context, the Diversity Condition seems like an extremely implausible requirement. As Heim (1992) herself concedes, "it doesn't seem right that one can never speak of wanting things one is convinced will happen or convinced won't happen" (p. 199). There's nothing incoherent about totally realistic desire ascriptions which are true in every doxastically accessible world:

(115) a. I live in Bolivia because I want to live in Bolivia. Iatridou (2000)

b. I want it to rain tomorrow (and I believe it will)

Conversely, there's nothing incoherent about totally unrealistic desire ascriptions which are false at every doxastically accessible world. [22]

(116) a. (Even though it's impossible,) I want this weekend to last forever.

b. I hope I get the job, but I'm certain that I won't.

As in §3.3, our alternative to the Diversity Condition is a new semantics for *want* where the modal base is directly restricted by a local context satisfying the complement's presuppositions:

(117) $[\![\text{want}]\!]^{w,f,g,\kappa} = \lambda x_e.\lambda p_{\langle s,t\rangle}.\forall w' \in \textsc{best}(x, w, f^{+\kappa}_{\textsc{dox}}, g) : p(w') = 1$

The Samaritan Inference fails for *want* for the same reasons it failed for *ought*. The desire ascription which serves as the premise of the Samaritan Inference can have a complement with non-trivial presuppositions, while the desire ascription that serves as the conclusion triggers no presuppositions and, thus, the premise quantifies over a smaller domain that the conclusion.

There is one problem our treatment of desire ascriptions leaves unsolved. When an attitude verb contains a presupposition trigger in its complement, like (118), it often generates two inferences: an inference that the attitude holder believes the presupposition,

[22]There may be rescue strategies that can save DC, but we won't explore them for the sake of brevity. For example, Grano and Phillips-Brown (2022) attempt to reconcile factual and counterfactual *want*-ascriptions with the DC by giving bouletic attitude verbs a modal base which is covertly conditional.

and the inference that the presupposition is actually true. Following Geurts (1998), we call these two inferences the *i-inference* and *e-inference*, respectively.

(118) Ann wants Bill to stop smoking.

$\rightsquigarrow$ Anne believes Bill smokes. i-inference

$\rightsquigarrow$ Bill smokes. e-inference

Our semantics does not predict the i-inference. Rather, our current semantics predicts something weaker; for example, in (118), it predicts that Anne believes Bill possibly smokes. It's hard to encode the i-inference in the truth-conditions without breaking things elsewhere. For example, if we simply added a presupposition requiring that f_{DOX} only contained worlds where the complement was true, then our semantics would predict that counterfactual *want*-ascriptions are infelicitous. One potential explanation is that the i-inference is generated by a pragmatic strengthening mechanism. However, for now, we will leave the question of how to properly derive the i-inference as an open question for future research.

3.5 Conclusion

This paper argued that the Samaritan Paradox arises from how deontic modals interact with presuppositions rather than from Upward Monotonicity (UM), as is often assumed. We also argued that the received solution to the Samaritan Paradox in natural language semantics, which posits that ⌜*S* ought A⌝ carries a diversity presupposition requiring that the domain contains both A-worlds and ¬A-worlds is *ad hoc* and subject to empirical counterexamples. We then developed a novel solution to the Samaritan Paradox according to which the domain of modal quantification for deontics is *restricted* by the commitments of their local context. This we showed provides a plausible solution to Åqvist's Paradox of Epistemic Obligation and the Paradox of Gentle Murder. Moreover, by appealing to the bounded semantics for epistemics proposed by Mandelkern (2019a, 2023), according to which the modal domains of epistemics are *bounded* by their local context, we can explain why epistemic, unlike deontics, license the Samaritan Inference. However, if Mandelkern is correct in his assessment that epistemics are sensitive to symmetric local contexts, this leaves us with an open questions: why are epistemics, but not deontics, sensitive to symmetric local contexts?

Chapter 4: Zvolenszky's Problem and Default Premise Semantics

> The problem, in a nutshell, is this: on Kratzer's (1981a; 1991) analysis, all sentences of the form 'If *p* then it must be that *p*' come out true. . . .
>
> Zvolenszky (2002)

4.1 Introduction

Many have observed that so-called *priority modals*, Portner's term for modals which "have to do with reasons for preferring one situation over another," interact with conditionals in a unique way (2009, p. 184). Theories of how modals interact with conditionals, like the Lewis-Kratzer restrictor theory, predict that ⌜if A, MODAL A⌝ is trivially true, a prediction borne out for epistemic modals:

(119) If Britney Spears drank a Pepsi, she must have drunk a Pepsi. epistemic

However, for priority modals, like deontic *ought* and bouletic *want*, this prediction is clearly wrong:[1]

[1] It's not important whether these conditionals are *actually* false. The important point is that they *can* be false.

(120) a. If Britney Spears drinks a Pepsi, she ought to drink a Pepsi. deontic

b. If Britney Spears drinks a Pepsi, she wants to drink a Pepsi. bouletic

Though first noticed by deontic logicians (van Fraassen, 1972; Jackson, 1985), this problem has recently received significant attention from linguists and philosophers of language where it's become known as Zvolenszky's problem or the "if p, ought p" problem (Zvolenszky, 2002; Cariani, 2013; Carr, 2014; Chung, 2020; Stojnić, 2024).

In this paper, we pursue a novel solution to "if p, ought p" problem that builds on Fabrizio Cariani's (2023) default premise semantics for anankastic conditionals. We argue that (120a) is false, because the fact that Britney Spears drinks a Pepsi is not itself a *reason* for Britney Spears *to drink* Pepsi. In contrast, the fact that Britney Spears drank a Pepsi is a *reason to believe* Britney Spears drank a Pepsi. In other words, while epistemics are sensitive to reasons for belief, priority modals are sensitive to reasons for action. In §4.3, we show how, following Cariani, this intuition can be fleshed out using tools from default logic (Reiter, 1980; Horty, 2007b, 2012). Along the way, we demonstrate that default premise semantics provides natural explanations of many problems plaguing the standard Kratzerian analysis of conditional *oughts* including problematic consequences of Strong Centering and Carr's (2014) "self-frustrating oughts."

4.1.1 The problem in Kratzer semantics

Kratzer (1981a,b) developed a premise semantics for modals and conditionals which is now standard in much of linguistics and philosophy of language. Her semantics is formulated relative to two parameters: a *conversational background* $f : W \to \mathcal{P}(\mathcal{P}(W))$ which maps a world to a set of propositions who intersection $\bigcap f(w)$ is called the *modal base*, and an *ordering source* $g : W \to \mathcal{P}(\mathcal{P}(W))$ which maps a world to a set of propositions which generates the following pre-order on the modal base:[2]

$$w \succsim_{g(w)} v \Leftrightarrow \{p \in g(w) : w \in p\} \supseteq \{p \in g(w) : v \in p\}$$

Given this definition, we can define the domain of *ought* as the set of undominated worlds in the modal base:[3]

$$\textsc{best}(w, f, g) := \{v \in f(w) : \neg\exists u \in \bigcap f(w)[u \succsim_{g(w)} v]\}$$

We can then analyze *ought* as a universal quantifier over this set of undominated worlds:

(121) $[\![\text{ought}]\!]^{w,f,g} = \lambda p_{\langle s,t\rangle}.\forall w' \in \textsc{best}(w, f, g) : p(w') = 1$

Following Lewis (1975), Kratzer (1986) argues that conditional antecedents restrict a (possibly covert) modal in their consequent.

(122) $[\![\textit{if}\,\text{A, B}]\!]^{w,f,g} = [\![\text{B}]\!]^{w,f^{+\text{A}},g}$ where $f^{+\text{A}} = f(w) \cap [\![\text{A}]\!]$

where we require $f^{+\mathbf{A}}$ to satisfy the usual Lewis-Stalnaker constraints on closeness:

- **Success**: $\bigcap f^{+\text{A}}(w) \subseteq \mathbf{A}$
- **Strong Centering**: $\bigcap f^{+\text{A}}(w) = \{w\}$ if $w \in \mathbf{A}$
- **Uniformity**: if $\bigcap f^{+\text{A}}(w) \subseteq \mathbf{C}$ and $\bigcap f^{+\text{C}}(w) \subseteq \mathbf{A}$, then $\bigcap f^{+\text{A}}(w) = \bigcap f^{+\text{C}}(w)$

The restrictor theory is attractive because it satisfies the common desideratum that conditional obligation "should be explicable in terms of the meanings of *if* and *should* construed independently" (Bonevac, 1998, p. 37). However, given these standard assumptions, it immediately follows that ⌜if A, ought A⌝ is always true.

Fact 4.1.1. ⌜if A, ought A⌝ is trivially true.

[2] Unlike in the previous two chapters, we do not define f as the modal base. This is done as a matter of convenience. Given our positive proposal, it's often easier to work directly with the conversational background than the set of worlds consistent with the conversational background.

[3] As in the previous two chapters, we will make the simplifying assumption that this pre-order is well-founded. That is, for any non-empty modal base $f(w) \neq \emptyset$, $\textsc{best}(w, f, g) \neq \emptyset$.

Proof. Follows immediately from the definition of *if* and Success. □

Let's first eliminate a few obvious ways one might attempt to solve the problem. Dropping Success as constraint on closeness may seem like a simple way of resolving the problem, but then we simply have a different problem, because Success is necessary for correctly deriving the triviality of (119). We might also posit that ⌜if A, ought A⌝ has a radically different logical form than ⌜if A, must have A⌝. For example, in deontic logic, it's common to find dyadic theories of conditional obligation.[4] Instead of attempting to capture the meaning of conditional obligations compositionally as arising from the interaction of *if* and *ought*, these theories take as primitive a dyadic conditional obligation operator $O(\cdot|\cdot)$. However, as Thomason (1981) and Kolodny and MacFarlane (2010) point out, given such a theory, conditional obligations are treated as idiomatic, but, whereas idioms are idiosyncratic to particular languages, nearly every language expresses conditional obligations with a conditional and a word expressing obligation.

More promising are the two linguistic solutions to the "if *p*, ought *p*" problem which have been pursued in the Kratzerian tradition: the first, inspired by the textbook solution to the Samaritan Paradox posits that *ought* presupposes its domain is diverse with respect to the truth of its prejacent, while the second posits that the antecedent of a conditional *ought* restricts a covert epistemic modal rather than the overt deontic modal. However, we believe that neither the former pragmatic nor the latter syntactic solution yields a satisfactory resolution of the puzzle.

4.1.1.1 The diversity solution

A prominent strategy for solving the "if *p*, ought *p*" problem draws on a traditional strategy for blocking the Samaritan Paradox (Prior, 1958). Prior observed that, since *ought* is mono-

[4]See, for example, the discussion in Chapter 10 of Chellas (1980).

tonic given a semantics where it's treated as a universal quantifier over accessible worlds, the presuppositions of its prejacent should be inherited under entailment. However, this is intuitively not the case:

(123) The old man who was mugged ought to be helped.
$\not\Rightarrow$ The old man ought to be mugged.

The received solution to the Samaritan Paradox is that deontic modals presuppose that their modal base contains both worlds where their prejacent is true and worlds where their prejacent is false (Heim, 1992; von Fintel, 1999; von Fintel and Heim, 2011).

Modal Base Satisfaction. An assertion of ⌜ought A⌝ in context c presupposes that the modal base $f(w)$ satisfies A's presuppositions for all $w \in c$.

Diversity Condition (DC). An assertion of ⌜ought A⌝ in a context c presupposes that $f(w) \cap [\![A]\!] \neq \emptyset$ and $f(w) \cap [\![\neg A]\!] \neq \emptyset$ for all $w \in c$.

Since presuppositions hold throughout the modal base by Modal Base Satisfaction, the Samaritan Inference fails, because the inferred *ought* does not satisfy the DC. Similarly, Frank (1997) proposed that ⌜if A, ought A⌝ is infelicitous because, when restricted to A-worlds, the domain of *ought* violates the DC.

As I argued in the previous chapter, one might question how explanatory the received solution really is, given that Modal Base Satisfaction and the DC are both *ad hoc* stipulations. However, even if we ignore this concern, serious empirical problems remain. Given the DC, ⌜if A, ought A⌝ is *always* defective. Zvolenszky (2002) dubs this the *flipside problem*, observing that the DC rules out the following felicitous assertions:

(124) a. The Dalai Lama never gets mad easily, so if the Dalai Lama is mad, then he should be mad. Zvolenszky (2002)

b. John knows the route well, so if John turns on Exit 49, then he should turn on

Exit 49. Carr (2014)

These conditionals are coherent and seemingly true: (124a) is true because, if the Dalai Lama is mad about something, that's a good reason to be mad about the thing in question, and (124b) is true because, if John takes Exit 49, that's a good reason to take Exit 49. Given that the DC predicts that these conditionals will always be unassertible, a solution based on the DC appears to be a non-starter.[5]

4.1.1.2 The double modal solution

The restrictor analysis posits that *if*-clauses always restrict a modal in the consequent of a conditional, but often there is no overt modal for the *if*-clause to restrict. Here the proponent of the restrictor analysis typically posits that the *if*-clause restricts a covert epistemic necessity modal. However, if we're already conjecturing that *if*-clauses restrict a covert epistemic necessity modal when no overt modal is present, then it's plausible to think that the *if*-clause may also restricts a covert epistemic necessity modal when an overt modal *is* present.

(125) a. $\Box_d$[antecedent][consequent]

b. $\Box_e$[antecedent][$\Box_d$ consequent]

Instead of always receiving the LF in (125a), the double modal solution posits that deontic conditionals may also receive the LF in (125b).

By positing a structural ambiguity, the double modal solution allows us to predict that many instances of ⌜if A, ought A⌝ are not trivially true.

(126) a. If Britney Spears drinks a Pepsi, she ought to drink a Pepsi.

b. must[if Britney Spears drinks Pepsi][ought Britney drink Pepsi]

[5]The proponent of the DC might appeal to a local accommodation mechanism here to avoid this criticism. For example, they might posit that the LF contains a presupposition accomodation operator à la Beaver and Krahmer (2001). However, as Chatain and Schlenker (2023) have recently observed, there is reason think that

Given the LF in (126b), the double modal analysis correctly predicts that (126a) is not trivially true. This is because there may be a world w' epistemically accessible from the actual world w where Britney Spears drinks a Pepsi, but the deontically best worlds relative to w' contain worlds where Britney doesn't drink Pepsi.

Geurts (2004) motivates the double modal analysis by noting that adverbial quantifiers often give rise to this sort of structural ambiguity. For example, he demonstrates that (127a) is ambiguous between a reading where the *if*-clause restricts the overt abverbial quantifier and one where it restricts a covert epistemic modal.

(127) a. If Beryl is in Paris, she often visits the Louvre

b. often[Beryl is in Paris][she visits the Louvre]

c. must[Beryl is in Paris][[often][she visits the Louvre]]

(127b) makes a claim about conditional frequency–it's true in a situation where Beryl visits the Louvre frequently on her trips to Paris, even if these trips are infrequent. In contrast, (127c) make a claim about unconditional frequency conditional on a supposition; so, this reading is true in a situation where we know that there are many cities where Beryl could be, and, moreover, we know, wherever she is, she visits that city's art museums frequently. This ambiguity can be resolved by moving the adverb to the sentence's periphery:

(128) a. Often if Beryl is in Paris she visits the Louvre.

b. If Beryl is in Paris, she visits the Louvre often.

Hence, Geurts argues that it's not unreasonable to think that deontic modals also generate a similar structural ambiguity.

The main problem with the double modal account is that it overgenerates. For example, the double modal analysis predicts that speakers should be able to access both trivial and non-trivial interpretations of "if Britney drinks a Pepsi, then she ought to drink a Pepsi." However, most speakers can only access the latter interpretation. Kratzer (2012) suggests that speakers might be unable to access the trivial reading because they "do not readily perceive trivial interpretations if non-trivial interpretations exist" (p. 107). However, as she concedes, this seems like an unattractive solution to the problem, because speakers have no problem accessing the trivial interpretation of a sentence like (129), even though there's also a non-trivial interpretation.

(129) I could not possibly work more than I do. Kratzer (2012)

Absent a principled explanation of why speakers are unable to access the trivial interpretation of (120a), we prefer to explore solutions that do not rely on double modalization.

4.2 Default premise semantics

We began the chapter by giving a high-level explanation of why ⌜if A, ought A⌝ is not trivially true. While the fact that A is always a reason to *believe* A, the fact that A is not always a reason *to do* A. So, ⌜if A, ought A⌝ is not trivially true, while ⌜if A, must have A⌝ is trivially true. Put a different way, epistemic modals and priority modals are sensitive to different types of reasons. Epistemic modals are sensitive to reasons for belief. Priority modals are sensitive to reasons for action. In this section, we will attempt to make this rough intuition precise. We take as primitive the notion of a reason for A. Following Scanlon (1998), we will adopt the rough working definition that a reason for A is "a consideration that counts in favor of" A (p. 17).

Non-monotonic logics, like default logic, give us a natural way of formally representing reasons. In default logic, we have rules of the form:

$$[\mathbf{A} \blacktriangleright \mathbf{C}]$$

Following Asher and Bonevac (1996), Morreau (1996) and Horty (2007b, 2012), we interpret such a rule as saying that A provides a *pro tanto* reason for doing C. For example, promising to get lunch with your friend Monica provides you with a *pro tanto* reason for getting lunch with Monica. This *pro tanto* reason may be overridden by other considerations. If you see a child drowning on your way to lunch with Monica, your reason for getting lunch, keeping the promise, is defeated by a more important consideration, helping the drowning child. Since adding new information may activate defaults which "defeat" lower priority defaults already active, adding information does not preserve entailment:

Monotonicity. If $\Gamma \models \psi$, then $\Gamma \cup \{\phi\} \models \psi$.

This is what makes default logic non-monotonic: while Monica making a promise entails that she must keep it, Monica making a promise and witnessing a child drowning does not entail that she must keep the promise.

To resolve a series of puzzles concerning anankastic conditionals, Cariani (2023) enriches Kratzer's semantics by allowing the ordering source to contain default rules. Cariani demonstrates that many puzzles involving anankastics are easily resolved if we allow the ordering source to contain default rules of the form: [**want A** ► **A**]. In other words, the ordering source contains a proposition which states that wanting to A is a *pro tanto* reason for doing A. Extant approaches to anankastics, in contrast, secure the same (if not worse) empirical coverage at much higher prices: either by sacrificing compositionality (Sæbø, 2001), positing non-standard LFs (von Fintel and Iatridou, 2005), or positing an ambiguity in the meaning of *want* (Condoravdi and Lauer, 2016). Like Cariani, we will pursue a strategy of solving the "if *p*, ought *p*" problem by encoding more information into the ordering source while keeping the LF as simple as possible and avoiding positing ambiguities.

4.2.1 A first-pass implementation

Given a default rule $\pi : \mathrm{A} \blacktriangleright \mathrm{C}$, we call π's antecedent, A, it's *precondition*, which we write as $\boldsymbol{prec}(\pi) = \mathrm{A}$, and we call π's consequent, B, it's *conclusion*, which we write $\boldsymbol{concl}(\pi) = \mathrm{C}$. Given a ordering source g of default rules, we calculate the modal domain by taking the intersection (assuming it's non-empty) of all the default rules whose preconditions are entailed by the modal base f.[6][7]

Definition 4.2.1. (Domain calculation (first-pass)).

Given a modal base f and a default ordering source g, we can define the following:

- $\boldsymbol{triggered}(f, g) := \{\pi \in g : \bigcap f \subseteq \boldsymbol{prec}(\pi)\}$
- if $\bigcap \boldsymbol{triggered}(f, g) \neq \emptyset$, then $\textsc{best}(w, f, g) := \bigcap \boldsymbol{concl}(\boldsymbol{triggered}(f, g))$

For the moment, we will ignore the additional complexities we must introduce to deal with the problem of conflicting defaults. We can already sketch how our semantics solves the "if p, ought p" problem by showing that (130) isn't a tautology.

(130) If Britney Spears drinks Pepsi, she ought to drink Pepsi.

For simplicity, let's suppose the ordering source contains only two default rules–Britney Spears drinking Pepsi gives her a *pro tanto* reason to brush her teeth, and Britney's contract saying she must drink Pepsi gives her a *pro tanto* reason to drink Pepsi.

$$g = \{[\textbf{drink-pepsi} \blacktriangleright \textbf{brush-teeth}], [\textbf{contract-says} \blacktriangleright \textbf{drink-pepsi}]\}$$

This means that, when the modal base is restricted to worlds where Britney Spears drinks Pepsi, the domain of *ought* is simply the set of worlds where Britney Spears brushes her teeth (assuming that the modal base does not contain the information that her contract

[6]For readability, I will the suppress the world arguments in what follows.

[7]The following account differs from Cariani (2023) who intersects the best-worlds with the modal base. This generates a theory which is conservative extension of Kratzer's semantics.

says drink Pepsi):

$$\textsc{best}(f^{\textbf{+drink-pepsi}}, g) = \bigcap \boldsymbol{concl}(\boldsymbol{triggered}(f^{\textbf{+drink-pepsi}}, g)) = \{\textbf{brush-teeth}\}$$

Since the worlds where Britney Spears brushes her teeth needn't only include worlds where she drinks Pepsi, (130) is false.

4.2.2 Conflicting defaults and moral dilemmas

We often have conflicting reasons for action. To return to our earlier example, since you made a promise to have lunch with Monica, you have a reason to go to lunch; however, if you pass a drowning child en route to lunch with Monica, you now have a reason not to go lunch. In default logic, this type of conflict is commonly depicted in an inheritance network called the Tweety Triangle, depicted in Figure 4.1. This diagram reflects a situation in which an agent has two conflicting epistemic reasons. On the one hand, Tweety is a bird so that's a *pro tanto* reason to believe Tweety can fly, but Tweety is a penguin so that's a *pro tanto* reason to believe he cannot fly. These cases are problematic for our first pass semantics, because it predicts that the domain of *ought* will be empty whenever a conflict arises between defaults.

In both examples, there's a default rule which clearly should take precedence. My reason for skipping lunch, saving a drowning child, should take precedence over my reason for getting lunch, keeping a promise, because saving drowning children is a more important duty than keeping lunch appointments. Similarly, my reason for believing Tweety cannot fly, that he's a penguin, should take precedence over my reason for believing Tweety can fly, that he's a bird, because the former is stronger evidence than the latter. Formally, we can implement this idea by assuming that there is priority ordering $>$ over defaults which is transitive and irreflexive (Horty, 2007a). While I'll largely ignore

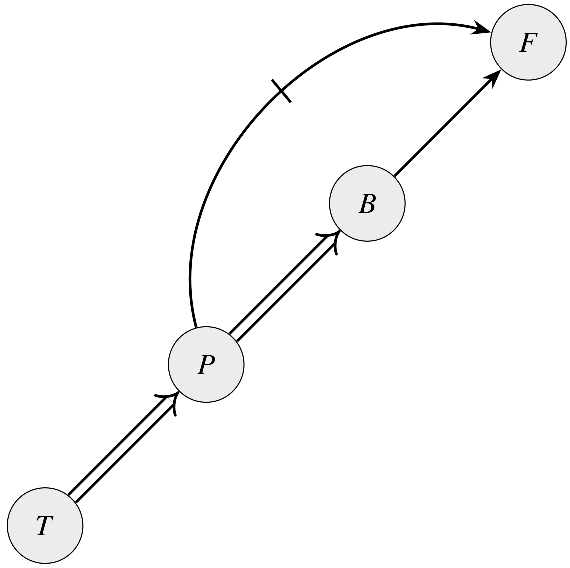

Figure 4.1: Tweety Triangle

the question of where in the semantics we should locate this priority ordering over defaults, I will note that we can straightforwardly treat this additional priority ordering over defaults as a secondary ordering source à la von Fintel and Iatridou (2008).

Definition 4.2.2. (Calculating the domain with tractable conflict).

Given a pair of a modal base f and a set of defaults g.

- $\textbf{\textit{triggered}}(f, g) = \{\pi \in g : \bigcap f \subseteq \textbf{\textit{prec}}(\pi)\}$
- given a set of defaults Π: $\textbf{\textit{conflicted}}(f, g, \Pi) = \{\pi \in \textbf{\textit{triggered}}(f, g) : \bigcap(\textbf{\textit{concl}}(\Pi) \cup \textbf{\textit{concl}}(\pi)) = \emptyset\}$
- $\textbf{\textit{effective}}(f, g) = \{\pi \in g : \neg\exists\Pi(\Pi \subseteq \textbf{\textit{triggered}}(f, g), \pi \in \textbf{\textit{conflicted}}(f, g, \Pi) \wedge \Pi \succ \pi)\}$
- if $\bigcap \textbf{\textit{effective}}(f, g) \neq \emptyset$, then $\textsc{best}(w, f, g) := \bigcap \textbf{\textit{concl}}(\textbf{\textit{effective}}(f, g))$

Introducing a priority ordering over defaults slightly complicates how we calculate the domain. Our dilemmas concerning lunch with Monica and Tweety are what we might call

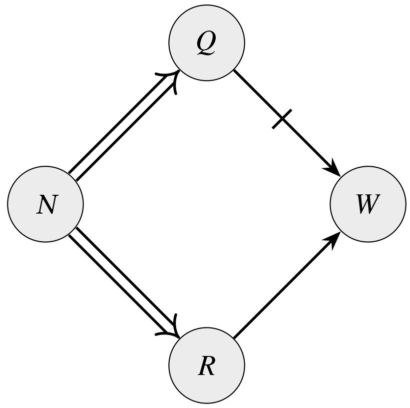

Figure 4.2: Nixon Diamond

tractable dilemmas. Even though two conflicting defaults are active, one clearly should take precedence. However, we often face *intractable dilemmas*, situations where two defaults are active and neither clearly takes precedence over the other. If two conflicting defaults are tied in priority, we want the domain to contain all the worlds consistent with one of the defaults. So, we construct the domain so that it contains all the worlds included in some maximum consistent subset of the conclusion of the effective defaults. A classic example of an intractable dilemma in the non-monotonic logic literature is the Nixon Diamond depicted in Figure 4.2. Nixon was a Quaker, which gives us a *pro tanto* reason to believe he opposes war; however, Nixon was also a Republican, which gives us a *pro tanto* reason to believe he supports war. If these two defaults are tied in priority, then the agent's domain should contain both worlds where Nixon is anti-war and worlds where Nixon is pro-war.

Definition 4.2.3. (Calculating the modal domain with intractable conflict)

- $\boldsymbol{maxicons}_X(Y) = \{Z \subseteq Y : X \cup Z \text{ is consistent and } \forall Z' \subseteq Z,$ $X \cup Z' \text{ inconsistent}\}$
- $\textsc{best}(w, f, g) = \{w : \exists Z \in \boldsymbol{maxicons}_f(\boldsymbol{concl}(\boldsymbol{effective}(f, g))), w \in Z\}$

Unlike Kratzer's original semantics, default premise semantics straightforwardly allows for the possibility of genuine moral dilemmas.[8] The semantics in Definition 4.2.3 corresponds to what Horty (2003) calls the *disjunctive account* of moral dilemmas according to which 'ought A' is true if and only if A is true in *all* maximally consistent subsets. Horty also considers a *conflict account* according to which 'ought A' is true if and only if A is true in *some* maximally consistent subset. Here I agree with von Fintel (2012) that the conflict account likely delivers truth-conditions for *ought* that are too weak.

(131) **Context:** my tax return and my review for *Linguistics & Philosophy* are both due, and I only have time to do one.
(Luckily,) I don't (really) have to do the taxes, because I (also) have to write the review. von Fintel (2012)

According to the conflict account, (131) is true because there's a maximal consistent subset where I just write my review for *Linguistics & Philosophy* and ignore my taxes. Yet, intuitively, this seems like the wrong verdict: I *have* to do my taxes.

4.2.2.1 Self-frustrating oughts

In addition to predicting that ⌜if A, ought A⌝ is trivially true, Carr (2014) observes that Kratzer's semantics also wrongly predicts that ⌜if A, ought not A⌝ is trivially false. She illustrates this with the Death in Damascus case from Gibbard and Harper (1978).

> **Death in Damascus.** If you are in the same city as Death tomorrow, then you'll die. Death has planned to be wherever he predicts you'll be, and he's very reliable in such predictions. Your options are to stay in Damascus or to go to Aleppo. But, as you know, if you stay in Damascus, then that's excellent

[8] von Fintel (2012) sketches a way of modeling moral dilemmas in Kratzer's semantics using multiple ordering sources. If we treat the priority ordering over defaults as a secondary ordering source, then we're proposing something in a very similar vein to what von Fintel is proposing.

evidence that Death will already be there. Similarly for going to Aleppo.

In this context, the following conditionals both seem true:

(132) a. If you're in Damascus, you shouldn't be in Damascus.

b. If you're in Aleppo, you shouldn't be in Aleppo.

While Kratzer's semantics predicts that the conditionals in (132) are both false, default premise semantics yields a natural explanation of why they're true. If you're in Damascus, you have a *pro tanto* reason to go to Aleppo; conversely, if you're in Aleppo, you have a *pro tanto* reason to go to Damascus. So, the conditional is sensitive to the following ordering source:

$$g(w) = \{[\textbf{in-alp} \blacktriangleright \textbf{in-dms}], [\textbf{in-dms} \blacktriangleright \textbf{in-alp}]\}$$

Given this ordering source, there are two maximally consistent sets: one where you go to Aleppo and one where you go to Damascus. When restricted to Damascus worlds, *should*'s domain only contains Aleppo worlds and vice versa when *should*'s domain is restricted to Aleppo worlds. Thus, we correctly predict the truth of both conditionals in (132).

4.2.2.2 The problem of impermissible antecedents

Given the Lewis-Kratzer restrictor analysis, we face the following problem: whenever a conditional *ought* asserts that the subject of the *ought*-ascription in its consequent did something impermissible, as in (133), it follows that they ought to have done what was impermissible.

(133) #If I murdered an innocent person, then I ought to have murdered them.

Since the modal base only contains worlds where I murdered an innocent person when restricted by the antecedent, it follows that I ought to have killed an innocent person. In default premise semantics, we do not face this problem. For example, suppose that the ordering source contains a default stating that you have a *pro tanto* reason not to murder innocent people:

$$g(w) = \{[\textbf{innocent} \;\blacktriangleright \neg\textbf{murder}]\}$$

Given this ordering source, default premise semantics predicts that (133) is trivially false, because the default $[\textbf{innocent} \;\blacktriangleright\; \neg\textbf{murder}] \in g$ is triggered; so, the domain of *ought* contains only worlds where I don't murder the innocent person.

4.2.3 Exploring the logic of default premise semantics

The current approach solves the "if *p*, ought *p*" problem by severing syntactic domain restriction from semantic domain restriction. This allows us to produce interpretations where ⌜if A, ought A⌝ is false, while preserving the simple restrictor analysis, but this also causes some basic principles of conditional logic to fail. For example, the following basic validity also no longer holds:[9]

Right Strengthening. $\ulcorner \text{if A, ought B} \urcorner \models \ulcorner \text{if A, ought } (\text{A} \wedge \text{B}) \urcorner$

This may seem like an unwanted result, but, in the context of conditional *oughts*, there appear to be clear cases where Right Strengthening seems to fail:[10]

(134) If John insults Mary, then he ought to apologize.

$\not\Rightarrow$ If John insults Mary, then he ought to both insult Mary and apologize.

Right Strengthening fails for the same reason ⌜if A, ought A⌝ is no longer trivially true: the fact that A isn't always a reason to do A.

[9] Thanks to Paolo Santorio (p.c.) for drawing my attention to this consequence of the semantics.

[10] Thanks to Chris Tucker (p.c.) for this example.

4.2.3.1 Failure of Conjunctive Sufficiency

Since Kratzer's semantics satisfies Strong Centering, the following is valid:

Conjunctive Sufficiency. ⌜A and C⌝ $\models$ ⌜if A, ought C⌝

The validity of Conjunctive Sufficiency for natural language conditionals is usually debated in the context of so-called "true-true" counterfactuals (Bigelow, 1976; McGlynn, 2012, a.o.) For example, in an early criticism of the Lewis-Stalnaker similarity-based semantics, Bigelow complained that, given Strong Centering, (135) is trivially true.

(135) If London is a large city, then Jupiter has twelve moons.

In the deontic context, Conjunctive Sufficiency licenses inferences which are arguably worse:

(136) John is drinking and John is driving.

$\not\Rightarrow$If John is drinking, he ought to be driving.

Here's the common diagnosis of the problem: the Lewis-Stalnaker semantics allows a conditional to be true even when there's no non-accidental "connection" between the truth of the antecedent and consequent (Bennett, 1974; Fine, 1975). In contrast, default premise semantics requires that there be a non-accidental connection between the antecedent and consequent in the form of a default rule. Hence, Conjunctive Sufficiency is invalid in default premise semantics, because there may be no default triggered by A which allows one to conclude C even if A and C are both true in the actual world.

There are, however, downstream consequences to invalidating Conjunctive Sufficiency. For example, Walters and Williams (2013) defend Conjunctive Sufficiency based on the observation, due to Nute (1980), that semantics which require a "connection" between the antecedent and consequent invalidate Weakening the Consequent:

Weakening the Consequent. If B $\models$ C, then ⌜if A, ought B⌝ $\models$ ⌜if A, ought C⌝

However, it's unclear that Weakening the Consequent is valid for deontic conditionals, because it generates conditional versions of Ross's paradox (Ross, 1944). For example, given Weakening of the Consequent, we would incorrectly predict that the following is a valid inference:[11]

(137) a. If you insult Susie, you should apologize.

b. $\not\Rightarrow$ If you insult Susie, you should either apologize or do it again.

4.2.3.2 Failure of Loop.

Given the Lewis-Stalnaker constraints on closeness, the following principle, which Kraus et al. (1990) call Loop, is valid in Kratzer's semantics for conditional *ought*s.[12]

Loop. $\ulcorner$if A, ought B$\urcorner \wedge \ulcorner$if B, ought C$\urcorner \wedge \ulcorner$if C, ought D$\urcorner \models \ulcorner$if A, ought C$\urcorner$

Santorio (2014, 2019) has recently observed that natural language counterfactuals often invalidate Loop, and Chung (2020) has observed that Santorio's counterexamples naturally generalize to conditional *ought*s. For example, consider the following variant of Santorio's original example:

> **Love Triangle.** Alice, Bob, and Carol are in a love triangle. Bob is pursuing Alice, Carol is pursuing Bob, and Alice is pursuing Carol. They all find their suitors annoying, but want to spend time with their beloved. There's a party planned and all three are invited.

Given this context, the following conditionals are all intuitively true:

[11] von Fintel (2012) argues that Ross's paradox should be given a pragmatic explanation where it arises because *ought* scoping over a disjunction licenses an implicature that both disjuncts are permissible. If this is correct, then Ross-type phenomena might not justify invalidating Weakening the Consequent.

[12] See Halpern (2013) for a proof that Loop is valid given the standard Lewis-Stalnaker constraints on closeness.

(138) a. If Alice goes to the party, Bob should go.

b. If Bob goes to the party, Carol should go.

c. If Carol goes to the party, Alice should go.

However, given Loop, we can then infer (139) which is intuitively false.

(139) #If Alice goes to the party, Carol should go.

This inference will not generally be valid in default premise semantics. Suppose we have the following ordering source:

(140) $g(w) = \{\textbf{alice} \blacktriangleright \textbf{bob}, \textbf{bob} \blacktriangleright \textbf{carol}, \textbf{carol} \blacktriangleright \textbf{alice}\}$

It's easy to verify that, relative to the ordering source in (140), the conditionals in (138) are all true. However, (139) is false, because the only default its antecedent triggers is **alice ▶ bob**; hence, it's domain only contain worlds where Bob goes to the party.

Moreover, the invalidity of Loop provides us with a way of choosing between different algorithms for calculating the domain of *ought*. Our algorithm, based on Cariani (2023), calculates BEST in a manner which is much simpler than alternative algorithms in many other systems of default logic. For example, Horty (2007b, 2012) imposes the constraint that the triggered set be stable where ***triggered***(f, g) is *stable* just in case it's a fixed-point of the ordering source:[13]

$$\boldsymbol{triggered}(f, g) = g$$

In other words, Horty's fixed-point algorithm requires that every default in the ordering source be triggered. Relative to the ordering source in (140), Horty's algorithm validates the Loop inference from (138) to (139), because, in his system, the antecedent of (139) will trigger every default in g.

$$\boldsymbol{triggered}(f^{+\textbf{alice}}, g)) = \{\textbf{alice} \blacktriangleright \textbf{bob}, \textbf{bob} \blacktriangleright \textbf{carol}, \textbf{carol} \blacktriangleright \textbf{alice}\}$$

[13] I am slightly simplifying Horty's system by ignoring the matter conflicting defaults. The idea remains more or less the same once we add a mechanism for dealing with conflicting defaults.

Hence, (139) will be evaluated with respect to the following domain:

$$\textsc{best}(f^{+\textbf{alice}}, g) = \bigcap \boldsymbol{concl}(\boldsymbol{triggered}(f^{+\textbf{alice}}, g)) = \{\textbf{alice} \wedge \textbf{bob} \wedge \textbf{carol}\}$$

So, given Horty's fixed-point algorithm, we predict (incorrectly) that the Loop inference from (138) to (139) is valid.

4.2.3.3 Failure of Agglomeration.

If we allow for the possibility of conflicting defaults, Agglomeration also fails, because, even if A is a *pro tanto* reason to do B and a *pro tanto* reason to do C, A may fail to be a *pro tanto* reason to do both B and C, because B and C may be inconsistent.

> **Agglomeration.** $\ulcorner$if A, ought B$\urcorner$, $\ulcorner$if A, ought C$\urcorner$ $\models$ $\ulcorner$if A, ought (B $\wedge$ C)$\urcorner$

Agglomeration is a very basic principle of conditional logic. For example, when discussing natural language counterfactuals, Hawthorne (2005) calls Agglomeration "overwhelmingly intuitive" (p. 397); so, this may seem like a problem. However, in our system, Agglomeration only fails when there are two triggered, conflicting defaults of equal priority. This means that Agglomeration can only fail in the context of an intractable dilemma. So, while Agglomeration fails, Consistent Agglomeration is still valid.[14]

> **Consistent Agglomeration.** When B and C are consistent, $\ulcorner$if A, ought B$\urcorner$, $\ulcorner$if A, ought C$\urcorner$ $\models$ $\ulcorner$if A, ought (B $\wedge$ C)$\urcorner$.

Since at least van Fraassen (1973), it's been widely acknowledged that Agglomeration must be weakened to Consistent Agglomeration, if we want a deontic logic that avoids predicting that $\ulcorner$ought (A $\wedge$ ¬A)$\urcorner$ is consistent.

[14] Jackson (1985) famously (or perhaps infamously) argued that even Consistent Agglomeration can fail. However, Cariani (2016) and Boylan (2023) have both recently given convincing arguments to the contrary.

4.3 Conclusion

In conclusion, the Lewis-Kratzer restrictor theory of how modals interface with conditionals incorrectly predict that the deontic conditional ⌜if A, ought A⌝ is trivially true; however, this theory correctly predicts that the epistemic conditional ⌜if A, must have A⌝ is trivially true. We propose a simple solution to the puzzle: deontic and epistemic conditionals are sensitive to different kinds of reasons. The fact that A is true is not always a reason for doing A; however, the fact that A is true is always a reason for believing A. We made this intuition precise with Cariani's (2023) default premise semantics, integrating a Kratzer-style (1981b; 2012) semantics for conditionals with default logic (Reiter, 1980; Horty, 2007b, 2012).

Chapter 5: Conclusion

This dissertation argued that several problems concerning the semantics of modality in natural language could be solved by positing that local contexts constrain the domain of modal quantification. Following Mandelkern (2019a, 2023), I argued that local contexts *bound* the admissible epistemic modal bases and, in Chapter 2, employed the resulting semantics to propose a novel solution to Karttunen's Problem concerning the strength of epistemic *must*. In Chapter 3, I argued that local contexts *restrict* the modal bases of deontic modals and bouletic attitude verbs. I used this fact to derive a novel solution to the Samaritan Paradox. However, my account leads to a proliferation of Zvolenszky's Problem. So, in Chapter 4, I proposed a novel solution to Zvolenszky's Problem which enriches the Kratzerian ordering source with default rules.

A major open question remains. Mandelkern's explanation of epistemic contradictions crucially relies on the fact that local contexts are computed symmetrically, while our account of the Samaritan Paradox crucially assumes that local contexts are computed asymmetrically. This raises an obvious question: how do speakers decide when to compute local contexts symmetrically versus asymmetrically? Schlenker (2009) posits that local contexts are asymmetric by default, but symmetric local contexts can be accessed at a processing cost. However, Schlenker's explanation suggests that symmetric interpretations should always be accessible, but dispreferred. However, in most contexts, speakers appear only able to access a symmetric or an assymmetric interpretation, not both. For example, recent experimental evidence from Kalomoiros and Schwarz (2021) suggests that disjunctions are always interpreted symmetrically, while recent experimental evidence from Mandelkern et al. (2020) suggests that speakers never access right-to-left filtering in conjunctions. This suggests that Schlenker's algorithm may not be the correct way to calculate local contexts, since it forces us to choose between local contexts which are uniformly symmetric or local contexts which are uniformly asymmetric.

Chapter A: Bounded semantics

A.1 The basic system

We work with the following basic language $\mathcal{L}$ built from a set of atomic propositions $\textsc{At} = \{A, B, \ldots\}$, negation, disjunction, conjunction, an epistemic possibility modality '$\Diamond$', and an epistemic necessity modality '$\Box$'.

$$\mathrm{A} ::= A \,|\, \neg \mathrm{A} \,|\, \mathrm{A} \vee \mathrm{B} \,|\, \mathrm{A} \wedge \mathrm{B} \,|\, \Diamond \mathrm{A} \,|\, \Box \mathrm{A}$$

We assume that negation, disjunction, and conjunction have the usual Boolean truth-conditions with an additional layer of bound conditions:

- $[\![\neg \mathrm{A}]\!]^{\kappa}$ is defined $\Leftrightarrow$ $[\![\mathrm{A}]\!]^{\kappa}$ is defined
- $[\![\mathrm{A} \wedge \mathrm{B}]\!]^{\kappa}$ is defined $\Leftrightarrow$ $[\![\mathrm{A}]\!]^{\kappa+\mathrm{B}}$ is defined and $[\![\mathrm{B}]\!]^{\kappa+\mathrm{A}}$ is defined
- $[\![\mathrm{A} \vee \mathrm{B}]\!]^{\kappa}$ is defined $\Leftrightarrow$ $[\![\mathrm{A}]\!]^{\kappa+\neg\mathrm{B}}$ is defined and $[\![\mathrm{B}]\!]^{\kappa+\neg\mathrm{A}}$ is defined
- $[\![Att_S(A)]\!]^{\kappa}$ is defined $\Leftrightarrow$ $\forall w \in \kappa : [\![\mathrm{A}]\!]^{Acc(S,w)}$ is defined

while the epistemic necessity modal '$\Box$' has the following local context-sensitive entry:

(141) a. $[\![\Box]\!]^{\kappa}$ is defined $\Leftrightarrow$ $\forall w : f_e(w) \subseteq \kappa \wedge \kappa \neq \emptyset$.

b. if defined, $[\![\Box \mathrm{A}]\!]^{w, f_e, g_e, \kappa} = 1 \Leftrightarrow \forall w' \in \textsc{dom}(w, f_e, g_e) : [\![\mathrm{A}]\!]^{w', f_e, g_e, \kappa} = 1$

c. where $\textsc{dom}(w, f_e, g_e) := f_e(w)$.[1]

[1]I assume that *must*'s ordering source is empty $g_e(w) = \emptyset$, because the strong locality constraint alone is sufficient to generate non-veridicality. However, everything that I say is also consistent with *must* having a non-empty ordering source.

Fact A.1.1. $\Box A \not\models A$

Proof. Suppose $w_@ \notin \kappa$ and $[\![A]\!]^{w_@} = 0$, but $f_e(w_@) \subseteq \kappa$ and κ is non-empty set of A-worlds. Then $[\![\Box A]\!]^{w_@} = 1$, but $[\![A]\!]^{w_@} = 0$. $\Box$

Fact A.1.2. $\neg A \wedge \Box A \models \bot$ and $\Box A \wedge \neg A \models \bot$.

Proof. Suppose $\neg A \wedge \Box A$ is true at w. Then $[\![\Box A]\!]^{w,f,g,\kappa+\neg A}$ must be defined; so, $f_e(w) \subseteq \kappa + \neg A$. So, if defined, $\Box A$ is false at w. Since the definedness conditions for conjunction are symmetric, it follows that $\Box A \wedge \neg A$ is also false at w. $\Box$

Fact A.1.3. $A \wedge \Box\neg A \models \bot$ and $\Box\neg A \wedge A \models \bot$.

Proof. Same proof strategy as Fact A.1.2. $\Box$

Fact A.1.4. $Att_S(\neg A \wedge \Box A) \models \bot$ and $Att_S(\Box A \wedge \neg A) \models \bot$.

Proof. Suppose $Att_S(\neg A \wedge \Box A)$ is defined and true. Then it must be that $[\![\neg A \wedge \Box A]\!]^{Acc(S,w)}$ is defined and true. However, from Fact A.1.2, we know that for an arbitrary local context, $\neg A \wedge \Box A$ is false if defined. So, we arrive at a contradiction. The same argument shows that $Att_S(\Box A \wedge \neg A)$ cannot be true and defined. $\Box$

A.2 Definites and indefinites

Following Mandelkern (2022), we assume that the indefinite has a witness bound and the definite has a familiarity bound:

- $[\![\text{an } x(A, C)]\!]^{w,g,\kappa}$ is defined $\Rightarrow$ if $\exists a \in D_e : [\![A \wedge C]\!]^{w,g[x\to a],\kappa} = 1$, then $[\![A \wedge C]\!]^{w,g,\kappa}$ is defined and true.
- $[\![\text{the } x(A, C)]\!]^{w,g,\kappa}$ is defined $\Rightarrow \forall\langle w', g'\rangle \in \kappa : [\![A]\!]^{w',g',\kappa}$ is defined and true.

Fact A.2.1. $\Box[\text{an } x(\text{A}, \text{C})] \not\models \text{the } x(\text{A}, \text{C})$

Proof. Suppose $[\![\Box[\text{an } x(\text{A}, \text{C})]]\!]^{w, f_e, \kappa}$ is defined and true. Then by the witness bound, $\exists a \in D_e : [\![\text{A} \wedge \text{C}]\!]^{w, g[x \to a], f_e, \kappa = f_e} = 1$, then $[\![\text{A} \wedge \text{C}]\!]^{w, g, f_e, \kappa = f_e}$ is defined and true. By familiarity, $[\![\text{the } x(\text{A}, \text{C})]\!]^{w, g, f_e, c}$ is defined only if $\forall \langle w', g' \rangle \in c : [\![\text{A}]\!]^{w', g', f_e, c}$ is defined. Now suppose $f_e \subset c$, and $(c \setminus f_e) \subseteq \overline{\mathbf{A}}$. Then the familiarity bound of the definite is not satisfied, despite the witness bound of the indefinite being satsified. $\Box$

Chapter B: Indirect restriction semantics

B.1 The basic system

We work with the following basic language $\mathcal{L}$ built from a set of atomic propositions $\texttt{At} = \{A, B, \ldots\}$, negation, disjunction, conjunction, and a deontic modality '$\bigcirc$'.

$$\mathrm{A} ::= A \mid \neg\mathrm{A} \mid \mathrm{A} \vee \mathrm{B} \mid \mathrm{A} \wedge \mathrm{B} \mid \bigcirc \mathrm{A}$$

We assume that negation, disjunction, and conjunction have the usual Boolean semantics plus the following asymmetric definedness conditions:

- $[\![\neg\mathrm{A}]\!]^{\kappa}$ is defined $\Leftrightarrow$ $[\![\mathrm{A}]\!]^{\kappa}$ is defined
- $[\![\mathrm{A} \wedge \mathrm{B}]\!]^{\kappa}$ is defined $\Leftrightarrow$ $[\![\mathrm{A}]\!]^{\kappa}$ is defined and $[\![\mathrm{B}]\!]^{\kappa+\mathrm{A}}$ is defined
- $[\![\mathrm{A} \vee \mathrm{B}]\!]^{\kappa}$ is defined $\Leftrightarrow$ $[\![\mathrm{A}]\!]^{\kappa}$ is defined and $[\![\mathrm{B}]\!]^{\kappa+\neg\mathrm{A}}$ is defined

Finally the deontic modal '$\bigcirc$' has the following local context-sensitive entry:

(142) a. $[\![\bigcirc\mathrm{A}]\!]^{w,f,g,\kappa} = 1 \Leftrightarrow \forall w' \in \textsc{best}(w, f^{+\kappa}, g) : [\![\mathrm{A}]\!]^{w,f,g,\kappa} = 1$

b. where $f^{+\kappa}(w) := f(w) \cap \kappa$, and

c. $\textsc{best}(w, f, g) := \{w' \in f(w) : \neg\exists w'' \in f(w).w'' \gtrsim_{g(w)} w'\}$.

Not that, $\kappa = \top$, $f^{+\kappa}(w) = f(w)$. Hence, the semantics in (144) collapses into the standard Kratzerian semantics.

In one sense, the semantics in (144) is non-monotonic, while in another sense the semantics in (144) is monotonic. If we think about preservation of *truth-at-world*, $\models$, it's easy to find counterexamples to Upward Montonicity (UM) (just pick A and B such that A $\models$ B, but such that A is evaluated with respect to a non-trivial κ while B is evaluated with respect to a trivial κ.) However, if we consider *truth-at-a-point*, or *diagonal consequence*, $\models_d$, UM is valid, because the domain of the two deontic modals are restricted by the same local context.

Fact B.1.1. If A $\models_d$ B, then $\bigcirc$A $\models_d$ $\bigcirc$B.

In other words, the semantics is monotonic if we hold fixed the semantic contribution of presupposition.

B.2 Local contexts of conditionals

To solve the Paradox of Gentle Murder, we show that Deontic Detachment fails in the current system. First, let's consider the material conditional A $\rightarrow_m$ C. It's easy to show that the material conditional has the following definedness conditions:[1]

- $[\![\mathrm{A} \rightarrow_m \mathrm{C}]\!]^{\kappa}$ is defined $\Leftrightarrow$ $[\![\mathrm{A}]\!]^{\kappa}$ is defined and $[\![\mathrm{C}]\!]^{\kappa+\mathrm{A}}$ is defined

Fact B.2.1. A, A $\rightarrow_m$ $\bigcirc$C $\not\models$ $\bigcirc$C

Proof. Assume A and A $\rightarrow_m$ $\bigcirc$C are true. Then A $\rightarrow_m$ $\bigcirc$C is true just in case $\bigcirc$C is true in its local context $\kappa + \mathbf{A}$. However, suppose $\textsc{best}(w, f^{\mathbf{A}}, g) \subset \textsc{best}(w, f, g)$, then it's easy to construct a model where $\bigcirc$A is false when evaluated against a trivial local context. □

Next, let's consider the strict conditional A $\rightarrow_s$ C, which has the following truth-conditions:

(143) $[\![\mathrm{A} \rightarrow_s \mathrm{C}]\!] = 1 \Leftrightarrow \forall w' \in f^{\mathrm{A}}(w) : [\![\mathrm{A}]\!] = 1$

[1]This can be proved from the definedness conditions of disjunction by observing that A $\rightarrow_m$ C $\Leftrightarrow$ $\neg$A $\vee$ C.

We can show that the local context of the consequent of a strict conditional is the union of all the antecedent worlds accessible from the initial local context:

- $[\![A \rightarrow_s C]\!]^\kappa$ is defined $\Leftrightarrow$ $[\![A]\!]^\kappa$ is defined and $[\![C]\!]^{\bigcup\{f^{+A}(w):w\in\kappa\}}$ is defined

Fact B.2.2. $A, A \rightarrow_s \bigcirc C \not\models \bigcirc C$

Proof. Same strategy as Fact B.2.1 □

Finally, let's consider the variable strict conditional $A \rightarrow_{vs} C$, which has the following truth-conditions:

(144) $[\![A \rightarrow_{vs} C]\!] = 1 \Leftrightarrow \forall w' \in \textsc{best}(w, f^{+A}, g) : [\![A]\!] = 1$

We can show that the local context of the consequent of a variably strict conditional is the union of all the closest antecedent worlds accessible from the initial local context:

- $[\![A \rightarrow_{vs} C]\!]^\kappa$ is defined $\Leftrightarrow$ $[\![A]\!]^\kappa$ is defined and $[\![C]\!]^{\bigcup\{\textsc{min}(w,f^{+A},g):w\in\kappa\}}$ is defined

Fact B.2.3. $A, A \rightarrow_{vs} \bigcirc C \not\models \bigcirc C$

Proof. Same strategy as Fact B.2.1 □

Chapter C: LFs with overt local contexts

C.1 The need for continuations

Local contexts, as Schlenker (2009, 2010a,b) conceives them, are purely pragmatic and calculated in parallel with truth conditions. However, since local contexts are sensitive to order and Logical Form (LF) encodes order, why not represent explicitly represent local contexts in the LF? In this appendix, we sketch how this can be done using the continuation-based framework developed by Barker (2022).

The easiest way to introduce overt local context variables at LF would be to treat all lexical items as predicates of local contexts. However, this naïve implementation leads to type mismatch. The entry for *ought* would be type $\langle\kappa, \langle\langle s, t\rangle, t\rangle\rangle$ and it's prejacent would be of type $\langle\kappa, \langle s, t\rangle\rangle$. We can't compose *ought* with its prejacent via PM, because both the codomain of *ought* and the codomain of its prejacent are of the wrong type (i.e., not of type t), and we can't use FA because the prejacent is not type κ. If it were allowed, we would like to momentarily ignore the presence of κ and simply compose *ought* directly with its prejacent using FA. It turns out we can formulate a version of FA like this, if we follow Barker (2022) and treat local contexts as continuations.

In programming languages, continuations play a role akin to local contexts in natural languages. Consider, for example, the expression $\ulcorner(1+2) * (3+4)\urcorner$. If we evaluate this expression left-to-right, we start by evaluating $\ulcorner(1+2)\urcorner$ and then we plug its value, 3, the *reduced expression* (redex), back into the original expression $\ulcorner(3) * (3+4)\urcorner$. We call the part of the expression which isn't the redex we're evaluating and which we can write as the function $\ulcorner \lambda x. x * (3+4) \urcorner$, the *continuation*. The continuation is the context in which the redex is evaluated and is a function from the value of the redex to the result of the entire computation.[1]

As we will see, the local contexts-as-continuations view has several attractive features. First, the view allows us to easily integrate standard algorithms for computing local contexts like the one proposed by Schlenker (2009, 2010a,b) into our compositional semantics. Moreover, whereas Karttunen, and the literature following him, typically assumes that local contexts are computed in a linear fashion from left-to-right, modeling local contexts as based on continuations allows us to compute local contexts based on *evaluation order*, which may differ from the linear order (Shan and Barker, 2006; Barker and Shan, 2014).

C.2 The formal system

Like Barker (2022), we will treat the first argument of a continuized expression as representing the local semantic context.

(145) $\boxed{\alpha} = \lambda\kappa.\kappa\alpha$ (continuized lexical items)

A continuation can be viewed as a generalization of the LIFT type-shifter introduced by Partee (1986). LIFT takes individual-type denotations and gives them generalized quantifier-type denotations.[2] For example, LIFT(John) = $\lambda\kappa.\kappa$(John) takes *John*, an expression of type e, and lifts it to an expression of type $(e \to t) \to t$. A continuation can be viewed as performing a similar operation: if α is type X, then continuizing $\kappa(\alpha)$ results in an expression of type $(X \to \sigma) \to \sigma$ where σ is the *answer type*–the type of end result of

[1]The presentation here will be somewhat light on details. For a more in-depth introduction to continuations, see chapter 5 of Friedman et al. (2008).

[2]This is why, for example, Barker (2002) used continuations to develop a compositional theory of quantification that does not have to rely on movement operations like Quantifier Raising (QR) or so-called "Cooper storage" (Cooper, 1983).

the entire computation. Following Frege, natural language semantics typically takes the answer type to be something of type t (i.e., a Boolean).[3]

Now that we can continuize lexical items, we need to answer the following question: how we do Function Application (FA) on two continuized lexical items? There are actually multiple ways of continuizing FA, but, for the moment, let's focus on the following:

> **Continuized Function Application.** If $[\beta\gamma]$ is an instance of function application combining to sub-LFs β and γ, then the continuized version $\boxed{[\beta\gamma]}$ is given by the following:

$$\boxed{[\beta\gamma]} = \lambda\kappa.\boxed{\gamma}(\boxed{\beta}(\lambda xy.\kappa[xy]))$$

This says that content is added to the local context only after it has been evaluated. This continuized version of FA solves the compositional problem we faced earlier, because it gives a means of doing FA "inside" a local context. Putting everything together, here's a sample derivation:

$$\begin{aligned}
\boxed{[\textbf{john smoked}]} &= \lambda\kappa.\boxed{\textbf{smoked}}(\boxed{\textbf{john}}(\lambda xy.\kappa[xy])) \\
&= \lambda\kappa.\boxed{\textbf{smoked}}((\lambda\kappa'.\kappa'\,\textbf{john})(\lambda xy.\kappa[xy])) \\
&= \lambda\kappa.\boxed{\textbf{smoked}}((\lambda xy.\kappa[xy])\,\textbf{john}) \\
&= \lambda\kappa.\boxed{\textbf{smoked}}(\lambda y.\kappa[\textbf{john}y]) \\
&= \lambda\kappa.(\lambda\kappa'.\kappa'\,\textbf{smoked})(\lambda y.\kappa[\textbf{john}y]) \\
&= \lambda\kappa.(\lambda y.\kappa[\textbf{john}y])\textbf{smoked} \\
&= \lambda\kappa.\kappa[\textbf{john smoked}]
\end{aligned} \tag{C.1}$$

The only difference between the continuized LF and the uncontinuized LF is the presence

[3] But, the answer type could, in principle, be something else like an object of type $\langle s, t\rangle$ (i.e., a set of world) if one adopts a dynamic semantics.

of κ, the local context, in the final line of the derivation. In fact, Plotkin (1975) proved that we can always recover the ordinary truth-conditions of the sentence, simply by applying the identity function $1_{\mathbf{id}} = \lambda\kappa.\kappa$ to the continuized result $(\lambda\kappa.\kappa[\textbf{john smoked}])1_{\mathbf{id}} =$ **john smoked**, a result Plotkin calls the Simulation Theorem.

C.2.1 Associating local contexts with entailments

To explain presupposition satisfaction, we need to associate local contexts with a set of entailments. In Barker's theory, these entailments are those that, for an arbitrary local context f, will be present no matter how the arguments of f are instantiated. First, we have to consider the information that will be excluded no matter what argument we instantiate. For example, if $f = \lambda p.\textbf{rain} \wedge p$, then $f p$ will entail **rain** no matter what p is instantiated. So, **rain** is an upper bound on the set of worlds compatible with the information in f. No matter what p we instantiate non-rain worlds will be excluded. Second, we have have to consider the information that will be included no matter what argument we instantiate. For example, if $f = \lambda p.\textbf{rain} \vee p$, then **rain** will entail $f p$ no matter what p is instantiated. So **rain** is a lower bound on the set of worlds compatible with the information in f. No matter what p we instantiate rain worlds will be included.

For an arbitrary local context f, least upper bound of f's semantic commitments $\lceil f \rceil$ and the greatest lower bound of f's semantic commitments $\lfloor f \rfloor$ as follows:

(146) $$\lceil f \rceil = \begin{cases} f & \text{if } f \text{ is type } st \\ \exists x_a.\lceil f x \rceil & \text{if } f \text{ is type } a \rightarrow b \end{cases}$$

(147) $$\lfloor f \rfloor = \begin{cases} f & \text{if } f \text{ is type } st \\ \forall x_a.\lfloor f x \rfloor & \text{if } f \text{ is type } a \rightarrow b \end{cases}$$

In other words, $\lceil f \rceil$ is the most informative way of saturating f's arguments, while $\lfloor f \rfloor$

is the least informative way of saturating f's arguments.

Definition C.2.1. The *commitments* of κ are $|\kappa| := \lceil\kappa\rceil \wedge \neg\lfloor\kappa\rfloor$.

We assume that the a presupposition p in a local context κ satisfied just in case $|\kappa| \longrightarrow p$; that is, just in case the commitments of κ entail p. We calculate the commitments by excluding all the worlds that do not satisfy the least upper bound, since these will be rejected no matter how the utterance is continued, and all the worlds that satisfy the greatest lower bound, since this information is redundant and can be safely ignored throughout the rest of the computation, because doing so won't affect the final semantic value.

Given these tools, we can now recover the usual commitments for local contexts familiar from Karttunen (1974) and subsequent dynamic theories like Heim (1982, 1983):

LF	Local Context at '[ ]'	$\lceil\kappa\rceil$	$\lfloor\kappa\rfloor$	$\|\kappa\|$
A and [ ]	$\lambda p.A \wedge p$	A	$\bot$	A
A or [ ]	$\lambda p.A \vee p$	$\top$	A	$\neg A$
If A then [ ]	$\lambda p.A \longrightarrow p$	$\top$	$\neg A$	A
Not []	$\lambda p.\neg p$	$\top$	$\bot$	$\top$

Barker's theory of local contexts resembles the theory of local context developed by Schlenker (2009, 2010a,b). Schlenker, recall, defines the local context of a proposition b occurring in a syntactic environment a_c as the smallest proposition p such that $a(p \wedge b')c' \Leftrightarrow ab'c'$ for all well-formed completions b' and c'. The main difference between Barker's continuation-based theory of local context and Schlenker's theory is that the former is a semantic theory, while the latter is a syntactic theory. In Schlenker's theory, one quantifies over syntactic completions, whereas, in Barker's theory, one quantifies over ways of saturating arguments.

While the commitments of local contexts have, thus far, arisen naturally from compositionally-driven mechanisms, we can also hard-wire them into lexical entries. For presupposition triggers, like the change-of-state verb *stop*, we can simply require that the commitments of the local context entail that the verb it modifies was previously true in the past:

(148) $[\![\text{stop}]\!] = \lambda\kappa.\kappa(\lambda px : |\kappa| \rightarrow \textsc{past}(px).\, \mathbf{stop}(px))$

In Barker's system, we can give factive attitude verbs entries where the commitments of their local context entails their complement:

(149) $[\![\text{know}]\!] = \lambda\kappa.\kappa(\lambda p : |\kappa| \rightarrow p.\, \mathbf{know}\, p)$

We can, likewise, give *ought* an entry restricted by the commitments of its local context, implementing the ideas we developed in Chapter 3.

(150) $[\![\text{ought}]\!]^{w,f,g} = \lambda\kappa.\kappa(\lambda p.\forall w' \in \textsc{best}_{\gtrsim_{g(w)}}(f(w) \cap |\kappa|) : p(w') = 1)$

It should be stressed here, as before, that we are hard-wiring indirect restriction into our semantics, it is not something that naturally falls out of Barker's semantic algorithm for local contexts. Figure A.1 shows a derivation of our solution to the Samaritan Paradox in Barker's system, and Figure A.2 shows a derivation of our solution to Åqvist's Paradox.

In the Barker-style compositional system we've adopted, Mandelkern's (2019a; 2023) bounded modality view that local contexts bound the admissible domains of local quantification for epistemics corresponds to the following entry:

(151) $[\![\text{might}]\!]^{w,f} = \lambda\kappa.\kappa(\lambda p_{\langle s,t\rangle} : |\kappa| \rightarrow f(w) \subseteq |\kappa|.\, \exists w' \in f(w) : p(w') = 1)$

We can implement symmetric local contexts by adopting a new symmetric rule for FA:[4]

Symmetric Continuized FA. $\boxed{[\beta\gamma]} = \lambda\kappa.\boxed{\beta}(\lambda x(\boxed{\gamma}(\lambda y.\kappa[xy]))$

The symmetric version of the continuized FA rule evaluates the right constituent first, then the left constituents. Mandelkern's satt projection rules face an explanatory problem

[4]As Barker (2022) points out, Barker (2002) uses the same composition rule in his theory of quantifiers to give a theory of displaced scope.

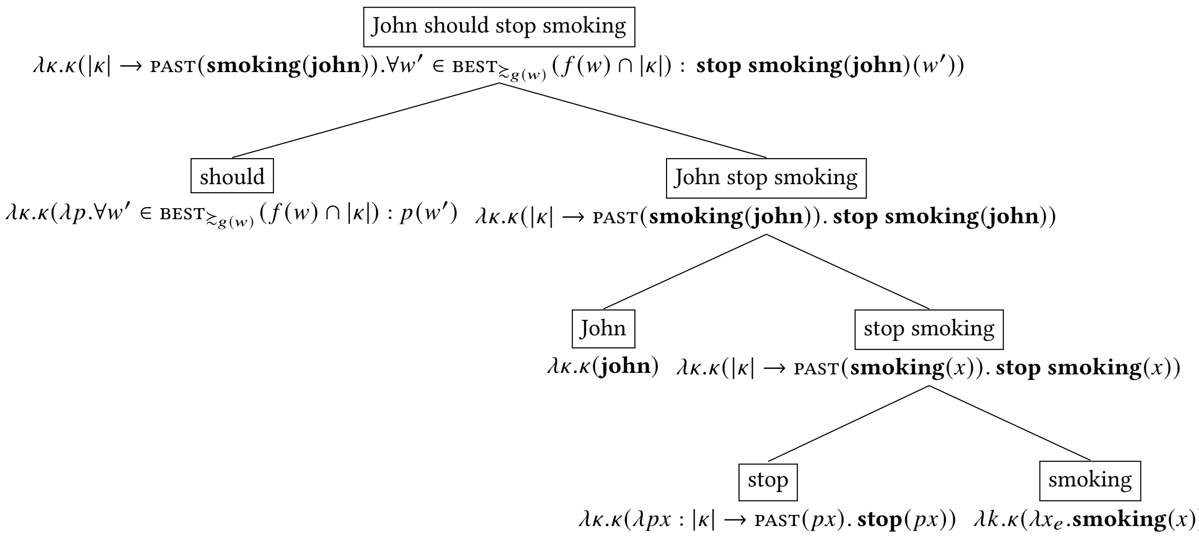

Figure C.1: Deriving *should+stop*

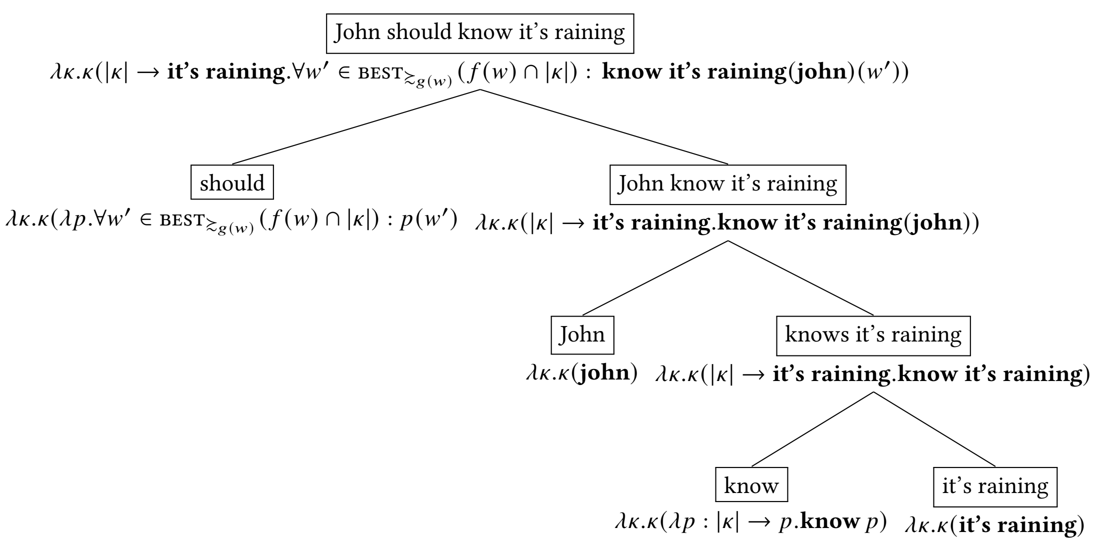

Figure C.2: Deriving *should+know*

resembling the complaint levied against the dynamic semantics proposed by Heim (1982, 1983). The rules Mandelkern postulates for satt projection are stipulative and *ad hoc*; however, as Barker (2022) observes, in the current setting, we only need to posit that *might* has the bound presupposition, and that things compose via the symmetric continuized version of FA. However, there is still a residual problem. Local contexts sometimes must be symmetric, but usually local contexts are asymmetric by default. However, Barker's system is completely unconstrained: we can apply the symmetric version of FA whenever we want. How exactly we should attempt to constrain Barker's system is an open question.

Chapter D: Default premise semantics

D.1 The basic semantics without conflict

Definition D.1.1. (Preconditions and conclusions). Given a default rule $\pi : \mathrm{A} \blacktriangleright \mathrm{C}$, we call π's antecedent, A, it's *precondition*, which we write as $\boldsymbol{prec}(\pi) = \mathrm{A}$, and we call π's consequent, B, it's *conclusion*, which we write $\boldsymbol{concl}(\pi) = \mathrm{C}$.

Definition D.1.2. (Domain calculation (first-pass)).

Given a modal base f and a default ordering source g, we can define the following:

- $\boldsymbol{triggered}(f, g) := \{\pi \in g : \bigcap f \subseteq \boldsymbol{prec}(\pi)\}$
- if $\bigcap \boldsymbol{triggered}(f, g) \neq \emptyset$, then $\textsc{best}(w, f, g) := \bigcap \boldsymbol{concl}(\boldsymbol{triggered}(f, g))$

Fact D.1.1. $\not\models$ ⌜if A, ought A⌝

Proof. Let $f = \{\top\}$ and $g = \{\mathbf{A} \blacktriangleright \neg\mathbf{A}\}$. Then $\textsc{best}(w, f^{+\mathrm{A}}, g) = \bigcap \boldsymbol{concl}(\boldsymbol{triggered}(f^{+\mathrm{A}}, g)) = \{\neg\mathbf{A}\}$. Hence, ⌜if A, ought A⌝ is false. □

Fact D.1.2. $\not\models$ ⌜if A, ought ¬A⌝

Proof. Let $f = \{\top\}$ and $g = \{\mathbf{A} \blacktriangleright \mathbf{A}\}$. Then $\textsc{best}(w, f^{+\mathrm{A}}, g) = \bigcap \boldsymbol{concl}(\boldsymbol{triggered}(f^{+\mathrm{A}}, g)) = \{\mathbf{A}\}$. Hence, ⌜if A, ought A⌝ is false. □

Fact D.1.3. ⌜A and B⌝ $\not\models$ ⌜if A, ought B⌝

Proof. Let $f = \{\top\}$ and $g = \{\mathbf{A} \blacktriangleright \neg\mathbf{B}\}$. Then $\textsc{best}(w, f^{+\mathrm{A}}, g) = \bigcap \boldsymbol{concl}(\boldsymbol{triggered}(f^{+\mathrm{A}}, g)) = \{\neg\mathbf{B}\}$. Hence, ⌜if A, ought B⌝ is false. □

Fact D.1.4. If B $\models$ C, then ⌜if A, ought B⌝ $\not\models$ ⌜if A, ought C⌝.

Proof. Let $f = \{\top\}$ and $g = \{\mathbf{A} \blacktriangleright (\mathbf{B} \wedge \neg\mathbf{C})\}$. Then $\textsc{best}(w, f^{+\mathrm{A}}, g) = \bigcap \boldsymbol{concl}(\boldsymbol{triggered}(f^{+\mathrm{A}}, g)) = \{\mathbf{B} \wedge \neg\mathbf{C}\}$. Hence, ⌜if A, ought B⌝ is true, but ⌜if A, ought C⌝ is false. □

D.2 Semantics with tractable conflict

Definition D.2.1. (Calculating the domain with tractable conflict).

Given a pair of a modal base f and a set of defaults g.

- $\boldsymbol{triggered}(f, g) = \{\pi \in g : \bigcap f \subseteq \boldsymbol{prec}(\pi) \wedge \bigcap(f \cup \boldsymbol{concl}(\pi)) \neq \emptyset\}$
- given a set of defaults Π: $\boldsymbol{conflicted}(f, g, \Pi) = \{\pi \in \boldsymbol{triggered}(f, g) : \bigcap(f + \boldsymbol{concl}(\Pi) + \boldsymbol{concl}(\pi)) = \emptyset\}$
- $\boldsymbol{effective}(f, g) = \{\pi \in g : \neg\exists\Pi(\Pi \subseteq \boldsymbol{active}(f, g), \pi \in \boldsymbol{conflicted}(f, g, \Pi) \wedge \Pi \succ \pi)\}$

D.3 Semantics with intractable conflict

Definition D.3.1. (Calculating the modal domain with intractable conflict)

- $\boldsymbol{maxicons}_X(Y) = \{Z \subseteq Y : X \cup Z$ is consistent and $\forall Z' \subsetneq Z$, $X \cup Z'$ inconsistent$\}$
- $\textsc{best}(w, f, g) = \{w : \exists Z \in \boldsymbol{maxicons}_{\cap f}(\boldsymbol{concl}(\boldsymbol{effective}(f, g))), w \in Z\}$

Bibliography

Adams, E. W. (1975). *The Logic of Conditionals: An Application of Probability to Deductive Logic*. Boston: Reidel.

Aikhenvald, A. Y. (2004). *Evidentiality*. Oxford University Press.

Alonso-Ovalle, L. and P. Menéndez-Benito (2010). Modal indefinites. *Natural Language Semantics 18*, 1–31.

Åqvist, L. (1967). Good samaritans, contrary-to-duty imperatives, and epistemic obligations. *Noûs 1*(4), 361–379.

Asher, N. (1987). A typology for attitude verbs and their anaphoric properties. *Linguistics and Philosophy 10*(2), 125–197.

Asher, N. and D. Bonevac (1996). Prima facie obligation. *Studia Logica 57*(1), 19–45.

Barker, C. (2002). Continuations and the nature of quantification. *Natural language semantics 10*(3), 211–242.

Barker, C. (2022). Composing local contexts. *Journal of Semantics 39*(2), 385–407.

Barker, C. and C.-c. Shan (2014). *Continuations and natural language*, Volume 53. Oxford Studies in Theoretical.

Beaver, D. and E. Krahmer (2001). A partial account of presupposition projection. *Journal of logic, language and information 10*, 147–182.

Bennett, J. (1974). Counterfactuals and possible worlds. *Canadian Journal of Philosophy 4*(2), 381–402.

Bhatt, R. (1998). Obligation and possession. *Papers from the UPenn/MIT roundtable on argument structure and aspect, MITWPL 32*, 21–40.

Bhatt, R. (1999). *Covert modality in non-finite contexts.* Ph. D. thesis, University of Pennsylvania.

Bigelow, J. C. (1976). If-then meets the possible worlds. *Philosophia 6*(2), 215–235.

Blumberg, K. and S. Goldstein (2023). Attitude verbs' local context. *Linguistics and Philosophy 46*(3), 483–507.

Blumberg, K. and B. Holguín (2019). Embedded attitudes. *Journal of Semantics 36*(3), 377–406.

Bonevac, D. (1998). Against conditional obligation. *Noûs 32*(1), 37–53.

Booth, R. (2022). *Underspecification in Modal Contexts.* Ph. D. thesis, Columbia University.

Borgonovo, C. and S. Cummins (2007). Tensed modals. *Coreference, modality and focus*, 1–18.

Boylan, D. (2023). Putting oughts together. *Semantics and Pragmatics 16*.

Burge, T. (1993). Content preservation. *Philosophical Review 102*(4), 457–488.

Burge, T. (1997). Interlocution, perception, and memory. *Philosophical Studies 86*(1), 21–47.

Cariani, F. (2013). 'Ought' and resolution semantics. *Noûs 47*(3), 534–558.

Cariani, F. (2016). Consequence and contrast in deontic semantics. *The Journal of Philosophy 113*(8), 396–416.

Cariani, F. (2023). Anankastic conditionals and the default theory of reasons. Manuscript, University of Maryland, College Park.

Carr, J. (2014). The *if p, ought p* problem. *Pacific Philosophical Quarterly 95*(4), 555–583.

Chatain, K. and P. Schlenker (2023). Local pragmatics redux: Presupposition accommodation and non-redundancy without covert operators. Manuscript Institut Jean Nicod, ENS.

Chellas, B. F. (1980). *Modal Logic: An Introduction.* New York: Cambridge University Press.

Chemla, E. and P. Schlenker (2012). Incremental vs. symmetric accounts of presupposition projection: an experimental approach. *Natural language semantics 20*, 177–226.

Chierchia, G. (1995). *Dynamics of meaning: Anaphora, presupposition, and the theory of grammar.* University of Chicago Press.

Chung, W. (2020). Intervention in deontic reasoning. *Semantics and Pragmatics 13*, 16–1.

Ciardelli, I. (2020). Indicative conditionals and graded information. *Journal of Philosophical Logic 49*, 509–549.

Condoravdi, C. (2002). Temporal interpretation of modals: Modals for the present and for the past. *The construction of meaning 5987.*

Condoravdi, C. and S. Lauer (2016). Anankastic conditionals are just conditionals. *Semantics and Pragmatics.*

Cooper, R. (1983). *Quantification and syntactic theory*, Volume 21. Springer.

Copley, B. (2003). So-called epistemic 'should'. *Snippets 9*, 7–8.

Copley, B. (2006). What *should* should mean? In *Language Under Uncertainty Workshop, Kyoto University.*

Degen, J. and J. Tonhauser (2022). Are there factive predicates? an empirical investigation. *Language 98*(3), 552–591.

Del Pinal, G. (2021). Oddness, modularity, and exhaustification. *Natural Language Semantics 29*, 115–158.

Del Pinal, G. (2021). Probabilistic semantics for epistemic modals: Normality assumptions, conditional epistemic spaces and the strength of must and might. *Linguistics and Philosophy 45*(4), 985–1026.

Fine, K. (1975). Critical notice of lewis, counterfactuals. *Mind 84*(335), 451–458.

von Fintel, K. (1998). The presupposition of subjunctive conditionals. *The interpretive* *tract 25*, 29–44.

von Fintel, K. (1999). NPI licensing, strawson entailment, and context dependency. *Journal of Semantics 16*(2), 97–148.

von Fintel, K. (2008). What is presupposition accommodation, again? *Philosophical perspectives 22*, 137–170.

von Fintel, K. (2012). The best we can (expect to) get? challenges to the classic semantics for deontic modals. Manuscript, MIT.

von Fintel, K. and A. S. Gillies (2008). Cia leaks. *Philosophical Review 117*(1), 77–98.

von Fintel, K. and A. S. Gillies (2010). Must . . . stay . . . strong! *Natural Language Semantics 18*(4), 351–383.

von Fintel, K. and A. S. Gillies (2021). Still going strong. *Natural Language Semantics 29*, 91–113.

von Fintel, K. and I. Heim (2011). Intensional semantics. Lecture notes, MIT.

von Fintel, K. and S. Iatridou (2005). What to do if you want to go to harlem: Anankastic conditionals and related matters. Manuscript, MIT.

von Fintel, K. and S. Iatridou (2008). How to say ought in foreign: The composition of weak necessity modals. *Time and modality*, 115–141.

Forrester, J. W. (1984). Gentle murder, or the adverbial samaritan. *Journal of Philosophy 81*(4), 193–197.

van Fraassen, B. C. (1968). Presupposition, implication, and self-reference. *The Journal of Philosophy 65*(5), 136–152.

van Fraassen, B. C. (1972). The logic of conditional obligation. *Journal of Philosophical Logic 1*(3/4), 417.

Frank, A. (1997). *On Context Dependence in Modal Constructions.* Ph. D. thesis, University of Stuttgart.

Frege, G. (1892). On sense and reference. *Translations from the philosophical writings of Gottlob Frege 2*, 56–85.

Friedman, D. P., M. Wand, and C. T. Haynes (2008). *Essentials of programming languages* (3rd ed.). MIT Press.

Geurts, B. (1996). Local satisfaction guaranteed: A presupposition theory and its problems. *Linguistics and Philosophy*, 259–294.

Geurts, B. (1998). Presuppositions and anaphors in attitude contexts. *Linguistics and Philosophy 21*(6), 545–601.

Geurts, B. (2004). On an ambiguity in quantified conditionals.

Geurts, B. and R. v. d. Sandt (2004). Interpreting focus again.

Giannakidou, A. and A. Mari (2016). Epistemic future and epistemic must: Nonveridicality, evidence, and partial knowledge. *Mood, aspect, modality revisited: New answers to old questions*, 75–117.

Gibbard, A. and W. L. Harper (1978). Counterfactuals and two kinds of expected utility. In A. Hooker, J. J. Leach, and E. F. McClennen (Eds.), *Foundations and Applications of Decision Theory*, pp. 125–162. D. Reidel.

Gillies, A. S. (2007). Counterfactual scorekeeping. *Linguistics and Philosophy 30*, 329–360.

Goodhue, D. (2017). Must φ is felicitous only if φ is not known. *Semantics and Pragmatics 10*, 14–EA.

Grano, T. and M. Phillips-Brown (2022). (counter)factual want ascriptions and conditional belief. *Journal of Philosophy 119*(12), 641–672.

Hacquard, V. (2006). *Aspects of modality.* Ph. D. thesis, Massachusetts Institute of Technology.

Hacquard, V. (2009). On the interaction of aspect and modal auxiliaries. *Linguistics and philosophy 32*, 279–315.

Hacquard, V. and A. Wellwood (2012). Embedding epistemic modals in english: A corpus-based study. *Semantics and Pragmatics 5*, 4–1.

Halpern, J. Y. (2013). From causal models to counterfactual structures. *Review of Symbolic Logic 6*(2), 305–322.

Hawthorne, J. (2005). Chance and counterfactuals 1. *Philosophy and Phenomenological Research 70*(2), 396–405.

Heim, I. (1982). *The Semantics of Definite and Indefinite Noun Phrases.* Ph. D. thesis, UMass Amherst.

Heim, I. (1983). On the projection problem for presuppositions. *Formal semantics–the essential readings*, 249–260.

Heim, I. (1992). Presupposition projection and the semantics of attitude verbs. *Journal of Semantics 9*(3), 183–221.

Hintikka, J. (1969). *Semantics for propositional attitudes.* Springer.

Holguín, B. (2020). Knowledge in the face of conspiracy conditionals. *Linguistics and Philosophy 44*(3), 737–771.

Homer, V. (2011). French modals and perfective: A case of aspectual coercion. In *Proceedings of the 28th West Coast Conference on Formal Linguistics*, pp. 106–114. Cascadilla Proceedings Project Somerville, MA.

Homer, V. (2021). Actualistic interpretations in french. *Semantics and Pragmatics 14*, 12–1.

Horty, J. (2007a). Defaults with priorities. *Journal of Philosophical Logic 36*(4), 367–413.

Horty, J. (2007b). Reasons as defaults. *Philosophers' Imprint 7*, 1–28.

Horty, J. F. (2003). Reasoning with moral conflicts. *Noûs 37*(4), 557–605.

Horty, J. F. (2012). *Reasons as Defaults.* Oxford, England: Oup Usa.

Iatridou, S. (2000). The grammatical ingredients of counterfactuality. *Linguistic inquiry 31*(2), 231–270.

Jackson, F. (1985). On the semantics and logic of obligation. *Mind 94*(374), 177–195.

Jackson, F. and R. Pargetter (1986). Oughts, options, and actualism. *Philosophical Review 95*(2), 233–255.

Kalomoiros, A. and F. Schwarz (2021). Presupposition projection from disjunction is symmetric. *Proceedings of the Linguistic Society of America 6*(1), 556–571.

Karttunen, L. (1971). Some observations on factivity. *Research on Language & Social Interaction 4*(1), 55–69.

Karttunen, L. (1972). Possible and must. In *Syntax and semantics volume 1*, pp. 1–20. Brill.

Karttunen, L. (1974). Presupposition and linguistic context. *Theoretical Linguistics 1*(1-3), 181—-193.

Karttunen, L. (1976). Discourse referents. In *Notes from the linguistic underground*, pp. 363–385. Brill.

Karttunen, L. (2016). Presupposition: What went wrong? In *Semantics and Linguistic Theory*, pp. 705–731.

Karttunen, L. and S. Peters (1979). Conventional lmplicature. In *Presupposition*, pp. 1–56. Brill.

Kaufmann, S. (2017). The limit assumption. *Semantics and Pragmatics 10*, 18–1.

Kibble, R. (1994). Dynamics of epistemic modality and anaphora. In *International workshop on computational semantics*, pp. 121–130. ITK, Tilburg.

Kiparsky, P. and C. Kiparsky (1971). Fact. In *Progress in linguistics.* Mouton.

Klinedinst, N. and D. Rothschild (2012). Connectives without truth tables. *Natural language semantics 20*, 137–175.

Kolodny, N. and J. MacFarlane (2010). Ifs and oughts. *Journal of Philosophy 107*(3), 115–143.

Kratzer, A. (1981a). The notional category of modality. *Formal semantics: The essential readings*, 289–323.

Kratzer, A. (1981b). Partition and revision: The semantics of counterfactuals. *Journal of Philosophical Logic 10*(2), 201–216.

Kratzer, A. (1986). Conditionals. *Chicago Linguistics Society 22*(2), 1–15.

Kratzer, A. (1991). Modality. *Semantics: An international handbook of contemporary research 7*, 639–650.

Kratzer, A. (2012). *Modals and conditionals: New and revised perspectives*, Volume 36. Oxford University Press.

Kraus, S., D. Lehmann, and M. Magidor (1990). Nonmonotonic reasoning, preferential models and cumulative logics. *Artificial Intelligence 44*(1-2), 167–207.

Kripke, S. A. (2009). Presupposition and anaphora: Remarks on the formulation of the projection problem. *Linguistic Inquiry 40*(3), 367–386.

Lackey, J. (1999). Testimonial knowledge and transmission. *Philosophical Quarterly 49*(197), 471–490.

Lackey, J. (2008). *Learning From Words: Testimony as a Source of Knowledge.* Oxford: Oxford University Press.

Lassiter, D. (2016). Must, knowledge, and (in)directness. *Natural Language Semantics 24*(2), 117–163.

Lewis, D. (1973). *Counterfactuals.* John Wiley & Sons.

Lewis, D. (1979). Scorekeeping in a language game. *Journal of philosophical logic 8*, 339–359.

Lewis, D. K. (1975). Adverbs of quantification. In E. L. Keenan (Ed.), *Formal Semantics of Natural Language*, pp. 3–15. Cambridge University Press.

Lewis, K. S. (2021). Anaphora and negation. *Philosophical Studies 178*(5), 1403–1440.

Mackay, J. (2019). Subjunctive conditionals' local contexts. *Linguistics and Philosophy 42*(3), 207–221.

Magri, G. (2009). A theory of individual-level predicates based on blind mandatory scalar implicatures. *Natural language semantics 17*, 245–297.

Mandelkern, M. (2016). Dissatisfaction theory. In *Semantics and linguistic theory*, Volume 26, pp. 391–416.

Mandelkern, M. (2019a). Bounded modality. *Philosophical Review 128*(1), 1–61.

Mandelkern, M. (2019b). What *must* adds. *Linguistics and Philosophy 42*(3), 225–266.

Mandelkern, M. (2022). Witnesses. *Linguistics and Philosophy 45*(5), 1091–1117.

Mandelkern, M. (2023). *Bounded meaning.* Oxford University Press.

Mandelkern, M., J. Zehr, J. Romoli, and F. Schwarz (2020). We've discovered that projection across conjunction is asymmetric (and it is!). *Linguistics and Philosophy 43*, 473–514.

Matthewson, L. (2015). Evidential restrictions on epistemic modals. In *Epistemic Indefinites.* Oxford University Press.

McGlynn, A. (2012). The problem of true-true counterfactuals. *Analysis 72*(2), 276–285.

Montague, R. (1969). On the nature of certain philosophical entities. *The Monist 53*(2), 159–194.

Morreau, M. (1996). Prima facie and seeming duties. *Studia Logica 57*(1), 47–71.

Ninan, D. (2014). Taste predicates and the acquaintance inference. *Semantics and Linguistic Theory 24*, 290–309.

Nute, D. (1980). *Topics in conditional logic*, Volume 20. Springer Science & Business Media.

Partee, B. H. (1986). Noun phrase interpretation and type-shifting principles.

Plotkin, G. D. (1975). Call-by-name, call-by-value and the λ-calculus. *Theoretical computer science 1*(2), 125–159.

Portner, P. (2009). *Modality.* Oxford University Press.

Pranav, A. and V. Hacquard (2013). Epistemics and attitudes. *Semantics and Pragmatics 6*(8), 1–59.

Prior, A. (1958). Escapism: The logical basis of ethics. In A. Melden (Ed.), *Essays in Moral Philosophy*, pp. 135–146.

Reiter, R. (1980). A logic for default reasoning. *Artificial Intelligence 13*(1-2), 81–137.

Roberts, C. (1989). Modal subordination and pronominal anaphora in discourse. *Linguistics and Philosophy 12*(6), 683–721.

Roelofsen, F. (2016). Information, issues, and live possibilities: might in conflicts, free choice, and questions. Presented at the 2016 Workshop on Situations, Information, and Semantic Content in Munchen.

Romoli, J. and M. Mandelkern (2018). Hierarchical structure and local contexts. In *Proceedings of Sinn und Bedeutung*, Volume 21, pp. 1017–1034.

Rooth, M. (1985). Association with focus.

Rooth, M. (1992). A theory of focus interpretation. *Natural language semantics 1*(1), 75–116.

Ross, A. (1944). Imperatives and logic. *Philosophy of Science 11*(1), 30–46.

Ruys, E. G. (2015). On the anaphoricity of too. *Linguistic Inquiry 46*(2), 343–361.

Sæbø, K. J. (2001). Necessary conditions in a natural language. *Audiatur vox sapientiae: A Festschrift for Arnim von Stechow*, 427–449.

Santorio, P. (2014). Filtering semantics for counterfactuals: Bridging causal models and premise semantics. In *Semantics and Linguistic Theory*, Volume 24, pp. 494–513.

Santorio, P. (2019). Interventions in premise semantics. *Philosophers' Imprint 19*(1), 1–27.

Scanlon, T. (1998). *What We Owe to Each Other.* Cambridge, Mass.: Belknap Press of Harvard University Press.

Schlenker, P. (2007). Anti-dynamics: Presupposition projection without dynamic semantics. *Journal of logic, language and information 16*, 325–356.

Schlenker, P. (2008). Presupposition projection: The new debate. In *Semantics and Lin-*

guistic Theory, pp. 655–693.

Schlenker, P. (2009). Local contexts. *Semantics and pragmatics 2*, 3–1.

Schlenker, P. (2010a). Local contexts and local meanings. *Philosophical Studies 151*, 115–142.

Schlenker, P. (2010b). Presuppositions and local contexts. *Mind 119*(474), 377–391.

Schroeder, M. (2011). *Ought*, agents, and actions. *Philosophical Review 120*(1), 1–41.

Schwarz, F. (2009). *Two types of definites in natural language.* University of Massachusetts Amherst.

Shan, C.-c. and C. Barker (2006). Explaining crossover and superiority as left-to-right evaluation. *Linguistics and Philosophy 29*, 91–134.

Sinnott-Armstrong, W. (1984). ?ought' conversationally implies ?can'. *Philosophical Review 93*(2), 249–261.

Soames, S. (1989). Presupposition. *Handbook of Philosophical Logic: Volume IV: Topics in the Philosophy of Language*, 553–616.

Stalnaker, R. (1974). Pragmatic presuppositions. In R. Stalnaker (Ed.), *Context and Content*, pp. 47–62. Oxford University Press.

Stalnaker, R. (1975). Indicative conditionals. In *Context and Content*, pp. 193–210. Oxford University Press.

Stalnaker, R. (1978). Assertion. In *Pragmatics*, pp. 315–332. Brill.

Stalnaker, R. (2002). Common ground. *Linguistics and Philosophy 25*(5-6), 701–721.

Stalnaker, R. (2014). *Context.* Oxford University Press.

Stalnaker, R. C. (1984). *Inquiry.* Cambridge University Press.

Stojnić, U. (2024). The inferential constraint and *if* ϕ *ought* ϕ problem. *Philosophical Studies.*

Strawson, P. F. (1950). On referring. *Mind 59*(235), 320–344.

Sudo, Y. (2014). Presupposition satisfaction in attitude contexts and modal subordination.

The art and craft of semantics: A Festschrift for Irene Heim 2, 175–199.

Swanson, E. (2006). *Interactions with context.* Ph. D. thesis, Massachusetts Institute of Technology.

Thomason, R. H. (1981). Deontic logic as founded on tense logic. In *New studies in deontic logic: Norms, actions, and the foundations of ethics*, pp. 165–176. Springer.

Tomberlin, J. E. (1981). Contrary-to-duty imperatives and conditional obligation. *Noûs*, 357–375.

van Fraassen, B. C. (1973). Values and the heart's command. *The Journal of Philosophy 70*(1), 5–19.

Veltman, F. (1985). *Logics for conditionals.* Ph. D. thesis, University of Amsterdam.

Veltman, F. (1996). Defaults in update semantics. *Journal of philosophical logic 25*, 221–261.

von Fintel, K. and S. Iatridou (2023). Prolegomena to a theory of x-marking. *Linguistics and Philosophy 46*(6), 1467–1510.

Walters, L. and R. Williams (2013). An argument for conjunction conditionalization. *Review of Symbolic Logic 6*(4), 573–588.

Werner, T. (2006). Future and non-future modal sentences. *Natural Language Semantics 14*(3), 235–255.

Willer, M. (2014). Dynamic thoughts on ifs and oughts. *Philosophers' Imprint 14*, 1–30.

Willer, M. (2016). Dynamic foundations for deontic logic. *Deontic Modality*, 324–354.

Willett, T. (1988). A cross-linguistic survey of the grammaticization of evidentiality. *Studies in Language. 12*(1), 51–97.

Williamson, T. (1996). Knowing and asserting. *Philosophical Review 105*(4), 489–523.

Williamson, T. (2000). *Knowledge and its Limits.* Oxford University Press, USA.

Yalcin, S. (2007). Epistemic modals. *Mind 116*(464), 983–1026.

Yalcin, S. (2016). Modalities of normality. *Deontic modality*, 230–255.

Zvolenszky, Z. (2002). Is a possible-worlds semantics of modality possible? a problem for kratzer's semantics. *Proceedings of Semantics and Linguistic Theory (SALT)*, 339–358.