Exploring Information

by

Samuel Carter

A Dissertation submitted to the

School of Graduate Studies

Rutgers, The State University of New Jersey

in partial fulfillment of the requirements

for the degree of

Doctor of Philosophy

Graduate Program in Philosophy

written under the direction of

Jeffrey C. King

&

Anthony Gillies

and approved by

New Brunswick, New Jersey

October, 2020

Abstract of the Dissertation

Exploring Information

by Samuel Carter

Dissertation Directors:

Jeffrey C. King
&
Anthony Gillies

Conditionals are familiar tools, playing an important role in both reasoning and communication. According to orthodoxy, conditionals can be divided into two categories: indicatives and subjunctives. This thesis sets out to assess how these two are alike and how they differ.

I argue that an austere account of this difference is available: conditionals, whether indicative or subjunctive, provide us with a way of investigating the status of the consequent under the supposition of the antecedent. In this way, both types of conditional are fundamentally tools for exploring information.

Acknowledgments

My greatest debt is to my advisors and the members of my committee.

Jeff King has been a constant guide through Rutgers. To a not insignificant extent, his encouragement is responsible for getting me to this point. He has taught me the importance of the details of the case and the value of reading closely. More than anything, he has introduced me to new and interesting questions with a regularity and energy that I will miss. It is because of Jeff that I know that now is always a good time to re-read *Demonstratives*.

Thony Gillies is the reason this thesis has the topic it does. He has been involved from its inception as a 3rd year paper to its completion in its current form, through many updates and revisions. Along the way he has been a source of support, advice and inspiration. The process of developing this thesis has at times felt like an exercise in belatedly drawing out the connections between topics he taught me in my first year in the programme.

Ernie Lepore has taught me about philosophy and philosophers in equal parts. His pressure to ensure that philosophical issues do not go overlooked has kept me honest, and has made both this thesis and my work in general far better than it could otherwise have been. I am immensely grateful to him for his energy and his deep commitment to his students. He will be pleased to know that, this sentence aside, the final version of the thesis is entirely λ-free.

Paul Pietroski arrived at Rutgers when I was in the third year of the programme. His generosity and dedication have been truly astounding. I am indebted to him for the pressure to re-evaluate assumptions across many different pieces of work that would have otherwise gone unconsidered. I could not have had better instruction in how to productively combine research in philosophy of language and linguistics without sacrificing the virtues of either. I hope that the present work does not fall too far short of that lesson.

Simon Charlow was my introduction to linguistics at Rutgers. His availability to offer answers and his patience regardless of how ill-formed the questions have been substantial, and I have gained a huge amount from conversations with him over the last five years. Many people could have taught me how to use formal tools to study language; Simon has taught me how not to misuse them.

I also owe much to members of the faculty who were not on my committee, of course. Andy Egan, Liz Camp, Michael Glanzberg, Brian McLaughlin, Jonathan Schaffer and Dean Zimmerman, amongst others, have all shown me more support than I could reasonably expected. Mercedes Diaz, more than anyone, is due thanks for the fact I am actually completing the programme.

I have been incredibly lucky to be immersed in the graduate community at Rutgers for 5 years. It couldn't have been better. During that time I have benefited daily from conversations with Austin Baker, Bob Beddor, David Black, Laura Callahan, Peter van Elswyck, Will Fleisher, Carolina Flores, Veronica Gomez-Sanchez, Jimmy Goodrich, Cameron Domenico Kirk-Giannini, Morgan Moyer, Lauren Richardson, Andrew Rubner, Eli Shupe, Una Stojnič and Isaac Wilhelm, amongst many others.

I owe especial thanks to Simon Goldstein. He has read much of the material here multiple times and in a variety of different states. His contribution is spread thickly throughout it. On my first visit to Rutgers, Simon woke me up at 5am to talk about my writing sample. From then on, he has been a continuous source of help and inspiration.

Other people outside of Rutgers have played an integral role along the way. John Hawthorne has been actively involved in this project at various different points. In fact, the observations which motivated the paper that the thesis grew out of the Q&A of a talk he gave at Oxford in 2016. Tim Williamson was my supervisor in Oxford before I came to Rutgers; without his commitment and belief in me, I wouldn't have. Daniel Altshuler's enthusiasm and hospitality have made the process fun and have taught me the value of collaboration. Chris Barker, Chris Connett, Cian Dorr, Jeremy Goodman, Phillip Koralus, Ofra Magidor, Salvador Mascarenhas, Duncan Pritchard, Jim Pryor, Craige Roberts and Cecilia Trifogli have all provided me with help and advice at various different points.

Some institutions also deserve acknowledgment: the Rutgers Philosophy of Language Reading Group, RuCCS, the Rutgers Linguistics Department, Baruirs Coffee Roasters on 41st & Roosevelt, the NY Philosophy of Language Workshop, the Corner Tavern on Easton, NASSLI & ESSLI, and the Rutgers Language Workshop. No thanks go to NJTransit.

Material from this thesis has been presented at two APAs. I am grateful to the audiences of those talks for their feedback. I also gained much from workshopping different sections of this thesis in the 3rd year seminar and the dissertation seminar, some of it multiple times. The participants in both provided me with many thoughtful and insightful comments. I owe them a great deal of thanks for their time and effort.

Material in chapters 1-2 includes overlap with work forthcoming in *Mind* as 'A Suppositional Theory of Conditionals'. I am grateful to the editors and referees for their contributions.

My family gave me the curiosity to start this process and the confidence to finish it. They have provided me with love and support, both in person and from a distance, throughout. I owe them my gratitude for bearing with me while I worked on it and my apologies that it is taking me even further away from them.

Most importantly, Natasha. You are my foundation and my home. You make me better. Thank you.

Contents

List of Figures

Chapter 0

Introduction

0.1 Plan

According to orthodoxy, conditionals come in two kinds: indicatives, like (1), and subjunctives, like (2):[1]

(1) If the butler did it, he used the candlestick.

(2) If the butler had done it, he'd have used the candlestick.

This thesis sets out to answer two questions:

(*i*) How are (1) and (2) alike?

(*ii*) How are (1) and (2) different?

Providing an answer to the former requires identifying a common core of conditional meaning. Providing an answer to the latter, in contrast, requires identifying a mechanism which explains the divergence in their behavior. In whatever way this is achieved, we would like our responses to each question to be consonant with one another.

An adequate combined answer to (*i*) and (*ii*) should be expected to satisfy some additional theoretical

[1] For classic discussion, see, e.g., Adams (1965), Lewis (1973), Stalnaker (1975), Slote (1978), Davis (1979), and Gibbard (1981), amongst others.

constraints. First, it should not posit unnecessary ambiguity. Stalnaker (1968, 1975, 2009*b*), Strawson (1986), and, more recently, Starr (2014*a*), defend Uniformity as a constraint on any minimally adequate theory of conditional:[2]

UNIFORMITY The semantic contribution of the conditional form is invariant across subjunctives and indicatives.

Uniformity requires that any difference between indicatives and subjunctives is not attributed to an ambiguity in the conditional form itself. According to Uniformity, '*if*' is univocal across subjunctives and indicatives.

Uniformity has a number of considerations in its favor, beyond mere parsimony. First, as Starr (2014*b*) observes, it is supported by cross-linguistic distribution. There are no attested examples of a language which employs two distinct particles playing the role played by '*if*' in English, one exclusively occurring in subjunctives and another occurring exclusively in indicatives. Second, were '*if*'-clauses in fact ambiguous, we would expect each conditional to admit of two distinct readings. Yet this is not what we find. In general, after other sources of ambiguity are accounted for, conditionals like (1) and (2) are associated with only a single reading each.[3]

As a second constraint on adequacy, we would like a theory to provide a plausible account of the discourse function of the conditional. Ramsey (1931) offers a canonical account of this function, writing:

> "If two people are arguing 'If p, then q?' [...], they are adding p hypothetically to their stock of knowledge and arguing on that basis about q"(Ramsey (1931, 247) .

[2] Here's Stalnaker's formulation of Uniformity:

> "The semantic rule that gives the truth conditions of the conditional as a function of the contextual parameter will be the same for both kinds of conditionals." (Stalnaker (2009*b*, 237, fn20))

Here's Strawson's:

> "[T]he least attractive thing that one could say about the difference between these two remarks is that [...] 'if ... then...' has a different meaning in one remark from the meaning which it has in the other." (Strawson (1986, 230))

.

[3] So-called 'sportscaster' conditionals seem to present an exception to this generalization. I discuss these kinds of examples in §1.2.2.

According to Ramsey, the conditional provides a means for interlocutors to explore their common information. It allows them to investigate properties of the state they would occupy if they were to incorporate the antecedent into that information, without having to accept the antecedent outright. This is vital to our ability to engage in activities like planning for contingencies, speculating about uncertainties and learning from our mistakes. A satisfactory theory should accommodate this, by vindicating the idea behind the Ramsey test.

0.2 Overview

I will address questions (*i*) and (*ii*) over the course of four chapters. The plan is as follows.

Chapter 1: Suppositional Theories of Conditionals

In Chapter 1, I introduce a proposal regarding conditional meaning. The contribution of supposition, at the level of discourse, and the contribution of '*if*'-clauses, at sub-sentential level, appear strikingly similar. (3) can be used to convey the same information as (4).

(3) Suppose the butler did it. Then he used the candlestick.

(4) Suppose the butler had done it. Then he would have used the candlestick.

A loose family of views (e.g., Mackie (1972), Edgington (1995), Barker (1995), DeRose and Grandy (1999), and Barnett (2006)) take this observation as the starting point for a theory of conditionals. According to views of this kind, '*if*'-clauses function to mark a complex speech act, one equivalent to sequentially supposing it's antecedent and then asserting its consequent.

Such views have two well-known shortcomings. First, they have been observed to make inadequate predictions concerning conditionals in embedded environments and to lack an adequate notion of validity. Second, they have also tended to appeal to an impoverished range of data—largely restricting themselves to the apparent similarity between (3) and (4). Yet, to develop a suppositional theory, we need both an adequate understanding of the effect of supposition in discourse and a full assessment of the extent to which these effects are reflected in the behavior of '*if*'-clauses.

I start by addressing the latter issue, presenting a range of further connections between supposition and '*if*"-clauses. The two are shown to interact in similar ways with subjunctive inference patterns (§1.2.1); counterfactuality (§1.2.2); epistemic contradictions (§1.2.3); and speaker ignorance (§1.2.4). I then turn to consider supposition itself (§1.3), identifying a number of logical properties we should expect an adequate account of supposition to capture.

On the basis of these data, I suggest that we need to revise the way we think about supposition. Introducing a supposition does not merely change the possibilities under consideration in a discourse (as suggested by, e.g., Stalnaker (1984), Yalcin (2007) and Green (2000), amongst others). It also changes the way we revise those possibilities when incorporating the effects of downstream supposition: after supposing ϕ, this information must be preserved through any later revisions. In concluding the chapter, I sketch an informal picture of the connection between '*if*'-clauses and supposition which suggests a potential explanation of the data surveyed.

Chapter 2: Information Dynamics

Chapter 2 implements the picture of conditionals sketched at the conclusion of Chapter 1. I start by introducing two different notions of information change: addition and revision (§2.2). Whereas the former is associated with changes to public information of the kind modelled by update semantics (Veltman (1996)), the latter is more commonly associated with changes to private information modelled by belief revision theory (Alchourrón et al. (1985)). I suggest that, in order to account for the difference between assertion and supposition, we need a theory which makes use of both.

To do this, I introduce a dynamic framework in which contexts keep track of both what information is currently under consideration and how that information is to be revised. By enriching our models beyond those employed in update semantics, this puts us in a position to model effects of supposition in addition to assertion.

I start by developing a semantics for a simple boolean language within the framework, which is shown to behave in a classical manner (§2.3.1). I then extend that language to include a supposition operation (§2.3.2) and dynamic strict conditional (§2.3.3) (as defended by, e.g., Dekker (1993), Gillies (2004, 2009) and Starr (2014*a*, ms)).

The conditional of natural language is given a decompositional analysis which combines both operators. According to the account, conditional constructions in natural language can be modeled as dynamic strict conditionals in which the antecedent occurs under the scope of the supposition operator. This proposal satisfies uniformity and vindicates the idea behind the Ramsey test.

To explain the divergent behavior of indicatives and subjunctives, the two are attributed different definedness conditions (§2.4). The former, unlike the latter, presuppose that their antecedent is epistemically possible. Combined with a notion of entailment which takes these definedness conditions into account (Strawson (1952), von Fintel (1997*a*, 1999)), this generates distinct logics for indicatives and subjunctives exhibiting the range of properties desired (§§2.5.1-4).

However, this approach also has a surprising consequence. It validates Collapse —the principle that indicatives and subjunctives are co-entailing. This consequence is discussed in the following chapter.

Chapter 3: Collapse

Chapter 3 abstracts away from the particular theory defended in Chapters 1 and 2 to consider the status of Collapse more generally. I argue that it is in fact motivated on independent grounds (§§3.2.1-3). If we are to preserve principles relating '*might*'-conditionals to the epistemic possibility of their antecedents and consequents, we will need to accept that indicatives and subjunctives are equivalent as long as their presuppositions are satisfied.

In the second half of the chapter, I turn to the philosophical consequences of Collapse. The principle is in tension with a popular picture of how indicatives and subjunctives differ, according to which the former, but not the latter, are information sensitive; unlike subjunctives, the acceptability of an indicative depends on the contextually salient information. Collapse, however, entails that in contexts which license both subjunctives must display the same sensitivity to contextual information as indicatives (§3.4).

Not only does Collapse conflict with an appealing theoretical picture, it also conflicts with widely reported intuitions about particular conditionals. Certain pairs of indicatives and their corresponding subjunctives elicit divergent judgments in context (Adams (1975)). For example: suppose I show you a coin and tell you that if I flip it and it lands heads, I will inform you of the outcome. If I don't flip it or if I flip it and it lands tails, I'll tell you nothing. Having heard nothing from me, it seems you should accept (5). However, you definitely should not accept (6).

(5) If the coin was flipped, it landed tails.

(6) If the coin had been flipped, it would have landed tails.

In the concluding section of the chapter (§3.5), I argue that these judgments can be explained by appeal to an independently plausible pragmatic principle, along with widely noted rules governing context shifting.

Originating with Heim (1991), a number of authors have noted that expressions differing only in their presuppositions can exhibit different pragmatic behavior. In particular, there appears to be pressure on speakers to employ expressions with stronger presupposition, where possible. As a result, the use of an expression with weaker presuppositions generates an implicature that truth-conditionally equivalent alternatives with stronger presuppositions are unlicensed.

When supplemented with a theory of covert context change, this phenomenon can be leveraged to explain the contrast between (5) and (6). The leaves us with an appealingly austere account of the difference between indicatives and subjunctives, on which the only differences between them arise from differences in their presuppositions.

Chapter 4: Causation & Revision

The first three chapters defend a view on which conditionals, whether indicative or subjunctive, are tools for exploring information. Evaluating a conditional involves revising a body of information with its antecedent and evaluating the consequent at the result. As observed in Chapter 3, this has the surprising consequence that the status of a subjunctive can sometimes depend on the body of information it is evaluated against.

Chapter 4 addresses the tension between this claim and another, well-recognized feature of subjunctives: they are sensitive to the causal structure of the world (Bennett (2003), Edgington (2004), Schaffer (2004), Kment (2006)). Judgments about subjunctives vary according to how the various parts of the world depend on one another.

Working within the framework presented in Chapter 2, I show that this tension is merely apparent. By establishing a mapping between causal models (Pearl (2000, 2009)) and entrenchment orderings over worlds (Grove (1988)), we can ensure that the revision operation of a context builds in information about what depends on what. Informally, the idea is that counterfactual revision will, as far as possible, leave matters causally upstream unchanged, while allowing matters causally downstream to vary freely. This provides us with a way of accommodating the sensitivity of subjunctives to causal structure, while retaining the information sensitive framework defended in the earlier chapters.

Chapter 1

Suppositional Theories of Conditionals

1.1 Introduction

Supposition and conditionals appear closely related. For example, (1) and (2) seem to play similar roles in discourse:

(1) Suppose that the butler did it. Then the gardener is innocent.

(2) If the butler did it, then the gardener is innocent.

Suppositional theories take this observation as a starting point, appealing to supposition to provide an account of the natural language conditional (e.g., (Mackie (1972)); (Edgington (1995)); (Barker (1995); amongst others). For example, here is J.L. Mackie:

> "The basic concept required for the interpretation of '*if*'-sentences is that of supposing [...] To assert 'If p, q' is to assert q within the scope of the supposition that p"(1972, 92-93).

This chapter considers the support for—and prospects of—a suppositional theory of conditionals. Existing theories have tended to suffer from two kinds of defect. First, they fail to offer empirical support for the connection between supposition and conditionals beyond examples like (1) and (2). In §1.2, I survey a wider range of linguistic data suggesting that the two are closely related, and supplement it, in §1.3, with further considerations about supposition itself. Second, they struggle to satisfy some basic desiderata in the logic of conditionals. In §1.4, I focus on two issues in particular, to do with accommodating the full range of environments in which conditionals occur and providing an adequate notion of validity.

1.2 Supposition and 'If'-Clauses

1.2.1 Conditional Inferences

Consider the following three inference patterns:[1]

(PR)	ϕ	$\models \psi \Rightarrow (\phi \wedge \psi)$	Preservation
(DA)	$\phi \vee \psi$	$\models \neg\phi \Rightarrow \psi$	Direct Argument
(CT)	$\phi \Rightarrow (\psi \Rightarrow \chi),\ \psi$	$\models \phi \Rightarrow \chi$	Conditional Telescoping

Preservation, the Direct Argument and Conditional Telescoping are often taken to be valid for indicative conditionals. Consider (3)-(5):[2]

(3) ✓ Ada is drinking red wine. (So) if she's eating fish, she's eating fish and drinking red wine.

(4) ✓ Claude is either in London or Paris. (So) if he's not in London, he's in Paris.

(5) ✓ If Lori is married to Kyle, then if she's married to Lyle, she's a bigamist. She's married to Lyle. (So) if she's married to Kyle, she's a bigamist.

Even those who ultimately deny the validity of one of the principles, (e.g., Stalnaker (1975), for the Direct Argument) typically accept that they appear to be valid.

[1](NB: throughout, $\rightsquigarrow$ is used for subjunctives, $\dashrightarrow$ for indicatives, $\Rightarrow$ for non-specific (i.e., subjunctive or indicative) conditionals and $\supset$ for the material conditional).

[2]I use ✓/✗ to indicate judgments about the acceptability of inferences. I use ?? to indicate markedness.

In contrast, the same inference patterns are standardly taken to be invalid for subjunctives:

(6) ✗ Ada is drinking red wine. (So) if she were eating fish, she'd be eating fish and drinking red wine.

(7) ✗ Claude is either in London or Paris. (So) if he weren't in London, he'd be in Paris.

(8) ✗ If Lori were married to Kyle, then if she were married to Lyle, she'd be a bigamist. She's married to Lyle. (So) if she were married to Kyle, she'd be a bigamist.

Counter-instances to (6)-(8) are easily identified. For example, circumstances in which Ada is drinking red wine and eating beef, but would be drinking white wine were she eating fish will constitute counter-instances to (6). Circumstances in which Claude is in London, but might be in Rome were he not will constitute counter-instances to (7). And circumstances in which Lori is married to Lyle, but would not be, were she to be married to Kyle will constitute counter-instances to (8).

Notably, however, embedding the rightmost premise under '*suppose*' leads each subjunctive inference pattern to improve considerably:

(9) ✓ Suppose Ada were drinking red wine. (Then) if she were eating fish, she'd be eating fish and drinking red wine.

(10) ✓ Suppose Claude were in London or Paris. (Then) if he weren't in London, he'd be in Paris.

(11) ✓ If Lori were married to Kyle, then if she were married to Lyle, she'd be a bigamist. Suppose she were married to Lyle. (Then) if she were married to Kyle, she'd be a bigamist.

That is, where the non-conditional premise is supposed — rather than asserted — subjunctive instances of (PR), (DA) and (CT) appear valid. Two brief observations are in order.

First, note that following supposition, the discourse particle '*then*' is preferred to '*so*' to indicate that one utterance stands in a consequence relation to another. In this respect, supposition patterns with '*if*'-clauses, which likewise license '*then*' (and not '*so*') in the matrix clause. Second, in addition to being embedded under suppose, the non-conditional premise occurs with an additional layer of past tense morphology in each of (9)-(11). We will return to the interaction between morphological marking and the entailment patterns in further detail in §2.6.1.

Assuming '*if*'-clauses behave like supposition, we would expect the two to have similar effects on our judgments about the validity of each inference pattern. To test this, we need to consider their deduction theorem equivalents:

(PR*)		$\models \phi \Rightarrow (\psi \Rightarrow (\phi \wedge \psi))$	CONDITIONAL PRESERVATION
(DA*)		$\models (\phi \vee \psi) \Rightarrow (\neg\phi \Rightarrow \psi)$	CONDITIONAL DIRECT ARGUMENT
(CT*)	$\phi \Rightarrow (\psi \Rightarrow \chi)$	$\models \psi \Rightarrow (\phi \Rightarrow \chi)$	CONDITIONAL CONDITIONAL TELESCOPING

(PR*), (DA*) and (CT*) are derived from the three original inference patterns by embedding the conclusion under a wide-scope '*if*'-clause containing the right-most premise. Put another way, they are each the result of an application of the Deduction Theorem to one of our starting patterns:

(DT) If $\Gamma, \phi \models \psi$, then $\Gamma \models \phi \Rightarrow \psi$. DEDUCTION THEOREM

Indicative instances of these three inference patterns are standardly taken to be valid. Indeed, the validity of the indicative instances follows from the acceptance, for indicatives, of (PR), (DA), (CT) and the DEDUCTION THEOREM.

If supposition and '*if*'-clauses behave alike, we would expect subjunctive instances of (PR*), (DA*), and (CT*) to be valid. Embedding the rightmost premise within an '*if*'-clause would have the same effect as embedding it under supposition. As (12)-(14) demonstrate, this prediction appears borne out.

(12) ✓ If Ada were drinking red wine, then if she were eating fish, she'd be eating fish and drinking red wine.

(13) ✓ If Claude were in London or Paris, then if he weren't in London he'd be in Paris.

(14) ✓ If Lori were married to Kyle, then if she were married to Lyle, she'd be a bigamist. (So) If Lori were married to Lyle, then if she were married to Kyle, she'd be a bigamist.

Summarizing, supposition and '*if*'-clauses have a similar effect on the acceptability of our three inference patterns. In their basic form, all three appear invalid for subjunctives. However, they improve substantially when the leftmost premise is embedded either (*i*) under supposition or (*ii*) in the antecedent of a subjunctive in which the conclusion is nested. This is provisional evidence that the relationship between the two extends beyond the similarity of constructions like (1) and (2).

1.2.2 Counterfactual Use

It is well known that subjunctives, unlike indicatives, support counterfactual uses (Lewis (1973), Stalnaker (1975), von Fintel (1999)). Indicatives are unacceptable in discourse contexts in which the negation of their antecedent is accepted. Yet subjunctives are typically fine in such environments. Thus, while (16) constitutes an acceptable bit of discourse, (15) does not.[3]

(15) The butler didn't do it. ??If he did it, he used the candlestick.

(16) The butler didn't do it. If he'd done it, he would've used the candlestick.

What has received less discussion is that this behavior disappears under supposition and in conditional consequents. That is, the discourse in (15) degrades substantially if the material in the first sentence is supposed, rather than asserted (as in (17)). And, it similarly degrades if it is embedded in a wide-scope '*if*'-clause (as in (18)).

(17) Suppose that the butler hadn't done it. ??If he'd done it, he would've used the candlestick.

[3]It is worth noting the existence of so-called 'sports-caster' conditionals, which lack the standard morphological marking associated with subjunctives yet permit counterfactual use. For example, having just watched an optimistic long-range effort sail over the bar, a soccer commentator can employ the indicative (‡):

(‡) If he passes, they score for sure.

Here, the negation of the antecedent is mutually accepted in the context. Thus the commentator's use of the conditional is counterfactual. Furthermore, in the context, it appears to receive an interpretation equivalent to the corresponding subjunctive 'If he had passed, they would have scored for sure.

I will argue, in Chapters 2-3, that the unavailability of counterfactual uses is the result of a presupposition of indicatives. Absent this pressupposition, the two types of conditional are equivalent. This suggests the interesting possibility that sports-caster conditionals like (‡) are the result of a form of presupposition cancellation Karttunen (1971), Abusch (2010), Abrusán (2016).

Insofar as such conditionals are *infelicitous* in contexts which entail their antecedent, I suggest that they should be identified as subjunctives with non-standard morphological marking, rather than indicatives permitting counterfactual use (see Chapter 2, §2.4 for further details).

(18) ??If the butler hadn't done it, then if he'd done it, then he would've used the candlestick.

Generalizing, when the antecedent of a subjunctive is inconsistent with (*i*) an earlier supposition, or (*ii*) the antecedent of an embedding subjunctive, subjunctives pattern with indicatives in being incompatible with counterfactual uses.

1.2.3 Epistemic Contradictions

Sequences of the form $\neg\phi; \Diamond\phi$ and $\Diamond\phi; \neg\phi$ give rise to infelicity.[4] Call such sequences Epistemic Contradictions. When their elements are conjoined, epistemic contradictions pattern with unembedded Moore sentences.

(19) ??It isn't raining but I don't know that.

(20) ??It isn't raining but it might be.

(19) and (20) are equally infelicitous. However, as noted by Yalcin (2007), the two pattern differently when embedded under supposition (as in (21)) and in the antecedents of subjunctives (as in (22)):

(21) Suppose it isn't raining but [I don't know that/??it might be].

(22) If it weren't raining but [I didn't know that/??it might've been], I'd have needlessly brought an umbrella .

By itself, this contrast constitutes a noteworthy similarity between supposition and '*if*'-clauses. Whereas both have an ameliorative effect on the felicity of Moore sentences within their scope, they inherit the infelicity of embedded epistemic contradictions.

However, this behavior generalizes beyond conjunction. When we consider epistemic contradictions in more complicated environments, further similarities between supposition and '*if*'-clauses emerge. For example, consider the discourses in (23.a-b), in which the first element of the epistemic contradiction

[4] Though Groenendijk and Stokhof (1991) dispute the latter.

occurs as a separate sentence, with the second embedded in the antecedent of a subjunctive. Such discourses appear either acceptable, or, at least, notable improvements on the epistemic contradictions in (19), (21) and (22).

(23) a. It isn't raining. If it might've been, I'd have brought an umbrella needlessly.

b. It might be raining. If it weren't, I'd have brought an umbrella needlessly.

The same is not true when the material in the first sentence is supposed, rather than asserted. The discourses in (24.a-b) is notably degraded in a way that (23.a-b) are not. Similarly where the material in the first sentence occurs within a wide-scope '*if*'-clause; (25.a-b) pattern with (24.a-b):

(24) a. Suppose that it weren't raining. ??If it might've been, I'd have brought an umbrella needlessly.

b. Suppose that it might've been raining. ??If it weren't, I'd have brought an umbrella needlessly.

(25) a. ?? If it weren't raining, then if it might've been, I'd have brought an umbrella needlessly.

b. ?? If it might've been raining, then if weren't, I'd have brought an umbrella needlessly.[5]

Summarizing, both supposition and '*if*'-clauses appear to have a similar effect on the acceptability of epistemic contradictions which occur partially within a subjunctive conditional. This pattern of behavior is somewhat surprising. On standard semantic accounts of epistemic contradicitons, (e.g., Veltman (1996), Gillies (2001), Yalcin (2007)) whereas the unacceptability of the first member of each pair is to be expected, the second members of each pair are predicted to be fine. Yet no such contrast appears to be present.[6]

[5]Note that (25.a) is the Import/Export equivalent of (22).

[6]It is worthwhile noting that discourses like (24) and (25) may improve somewhat with the addition of '*neverthless*' in the embedded RH-conjunct of the epistemic contradiction:

(25′) Suppose that it weren't raining. If it nevertheless might've been, I'd have brought an umbrella needlessly.

Notice, however, that the same effect (to the extent it exists) can be observed in straight epistemic contradictions under supposition and in subjunctive antecedents:

(21′) Suppose it isn't raining, but it nevertheless might be.

1.2.4 Evidential Weakness

Finally, the requirements governing the introduction of supposition and '*if*'-clauses appear similarly weak. It is notably marked to assert $\Diamond\phi$ and ϕ sequentially.

(26) ?? Maybe the butler did it. He did. So the gardener is innocent.

Assuming there is no sudden revelation between the first two utterances, (26) does not constitute an acceptable bit of discourse for a single speaker to engage in. It is hardly mysterious why. Someone who claims that the butler might have done it thereby implicates that they are not in a position to make the unqualified claim that he did it. Thus, absent acquisition of new information, sequential assertion of the first and second sentence is difficult to reconcile with the assumption of speaker co-operativity (Grice (1967)). A speaker who satisfies the requirements of quality on the latter utterance will violate the requirements of quantity on the former, and *vice versa*.

The same is not true when the same material is supposed, or occurs embedded within an '*if*'-clause.

(27) Maybe the butler did it. Suppose he did. Then the gardener is innocent.

(28) Maybe the butler did it. If he did, then the gardener is innocent.

Both (27) and (28) are perfectly acceptable bits of discourse. Summarizing, the evidential requirements on both supposition and conditionals appear weak when compared to assertion. Unlike assertion, someone who introduces a supposition does not thereby represent themselves as being in a strong evidential position regarding the material supposed. Equally, someone who employs a conditional need not be have any particular epistemic attitude towards its antecedent.[7]

(22′) If it isn't raining, but it nevertheless might be, then I'd have needlessly brought an umbrella.

[7] As noted we noted above, indicatives are generally infelicitous where there antecedent is ruled out. However, it is important to distinguish what is ruled out by the discourse context from what is ruled out by individual interlocutors (Stalnaker (2014)). As long as this distinction is enforced, we have no reason to exclude the possibility that a speaker may utter an indicative despite knowing its antecedent to be false.

1.2.5 Summary

We have identified a range of further connections between supposition and '*if*'-clauses. Both (*i*) have a common effect on a variety of inference patterns involving subjunctive conditionals; (*ii*) block the availability of counterfactual uses of subjunctives; (*iii*) interact in similar ways with epistemic contradictions; and (*iv*) impose weaker requirements than assertion on the evidential state of the speaker. These observations are significant in at least two ways: first, they provide substantial additional support for giving an account of conditionals in terms of supposition, beyond the basic similarity between sentences like (1) and (2). Second, they serve as a guide to developing such a theory, providing important evidence both about the behavior of supposition and about how it is connected to the behavior of conditionals.

In §1.3, we will look more generally at the properties exhibited by supposition in discourse. In §1.4, we will turn to consider previous suppositional theories, assessing both their merits and their potential shortcomings. Finally, in §1.5, I will sketch an outline of alternative kind of theory, to be developed in more detail in Chapter 2.

1.3 Supposition

It is important to distinguish between the mental act of supposing and the public act of introducing a supposition in a discourse.[8] To this end, we will use 'cognitive supposition' for the former and 'verbal supposition' for the latter. An agent can engage in cognitive supposition without making a corresponding verbal supposition. And equally (*pace* Green (2000)), an agent can make a verbal supposition without engaging in being in the cognitive state.

[8]In fact, it may be necessary to distinguish between multiple mental states which are picked out by the same verb. As Green (2000) notes, there is is an attitude ascribed by sentences like (†.a) which appears to be interchangeable with belief and which requires the subject to think the content likely. This appears importantly different from the attitude ascribed by ((†.b), which is not interchangeable with belief and which does not require the subject to ascribe any likelihood to the content.

(†) a. Having considered the alternatives, Olga supposed that death was inevitable.
b. To establish her result, Maya first supposed it's negation.

It is the latter for which I reserve the label 'cognitive supposition'. I will return to this distinction in greater detail in §2.6.2

Our primary focus (both in this chapter and throughout) will be on verbal supposition. Accordingly, where context suffices to establish the intended sense, I will continue refer to it simply as supposition. We will briefly return to the relationship between the mental act and the public contribution to a discourse towards the end of Chapter 2.

§§1.2.1-4 surveyed a range of properties of verbal supposition, ones which it shares with '*if*'-clauses. In this section we will further individuate the target phenomenon, identifying a number of other properties a theory should account for and looking at how it interacts with more general features of discourse structure.

1.3.1 The Structure of Supposition

Productivity: Supposition is productive, insofar as information supposed can be employed to draw further inferences. If ϕ entails ψ, then after supposing ϕ, it is acceptable to conclude that ψ. Inferences like the one in (29.a) appear generally safe. As a limiting case, introducing a supposition establishes its own content; supposing ϕ results in a discourse context in which ϕ is accepted.

(29) a. ✓Suppose that Philippe were in Paris. Then he'd be in France.

b. Suppose that John's parents have two sons. Then he might invite his brother to the party.

As a corollary, material introduced via supposition can satisfy the presuppositions of downstream utterances. The first sentence of the discourse of (29.b) suffices to ensure that the existence and uniqueness presuppositions of the the definite in the second are fulfilled. This is as we would expect if presupposition satisfaction is correctly understood in terms of requirements on what is accepted in a context.

Accumulativity: Suppositions accumulate through a discourse. After supposing ϕ and ψ sequentially, it is acceptable to conclude that $\phi \wedge \psi$. That is, inferences like the one in (30) are generally safe.

(30) ✓ Suppose the Mets score one run in the first inning. Suppose that they score two in the second. Then they'd have scored three runs in their first two innings.

More generally, the effects of supposition are persistent through further contributions to the discourse. After a supposition is introduced, its content remains accepted until the supposition is retracted, regardless of what other material is asserted or supposed.

Regularity: Changes to what information is supposed bring about changes in the context that results in a regular manner. If, after supposing ϕ, it is acceptable to conclude both ψ and χ, then after supposing $\phi \wedge \psi$ it is acceptable to conclude χ. If it is safe to draw the inference in (31.a), then it appears safe to draw the inference in (31.b).

(31) a. Suppose that Berlin had fallen to the Russians. Then Vienna and Hamburg would have too.

b. Suppose that Berlin and Vienna had fallen to the Russians. Then Hamburg would have too.

Regularity imposes a constraint on how sensitive what is accepted in a discourse can be to what is supposed. If a supposition is sufficient to get the discourse context into a particular state, then strengthening that supposition with some material accepted in that state will not change its effect on the discourse context.

Conservativity: The changes brought about via supposition are conservative. If ϕ is accepted, but ψ is not, then after supposing $\neg\psi$ it is acceptable to conclude $\phi \wedge \neg\psi$. Inferences like the one in discourse (32) are generally safe.

(32) ✓ Philippe is in Paris. Maybe Marie is too... Suppose that she were. Then Philippe and Marie would both be in Paris.

Conservativity imposes a constraint on how much introducing a supposition in a discourse can change what is accepted. If the material being supposed is compatible with what is already accepted, then anything accepted prior to the supposition will continue to be accepted after it.

1.3.2 Retraction

Before proceeding, it is important to highlight a central feature of supposition which the present discussion fails to address. The effects of a supposition are persistent; they endure beyond its syntactic scope. However, they are not irreversible.

The discourse in (33.a) can be felicitously extended with (33.b). For this to occur, the former supposition's effect need to be retracted: in the discourse context resulting from the latter supposition, the material introduced by the former must no longer be accepted. Each of the patterns observed in the previous section holds only under the assumption that no retraction occurs.

(33) a. Suppose that it's raining. Then the park will be wet.

b. …But suppose that it isn't. Then the park will be dry.

b′. …Suppose that we go for a picnic. Then we'll be miserable.

That (33.a-b) are presented as contrasting appears important in triggering retraction. Where (33.a) is followed by (33.b′) instead, no retraction is triggered; it is accepted that we will be miserable in the posterior context precisely because we preserve both the supposition that we will go for a picnic and the supposition that it is raining. Crucially, the relevant difference here appears to have more to do with how the two utterances cohere than simply whether the two suppositions are in tension. Even if the initial supposition that it is raining results in a context in which it is accepted that we won't go for a picnic, proceeding to suppose that we will go for one does not suffice to trigger retraction. Rather, it will result in a new context in which it is accepted both that it is raining and that we go for a picnic.

In order to properly understand a discourse, interlocutors need to know not just what is being said, but how what is being said is related to what has been said before. A number of authors have argued that these kinds of relationships between utterances also play an essential role in determining how the discourse develops over time. That is, the properties of the discourse context in which an utterance ends up being evaluated depend not just on the sequence of utterances which preceded it, but also how those utterances are structured (Hobbs (1985), Roberts (1996, 2012), Kehler (2002), Asher and Lascarides (2003)).

For example, the domain of quantification of a modal expression depends, in part, on the preceding discourse (Kratzer (1977, 1979)). In cases of so-called modal subordination, the prejacent of one modal anaphorically supplies a restriction on the domain of later modals (Roberts (1989)). Asher and McCready (2007) and Stojnić (2017, 2019) have argued that, in addition to establishing a coherent discourse, the relationship between utterances has an important role to play in determining the availability of modal subordination.

(34) a. A wolf might walk in. It would eat you.

b. A wolf might walk in. But it probably won't.

The discourses (34.a-b) differ regarding what restrictions are placed on the domains of the modals in their second sentence. In the former, the modal is subordinated to the prejacent of the preceding sentence; it quantifies only over scenarios in which a wolf enters. In the latter, by contrast, there is no subordination,(on pain of incoherence). The modal quantifies unrestrictedly over both wolf-entrance and wolf-absent scenarios.

According to Asher & McCready and Stojnić, the availability of subordination co-varies with the discourse relations between utterances. Where one utterance elaborates on another (as, they claim, occurs in (34.a)), modals in the former can be anaphorically restricted by material in the latter. Where two utterances stand in contrast, however, (as, they claim, occurs in (34.b)), subordination will be unavailable.

This proposal suggests a tentative but somewhat attractive hypothesis: as with subordination, so with supposition. In the discourse formed of (33.a-b), the second pair of sentences is presented as a contrast with (33.a). The same is not true of (33.a-b′). Instead, (33.b′) supplies additional comment on the situation described by the (33.a), providing an elaboration. Perhaps the discourse relations between utterances have an important role to play in determining whether a supposition remains in force.[9]

On a view of this kind, the introduction of a supposition occurs in a very different way to its retraction. Whereas the former is overtly marked by a single utterance (i.e., of the form ⌜*Suppose that* ϕ ⌝), the latter depends on global features of the discourse to do with how utterances are related to one another. We will return to the relationship between discourse structure and retraction again in the discussion of mood in §2.6.1.

1.4 Suppositional Theories: A Brief Overview

A theory of conditionals which makes essential appeal to supposition has been defended (in various, closely related, forms) by a number of authors, including Mackie (1972), Edgington (1986, 1995), Barker (1995), DeRose and Grandy (1999), and Barnett (2006). What is common to variants of the theory is a commitment to the claim that, in uttering a sentence of the form ⌜*if* ϕ, ψ⌝ (where ϕ and ψ are clauses with declarative mood) an agent performs a speech act equivalent (in some respect) to sequentially: (*i*) supposing that ϕ, and (*ii*) asserting that ψ. Interpreting such theories charitably, we should also assume that the speech act incorporates the effect of 'undoing' the supposition that ϕ, in order to ensure that the effect is limited to the syntactic scope of the '*if*'-clause.

Call this the Speech-Act Suppositional Theory. Below are formulations of the theory by three of its proponents:

"[The Speech Act Suppositional Theory] explains conditionals in terms of what would

[9] Note that, even if they are similarly sensitive to discourse structure, the prospects for explaining the full range of phenomena associated with supposition as a special case of subordination are not good. Like supposition, subordination involves an inter-sentential restriction on some collection of possibilities. However, whereas supposition imposes a restriction on the possibilities under consideration in the context (i.e., information state of the context), subordination is strictly a restriction on the domain of modal expressions. More importantly, as we saw in §1.2.1, unlike subordination, supposition does not just modify the possibilities under consideration. It also changes the effect of future modifications to that collection of possibilities.

> probably be classified as a complex illocutionary speech act, the framing of a supposition and putting something forward within its scope." (Mackie (1972, 100))

> "To assert or believe ⌜*if* ϕ, ψ⌝ is to assert (believe) ψ within the scope of the supposition, or assumption, that ϕ." (Edgington (1986, 5))

> "The pragmatic theory of '*if*' states that utterance of ⌜*if* ϕ, ψ⌝ is an assertion of ψ grounded on supposition of ϕ where [the speaker] implicates via the presence of '*if* ϕ' that their assertion of ψ is so grounded." (Barker (1995, 188))

Neither Mackie or Edgington provides a non-metaphorical gloss of their talk of one speech act occurring within the scope of another. However, the intended position appears relatively clear. Performing a speech act of supposing that ϕ results in a new discourse context (which differs from the discourse context which would result from asserting that ϕ). An assertion of ψ in this new context may differ (in its felicity, its illocutionary effects, etc.) from an assertion of ψ in the prior context. The Speech-Act Suppositional Theory says that the primary conventionally determined contribution of an '*if*'-clause is not to determine a particular content in combination with the matrix clause. Rather, it is to indicate that the speaker is performing a particular speech act, equivalent to asserting the consequent in the discourse context created by supposing the antecedent.

Speech-Act Suppositional Theories vary in their explanation of how '*if*'-clauses come to conventionally indicate the speech act performed by the sentence in which they occur. For example, Barker (1995) suggests that the clause introduces a conventional implicature. That is, that a complex speech act of the appropriate type is being performed is part of the not-at-issue content of a conditional. Alternatively, Barnett (2006) suggests that '*if*'-clauses express that the speaker is in a particular kind of cognitive state, namely, one of supposing the clause's content. Declarative consequents, correspondingly, expresses an attitude of belief. It will, presumably, need to be a corollary of this view that the discourse effect of a sequence like (1) is exhausted by what its component utterances convey about the attitudes of the speaker. While they diverge on important issues, the differences between these versions of the theory can be largely set aside for present purposes.

The Speech-Act Suppositional theory has a number of appealing features. Most obviously, it accounts for the apparent interchangeability of (1) and (2). To utter (2) just is, according to the theory, to perform a speech act equivalent to the speech acts performed by an utterance of the two sentences in (1).

It also does well accounting for the full range of clause types which can occur in conditional consequents. Conditionals happily embed both interrogatives (as in (35)) and imperatives (as in (36)):

(35) If the C train is running on the express track, will it stop at Columbus Circle?

(36) If the C train is running on the express track, change at 42nd St!

Many standard accounts of conditionals have difficulty compositionally accounting for the meaning such constructions. As Edgington (1995), Barker (1995) and Barnett (2006) note, however, Speech-Act Suppositional Theories seem comparatively well positioned. Just like assertions, both questions and commands can occur in the kind of discourse context that results from a supposition. Accordingly, on the assumption that clause-types encode (or at least constrain) the illocutionary acts performed in uttering them (Sadock and Zwicky (1985), Stenius (1967), Portner (1997)), the Speech Act theory can extend their account of conditionals with declarative consequents to those with interrogatives and imperatives. That is, (35) and (36) will have the discourse effect of asking a question and issuing a command, respectively, in the context that results from supposing the C train is running express.[10]

However, the theory also faces substantial, well-known challenges. I will consider two of the most

[10] As Edgington puts it:

> "Any kind of speech act can be performed unconditionally or conditionally. There are conditional questions, commands, promises, agreements, offers, etc., as well as conditional assertions. [...] It is overwhelmingly plausible that the clause, 'if he phones', does the same job in conditional statements, commands, questions, promises, expressions of wish, etc., and hence that a theory of conditionals should be applicable to more than conditional statements." (Edgington (1995, 177))

And, similarly, DeRose & Grandy:

> "[O]ne naturally wonders why one couldn't conditionally perform speech acts other than assertion by means of '*if*' sentences with indicative '*if*' clauses, but where the other clause, taken by itself, would execute the other speech act in question. And, when we look to natural language, we find constructions which seem to do just that. [...] That we can issue conditional warnings, questions, and commands in such a way is unsurprising given our view. " (DeRose and Grandy (1999, 410))

significant ones in turn.

1.4.1 Embeddability

Conditionals can occur felicitously in a variety of sub-sentential environments (see Kolbel (2000), in particular, for a succinct overview). Amongst other examples, they can be conjoined (e.g., (37)), disjoined (e.g., (38)), embedded under negation (e.g., (39)) and in conditional consequents (e.g., (40)):

(37) If Lea rolls a six, she'll win, and if Lea wins, Caroline will lose.

(38) Either if Lea rolls six, she'll win or if Lea rolls a six, she'll tie.

(39) It isn't the case that if Lea rolls a six, she'll win.

(40) If Caroline rolls a five, then Lea will win if she rolls a six.

This behavior generates a problem for the Speech-Act Suppositional Theory. It is typically assumed that speech acts cannot be performed using a clause in an embedded environment (Austin (1962), Searle and Vanderveken (1985), Murray and Starr (2018, Forthcoming); though cf. Krifka (2001, 2004)). The proponent of the theory must, accordingly, provide an account of (*i*) what the conditional contributes to the content of a sentence when it occurs in an embedded environment, and (*ii*) what the role of the '*if*'-clause is in such environments, if not to mark a speech act.

Theorists have tended to adopt one of three different kinds of strategies to address this concern, which we can broadly categorize as denial, coercion and narrow-scoping. None appear fully satisfactory.

Denial: Some speech act suppositional theorists simply deny that sentences in which a conditional occurs in a certain embedded environment are linguistically possible. For example, Edgington (1995, 172) claims that disjunctions of conditionals are 'virtually uninstantiated' and ungrammatical (cf. Appiah (1985)). Similarly, Barnett (2006, 548), while not denying that such construction are grammatical, tells us that they are 'not meaningful statements of anything at all'.

The problem with this kind of response is that it is wrong. (38) is a perfectly grammatical sentence of English. And it is equally clear what it means: it means that either if Lea rolls a six, she'll win, or if she rolls a six, she'll tie. It should be accepted by someone who accepts that one of the two conditionals is true, and rejected by anyone who accepts that both are false.

Coercion: Mackie (1972, 103) opts for split view, on which '*if*' functions as a speech act markers when unembedded, but denotes the material conditional in (certain) embedded environments. Rather than positing ambiguity, he suggests that when it is embedded, its meaning is coerced into that of a truth-functional operator.

Aside from being *ad hoc*, Mackie's proposal also faces more significant challenges. In particular, it is unclear how it can account for the acceptability of certain inferences. The inference from (38) and (39) to the conclusion that if Lea rolls a six, she'll tie appears impeccable. Yet it is hard to see how Mackie could account for this. It is integral to the proposal that accepting ψ under the supposition that ϕ is not equivalent to accepting the material conditional $\phi \supset \psi$. But accepting the disjunction of two material conditionals along with the negation of the first, need not lead one to accept the consequent of the second under the supposition of its antecedent. Similarly, it appears that someone who accepts (40) and accepts that Caroline will roll a five should accept that Lea will win if she rolls a six. But on Mackie's view, the latter, unembedded conditional does not follow from the material conditional it would be coerced into expressing when embedded. In general, taking '*if*'-clauses to have a different semantic function depending on their environment threatens our ability to account for inferences in which a conditional occurs embedded in the premises, but unembedded in the conclusion.

Narrow-scoping: A number of theorists have argued that apparent embeddings of conditionals are in fact no such thing. What appears at surface level to be a conditional scoping below another operator, should instead be treated as an instance of that operator applying to the consequent of the conditional.

Thus, for example, Edgington (1995, 173) and Barnett (2006, 546) both suggest that negations of conditionals should be interpreted as conditionals with negated consequents. Setting aside questions of how the negation comes to be interpreted under the scope of the antecedent, this strategy yields counter-intuitive predictions. Not winning is tieing or losing. So (39) is predicted to be equivalent to (41):

(41) If Lea rolls a six, she'll either tie or lose.

But this seems wrong. Assuming that Lea wins iff she rolls a higher number than any other player, and no player has yet rolled, we should accept (39) (afterall, someone else might also roll a six). However, we should deny (41) (afterall, they might not).

1.4.2 Validity

A minimally adequate theory of conditionals ought to be able to be supplemented with a notion of validity in order to generate predictions about how conditionals interact with other logical vocabulary. The status of inference patterns relating conditionals and expressions such as negation, disjunction and conjunction are amongst the core subject matter of the study of conditionals. Any satisfactory theory should, at least in principle, be capable of adjudicating questions of these kinds.

The problem is that, under the Speech-Act Suppositional Theory, the expressions belong to fundamentally different semantic categories. On the theory's standard version, negation, disjunction, conjunction, etc. are assigned their classical, truth-functional meaning. Accordingly, their logical properties are to be understood in terms of their effect on a sentence's truth-conditions. In contrast, any logical properties ascribed to the conditional will arise from its effect on the speech acts which can be performed by an utterance of a sentence, rather than on the truth-conditions of that sentence. Indeed, on the standard version of the theory, sentences with a conditional at widest scope cannot be attributed truth conditions at all.

Accordingly, any notion of validity capable of making predictions about the interaction of conditionals with other logical vocabulary needs to be formulated as a relation between (classes of) speech acts. It is not at all clear what such a notion of entailment will look like. In particular, since utterances of conditionals lack truth conditions on the Speech-Act Suppositional Theory, it cannot amount to mere relation of truth preservation between the assertion of the premises and the assertion of the conclusion.

Response to the second problem have typically taken one of two forms: they offer either a probabilistic account of validity, or one which appeals to the (non-quantitative) cognitive attitudes of agents. We can briefly consider the prospecticts of each in turn.

Probabilistic validity: Edgington (1995), following Adams (1965, 1970, 1975), suggests we classify of an argument as valid iff the sum of the uncertainty of its premises is at least as great as the uncertainty of its conclusion (where the uncertainty of ϕ is 1 minus the probability of ϕ). Supplementing this with the proposal that indicative conditionals be ascribed a probability equal to the conditional probability of their consequent on their antecedent, this gives us a response to the concern about how to evaluate arguments.

One point to note is that on the proposal, both (PR) and (DA) will come out as invalid for indicatives. An agent's uncertainty about $\phi \wedge \psi$ conditional on ψ can exceed her unconditional uncertainty about ϕ. And her uncertainty about ψ conditional on ϕ can exceed her uncertainty about the material conditional.

This issue is not as significant as it might seem. In arguing that the inference patterns were valid for indicatives, we tacitly limited ourselves to considering cases where the premises were certain. And the probabilistic proposal, as Edgington (1995) notes, can capture this: in the limiting case where the premises recieve probability 1, the conclusions do to. If we consider instead whether having high confidence in the premises requires having high confidence in the conclusion, the inference patterns no longer appear good.

A more substantial question for the probabilistic proposal is what motivates assigning the conditional a probability equal to the conditional probability of its consequent given its antecedent in the first place. Since '*if*'-clauses are treated as markers of speech acts, the conditional's probability cannot be identified with the probability of some conditional content it expresses.[11] So on what basis do we assign it a probability?

Adams (1965) and Edgington (1995) have a response. What matters, they suggest, is that arguments preserve assertability. And what is a claim's assertability? For non-conditional claims, according to them, its degrees of assertability is simply its probability. For conditional claims, it is the conditional probability of the consequent given the antecedent. So the requirement that arguments preserve assertability simply collapses into the probabilistic notion of validity.[12]

Regardless of what determines the assertability of conditionals, there are reasons to think that Adams and Edgington's claim that assertability goes by probability cannot be right in general. It is never appropriate to assert that one has lost the lottery on the basis of the probability alone, no matter how great that probability is (DeRose (1996), Williamson (2000), Hawthorne (2003)). Yet there are plenty of claims whose probability falls short certainty, and yet which meet the standard for assertability. If assertability (for non-conditional claims) were proportional to probability, we would be at a loss to reconcile these two observations (cf. Dudman (1992, 204-207), DeRose (2010, fn11), Carter (ms, §4.2)).[13]

A second worry for the probabilistic proposal is that it does not readily extend to subjunctive conditionals. As we noted above, subjunctives permit use in counterfactual discourse contexts—contexts in which their antecedent has probability 0. Yet, if it is determined according to bayes theorem, in such contexts the conditional probability of the consequent on the antecedent will be undefined.

Alternative probability measures are available, which allow for definedness in cases where the antecedent has probability 0 by taking conditional probabilities as primitive (Rényi (1955), Popper

[11] Indeed, given the various triviality results in the vicinity (Lewis (1976), Bradley (2000), Hájek (2011), Fitelson (2015), Russell and Hawthorne (2016)), this is treated as a significant advantage of the theory by its proponents.

[12] Similar claims can be found in Lewis (1976) and DeRose and Grandy (1999)

[13] On, that is, the seemingly trivial assumption that whether something is assertable *simpliciter* is dependent on its degree of assertability.

(1959)). Nevertheless there remains a question of how to constrain such measure so that the identification of the probability of a subjunctive with the conditional probability of the consequent on the antecedent yields the correct result (Edgington (1995)).

Cognitive validity: An alternative entertained by Barnett (2006) is to offer a notion of validity in terms of the attitudes of rational agents. That is, he suggests, 'an argument is valid just in case whoever accepts its premises is rationally required to accept its conclusion' (551). This needs to be supplemented with an account of what it is to accept a conditional, since on the speech act theory they are not the kinds of things which agents bear attitudes to directly. Accordingly, Barnett proposes that an agent accepts a claim (conditional or non-conditional) iff she would accept the content of anything it states if she were supposing the content of any supposition it introduces. Since Barnett takes conditionals to state their consequents while introducing their antecedent as a supposition, this gives the result that an agent accepts that if the butler did it, then he used the candlestick, iff, were she to suppose that the butler did it, she would accept that he used the candlestick.

One worry is that the notion of acceptance is unappealingly *ad hoc*. In order to get a plausible result, Barnett must stipulate an account of acceptance to match up directly with the various effects of the speech act expressed by a conditional on his view. As a result, it doesn't deal well with acceptance conditions of complex sentences involving conditionals. For example, both conjunctions like (37) and disjunctions like (38) introduce two suppositions and make two statements. Yet in neither case would a rational agent accept the sentence only if they accepted the content of both statements under both suppositions.

A second worry is that, by tying validity to rationality, the account will inevitably give up too many inferences. For example, preserving conjunction introduction might be thought a non-negotiable condition on a non-standard account of validity. Yet a lesson of the lottery (Kyburg (1961)) and the preface (Makinson (1965)) would seem to be that a rational agent can accept each of a set of claims without accepting their conjunction.

1.5 Summary

Speech-act based theories offer a simple picture of the relationship between supposition and '*if*'-clauses. At a level of abstraction, the two simply offer different ways of bringing about the same kinds of changes to a discourse. However, as we have seen, such theories also face significant challenges. While responses to these challenges are available, questions about their success suggest it is worthwhile investigating other options.

We will start by attempting to better characterize the way that (verbal) supposition differs from bare assertion. Here, the pattern of behavior considered in §1.2 is particularly suggestive. Consider the contrast between (6)-(8) and (9)-(11), repeated below:

(6) ✗ Ada is drinking red wine. (So) if she were eating fish, she'd be eating fish and drinking red wine.

(7) ✗ Claude is either in London or Paris. (So) if he weren't in London, he'd be in Paris.

(8) ✗ If Lori were married to Kyle, then if she were married to Lyle, she'd be a bigamist. She's married to Lyle. (So) if she were married to Kyle, she'd be a bigamist.

(9) ✓ Suppose Ada were drinking red wine. (Then) if she were eating fish, she'd be eating fish and drinking red wine.

(10) ✓ Suppose Claude were in London or Paris. (Then) if he weren't in London, he'd be in Paris.

(11) ✓ If Lori were married to Kyle, then if she were married to Lyle, she'd be a bigamist. Suppose she were married to Lyle. (Then) if she were married to Kyle, she'd be a bigamist.

Intuitively, the suppositional variants differ in two respects: first, not everything accepted in the prior context needs to continue to be accepted in the context resulting from supposition. For example, if it was previously accepted that Ada is drinking white wine (and white wine only), supposing she is drinking red can lead us to give up that information. Second, after the supposition is introduced, the supposed information must be preserved when evaluating downstream subjunctives.

This offers the outline of an explanation of why the three inference patterns improve under supposition. For example, having supposed that Ada is drinking red wine, we enter a new context in which this information is accepted. This information is then required to be held fixed while evaluating the conditional in the conclusion. Hence, it is guaranteed that at any state that results from incorporating the information that Ada is eating fish into our context will continue to incorporate the information that she is drinking red wine. Our judgments about (10) and (11) can be accounted for in the same way. Information introduced by assertion, on the other hand, is not required to be preserved. Hence, (6)-(8) are not judged generally valid.

This rough picture generalizes to other kinds of example. Consider the contrast in the availability of counterfactual use in (16)-(17):

(16) The butler didn't do it. If he'd done it, he would've used the candlestick.

(17) Suppose that the butler hadn't done it. ??If he'd done it, he would've used the candlestick.

After a bare assertion that the butler is innocent, the subjunctive allows us to evaluate its consequent at some (minimally different) possibilities in which he was guilty. However, supposing that he is innocent requires us to hold the information that he didn't do it fixed. Accordingly, downstream subjunctives whose antecedents entail his guilt will be expected to impose conflicting constraints.

Likewise for epistemic contradictions. After the assertion that it isn't raining, the subjunctive in (23.a) allows us to evaluate its consequent at some (minimally different) possibilities at which it is accepted that he might have been.

(23) a. It isn't raining. If it might've been, I'd have brought an umbrella needlessly.

(24) a. Suppose that it weren't raining. ??If it might've been, I'd have brought an umbrella needlessly.

In contrast, if it is supposed that it is raining, instead (as it is in (24.a)), this information must be held fixed. Yet there are no possibilities at which it is both accepted that it is raining and accepted that it might not be. Hence, the latter discourse is predicted to be infelicitous.

The remainder of the paper develops a new suppositional theory of conditionals (both indicative and subjunctive) which implements this rough picture to account for the data. Informally, the idea is as that supposition and assertion correspond to different ways of changing a body of information. The effects of a supposition on a discourse context can be separated into two parts: (i) it induces a minimal revision to the possibilities under consideration, to incorporate the supposed information; and (ii) it modifies what will count as a minimal revision in the future, ensuring that the information supposed will be preserved.

'*if*'-clauses trigger the same kind of change as supposition at a local level. The primary difference is that the effects of the latter are restricted to their syntactic scope. In a slogan: supposition is discourse-level '*if*'; '*if*' is sentence-level supposition.

In the next chapter, I develop this sketch in a more formal setting. We will start by considering different notions of information change. From there, I offer a dynamic semantics for a language containing both a conditional and a supposition operator, and show how it accounts for the desiderata introduced in the preceding sections. The resulting theory preserves many of the philosophical advantages of the speech act theory, while avoiding challenges to do with embeddability and validity.

Chapter 2

Information Dynamics

2.1 Introduction

The previous chapter introduced two ideas. First, that assertion and supposition are associated with fundamentally different ways of modifying a body of information. Second, that the kind of modification triggered by '*if*'-clauses is the kind associated with supposition rather than the kind associated assertion. In this chapter, we will develop each of these ideas in a more rigorous setting.

We will start, in §2.2, by surveying different types of information change. In particular, we will look at two ways of incorporating new information into old: addition and revision. The former is familiar from theories of conversation dynamics (in particular, Veltman (1996)) while the latter is most closely associated with theories of belief dynamics (in particular, Alchourrón et al. (1985)). However, as we will see, these differences in traditional application are superficial. Both can be characterized in terms of different ways of moving through an information structure.

In §2.3, we will employ these two types of information change to give a dynamic semantics for a formal language containing both supposition and conditional operators. This semantics aims to implement the informal picture of the relationship between supposition and '*if*'-clauses with which the previous chapter concluded. In §2.4, we will extend the language from the previous section to model indicatives and subjunctives. I propose an austere account of the distinction, on which differences in their behavior are attributed to differences in their pressuppositions. In §2.5, this is shown to be sufficient to account for the phenomena discussed in Chapter 1, while avoiding issues with embeddability and validity. In §2.6, we will step back to consider philosophical aspects of view. While different interpretations of the formal system are available, I suggest that we can think of it as offering a different species of suppositional theory, one on which the connection between '*if*'-clauses and verbal supposition occurs at the level of the cognitive acts employed in processing them, rather than the speech acts performed by uttering them.

2.2 Information

An information state corresponds to a view on the world. Each information state rules some ways the world could be out and leaves other ways the world could be open. Anything that represents the world as being a particular way can be represented by an information state. Maps and minds, movies and manuscripts, myths and mime: all belong to kinds of things we can use information states model. We associate each kind of representational item with an information state which leaves a way the world could be open iff the item does not represents the world as not being that way.

Information states can be ordered. Where s and s' are information states, $s \leq s'$ iff every way the world could be which is ruled out by s' is also ruled out by s. In this case, we say that s is at least as strong as s' or that s incorporates s'.

An information structure is a collection of information states which exhaust the ways of representing the world.

Definition 1.

An information structure $\langle \mathbb{S}, \leq \rangle$ is a pair comprising a set of information states $\mathbb{S} = \{s, s', ...\}$ and partial order, $\leq$, which form a complemented distributed bounded lattice. That is:

- $\mathbb{S}$ has a least element, $\bot$, and greatest element, $\top$;

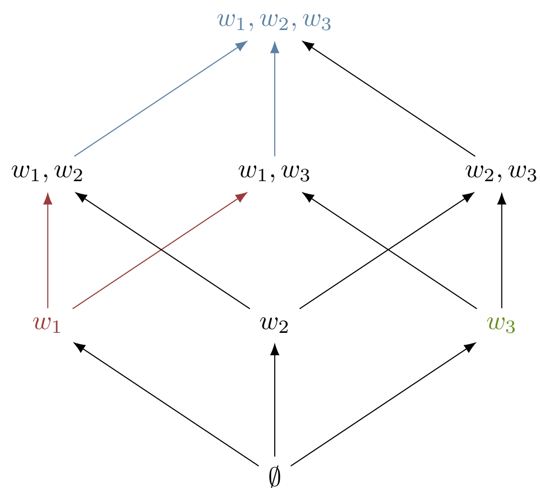

Figure 2.1: An information structure.

- For any $X \subseteq \mathbb{S}$, X has a least element, $\bigwedge(X)$, and greatest element, $\bigvee(X)$;
- For any $s, s', s'' \in \mathbb{S}$: $s \wedge (s' \vee s'') = (s \wedge s') \vee (s \wedge s'')$;
- Every $s \in \mathbb{S}$ has a unique complement, $\overline{s}$: a state such that $s \wedge \overline{s} = \bot$ and $s \vee \overline{s} = \top$.

$s \wedge s'$ is the meet of s and s': the greatest element of $\mathbb{S}$ which is at least as strong as s and at least as strong as s'. $s \vee s'$ is the join of s and s': the least element of $\mathbb{S}$ which no stronger than s and no stronger than s'.

A simple example of an information structure is the powerset algebra. Where $\mathcal{W}$ is a set of worlds, $\langle \mathcal{P}(\mathcal{W}), \subseteq \rangle$ is an information structure. Its greatest element is $\mathcal{W}$ and its least is $\emptyset$. Meet and join correspond to intersection and union over sets of worlds, respectively: that is, $s \wedge s' = s \cap s'$ and $s \vee s' = s \cup s'$. Complementation corresponds to complementation: that is, $\overline{s} = \mathcal{W}/s$. **Figure 2.1** depicts an powerset information structure based on the set of worlds $\{w_1, w_2, w_3\}$. The meet and join of $\{w_1, w_2\}$ with $\{w_1, w_3\}$ are designated in red and blue, respectively, and the complement of $\{w_1, w_2\}$ in green. In §2.3 onwards, we will focus on powerset information structures for ease of exposition. However, for the remainder of this section, we will continue to think about information structures in general terms.

2.2.1 Information Change

We can now ask how an information structure constrains the way we transition between information states. Or, alternatively put: what are the rules governing information change?

A transition rule is an operation which maps a pair of information states to a new information state. Each transition rule tells us if we start in one state, s, and receive the information represented by a second state, s', what state we should end up in. Thus, a transition rule can be thought of as describing a particular way to move through an information structure.

The addition operation, $+$, is the simplest transition rule we will consider. As a rule for discourse update in conversations, it is entertained by Veltman (1996) (for boolean languages) and by Stalnaker (1978) (generally).

Definition 2 (Addition).

$+$ is an addition operation iff $+$ satsfies $(+_1)$-$(+_2)$:

$(+_1)$ $s + s' \leq s$ and $s + s' \leq s'$

$(+_2)$ there is no s'' s.t. $s'' < s + s'$, $s'' \leq s$, and $s'' \leq s$

$s + s'$ is the weakest body of information which is at least as strong as both s and s'. It corresponds directly to the meet operation in the information structure: that is, $s+s' = s \wedge s'$. Accordingly, addition shares the same properties: it is commutative ($s+s' = s'+s$), associative ($s+(s'+s'') = (s+s')+s''$) and idempotent ($s + s = s$).

Addition is a maximally conservative way of modifying information. $s + s'$ retains any information incorporated by s and any information incorporated by s'. Adding new information never leads to information loss: it returns a state that is at least as strong as both the old state and the incoming state.

This characteristic of addition seems unequivocally desirable when the old information and the incoming information are compatible. Things are less clear, though, when the two conflict. In this case, addition takes one directly to the absurd state, $\perp$. Yet, this is not the only option available.

Sometimes acquiring incompatible information can lead us to give up part of our old information.

The revision operation, $*$, provides us with a transitions rule accommodating this idea. It is the information-theoretic analogue of the operation on closed sets of sentences defined in Alchourrón et al. (1985).[1]

Definition 3 (Revision).

$*$ is a revision operation iff $*$ satisfies $(*_1)$-$(*_4)$:

$(*_1)$	$s * s' = s \wedge s'$	*if* $s \wedge s' \neq \perp$;
$(*_2)$	$s * s' \leq s'$	
$(*_3)$	$\perp < s * s'$	*if* $s' \neq \perp$;
$(*_4)$	$s * (s' \wedge s'') = (s * s') \wedge s''$	*if* $(s * s') \wedge s'' \neq \perp$.

The first constraint says that revision should be minimal: as long as the incoming information is compatible with the old information, revision returns the weakest state which incorporates both. The second constraint says that revision should be successful: revising with some incoming information should return a state which is at least as strong as the information revised with. The third constraint says that revision should be consistent: revising with incoming information should not return an absurd state, as long as the incoming information is non-absurd. The fourth constraint says that revision should be orderly: revising with the meet of two states should return the meet of the second with the state that results from revising with the first, as long as the latter is non-absurd. Given these constraints, it is east to see that revision is neither commutative ($s * s' \neq s' * s$) nor associative ($s * (s' * s'') \neq (s * s') * s''$). However, it is idempotent ($s * s = s$).

Revision is often charactized as capturing a notion of minimal change (though cf. Rott (2000)). As long as the incoming information is compatible with the old information, it returns a state which is at least as strong as both. However, where the two conflict, revising with incoming information can lead one to give up old information.

We will sometimes want to consider operations which are less conservative than full revision. Say

[1] See Grove (1988) and Stalnaker (2009*a*) in particular for discussion.

that $*$ is a quasi-revision operation iff it satisfies ($*_2$)-($*_4$).Crucially, unlike revision operations, a quasi-revision operation may lead us to give up old information even when the incoming information is compatible with it.[2] Every revision operation is a quasi-revision operation, but not *vice versa*.

While importantly different, revision and addition are related. In fact, we can define revision in terms of addition and a second operation. The subtraction operation, $-$, provides a rule-governed way of giving up information. In this respect, it can be thought of as the opposite of addition.

Definition 4 (Subtraction).

$-$ is a subtraction operation iff $-$ satisfies ($-_1$)-($-_5$):

($-_1$)	$s - s' = s$	*if* $s \wedge s' \neq s$;
($-_2$)	$s \leq s - s'$	
($-_3$)	$s' < (s - s')$	*if* $s' \neq \top$;
($-_4$)	$s - (s \wedge s') \leq (s - s') \vee (s - s'')$	
($-_5$)	$s - s' \leq s - (s' \wedge s'')$	*if* $s - (s' \wedge s'') \not\leq s'$,

The first constraint says that subtraction should be minimal: as long as the old information is not at least as strong as the information to be subtracted, subtraction is idle. The second constraint says that subtraction should be conservative: subtracting some information should always return a state weaker than the old state. The third constraint says that subtraction should be successful: subtraction should return a state which strictly weaker the information subtracted, as long as that information is non-trivial. The fourth and fifth constraints say that subtraction should be orderly: subtracting the meet of two states should return a state at least as strong as subtracting as the join of subtracting with each individually. And subtracting the meet of two states should not return a state stronger than the one returned by subtracting either state by itself, unless that state is strictly weaker than the result of subtracting the meet.

[2]Quasi-Revision operations have recently received some interest in epistemology(Boylan and Schultheis (2019), Goodman and Salow (2018)). For example, suppose that on the basis of a quick glance, one judges a distant tree to be between 650cm and 750cm. Upon learning that the tree is less than 660cm high, it seems it would be reasonable to give up one's original judgment regarding the minimum height of the tree. If one thinks that one should believe only what one knows, this can be motivated by the existence of a margin-for-error principle governing knowledge (Williamson (2000)). Yet this kind of revision would violate minimality.

As first observed in Levi (1977) for the corresponding operations on of closed sets of sentences, revision can be defined in terms of subtraction and addition.[3]

Levi Identity $s * s' = (s - \overline{s}') + s'$

Revising with incoming information is equivalent to subtracting its complement and then adding the information itself (for some subtraction operation).

2.3 Semantics

Where $\{\mathsf{A}, \mathsf{B}, ..., \bot, \top\}$ is a set of atomic sentences, L_0 is a boolean language.

Definition 5.

- $\{\mathsf{A}, \mathsf{B}, ..., \bot, \top\} \subseteq \mathsf{L}_0$;
- If $\phi, \psi \in \mathsf{L}_0$, then $\neg\phi, \phi \wedge \psi \in \mathsf{L}_0$;
- Nothing else is a member of L_0.

We define disjunction and the material conditional stipulatively in terms of negation and conjunction, so that $\phi \vee \psi \equiv_{def} \neg(\neg\phi \wedge \neg\psi)$ and $\phi \supset \psi \equiv_{def} \neg(\phi \wedge \neg\psi)$.

Let worlds $w, v, ... \in \mathcal{W}$ be functions from atomic sentences to truth values. Every world maps $\top$ to 1 and $\bot$ to 0. $\langle \mathcal{P}(\mathcal{W}), \subseteq \rangle$ is the powerset information structure of based on the set of worlds $\mathcal{W}$. We can define a static interpretation function, $[\![\cdot]\!]$, which maps sentences of L_0 to states in the information structure.

[3] I use 'subtraction' and 'addition', respectively, as names for the contraction and expansion operations of classical AGM on closed sets of sentence, since standared names of the latter have misleading connotations in the information structure setting.

Definition 6 (Static Semantics).

$$\begin{array}{llll} (i) & [\![\mathsf{A}]\!] & = & \{w \mid w(\mathsf{A}) = 1\} \\ (ii) & [\![\neg\phi]\!] & = & \overline{[\![\phi]\!]} \\ (iii) & [\![\phi \wedge \psi]\!] & = & [\![\phi]\!] \wedge [\![\psi]\!] \end{array}$$

We can think of the static denotation of a sentence as the information which the sentence semantically conveys. An atomic sentence denotes the information state comprising the worlds at which it is true. The negation of a sentence denotes the complement of the denotation of the sentence negated. The conjunction of two sentences denotes the meet of the denotation of the sentences conjoined. Since states in the structure are simply sets of worlds, the interpretation function simply maps sentences to their coarse-grained propositional content.

2.3.1 Dynamic Semantics: Basics

In static semantics, the meaning of a sentence is a proposition—a set of points of evaluation. In dynamic semantics, the meaning of a sentence is a context change potential (CCP) — a function from points of evaluation to new points of evaluation. Points of evaluation in a dynamic framework represent discourse contexts. By identifying meanings with CCPs, we are thus able to model both how an utterance's evaluation depends on context, and how context depends on what is uttered.

A dynamic theory's choice of points of evaluation will reflect the type of changes to context it aims to model. Chapter 1 suggested that supposition has a dual effect: (i) it changes what possibilities are under consideration; and (ii) it changes how later suppositions change what possibilities under consideration. Accordingly, we will identify contexts with pairs, containing an information state, and a rule for transitioning between them.

Definition 7 (Context).

A context σ is a pair $\langle s_\sigma, *_\sigma \rangle$, such that:

- s_σ is an information state;
- $*_\sigma$ is a quasi-revision operation.

Each context comprises a view on the world (an information state) and a way of changing that view on the world (a (quasi-)revision operation). Where $*_\sigma$ is a full revision operation (i.e., it satisfies ($*_1$)-($*_4$)), we will say that σ is proper. Otherwise, we will say that σ is adequate. The underlying idea, to be reflected in our account of entailment, is that every conversation starts in a proper context, but that update may result in a context which is merely adequate. Where s_σ is empty, we say that σ is absurd. Observe that there is more than one absurd context (one for every quasi-revision operation).

We define a dynamic interpretation function, $[\cdot]$, which maps sentences of L_0 to CCPs.

Definition 8 (Basic Dynamic Semantics).

$$\begin{array}{llcl} (i) & \sigma[\mathsf{A}] & = & \langle s_\sigma + [\![\mathsf{A}]\!], *_\sigma\rangle \\ (ii) & \sigma[\neg\phi] & = & \langle s_\sigma / s_{\sigma[\phi]}, *_\sigma\rangle \\ (iii) & \sigma[\phi\wedge\psi] & = & \sigma[\phi][\psi] \end{array}$$

Update with A returns the context resulting from adding its static content to the information state of the input context. Update with $\neg\phi$ returns the context resulting from eliminating any worlds from the information state of the input which would survive update with ϕ. Update with $\phi\wedge\psi$ returns the context resulting from sequential update with ϕ and ψ.

Across L_0, update leaves the revision operation of the context unchanged. In fact, we can observe something stronger: within the boolean fragment, update is purely additive. That is:

$$\sigma[\phi] = \langle s_\sigma + [\![\phi]\!], *_\sigma\rangle$$

For $\phi \in \mathsf{L}_0$, update with ϕ has the effect of adding $[\![\phi]\!]$ to the information state of the input context. This reflects the informal idea, mentioned in Chapter 1, that assertion simply amounts to eliminating possibilities from consideration which are incompatible with the information conveyed by the sentence asserted (cf., in particular, Stalnaker (1978, 2002, 2014)).

We define a notion of acceptance and entailment in terms of information preservation.

Definition 9 (Acceptance and Entailment).

(i) $\sigma \models \phi$ iff $s_\sigma = s_{\sigma[\phi]}$

(ii) $\psi_i, ..., \psi_j \models \phi$ iff for all proper σ either:

- $\sigma[\psi_i], ..., [\psi_j] \models \phi$; or
- $\sigma[\psi_i], ..., [\psi_j][\phi]$ is undefined.

A context accepts ϕ iff updating with ϕ leaves the information state of the context unchanged. A sequence of premises entail a conclusion iff the context which results from updating with the premises sequentially accepts the conclusion (where sequential update with the premises and conclusion is defined).[4] Observe that, under **Def.9**, the logic of L_0 is classical.

The second clause of the definition encodes a version of Strawson entailment (Strawson (1952), von Fintel (1997*a*, 1999)): in evaluating the validity of an argument we consider only those contexts at which the premises and conclusion are defined (in their local context). This clause is idle over L_0, where every CCP is total over the domain of contexts (both proper and adequate). However, it will allow us to model the effect of presuppositions on judgments about validity (Chapter 3, §3.1.1 in particular, for discussion of the features of Strawson entailment).

2.3.2 Supposition

Our next step is to enrich the theory to model the effects of supposition. To do so, we introduce a monadic operator, *Sup*. Let L_1 be the language that results from embedding sentences of the boolean fragment under *Sup*. We call L_1 the suppositional fragment of our final language.

Definition 10.

- If $\phi \in \mathsf{L}_0$, then $Sup(\phi) \in \mathsf{L}_1$;

[4] Thus, our preferred entailment relation is an update-to-test relation (cf. Veltman (1996), van Benthem (1996))

- Nothing else is a member of L_1.

On the informal picture sketched in Chapter 1, what is supposed must be held fixed through future suppositional changes to the discourse context. After supposing ϕ, any future revision to the possibilities under consideration should return an information state which incorporates the information conveyed by ϕ.

To implement this idea, we need to define an update operation on (quasi-)revision operations.

Definition 11 (Quasi-Revision Update).

For any quasi-revision operation $*$ and information states s, s', and s'':

$$s *_{|s''} s' = s * (s' \wedge s'')$$

Intuitively, $*_{|s}$ is the operation just like $*$, but which preserves the information s. That is, it only ever returns an information state which is stronger than s. $s *_{|s''} s'$ is the $*$-revision of s which returns a state incorporating s' and s''. If $*$ is a full-revision operation, then $*_{|s}$ is a quasi-revision operation. However, $*_{|s}$ may fail to be a full revision operation. Counterinstances to $(*_1)$ will occur when s is compatible with s' (i.e., $s \wedge s' \neq \bot$), but not with $s \wedge s''$ (i.e., $s \wedge s'' = \bot$). In this case, $s *_{|s''} s'$ will not be purely additive; that is, it will not be equal to $s \wedge s'$.

We are now in a position to extend our dynamic interpretation function to L_1.

Definition 12 (Supposition).

$$\sigma[Sup(\phi)] = \langle s_\sigma * [\![\phi]\!], *_{\sigma|[\![\phi]\!]} \rangle$$

$Sup(\phi)$ has a dual effect on a context, σ: first, it revises its information state, s_σ, with the information conveyed by ϕ. Second, it replaces its (quasi-)revision operation, $*_\sigma$, with the operation just like it, but which preserves the information conveyed by ϕ. Informally, $Sup(\phi)$ has the effect of (minimally) changing the set of possibilities under consideration so that they incorporate $[\![\phi]\!]$, and ensuring that the result of any further suppositional changes also entail this information.

The supposition operation defined has all of the features we identified as desiderata.

If $\phi \models \psi$, then $Sup(\phi) \models \psi$. PRODUCTIVITY

$Sup(\phi)$,..., $Sup(\psi) \models \phi \wedge \psi$. ACCUMULATIVITY

If $Sup(\phi) \models \psi \wedge \chi$, then $Sup(\phi \wedge \psi) \models \chi$ REGULARITY

If $\sigma \models \phi$ but $\sigma \not\models \neg\psi$, then $\sigma[Sup(\psi)] \models \phi$. CONSERVATIVITY

First, supposition is productive. Supposing ϕ returns a context which accepts anything that is entailed by ϕ. Productivity follows directly from $(*_2)$. The result of revision with an information state will incorporate that information state. Second, supposition is accumulative. Supposing ϕ and supposing ψ sequentially (with an arbitrary series of intervening updates) returns a context which accepts $\phi \wedge \psi$. Accumulativity follows from $(*_2)$ and the definition of update to revision. After supposing some information state, any later supposition will return a state which incorporates that information. Third, supposition is regular. If supposing ϕ returns a context which accepts $\psi \wedge \chi$, then supposing $\phi \wedge \psi$ will return a context which accepts χ. Regularity directly follows from $(*_4)$. Finally, supposition is conservative. If a (proper) context accepts ϕ but does not accept $\neg\psi$, then supposing ψ will return a context which continues to accept ϕ. Conservativity follows directly from $(*_1)$.

2.3.3 Conditionals

Finally, we need to add a conditional to our language. To do so, we start by introducing a binary operator $\rightarrow$. Let L_2 be the language that results from adding $\rightarrow$ in a limited way. We will call L_2 the conditional fragment of our language.

Definition 13.

- If $\phi \in \mathsf{L}_0 \cup \mathsf{L}_1$, and $\psi \in \mathsf{L}_0 \cup \mathsf{L}_2$, then $\phi \rightarrow \psi \in \mathsf{L}_2$;
- Nothing else is a member of L_2.

$\phi \to \psi$ expresses a generalisation of the dynamic strict conditional defended by Dekker (1993), Gillies (2004, 2009) and Starr (2014*b*,*c*, ms), amongst others.

Definition 14.

$$\sigma[\phi \to \psi] = \begin{cases} \sigma, & \textit{if } \sigma[\phi] \models \psi; \\ \langle \emptyset, *_\sigma \rangle, & \textit{otherwise.} \end{cases}$$

$\phi \to \psi$ checks whether ψ is accepted by $\sigma[\phi]$. If so, it returns σ; if not, it returns an absurd state, $\langle \emptyset, *_\sigma \rangle$. Stated informally, $\phi \to \psi$ induces a test which passes iff after update with ϕ, ψ conveys no new information.

The dynamic strict conditional offers a simple implementation of the idea that, to evaluate a conditional, we first update the context with its antecedent and then evaluate its consequent at the result. Crucially, however, Chapter 1 suggested that we update the context not as we would for an assertion of the antecedent, but rather as we would if it were supposed. Accordingly, we will not identify the natural language conditional, $\Rightarrow$, with the dynamic strict conditional directly. Rather, we will take it to correspond to a complex sentence within our language.

Definition 15.

$$(i.) \quad \phi \Rightarrow \psi \quad =_{def} \quad Sup(\phi) \to \psi$$

$$(ii.) \quad \sigma[Sup(\phi) \to \psi] \quad = \quad \begin{cases} \sigma, \text{ if } \sigma[Sup(\phi)] \models \psi \\ \langle \emptyset, *_\sigma \rangle, \text{ otherwise.} \end{cases}$$

$\phi \Rightarrow \psi$ is introduced to be definitionally equivalent to the strict conditional in which its antecedent occurs under *Sup*; that is $\phi \Rightarrow \psi = +def Sup(\phi) \to \psi$. Stated informally, $Sup(\phi) \to \psi$ induces a test which decomposes into two stages: first, it finds the result of updating σ with $Sup(\psi)$. This update revises s_σ with the information conveyed by ϕ, and replaces $*_\sigma$ with its ϕ-preserving variant. Second, it checks that the resulting context accepts ψ. If so, it returns σ; if not, it returns an absurd state.

Our complete language $\mathsf{L} = \mathsf{L}_0 \cup \mathsf{L}_1 \cup \mathsf{L}_2$ is the union of the boolean fragment, suppositional fragment and conditional fragment. Within this language, every sentence is associated with instructions of a certain kind: instructions on how to change the context. The conditional, in particular, encodes instructions to either keep the context the same, or return an absurd context, depending on the status of the consequent after the antecedent is supposed. It is easy to see, on this picture how (1) and (2) (repeated below) are alike. The latter will be accepted at a context iff the former constitutes a valid argument.

(1) Suppose that the butler did it. Then the gardener is innocent.

(2) If the butler did it, then the gardener is innocent.

Because every sentence of the language is associated with a meaning of the same kind, unlike Speech Act Suppositional Theories, it is easy to satisfy the requirements of embedding and validity. The language is compositional—conditionals freely embed under negation, in conjunction and disjunction and in the consequents of conditionals. And, as we have seen, it is easy to define a non-standard notion of validity in terms of update and acceptance, which neverthless remains classical over the boolean fragment of the language.

Finally, observe that the account satisfies Uniformity. Both indicatives and subjunctives are associated with the same basic meaning: both test whether supposing the antecedent returns a context where the antecedent is accepted. The introduction argued that Uniformity is a condition on any minimally adequate theory of indicatives and subjunctives. Nevertheless, in virtue of satisfying it we are left with a clear explanatory burden: why is it that the two types of conditional exhibit such different behavior?

2.4 Indicatives & Subjunctives

Indicatives and subjunctives differ in morphological marking. The antecedent of (43) contains a layer of past perfective morphology lacking in (42).

(42) If the butler did it, he used the candlestick.

(43) If the butler had done it, he would have used the candlestick.

However, the morphological marking in the antecedent of (43) does not appear to encode its standard past-tense meaning (Starr (2014*a*, 1026)). Temporal modifiers (such as, e.g., '*tomorrow*') which are typically unacceptable under past tense morphology can occur freely in subjunctive antecedents. For example, whereas (44.a) is marked, (44.b) is a felicitous continuation of the discourse.

(44) Sherlock arrived before the witnesses departed.

a. ?? If he arrived tomorrow, he was too late.

b. If he had arrived tomorrow, he would have been too late.

This has motivated a number of theorists to argue that the additional layer of morphological marking has a dedicated semantic contribution in conditionals (Iatridou (2000), Schulz (2007), Starr (2014*a*), amongst others; for dissent cf., Ippolito (2003), Arregui (2007, 2009) a.o.). It is its distinct morphological marking, then, rather than any ambiguity in the meaning of '*if*', to which the subjunctive's contrast with the corresponding indicative is to be attributed.

In this section, I defend an account of this kind. I propose that the variation in morphological marking in conditional antecedents encodes a difference in presupposition. Stated simply, indicatives presuppose the possibility of their antecedent; subjunctives don't.

On this view, while the two types of conditional agree in their at-issue meaning (§2.3.3), the contexts in which they are licensed differ. Crucially, our adopted notion of validity (§2.3.1) makes the validity of an argument sensitive to the licensing conditions of its constituent sentences. As a result, the two types of conditional are predicted to have different logics and to interact in importantly different ways with supposition. We will see in §2.5 that this allows us to accommodate the patterns of behavior observed in Chapter 1.

The proposed account of indicative/subjunctive morphology is closely related to one previously defended by Han (1998) and Sampanis (2012). According to Han and Sampanis, subjunctive inflection encodes a [-REALIS] feature, interpreted as indicating that the situation described by the clause in which it occurs is 'unrealized'. The present proposal can be seen as implementing this idea: subjunctive mood in a clause encodes the absence of a presupposition that it is compatible with the discourse environment.

2.4.1 Indicatives

As noted in Chapter 1, indicative conditionals do not permit counterfactual use (§1.2.2). Following Stalnaker (1975), Karttunen and Peters (1979), von Fintel (1997*b*), Gillies (2009, 2020) and Starr (2014*a*) (amongst others), say that indicative conditionals presuppose the possibility of their antecedent. Using $\dashrightarrow$ as notational shorthand for the indicative:

Definition 16.

$$\sigma[\phi \dashrightarrow \psi] = \begin{cases} \sigma[Sup(\phi) \rightarrow \psi], & \textit{if } \sigma \not\models \neg\phi \\ undefined, & \textit{otherwise.} \end{cases}$$

$\sigma[\phi \dashrightarrow \psi]$ is defined only if σ is compatible with the information conveyed by ϕ. In this case, it applies the test induced by $Sup(\phi) \rightarrow \psi$.

Supposition has an important property. Where σ is a proper context compatible with ϕ, supposing ϕ has the same effect on the information state of the context as updating with ϕ directly. That is:

Fact 1.

For all proper σ and $\phi \in \mathsf{L}_0$:

If $\sigma \not\models \neg\phi$, then $s_{\sigma[\phi]} = s_{\sigma[Sup(\phi)]}$.

Fact 1 follows directly from the observations that (i) if $\sigma \not\models \neg\phi$, then $s \wedge [\![\phi]\!] \neq \bot$ and (ii) the revision operation of any proper context obeys $(*_1)$. Revising with a compatible information state has the same effect as adding it. Thus, in contexts compatible with ϕ, the only difference between $Sup(\phi)$ and ϕ is their effect on later revisions.

It follows that the indicative is equivalent to the plain dynamic strict conditional over the closure of L_0 under $\dashrightarrow$.

Fact 2.

Where $\phi \dashrightarrow \psi$ belongs to the closure of L_0 under $\dashrightarrow$:

$$\phi \dashrightarrow \psi \mathrel{=\!\mid\!\models} \phi \rightarrow \psi \qquad . \quad \cdot$$

A full proof of **Fact 2** can be found in **Appendix A**. We start by showing that for all proper σ at which both are defined, σ accepts $\phi \dashrightarrow \psi$ iff it accepts $\phi \rightarrow \psi$, as long as $\phi, \psi \in \mathsf{L}_0$. The latter is defined everywhere. So, suppose that $\phi \dashrightarrow \psi$ is defined at σ. Then, by **Def.16**, it follows that σ does not accept $\neg\phi$. So, by **Fact 1**, $s_{\sigma[Sup(\phi)]}$ just is $s_{\sigma[\phi]}$: the two updates produce the same information state. But observe that σ accepts $\phi \rightarrow \psi$ iff $s_{\sigma[\phi]} = s_{\sigma[\phi][\psi]}$, and accepts $\phi \dashrightarrow \psi$ iff $s_{\sigma[Sup(\phi)]} = s_{\sigma[Sup(\phi)][\psi]}$. So, since $\psi \in \mathsf{L}_0$ and update with boolean expressions is revision operation insensitive, it follows that σ accepts the former iff it accepts the latter. We can then proceed to prove **Fact 2** by induction.

This is a comforting result if one is sympathetic to the thought that logic generated by the dynamic strict conditional is appropriate for indicatives. Our suppositional theory, supplemented with the idea that indicatives presuppose the possibility of their antecedent, coincides with the dynamic strict account (Gillies (2004, 2009), Gillies (2009)).

The dynamic strict account of the indicative has some nice features. The indicative is predicted to be equivalent to the material conditional; that is, $\phi \dashrightarrow \psi \dashv\models \phi \supset \psi$. Yet, since the two are associated with different CCPs and have different definedness conditions, they behave differently in embedded environments. As shown by, e.g., Rothschild (2014) and Starr (ms), this allows for the possibility that they may interact differently with judgments about probability.

(MP)	$\phi \Rightarrow \psi, \phi$	$\models$	ψ	Modus Ponens
(IE)	$(\phi \wedge \psi) \Rightarrow \chi$	$\dashv\models$	$\phi \Rightarrow (\phi \Rightarrow \chi)$	Import/Export

It also satisfies both *Modus Ponens* and Import/Export. Yet, *contra* Gibbard (1981), it does so without collapsing into the material conditional. As pointed out by Gillies (2009), the two are logically equivalent but remain semantically distinct. Finally, it satisfies the deduction theorem. If some sequence of premises, followed by ϕ, together entail ψ, then that sequence of premises alone is sufficient to establish $\phi \dashrightarrow \psi$. As we suggested in Chapter 1, this seems to be the right result.

2.4.2 Subjunctives

Unlike indicatives, subjunctives do permit counterfactual uses. Following Stalnaker (1975) and von Fintel (1997*b*), say that subjunctive conditionals have trivial presuppositions: they do not introduce definedness conditions beyond those of their constituents.

Definition 17.

$$\sigma[\phi \rightsquigarrow \psi] = \sigma[Sup(\phi) \rightarrow \psi]$$

Unlike indicatives, subjunctives are defined in contexts which accept the negation of their antecedent. In such contexts, supposition of the antecedent returns the revision of the input information state with the antecedent (along with a revision operation which preserves the information conveyed by the antecedent). Clearly, this information state will include 'counterfactual' possibilities—possibilities not included in the information state of the input context. The test imposed by the subjunctive passes iff this information state is left unchanged after update with the consequent.

Subjunctives are not strawson equivalent to the corresponding dynamic strict conditional. Where σ accepts $\neg\phi$, $\phi \rightarrow \psi$ will be trivially accepted. Yet $\phi \rightsquigarrow \psi$ need not be. Accordingly, they lack many of the properties of indicatives.

While *Modus Ponens* remains valid for subjunctives with boolean consequents, there will be counterinstances involving embedded subjunctives. Supposing ϕ can result in a context at which a subjunctive is accepted, without that subjunctive being accepted at a context which results from updating with ϕ.[5] In fact, this seems like the right prediction. Unlike its indicative counterpart, the argument in (45) appears bad.

(45) ✗ a. If the match had been wet when it was struck, then if it had lit when it was struck, then it would have lit when it was wet.

b. The match was wet when it was struck.

c. So if it had lit when it was struck, then it would have lit when wet.

(45.a) appears unimpeachable—indeed, according to our current semantics, it has the status of a logical truth. Yet, upon learning (45.b), there remains good reason to resist accepting (45.c). Afterall, we might reasonably insist that if the match had lit when it was struck, then it would have been dry.

The current semantics explains this judgment. We accept (45.a), since in evaluating the embedded subjunctive, we hold the information that the match was wet fixed. Asserting (45.b) adds the information that the match was wet (and struck) to the information state. However, it does not require this information to be held fixed. Hence, in evaluating the conclusion, we are permitted to adopt an information state at which the information that the match is wet is relinquished upon revising with the information that the match was struck and lit.

What about Import/Export? It turns out that the status of Import/Export for subjunctives turns on the properties of iterated revision. Darwiche and Pearl (1997) propose canonical constraints on sequential revision. These are captured, in the present framework, by $(*_5)$-$(*_7)$:

[5]For example, let $s_\sigma = \{w_{\mathsf{Ab}}\}$ and suppose that $s_\sigma * [\![\mathsf{B}]\!] = \{w_{\mathsf{aB}}\}$. It is guaranteed that $\sigma \models \mathsf{A} \rightsquigarrow (\mathsf{B} \rightsquigarrow (\mathsf{A} \wedge \mathsf{B}))$. Yet $\sigma[\mathsf{A}] \not\models \rightsquigarrow (\mathsf{B} \rightsquigarrow (\mathsf{A} \wedge \mathsf{B}))$ (where: $w_{\mathsf{Ab}}(\mathsf{A}) = w_{\mathsf{aB}}(B) = 1 - (w_{\mathsf{Ab}}(\mathsf{B}) = 1 - w_{\mathsf{aB}}(\mathsf{A})))$.

$(*_5)\quad (s * s') * s'' = s * s'\qquad \textit{if}\ s'' \wedge s' = s''\ \text{or}\ s'' \wedge s' = \bot;$

$(*_6)\quad (s * s'') * s' \leq s'\qquad \textit{if}\ s * s' \leq s''$

$(*_7)\quad ((s * s'') * s') \wedge s'' \neq \bot\qquad \textit{if}\ (s * s') \wedge s'' \neq \bot$

As demonstrated in **Appendix A**, Import/Export will be valid for subjunctives as long as the quasi-revision operation of every proper context satisfies Darwiche and Pearl's constraints on iterated revision. In fact, this is stronger than necessary, since $(*_5)$ by itself suffices.

Finally, the deduction theorem is invalid over the closure of L_0 under $\rightsquigarrow$. While a subjunctive may be entailed by a sequence of premises followed by ϕ, the same sequence of premises need not entail the nested subjunctive in which ϕ takes wide-scope. Again, the difference comes down to the fact that in the former case, ϕ is merely asserted, whereas in the latter, it occurs under supposition in the antecedent. Since, unlike indicatives, subjunctives allow for counterfactual revision, this can change the evaluation of downstream material.

While subjunctives are not equivalent to the dynamic strict conditional, dynamic strict behavior can be recovered under supposition. That is, our subjunctive exhibits an important property:

> **Fact 3.**
>
> If $\phi_i, ..., \phi_j \models \psi \rightarrow \chi$, then $Sup(\phi_i), ..., Sup(\phi_j) \models \psi \rightsquigarrow \chi$.
>
> (Where $\phi_i, ..., \phi_j \in \mathsf{L}_0$.)

The subjunctive behaves like a dynamic strict conditional in arguments as long as the premises are supposed rather than asserted. Proof is provided in **Appendix A**.

This allows us to recover valid variants of a range of arguments which are otherwise invalid for the subjunctive. Supposing a material conditional allows one to infer the corresponding subjunctive. A version of *Modus Ponens* is valid in which the non-conditional premise occurs under supposition. And the deduction theorem can be preserved in the variant form in which the premises are supposed. In each case, this appears to match with our judgments about the effect of replacing assertion with supposition in the respective arguments.

2.5 Results

In this section, I survey how the account developed in §§2.2-2.4 accommodates the phenomena surveyed in Chapter 1. More technical material is delayed to **Appendix A**.

2.5.1 Inference Patterns

In §1.2.1, we observed that indicatives and subjunctives interact differently with the inference patterns of Preservation, the Direct Argument and Conditional Telescoping.

$$\begin{array}{lrl} (\underset{\dashrightarrow}{\text{PR}}) & \phi & \models \psi \dashrightarrow (\phi \wedge \psi) \\ (\underset{\dashrightarrow}{\text{DA}}) & \phi \vee \psi & \models \neg\phi \dashrightarrow \psi \\ (\underset{\dashrightarrow}{\text{CT}}) & \phi \dashrightarrow (\psi \dashrightarrow \chi),\ \psi & \models \phi \dashrightarrow \chi \end{array}$$

Preservation, the Direct Argument and Conditional Telescoping are each valid within the fragment of L_0 closed under $\dashrightarrow$. This follows directly from **Fact 2** and the observation that they are valid for the strict conditional. Furthermore, since the Deduction Theorem is valid for indicatives, we know that their conditional variants will be likewise valid.

$$\begin{array}{lrl} (\underset{\dashrightarrow}{\text{PR}}^*) & & \models \phi \dashrightarrow (\psi \dashrightarrow (\phi \wedge \psi)) \\ (\underset{\dashrightarrow}{\text{DA}}^*) & & \models (\phi \vee \psi) \dashrightarrow (\neg\phi \dashrightarrow \psi) \\ (\underset{\dashrightarrow}{\text{CT}}^*) & \phi \dashrightarrow (\psi \dashrightarrow \chi) & \models \psi \dashrightarrow (\phi \dashrightarrow \chi) \end{array}$$

Hence, the current proposal is immediately able to accommodate the judgments about indicatives discussed.

Subjunctives behave differently. Preservation, the Direct Argument and Conditional Telescoping are invalid for the subjunctive in the closure of L_0 under $\rightsquigarrow$. (i')-(iii') are invalid:

$$\begin{array}{lrl} (\underset{\rightsquigarrow}{\text{PR}}) & \phi & \models \psi \rightsquigarrow (\phi \wedge \psi) \\ (\underset{\rightsquigarrow}{\text{DA}}) & \phi \vee \psi & \models \neg\phi \rightsquigarrow \psi \\ (\underset{\rightsquigarrow}{\text{CT}}) & \phi \rightsquigarrow (\psi \rightsquigarrow \chi),\ \psi & \models \phi \rightsquigarrow \chi \end{array}$$

It is easy to see why. Since the antecedent of the subjunctive conclusion need not be compatible with the context that results from update with the premises, there is no guarantee that the premises will continue to be accepted at the context at which the consequent is evaluated. For example, consider ($\underset{\leadsto}{\mathrm{PR}}$) (the subjunctive instance of Preservation). If update with ϕ returns a context which rules out ψ, then supposing ψ can return a new context which fails to support ϕ. Hence, update with ϕ does not always result in a context at which $\psi \leadsto (\phi \wedge \psi)$ is accepted.

Each inference pattern can, nevertheless, be made valid if the rightmost premise is embedded under supposition. That is, if we treat non-conditional premises as supposed, rather than asserted, validity is recovered.

$$
\begin{array}{lrl}
(\mathrm{S}\underset{\leadsto}{\mathrm{PR}}) & \mathit{Sup}(\phi) & \models \psi \leadsto (\phi \wedge \psi) \\
(\mathrm{S}\underset{\leadsto}{\mathrm{DA}}) & \mathit{Sup}(\phi \vee \psi) & \models \neg\phi \leadsto \psi \\
(\mathrm{S}\underset{\leadsto}{\mathrm{CT}}) & \phi \leadsto (\psi \leadsto \chi),\ \mathit{Sup}(\psi) & \models \phi \leadsto \chi
\end{array}
$$

($\mathrm{S}\underset{\leadsto}{\mathrm{PR}}$)-($\mathrm{S}\underset{\leadsto}{\mathrm{CT}}$) are each valid. For the former two, this follows directly from **Fact 3**. The latter requires the additional assumption that the revision operation of a proper context satisfies the iterated revision conditions, ($*_5$)-($*_7$). Proofs are provided in **Appendix A**. However, an informal gloss is easily available: information introduced via supposition must be preserved by revisions with the antecedent of the subjunctive conclusion. Accordingly, it is guaranteed to be accepted at the context at which consequent is evaluated. For example, consider ($\mathrm{S}\underset{\leadsto}{\mathrm{PR}}$) (the suppositional variant of Preservation). After supposition of ϕ, supposing ψ will return a context which supports both ϕ and ψ. Accordingly, the conditional $\psi \leadsto (\phi \wedge \psi)$ is guaranteed to be accepted.

Since the deduction theorem is valid for $\rightarrow$, the same reasoning also accounts for the conditional variants.

$$
\begin{array}{lrl}
(\underset{\leadsto}{\mathrm{PR}}^*) & & \models \phi \Rightarrow (\psi \leadsto (\phi \wedge \psi)) \\
(\underset{\leadsto}{\mathrm{DA}}^*) & & \models (\phi \vee \psi) \leadsto (\neg\phi \Rightarrow \psi) \\
(\underset{\leadsto}{\mathrm{CT}}^*) & \phi \leadsto (\psi \leadsto \chi) & \models \psi \leadsto (\phi \leadsto \chi)
\end{array}
$$

($\underset{\leadsto}{\mathrm{PR}}^*$)-($\underset{\leadsto}{\mathrm{CT}}^*$) are each valid. By the deduction theorem for the strict conditional, if we know that $\Gamma, \mathit{Sup}(\phi) \models \psi$, it follows that $\Gamma \models \mathit{Sup}(\phi) \rightarrow \psi$. But the subjunctive is just equivalent to $\mathit{Sup}(\phi) \rightarrow \psi$. So from the fact that the suppositional variants of each inference pattern are valid, it follows that the conditional variants are too. Accordingly, we are able accommodate the full range of judgments about the validity/invalidity of inferences discussed in Chapter 1.

2.5.2 Counterfactual Use

Indicatives presuppose that their antecedent is compatible with the context in which they are used. Accordingly, counterfactual uses are predicted to be marked. Subjunctives, which lack this presupposition, are in general predicted to be compatible with counterfactual use. However, as we saw, the availability of counterfactual uses of subjunctives disappears when incompatible information is introduced via supposition, or in a wide scope '*if*'-clause, as in (17)-(18) (repeated).

(17) Suppose that the butler hadn't done it. ??If he'd done it, he would've used the candlestick.

(18) ??If the butler hadn't done it, then if he'd done it, then he would've used the candlestick.

There is a simple explanation of this on the present account. Information introduced via supposition is preserved by any later suppositions. Where ϕ and ψ are inconsistent, successive updates with $Sup(\phi)$ followed by $Sup(\psi)$ will return an absurd context. Since they are inconsistent, no non-absurd information state accepts both.

Accordingly, evaluating a subjunctive in an environment that results from supposing something inconsistent with its antecedent will require evaluating its consequent at an absurd context. In general, we should expect this to result in infelicity. Since absurd contexts accept everything, such uses of a subjunctive will be uninformative.[6]

[6] A reasonable source of concern is that I have offered a different kind of explanation of the infelicity of counterfactual uses of indicatives and counterfactual uses of subjunctives under supposition. The former is predicted to be bad due to presupposition failure, whereas the latter is predicted to be bad due to triviality. Isn't this at least somewhat unsatisfying?

It is easy to avoid this shortcoming by modifying our account of subjunctives in a minimally invasive way. Rather than treating them as being everywhere defined, we could instead take $\phi \rightsquigarrow \psi$ to be defined in σ only if $s_{\sigma[Sup(\phi)]} \neq \emptyset$. This will preserve the acceptability counterfactual uses in general, while predicting a similar source of infelicity in cases where their antecedent conflicts with prior supposition (or is inconsistent itself).

2.5.3 Epistemic Contradictions

In §1.2.3, we saw that epistemic contradictions (sequences comprising a '*might*' claim along with the negation of its prejacent) interact in similar ways with '*if*'-clauses and supposition. Not only are their conjunctions marked in both environments, they also exhibit more complicated and less widely discussed behavior. Discourses like (24.a) and the corresponding conditionals like (25.a) appear equally infelicitous.

(24) a. Suppose that it weren't raining. ??If it might've been, I'd have brought an umbrella needlessly.

(25) a. ?? If it weren't raining, then if it might've been, I'd have brought an umbrella needlessly.

The framework developed in §§2.3-2.4 does not extend to epistemic modals. Accordingly, it is not equipped to make predictions about epistemic contradictions or their interaction with '*if*'-clauses/supposition.

It is not complicated to extend the semantics to a language including modals. We introduce them via stipulative definition, so that $\Box\phi \equiv_{def} \top \Rightarrow \phi$:

Definition 18 (Repeated).

$$\sigma[\Box\phi] = \begin{cases} \sigma, & \textit{if } \sigma \models \phi; \\ \langle \emptyset, *_\sigma \rangle, & \textit{otherwise.} \end{cases}$$

The more significant issue is that our framework does not provide us with a way of embedding expressions under supposition unless they associated with propositional content. Supposition changes the information state of a context via revision and revision is defined in terms of the information conveyed by an expression. Yet, under their standard dynamic treatment, modal claims do not convey information.

Appendix B provides a solution to this issue by complicating the framework. Relative to a revision operation, $*$, we can associate every expression in the language with an informational content: the information states which accept it when paired with $*$. We then generalize our notions of revision to apply, not to individual states, but to sets of information states.

This allows us to extend our supposition operation to the modal fragment of the language, L_2, containing $\Box$, $\Diamond$ and $\Rightarrow$, giving us the full range of embedding behavior. As demonstrated in the appendix, once extended in this way, the framework easily accounts for the behavior of epistemic contradictions, both in the simple cases discussed in Yalcin (2007), and the more complicated cases like (24.a-b) and (25.a-b).

2.5.4 Evidential Weakness

Unlike (26), (27)-(28) are perfectly acceptable bits of discourse. One can suppose some information or employ it in the antecedent of a conditional, even if one is not in a position to assert it.

(26) ?? Maybe the butler did it. He did. So the gardener is innocent.

(27) Maybe the butler did it. Suppose he did. Then the gardener is innocent.

(28) Maybe the butler did it. If he did, then the gardener is innocent.

It is easy to see why this would be the case for conditionals. On dynamic proposals like the present, a conditional does not convey new information. Instead, it provides a way for interlocutors to explore the information they already possess, via the test it induces. Indicatives and subjunctives check the properties of an information state under different kinds of modification: addition, for the former; revision, for the latter. Since, unlike assertion, the antecedent is not permanently incorporated into the discourse context, there is no expectation on the speaker to have a special epistemic standing with respect to it.

Supposition and conditionals behave alike in our framework. The primary difference between the two is that whereas the effects of '*if*'-clauses are limited to their syntactic scope, the effects of supposition are persistent. Nevertheless, they are not permanent. As discussed in Chapter 1, suppositions can be withdrawn in ways that assertion cannot. Like conditionals, supposition provides a tool for exploring the information at a context, without imposing any irreversible changes on it. Thus, as with conditional antecedents, we should not expect suppositions to be held to the same standard as assertions.

2.6 Supposition, Reconsidered

In this concluding section, we return to consider the properties of supposition: how it interacts with different kinds of morphological marking, and how verbal and cognitive supposition are related.

2.6.1 Supposition and Mood

§2.4 proposed that the difference in morphological marking between indicative and subjunctive antecedents semantically encoded a difference in presupposition. Verbal suppositions can occur with either kind of morphology, as (46.a-b) demonstrate.

(46) a. Suppose the butler did it.

b. Suppose the butler had done it.

(47) The butler didn't do it. Suppose he [had/??did]. Then there'd be blood on the candlestick.

An appealingly simple hypothesis is that the morphological marking found in (46.a) and (46.b) has precisely the same contribution as it does in indicative/subjunctive-antecedents. The contrast in discourses like (47) supports this hypothesis. Here, counterfactual use of the supposition is acceptable only if the subordinate clause has an additional level of past tense marking.

As long as it occurs in a non-counterfactual environment, the simple hypothesis predicts that the effect of supposition on downstream conditionals and later suppositions will be independent of its morphological marking. In particular, information supposed will be required to be held fixed in evaluating later conditionals, regardless of whether it has an additional layer of past tense or not. As (48)-(49) demonstrate, this prediction seems correct.

(48) a. ✓ Suppose the Mets outscored the Cubs.

b. ... Then, if the Cubs had scored 12, the Mets would have scored 13.

(49) a. ✗ The Mets outscored the Cubs.

b. ... So, if the Cubs had scored 12, the Mets would have scored 13.

In the discourse context generated by (48.a), (48.b) is judged true. Despite lacking an additional layer of past tense marking, the information conveyed by the complement clause of the former appears required to be preserved when evaluating the latter. In contrast, the same subjunctive can naturally be judged false if it occurs in a discourse context following (49.a) instead.

The relative parsimony of this hypothesis give us *prima facie* reason to accept it. However, if we are to do so, we will require some explanation of why the inferences in (9)-(11) appear easier to reject when the supposition is stripped of past-tense morphology. If additional past-tense morphology merely has an effect on presuppositions, we would expect (50)-(52) to be as good as their 'subjunctive' variants (discussed in §1.2.1).

(50) ? Suppose Ada is drinking red wine. (Then) if she were eating fish, she'd be eating fish and drinking red wine.

(51) ? Suppose Claude is in London or Paris. (Then) if he weren't in London, he'd be in Paris.

(52) ? If Lori were married to Kyle, then if she were married to Lyle, she'd be a bigamist. Suppose she's married to Lyle. (Then) if she were married to Kyle, she'd be a bigamist.

To account for this contrast, I propose, we need to recognize the additional effect that morphological marking can have on discourse structure. A discourse is not mere collection of utterances. Understanding a discourse requires understanding the relations between distinct utterances. Grammatical mood can play a role in guiding this process. In particular, a shift between 'indicative' anf 'subjunctive' morphology often indicates that two claims are being presented as contrasting.

Whereas, in its discourse context, (53.a) is most naturally heard as introducing an incompatible alternative to the possibility of Bob attending and us all drinking wine, this reading is notably less prominent for (53.b). The most natural interpretation of the latter (but not of the former) implies that, if both Mary and Bob come, we'll drink wine and do shots.

(53) If Bob comes to the party, we'll drink wine.

a. ...If Mary were to come, we'd do shots.

b. ...If Mary comes, we'll do shots.

If, as suggested in §1.3.2, contrast can trigger withdrawal of suppositions, this would provide an explanation of why the inferences in (50)-(52) are degraded. The change in morphological marking between the supposition and the conditional indicates that they are being presented as contrasting, which in turn triggers withdrawal of the supposition. Yet withdrawing the downstream effect of supposition prior to evaluating the final subjunctive will result in the inference no longer being valid. Clearly, much more needs to be done to investigate the relation between supposition, mood and discourse structure. The brief discussion in this section has aimed merely to demonstrate the viability of an approach which can allow us to preserve a simple, univocal account of both supposition and 'indicative'/'subjunctive' morphology.

2.6.2 Supposition and Cognition

Above, we distinguished between verbal supposition and cognitive supposition. The former is a public act, resulting in changes the state of a conversation. The latter is a mental act, resulting in changes in an individual's state of mind.

To investigate the relationship between the two, we can start by comparing sentences like (54) and (55).

(54) Suppose that the gardener witnessed the murder.

(55) Sherlock supposed that the gardener witnessed the murder.

According to Dorr and Hawthorne (2013), '*suppose*' is lexically ambiguous across (54) and (55). In (55), it is as an attitude verb, relating the subject to a propositional content. In (54), by contrast, it is a *sui generous* speech act marker, which triggers a change in the context akin to that brought about by '*if*'. A view of this kind is also implicit in many Speech Act Suppositional Theories, such as those of Barnett (2006) and Barker (1995).

In favor of ambiguity, Dorr & Hawthorne cite apparent evidence that '*suppose*' does not combine with subjunctive mood when employed in its attitude ascription use. Whereas both forms of morohplogical marking are available in (56.a), they claim that only indicative is permitted in (56.b-c). This provides *prima facie* evidence, they suggest, that the two involve lexically distinct expressions: one, occurring in verbal supposition, which permits subjunctive complements clauses, and another, occuring in attitude ascriptions, which does not.

(56) a. Suppose that Gore [was/were/had been] president. How would things have been different?

b. He supposed that Gore [was/??were/??had been] president.

c. He is supposing that Gore [was/??were/??had been] president.

The subjunctive variants of (56.b-c) are indisputably marked in some respect. However, it is far from clear that this is due to a general prohibition on the presence of subjunctive mood in the complement clause of '*suppose*' whenever it occurs as an attitude verb. Crucially, in more complex constructions, such as (57.a-b), the same morphological marking appears perfectly acceptable within an attitude ascription.

(57) a. Having supposed that Gore [were/had been] president, he went on describe how things would have been different.

b. After supposing that Gore [were/had been] president, he will go on to describe how things would have been different.

This improvement is in need of explanation. Why is subjunctive mood marked in (56.b-c) if it is acceptable in (57.a-b)?

Interestingly, there is relatively strong evidence that '*suppose*' is in fact ambiguous between two attitude verbs: an atelic one and a telic one. The existence of the former is supported by sentences like (58.a), while the existence of the latter is supported by sentences like (58.b).

(58) a. [For years/until being corrected], he supposed that cows drank milk.

b. [To derive a contradiction/at the start of the proof], he supposed that the conjecture were true.

In constructions like (58.a), '*suppose*' combines with standard markers of atelic verbs and appears to denote a state roughly equivalent to belief.[7] Both '*for*'-PPs and '*until*' are typically assumed to require that the main verb of the clause they combine with is atelic (see Filip (2012) for a survey of these aspectual tests in English). In constructions like (58.b), by contrast, it combines with standard markers of telic verbs, and clearly does not denote a state equivalent to belief.[8]

These data motivate distinguishing between two readings of '*suppose*' in attitude ascriptions: the atelic reading (present in (58.a)), which denotes a state of endorsing a proposition and the telic reading (present in (58.b)), which denotes an action of adopting a proposition for the sake of argument.

Dorr & Hawthorne's data suggest that the former, atelic reading does not combine with subjunctive complement clauses (and hence should be distinguished from the verb that occurs in verbal suppositions like (56.a)). However, the latter, telic reading appears to do so freely.

There is good reason to think that constructions like those in (57.a-b) require their subordinate clause to denote an event with a result state. In each case, the reference time of the main clause is required to follow the completion of the event denoted by complement of '*having*'/'*after*' (in this case, the event of supposing Gore to be President). Accordingly, only the telic reading of the verb would be expected to be permitted in such environments, since atelic verbs denote events lacking an associated result state (Dowty (1979)). So, (57.a-b) provide us with strong evidence that the telic

[7]Note the oddity of continuing either variant of (58.a) with '... but he never believed it'.
[8]Continuing either variant of (58.b) with '...but he didn't believe it' is perfectly acceptable.

use of '*suppose*', unlike its atelic counterpart, does take a subjunctive complement.[9]

With this distinction in hand, a simpler account of the relationship between verbal supposition and attitude ascriptions is available. Rather than concluding that verbal suppositions like (56.a) involve a *sui generis* speech act marker on the grounds that they allow for subjunctive mood, we can instead simply treat them as imperatives headed by the telic variant of the verb. On the resulting picture, verbal and cognitive supposition are closely connected. Utterances like of sentences like (54) can be understood as belonging to a sub-category of the more general illocutionary act of commanding. That is, verbal supposition simply amounts to an instruction to engage in the act of cognitive supposition (i.e., adopting a proposition for the sake of argument).

This suggests a very different relationship between verbal supposition conditionals. Speech act suppositional theories treat '*if*'-clauses as indicating the performance of a particular kind of complex speech act, one built up out of verbal supposition and assertion. In doing so, they hope to explain why conditionals like (2) elicit the same judgments as discourses like (1).

(1) Suppose that the butler did it. Then the gardener is innocent.

(2) If the butler did it, then the gardener is innocent.

In contrast, there is no need to appeal to any similarity in illocutionary force on the present proposal. Instead, all we need is that the two are related at the level of cognition. To explain our convergent judgments, it suffices that in evaluating (2) an agent is required to evaluate its consequent after performing the same cognitive act towards its antecedent as is elicited by the imperative in (1). That is, a conditional is accepted by interlocutors just in case, upon supposing its antecedent, they would accept its consequent.

[9] Indeed, note that this was already evident from (58.b), in which the complement clause is has subjunctive morphology.

2.7 Summary

This chapter defended a suppositional theory of the natural language conditional within a dynamic framework. This theory satisfies Uniformity. Both indicatives and subjunctives are associated with a common 'conditional meaning'.

Importantly, by (i) allowing their presuppositions to differ and (ii) employing a notion of validity sensitive to definedness, it is nevertheless possible to obtain different logical properties for each. Whereas our indicative validates *modus ponens*, the deduction theorem and is bounded from above and below by the material conditional, our subjunctive exhibits none of these properties.

Despite their differing logics, indicatives and subjunctives remain closely related. In particular, our framework validates Collapse:

Collapse $\quad \phi \dashrightarrow \psi \dashv\vDash \phi \rightsquigarrow \psi$

Collapse says that indicatives and subjunctives are Strawson equivalent. Where defined, one is accepted iff the other is. Collapse is not an accidental feature of the framework. Rather, it is an essential feature of the strategy employed for distinguishing the two types of conditional. Since indicatives and subjunctives differ only in their presupposition, and, to establish validity, we look only at contexts where those presuppositions are satisfied, they are guaranteed to be equivalent.

The next section takes a step back, abstracting away from the current theory to argue that Collapse is, in fact, independently motivated. Rather than being a defect of the framework employed here, it is instead a desirable result, one that any adequate account the relationship between indicatives and subjunctives will need to vindicate.

Appendix A

We start by proving **Fact 2**. That is, where defined, indicatives induce the same update on context as the dynamic strict conditional. Let $\mathsf{L}_{\dashrightarrow}$ be the closure of L_0 under $\dashrightarrow$.

Fact 2 1.

Where $\phi \dashrightarrow \psi \in \mathsf{L}_{\dashrightarrow}$:

$$\sigma[\phi \to \psi] = \sigma[\phi \dashrightarrow \psi] \qquad (\textit{if defined}).$$

Proof. First, note that $\phi \to \psi$ and $\phi \dashrightarrow \psi$ are both tests. Thus, where defined, $\sigma \models \phi \dashrightarrow \psi$ iff $\sigma[\phi \dashrightarrow \psi] = \sigma$; otherwise $\sigma[\phi \dashrightarrow \psi] = \langle \emptyset, *_\sigma \rangle$ (and similarly for $\phi \to \psi$). It follows that $\sigma[\phi \to \psi] = \sigma[\phi \dashrightarrow \psi]$(where defined) iff for all σ on which $\phi \dashrightarrow \psi$ is defined, $\sigma \models \phi \to \psi$ iff $\sigma \models \phi \dashrightarrow \psi$. That is, to prove that their denotations coincide where defined, it suffices to demonstrate that they are strawson co-entailing.

Suppose $\phi, \psi \in \mathsf{L}_0$. Let σ be an arbitrary proper context s.t. $\sigma[\phi \dashrightarrow \psi]$ is defined; that is, $\sigma \not\models \neg\phi$. It follows from $(*_1)$ that $s_{\sigma[Sup(\phi)][\psi]} = s_{\sigma[\phi][\psi]}$. Thus, $\sigma \models \phi \to \psi$ iff $\sigma \models \phi \dashrightarrow \psi$.

Let $\sigma^{\downharpoonleft *} = \langle s_\sigma, * \rangle$. That is, $\sigma^{\downharpoonleft *}$ is the context which replaces the revision operation of σ with $*$. Observe that where $\phi, \psi \in$L$_0$, for any $*$, $\sigma \models \phi \to \psi$ iff $\sigma^{\downharpoonleft *} \models \phi \to \psi$. That is, whether $\phi \to \psi$ is accepted at σ is insensitive to the quasi-revision operation of the context. Putting the two observations together, it follows that, for all $*$: $\sigma \models \phi \to \psi$ iff $\sigma^{\downharpoonleft *} \models \phi \dashrightarrow \psi$.

Suppose, for induction, that where $\phi \in \mathsf{L}_0$, $\psi \in \mathsf{L}_{\dashrightarrow}$, for all $*$: $\sigma \models \phi \to \psi$ iff $\sigma^{\downharpoonleft *} \models \phi \dashrightarrow \psi$. Suppose $\chi \in \mathsf{L}_0$. We will show that $\sigma \models \chi \to (\phi \dashrightarrow \psi)$ iff $\sigma \models \chi \dashrightarrow (\phi \dashrightarrow \psi)$ (if defined).

Let σ be an arbitrary context such that $\sigma[\chi \dashrightarrow (\phi \dashrightarrow \psi)]$ is defined. Thus, $\sigma \not\models \neg\phi$. It follows from $(*_1)$ that $s_{\sigma[\chi]} = s_{\sigma[Sup(\chi)]}$. So $\sigma[Sup(\chi)] = (\sigma[\chi])^{\downharpoonleft *_{|[\![\phi]\!]}}$.

Since, for all $*$, $\sigma \models \phi \to \psi$ iff $\sigma^{\downharpoonleft *} \models \phi \to \psi$ it follows that $\sigma[\chi] \models \phi \to \psi$ iff $(\sigma[\chi])^{\downharpoonleft *_{|[\![\phi]\!]}} \models \phi \dashrightarrow \psi$. But trivially, $\sigma[\chi] \models \phi \to \psi$ iff $\sigma[\chi] \models \phi \dashrightarrow \psi$. Thus $\sigma[\chi] \models \phi \dashrightarrow \psi$ iff $\sigma[Sup(\chi)] \models \phi \dashrightarrow \psi$. So

$\sigma \models \chi \rightarrow (\phi \dashrightarrow \psi)$ iff $\sigma \models \chi \dashrightarrow (\phi \dashrightarrow \psi)$. QED. $\square$.

Next, we prove that Import/Export is valid given the iterated revision constraints on proper contexts.

Fact 4.

Given the requirement that, for all proper σ, $*_\sigma$ satisfies $(*_5) - (*_7)$):

$$Sup(\phi \wedge \psi) \rightarrow \chi \;=\!\!|\!\!\models Sup(\phi) \rightarrow (Sup(\psi) \rightarrow \chi)\ .$$

Proof. Observe that $Sup(\phi \wedge \psi) \rightarrow \chi =\!\!|\!\!\models Sup(\phi) \rightarrow (Sup(\psi) \rightarrow \chi)$ iff for all proper σ: $\sigma[Sup(\phi \wedge \psi)] = \sigma[Sup(\phi)][Sup(\psi)]$. That is, sequential supposition of a pair of sentences is has the same effect on context as supposition of their conjunction. For an arbitrary choice of context $\sigma = \langle s, * \rangle$, $\sigma[Sup(\phi \wedge \psi)] = \langle s * (\llbracket\phi\rrbracket \wedge \llbracket\psi\rrbracket), *_{|\llbracket\phi\rrbracket \wedge \llbracket\psi\rrbracket} \rangle$. In comparison, $\sigma[Sup(\phi)][Sup(\psi)] = \langle s * \llbracket\phi\rrbracket) *_{|\llbracket\phi\rrbracket} \llbracket\psi\rrbracket,\ *_{|\llbracket\phi\rrbracket|\llbracket\psi\rrbracket} \rangle$. First, note that by the definition of update on quasi-revision operations, we know that $*_{\llbracket\phi \wedge \psi\rrbracket} = *_{|\llbracket\phi\rrbracket|\llbracket\psi\rrbracket}$. That is, sequential update of a revision operation with two information states has the same effect as update with the meet of the two states.

Next, note that $(s * \llbracket\phi\rrbracket) *_{|\llbracket\phi\rrbracket} \llbracket\psi\rrbracket = (s * \llbracket\phi\rrbracket) * (\llbracket\phi\rrbracket \wedge \llbracket\psi\rrbracket)$. Yet, by $(*_5)$, $(s * \llbracket\phi\rrbracket) * (\llbracket\phi\rrbracket \wedge \llbracket\psi\rrbracket)$ is the same as $s * \llbracket\phi\rrbracket \wedge \llbracket\psi\rrbracket$.

Putting the two together, it follows that $\sigma[Sup(\phi \wedge \psi)] = \sigma[Sup(\phi)][Sup(\psi)]$. Yet σ was arbitrary. So $Sup(\phi \wedge \psi) \rightarrow \chi =\!\!|\!\!\models Sup(\phi) \rightarrow (Sup(\psi) \rightarrow \chi)$. QED. $\square$.

Finally, we aim to show that that (Pres) and (DA) are valid under supposition, and that (CT) is valid under supposition given the iterated revision constraints.

We start by establishing **Fact 3**: that subjunctives behave like the dynamic strict conditional under supposition.

Fact 3 1.

If $\phi_i, ..., \phi_j \models \psi \rightarrow \chi$, then $Sup(\phi_i), ..., Sup(\phi_j) \models \psi \rightsquigarrow \chi$.

(Where $\phi_i, ..., \phi_j, \psi, \chi \in \mathsf{L_0}$.)

Proof. First, observe that $Sup(\phi_i), ..., Sup(\phi_j) \models \psi \rightsquigarrow \chi$ iff for every proper context σ, sequential update with the premises followed by supposition of the antecedent returns a context which accepts the consequent; that is, $\sigma[Sup(\phi_i)], ..., [Sup(\phi_j)][Sup(\psi)] \models \chi$.

We know, by ($*_2$), that for any σ, the supposing a series of claims returns an information state which is at least as strong as the meet of their static content. That is, $s_{\sigma[Sup(\phi_i)],...,[Sup(\phi_j)][Sup(\psi)]} \in [\![\phi_i]\!] \wedge ... \wedge [\![\phi_j]\!] \wedge [\![\psi]\!]$.

Yet, note that if $\phi_i, ..., \phi_j \models \psi \to \chi$, then for any information state $s \in [\![\phi_i]\!] \wedge ... \wedge [\![\phi_j]\!] \wedge [\![\psi]\!]$, we can conclude that $s \in [\![\chi]\!]$. Afterall, let σ be a context pairing s with an arbitrary revision operation. Since the inference is valid, we know that $s_{\sigma[\phi_i],...,[\phi_j][\psi]} = s_{\sigma[\phi_i],...,[\phi_j][\psi][\chi]}$. But $s_{\sigma[\phi_i],...,[\phi_j][\psi]}$ just is s, since it is at least as strong as the informational content of each sentence. So, we can conclude that $s \in [\![\chi]\!]$.

Putting this together, it follows that, for any σ, $s_{\sigma[Sup(\phi_i)],...,[Sup(\phi_j)][Sup(\psi)]} \in [\![\chi]\!]$. So, since each of the claims belongs to L_0, supposing $\phi_i, ..., \phi_j, \psi$ will return an information state which accepts χ. Or, equivalently, $Sup(\phi_i), ..., Sup(\phi_j) \models \psi \rightsquigarrow \chi$. QED. $\square$.

Fact 5.

(i)-(ii) are valid. (iii) is valid given the quasi-revision operations of proper contexts satisfy ($*_5$)-($*_7$):

$$\begin{array}{llcl}
(i) & Sup(\phi) & \models & Sup(\psi) \to (\phi \wedge \psi) \\
(ii) & Sup(\phi \vee \psi) & \models & Sup(\neg\phi) \to \psi; \\
(iii) & Sup(\phi) \to (Sup(\psi) \to \chi), Sup(\psi) & \models & Sup(\phi) \to \chi
\end{array}$$

Proof. (i)-(ii) follow immediately from **Fact 3**.

Turning to (iii), observe that $Sup(\phi) \to (Sup(\psi) \to \chi), Sup(\psi) \models Sup(\phi) \to \chi$ iff for all proper σ, if $\sigma \models Sup(\phi) \to (Sup(\psi) \to \chi)$, then $\sigma[Sup(\psi)][Sup(\phi)] \models \chi$. For an arbitrary proper $\sigma = \langle s, * \rangle$, where $*$ satisfies ($*_5$)-($*_7$), suppose that $\sigma \models Sup(\phi) \to (Sup(\psi) \to \chi)$. Then $\sigma[Sup(\phi)][Sup(\psi)] \models \chi$.

So, it suffices to demonstrate that $\sigma[Sup(\phi)][Sup(\psi)] = \sigma[Sup(\psi)][Sup(\phi)]$.

First, we know that $*_{\sigma[Sup(\phi)][Sup(\psi)]} = *_{|[\![\phi]\!]|[\![\psi]\!]}$ and $*_{\sigma[Sup(\psi)][Sup(\phi)]} = *_{|[\![\psi]\!]|[\![\phi]\!]}$. By the definition of $*_{|[\![\phi]\!]}$, we know that $*_{|s|s'} = *_{|s'|s}$. So $*_{\sigma[Sup(\phi)][Sup(\psi)]} = *_{\sigma[Sup(\psi)][Sup(\phi)]}$.

Next, note that $s_{\sigma[Sup(\phi)][Sup(\psi)]} = s_{\sigma[Sup(\phi)]} *_{|[\![\phi]\!]} [\![\psi]\!]$. Correspondingly, $s_{\sigma[Sup(\psi)][Sup(\phi)]} = s_{\sigma[Sup(\psi)]} *_{|[\![\psi]\!]} [\![\phi]\!]$. But we know, by the definition of $*_{|s}$ that $s_{\sigma[Sup(\phi)]} *_{|[\![\phi]\!]}, [\![\psi]\!]) = (s_\sigma * [\![\phi]\!]) * ([\![\phi]\!] \wedge [\![\psi]\!])$. Equally, we know that $s_{\sigma[Sup(\psi)]} *_{|[\![\psi]\!]}, [\![\phi]\!]) = (s_\sigma * [\![\psi]\!]) * ([\![\psi]\!] \wedge [\![\phi]\!])$. But, by $(*_5)$, $(s * s') * s' \wedge s'' = s * (s' \wedge s'') = (s * s'') * (s'' \wedge s')$. So $s_{\sigma[Sup(\phi)][Sup(\psi)]} = s_{\sigma[Sup(\psi)][Sup(\phi)]}$

Putting the two together, it follows that $\sigma[Sup(\phi)][Sup(\psi)] = s_{\sigma[Sup(\psi)][Sup(\phi)]}$. But σ was arbitrary. So $Sup(\phi) \to (Sup(\psi) \to \chi), Sup(\psi) \models Sup(\phi) \to \chi$. QED. $\square$.

Appendix B

This appendix shows how to extend the framework to allow for embedding of arbitrarily sentences containing modals and conditionals under Sup. As discussed in §2.5.3, we introduce modals via stipulative definition, so that $\Box\phi \equiv_{def} \top \Rightarrow \phi$ and $\Diamond\phi \equiv_{def} \neg(\top \Rightarrow \neg\phi)$. $\Box\phi$ denotes the test which passes at σ iff σ accepts ϕ. $\Diamond$ is the dual of $\Box$; $\Diamond\phi$ denotes the test that passes at σ iff σ does not accept $\neg\phi$.

We define a generalized revision operation. A revision operation takes a pair of information states to a new information state (cf. **Def.3**). A generalized revision operation derived from it takes an information state and a set of information states to a new information state.

Definition 19.

Where $*$ is a quasi-revision operation, $\circledast$ is a derivative, generalized quasi-revision operation iff it satisfies ($\circledast_1$)-($\circledast_2$), for any set of information states S:

($\circledast_1$) $\quad s \circledast S \in \{s * s' |\ s' \in S\}$;

($\circledast_2$) $\quad$ For all $s' \in S$: if $s \circledast S \leq s * s'$, then $s \circledast S = s * s'$.

The first constraint says that generalized revision should be an instance of revision: revising with a set of information states returns the result of revising with some state in that set. The second constraint says that generalized revision should be minimal: revising with a set of information states should not return a state strictly stronger than the result of revising with some member of that set.

We define an update operation on generalized revision.

Definition 20.

$s \circledast_{|S} S' = s \circledast \{s \wedge s' | s \in S, s' \in S'\}$

$\circledast_{|S}$ is the generalized revision operation just like $\circledast$ but which preserves the information in S. The $\circledast_{|S}$-revision of s with S' just returns the $\circledast$-revision of s with the set containing the meet of each state in S with each state in S'.

We make two modifications to the framework introduced in Chapter 2. First, extend the language to allow for embedding of arbitrarily complex modal sentences under Sup.

Definition 21.

We define a pair of languages, L_0^+ and L_1^+ simultaneously.

L_0^+:

- $\{\mathsf{A}, \mathsf{B}, ..., \bot, \top\} \subseteq \mathsf{L}_0^+$;
- If $\phi, \psi \in \mathsf{L}_0^+$, then $\neg\phi, \phi \wedge \psi \in \mathsf{L}_0^+$;
- If $\phi \in \mathsf{L}_0^+ \cup \mathsf{L}_1^+$ and $\psi \in L_0^+$, then $\phi \to \psi \in \mathsf{L}_0^+$.

L_1^+:

- If $\phi \in \mathsf{L}_0^+$, then $Sup(\phi) \in \mathsf{L}_1^+$.

Let $\mathsf{L}^+ = \mathsf{L}_0^+ \cup \mathsf{L}_1^+$. Unlike L, L^+ permits arbitrarily complex embeddings of modal expressions under Sup.

Second, we enrich the objects playing the role of points of evaluation.

Definition 22 (Enriched Contexts).

An enriched context π is a pair $\langle s_\pi, \circledast_\pi \rangle$, such that:

- s_π is an information state;
- $\circledast_\pi$ is a generalized quasi-revision operation.

Where $\phi \in L_0^+$, we define update as above.

Definition 23.

$$\begin{array}{llll} (i) & \pi[\mathsf{A}] & = & \langle s_\pi + [\![\mathsf{A}]\!], \circledast_\pi \rangle \\ (ii) & \pi[\neg\phi] & = & \langle s_\pi / s_{\pi[\phi]}, \circledast_\pi \rangle \\ (iii) & \pi[\phi \wedge \psi] & = & \pi[\phi][\psi] \\ (iv) & \pi[\phi \rightarrow \psi] & = & \begin{cases} \pi, & \textit{if } \pi[\phi] \models \psi; \\ \langle \emptyset, \circledast_\pi \rangle, & \textit{otherwise.} \end{cases} \end{array}$$

For any generalized quasi-revision operation $\circledast$, $\|\cdot\|^{\circledast}$ is the function mapping expressions to the set of information states that accept them when paired with $\circledast$. That is:

$$\|\phi\|^{\circledast} = \{s |\ \langle s, \circledast \rangle \models \phi\}$$

Call $\|\phi\|^{\circledast}$ the informational content of ϕ at $\circledast$. The informational content of a boolean expression is simply the downset of its informational content. That is:

$$\|\phi\|^{\circledast} = \{s \in \mathbb{S} |\ s \leq [\![\phi]\!]\} \qquad \textit{if} \quad \phi \in \mathsf{L_0}.$$

However, where ϕ is non-boolean, the informational content of ϕ may not correspond to a downset at all.

We define supposition in terms of generalized revision with the informational content of the expression supposed. For example, $\| \Diamond\phi \|^{\circledast} = \{s | \exists s' \in \| \phi \|^{\circledast} \colon s \geq s'\}$. That is, the informational content of $\Diamond\phi$ will contain all the information states for which it is possible to find a stronger state within the informational content of ϕ.

Definition 24.

$$\pi[Sup(\phi)] = \langle s_\pi \circledast \| \phi \|^{\circledast_\pi}, \circledast_{\pi|(\|\phi\|^{\circledast_\pi})} \rangle$$

That is, supposing ϕ has a dual effect on a context π: first, it revises the information state of the context with the informational content of ϕ at $\circledast_\pi$. Second, it updates the generalized quasirevision operation of the input context with the informational content of ϕ at $\circledast_\pi$.

Our enriched framework is a conservative extension of the old framework. That is, update with $Sup(\phi)$ has the same effect as long as ϕ belongs to the boolean fragment of the language.

Fact 6.

Where σ is a context and π an enriched context such that (i) $s_\sigma = s_\pi$ and (ii) $\circledast_\pi$ is a generalized quasi-revision operation derived from $*_\sigma$:

$$s_{\sigma[Sup(\phi)]} = s_{\pi[Sup(\phi)]} \qquad\qquad \textit{if } \phi \in L_0$$

Proof: If $\phi \in \mathsf{L}_0$, then $\| \phi \|^{\circledast_\pi} = \downarrow [\![\phi]\!]$, the downset of the the static content of ϕ. The static content of s_π is either stronger than s_π, weaker than it, or incomparable. Suppose that $s_\pi \leq [\![\phi]\!]$. By $(*_1)$ and $(\circledast_2)$, $s_\pi \circledast_\pi \| \phi \|^{\circledast_\pi} = s_\pi = s_\sigma * [\![\phi]\!]$. Suppose that $s_\pi \geq [\![\phi]\!]$. By $(*_1)$ and $(\circledast_2)$, $s_\pi \circledast_\pi \| \phi \|^{\circledast_\pi} = s_\pi \wedge [\![\phi]\!] = s_\sigma * [\![\phi]\!]$. Suppose, finally that s_π and $[\![\phi]\!]$ are incomparable. By $(*_4)$, for every $s' < [\![\phi]\!]$, $s_\pi * s' \leq s_\pi * [\![\phi]\!]$. So $s_\pi \circledast_\pi \| \phi \|^{\circledast_\pi} = s_\sigma * [\![\phi]\!]$. But observe that $s_{\sigma[Sup(\phi)]} = s_\sigma * [\![\phi]\!]$ and $\pi[Sup(\phi)] = s_\pi \circledast_\pi$. QED.

Our enriched framework also predicts that sequential supposition of an epistemic contradiction returns the absurd state.

Fact 7.

(i) $\quad Sup(\neg\phi \wedge \Diamond\phi) \models \bot$

(ii) $\quad Sup(\Diamond\phi \wedge \neg\phi) \models \bot$

(iii) $\quad Sup(\neg\phi), Sup(\Diamond\phi) \models \bot$

(iv) $\quad Sup(\Diamond\phi), Sup(\wedge\neg\phi) \models \bot$

Proof. (i)-(ii) follow immediately from the observation that $\| \neg\phi \wedge \Diamond\phi \|^{\circledast} = \| \Diamond\phi \wedge \neg\phi \|^{\circledast} = \{\emptyset\}$. Tat is, epistemic contradictions are accepted only at the absurd state. (iii)-(iv) follow from the observation that $\{s \wedge s' | \; s \in \parallel \neg\phi \parallel^{\circledast}, s' \in \parallel \Diamond\phi \parallel^{\circledast}\}$ is the singleton of the absurd state. For any s, $s \circledast_{|\|\Diamond\phi\|^{\circledast}} \parallel \neg\phi \parallel^{\circledast} = s \circledast_{|\|\neg\phi\|^{\circledast}} \parallel \Diamond\phi \parallel^{\circledast}$. Both are equivalent, by the definition of updated generalised revision operations, to $s \circledast \{s \wedge s' | \; s \in \parallel \neg\phi \parallel^{\circledast}, s' \in \parallel \Diamond\phi \parallel^{\circledast}\}$.

It follows directly from **Fact 7** that each of the discourses in Chapter 1, §2.3 are absurd in the following sense: the information state at which the (nested) consequent of the conclusion is evaluated is guaranteed to be empty. Accordingly, the indicative instances are guaranteed to be underfined. We can generalize this to a uniform prediction of infelicity by adopting the proposal in §2.5.2 that $\phi \rightsquigarrow \psi$ is defined at a context iff update with $Sup(\phi)$ does not return an absurd context.

Chapter 3

Collapse

3.1 Indicatives & Subjunctives

Consider again the following pair of conditionals (repeated from the Introduction):

(59) If the butler did it, he used the candlestick.

(60) If the butler had done it, he would've used the candlestick.

According to orthodoxy, (59) and (60) differ in meaning.[1] As we observed above, one way the difference between the two is manifest involves the contexts at which they are licensed (§1.2.2). Unlike subjunctives, indicatives are unacceptable in counterfactual environments—contexts in which their antecedent has been ruled out. Thus, while (16) constitutes an acceptable bit of discourse, (15) does not.

(15) The butler didn't do it. ??If he did, he used the candlestick.

(16) The butler didn't do it. If he had, he would've used the candlestick.

[1]For classic discussion, see, e.g., Adams (1965), Lewis (1973), Stalnaker (1975), Slote (1978), Davis (1979), and Gibbard (1981), amongst others.

At this point, it will be helpful to introduce some general terminology. A discourse context, c, is the kind of real-world environment in which a conversation takes place. The previous chapter proposed (partially) modeling discourse contexts in terms of a pair of set-theoretic objects: pairs comprising an information state and a transition rule. In this chapter, we abstract away from the particular details of that model, exploring the relationship between indicatives and subjunctives in a more neutral setting.

We can talk about the presuppositions of various expressions using the notion of **truth-in-a-context**.[2] Where ϕ is true-in-c, we'll write $c \models \phi$.[3] We'll say that ϕ is licensed at c iff its presuppositions are true there.

As argued above, there is then a simple (and seemingly popular) story to be told about the behavior in (15)-(16). Indicatives presuppose their antecedent to be epistemically possible in the context in which they are used.

Indicative Licensing: $\phi \dashrightarrow \psi$ is licensed at c only if $c \models \Diamond\phi$.

Indicative Licensing has been defended by Stalnaker (1975), Karttunen and Peters (1979), von Fintel (1997*b*), Gillies (2009, forthcoming) and Starr (2014*a*, ms) amongst others. On the assumption that $\Diamond\phi$ is true-in-c only if ϕ hasn't been ruled out at c, Indicative Licensing coincides with the proposal in §2.4.1 and explains the infelicity of (15).

Subjunctives are standardly assumed to be licensed in both counterfactual and non-counterfactual environments.[4,5] §2.4.2 followed this assumption, attributing them trivial presuppositions which are

[2] The notion of truth-in-a-context is originally due to Kaplan (1989). I employ it here without any particular commitment to contexts being the kinds of things Kaplan says contexts are. In particular, if the truth-in-a-context of an expression containing one or more epistemic modals is sensitive to some body of information, then contexts will need to be the kind of thing which can determine one of those.

[3] I do not intend to insist that the *truth* part of *truth-in-a-context* be taken too seriously. Some approaches to the semantics of conditionals and modals distinguish between a technical and a philosophical notion of truth (see, e.g., Yalcin (2011)). In broad terms, the former models the positive evaluative status of a sentence in a context, while the latter characterizes the accuracy of those sentences belonging to a factual fragment of the language. Insofar as this distinction is taken seriously, it is the former which is important for present purposes.

[4] Note that, in the sense used here, for c to be a counterfactual context for ϕ does not merely require that ϕ is false in c; rather, it requires that $\Box\neg\phi$ is true in c.

[5] Lakoff (1970) and Portner (1992) defend the stronger condition that subjunctives presuppose their antecedent is epistemically impossible. They consequently require an explanation to of the apparently felicitous non-counterfactual uses.

everywhere satisfied. As such, it adhered to von Fintel (1997*b*)'s observation that indicatives and subjunctives are in non-complementary distribution—i.e., that there exist contexts at which both an indicative and its corresponding subjunctive are licensed.[6] These contexts will simply be those at which the two conditionals' common antecedent is epistemically possible.

If indicatives and subjunctives are in non-complementary distribution, a number of interesting issues can be raised. In particular, we can non-trivially ask: what entailment relations hold between the two types of conditional at those contexts which license both?

This chapter takes up that question. I will argue that under certain apparently well-motivated assumptions about the logic of indicatives and subjunctives, they are equivalent in the appropriate sense. Where licensed, each entails the other. Defending this principle, which I will call COLLAPSE, is the primary goal of the chapter.

COLLAPSE $\phi \dashrightarrow \psi \mathrel{=\!\!|\!\models} \phi \rightsquigarrow \psi.$

§3.2 introduces three inference patterns, each of which has substantial appeal. When taken together, however, it is shown that they give rise to COLLAPSE. §3.3 considers the support for these three principles in greater detail, and looks at some motivations which can be offered for rejecting them. After arguing that there is no way of avoiding COLLAPSE without incurring costs, §3.4 considers the philosophical implications of COLLAPSE, seeking to show that the principle is simultaneously more and less radical than it at first seems. Finally, §3.5 aims demonstrate that the reasons others have wanted to deny COLLAPSE can be accounted for by appeal to an independently plausible pragmatic principle and a popular story about how context changes occur during discourse.

[6] The *locus classicus* here, due to Anderson (1951), concerns the things doctors tend to say about potential arsenic poisoning victims. The same point can be made by the observation that (†) is a perfectly fine piece of discourse.

(†) Maybe the butler did it. If he had done, he'd have used the candlestick.

See §3.5.3 for further discussion.

3.1.1 Strawson Entailment

Since we are dealing with expressions some of which have non-trivial presuppositions, we want a notion of entailment which takes this into account. In particular, we need to know how to treat contexts at which one (or more) of the premises/conclusion is unlicensed when determining whether an inference pattern is valid.

Following the strategy of Chapter 2, we will reason about inference patterns in terms of Strawson Entailment (Strawson (1952), von Fintel (1997*b*, 1999)). Informally, the idea is that in evaluating a particular inference we should consider all and only those contexts at which both the premises and the conclusion are licensed. It is valid iff, within this restricted domain, there is no context at which the premises are true but the conclusion false.

Strawson Entailment

$\Gamma \models \phi$ iff in all contexts at which (the elements of) Γ and ϕ are licensed, ϕ is true if $\bigwedge \Gamma$ is.

In thinking about Strawson entailment, it will be useful to have a function, δ, which maps an expression to the set of its presuppositions—the expressions which must be true-in-a-context for it to be licensed there. Thus, according to Indicative Licensing, $\Diamond\phi \in \delta(\phi \dashrightarrow \psi)$. Where Γ is a set of sentences, we will adopt the notational convention that $\delta(\Gamma) = \bigcup_{\phi\in\Gamma} \delta(\phi)$. It is then easy to define Strawson entailment in terms of classical entailment and δ:

Fact 8.

Γ Strawson entails ϕ iff $\Gamma \cup \delta(\Gamma \cup \{\phi\})$ classically entails ϕ.[7]

Importantly, some inference rules which are valid for classical entailment fail for Strawson entailment. In particular, Strawson entailment does not vindicate Cut (Smiley (1967), see also Cariani and Goldstein (forthcoming)).

Cut If $\Gamma \models \phi$ and $\Delta, \phi \models \psi$, then $\Delta, \Gamma \models \psi$.

[7] Note that on the assumption that an expression is licensed in c only if the presuppositions of its presuppositions are true-in-c, then $\psi \in \delta(\phi)$ implies $\delta(\psi) \subseteq \delta(\phi)$. That is, δ is identical to its own transitive closure.

To see why, consider the case in which presuppositions of ϕ are not entailed by the presuppositions of (the elements of) $\Gamma \cup \Delta \cup \{\psi\}$. Then the contexts considered in evaluating whether Γ Strawson entails ϕ and in evaluating whether Δ, ϕ Strawson entail χ, will be a strict subset of those considered in evaluating whether Δ, Γ Strawson entail ψ. Accordingly, that the former two inferences are Strawson valid does not guarantee that there is no context at which the elements of Δ, Γ are licensed and true, yet ψ is licensed but false.

Transitivity of Entailment (that is, if $\Gamma \models \phi$ and $\phi \models \psi$, then $\Gamma \models \psi$) is a limiting instance of Cut. It also fails in its full generality, for the same reasons. However, Strawson entailment does preserve Cut (and, hence, the Transitivity of Entailment) in a restricted form. The rule is valid in the special case in which the presuppositions of (the elements of) Γ, Δ, and ψ, along with the former two themselves, are at least as strong as the presuppositions of ϕ. Call the resulting principle Strawson Cut.

Strawson Cut Suppose that $\delta(\Gamma \cup \Delta \cup \{\psi\}), \Gamma, \Delta \models \bigwedge \delta(\phi)$.
Then, if $\Gamma \models \phi$ and $\Delta, \phi \models \psi$, then $\Delta, \Gamma \models \psi$.

Strawson Cut says that, if Γ and Δ, along with the presuppositions of (the elements of) $\Gamma \cup \Delta \cup \{\psi\}$, entail each of the presuppositions of ϕ, then the relevant instance of Cut will be Strawson validity preserving. It is this restricted rule which we will rely on below.

In addition to invalidating some classical inference rules, Strawson entailment also validates some novel rules. Of particular importance in what follows is the rule I will call Reduction.

Reduction If $\Gamma, \phi \models \psi$ and $\phi, \delta(\phi) \in \delta(\Gamma \cup \{\psi\})$, then $\Gamma \models \psi$.

Reduction says that if Γ and ϕ Strawson entail ψ, but ϕ and its presuppositions are presuppositions of ψ and (the elements of) Γ, then Γ Strawson entails ψ by itself. To see why, note that since ϕ and its presuppositions are amongst the presuppositions of ψ and (the elements of) Γ, in evaluating the latter entailment we will restrict our attention to only those contexts in which ϕ is true and licensed. But it is established that in all such contexts if the elements of Γ are true and licensed, then ψ is true, if licensed.

The relationship that holds between the presuppositions of complex expressions and the presuppositions of their parts is subject to complex and unresolved questions (for discussion see, e.g., Karttunen (1973, 1974), Heim (1983, 1990),Geurts (1999), Beaver (1993, 2001), Schlenker (2007, 2008, 2009), Rothschild (2008)). However, for present purposes, difficult cases can be set aside. Instead, we can restrict our attention to matters on which there is a large degree of consensus. I follow Karttunen (1973, 1974) in adopting three assumptions: First, that negation and modals are transparent to presuppositions. That is, $\delta(\Diamond\phi) = \delta(\neg\phi) = \delta(\phi)$. Second, that conjunctions inherit the presuppositions of their left-hand conjunct, along with the presuppositions of the right-hand conjunct conditional on the left.[8] That is, $\delta(\phi \wedge \psi) = \delta(\phi) \cup \{\phi \supset \chi \mid \chi \in \delta(\psi)\}$. Finally, that the presuppositions of a conditional include at least the presuppositions of its antecedent along with the presuppositions of its consequent conditional on its antecedent. That is, $\delta(\phi \dashrightarrow \psi) \subseteq \delta(\phi) \cup \{\phi \supset \chi \mid \chi \in \delta(\psi)\}$ (*mutatis mutandis* for subjunctives).[9] Wherever these assumptions play a role in the argument below, the role they play will be noted.

3.2 Constructing Collapse

3.2.1 If/And

First, consider the following pair of inference patterns:

If/And	i. $\Diamond\phi, \phi \dashrightarrow \Diamond\psi \models \Diamond(\phi \wedge \psi)$	Indicative
	ii. $\Diamond\phi, \phi \rightsquigarrow \Diamond\psi \models \Diamond(\phi \wedge \psi)$	Subjunctive

If/And says that, given the epistemic possibility of its antecedent, a $\Diamond$-conditional entails the epistemic possibility of its antecedent and consequent conjoined. Both indicative and subjunctive variants of this principle look to be in good standing.

Take the indicative case first. An individual who argues from (61.a) to (61.c) reasons impeccably (indeed, to the point of sounding boring).

[8] For recent discussion, see Chemla and Schlenker (2012), and Mandelkern et al. (2017).

[9] For subjunctives, it is plausible that we also want the stronger condition: $\delta(\phi \dashrightarrow \psi) \subseteq \delta(\phi) \cup \{\phi \rightsquigarrow \chi \mid \chi \in \delta(\psi)\}$.

(61) a. Maybe the butler did it.

b. If he did it, maybe he used the candlestick.

c. So, maybe the butler did it using the candlestick.

It is hard to see how an inference of this form could fail. Along with arguments like (61.a-c), indicative IF/AND draws support from the oddity of accepting its premises while simultaneously accepting the negation of its conclusion. Since neither '*maybe*' nor conditionals embed happily under sentential negation, we can see this most easily by considering other downward monotonic environments that do embed them (such as, e.g., the scope of '*no-one*').

(62) a. Anyone might buy a lottery ticket.

b. Anyone might win the lottery, if they buy a ticket.

c. No-one might buy a winning lottery ticket.

An individual who accepts (62.a-c) has reasoned sub-optimally (indeed, to the point of sounding unintelligible). Turning to the subjunctive case, analogous considerations appear to mitigate equally strongly in its favor.

(63) a. Maybe the butler did it.

b. If he'd done it, maybe he'd have used the candlestick.

c. So, maybe the butler did it using the candlestick.

The reasoning in (63.a-c) seems just as good as its indicative counterpart. Likewise, (64.a-c) seem no less inconsistent.

(64) a. Anyone might buy a lottery ticket.

b. Anyone might win the lottery, if they were to buy a ticket.

c. No-one might buy a winning lottery ticket.

These observations are not new. The validity of If/And has been previously advocated in Gillies (2010, forthcoming) (for indicatives) and Gillies (2007) and Goldstein (2018) (for subjunctives), though cf. Ciardelli (2020) for recent dissent.

Note that I will take it for granted that the examples above involve a consequent embedded modal (rather than syntactically less plausible wide-scoping). There is good reason to think that, in this position, '*might*' admits both an epistemic and circumstantial reading (see, e.g., Lewis (1973, 1986), Stalnaker (1981, 1984), DeRose (1994, 1999), Bennett (2003) and Asher and McCready (2007), for discussion). However, the latter reading of (62.b)-(64.b) would appear to be highly unnatural, if it is available at all. Moreover, there is clearly no circumstantial reading of '*maybe*' when it occurs in the same position (e.g., (61.b)-(63.b)).[10]

3.2.2 And/If.

Next, consider the following further pair of inference patterns:

And/If i. $\Diamond(\phi \wedge \psi) \models \phi \dashrightarrow \Diamond\psi$ Indicative

ii. $\Diamond(\phi \wedge \psi) \models \phi \rightsquigarrow \Diamond\psi$ Subjunctive

And/If says that a $\Diamond$-conditional is entailed by the epistemic possibility of its antecedent and consequent conjoined. The two variants are not quite the converses of their If/And counterparts, since the latter included the epistemic possibility of the antecedent as a premise. However, this

[10] I take the strongest argument in favor of the availability of an epistemic reading of '*might*'-subjunctives to be following (closely related to arguments proposed by Stalnaker (1981) and DeRose (1994, 1999)).

Suppose I am unsure whether it is raining. I'm convinced that if it is, then had we gone to the park it would have been miserable. I'm also convinced that if it isn't, then had we gone to the park, it wouldn't have been miserable. Nevertheless, as long as I don't know, it seems that I'm in a position to assert (‡):

(‡) If we had gone to the park, it might have been miserable, but then again, it might not have.

Insofar as, upon learning that the weather was fine, the assertability of (‡) decreases dramatically, the modals in the consequent display precisely the sensitivity to speaker's information we would expect if they were epistemic in flavor.

Note also that, insofar as '*maybe*' only permits epistemic readings, the difficulty in hearing (∗.a-b) as ascribing anything but an unusually inconsistent state of mind to an otherwise renowned detective is evidence in favor of the robustness of the epistemic reading of 'might':

(∗) a. Sherlock thinks that if the butler had done it, maybe he'd have used the revolver.
b. However, he doubts that if the butler had done it, he might have used the revolver.

difference is superficial. The possibility of the antecedent is entailed by the possibility of its conjunction with the consequent.

And/If also looks to be in good standing, for both variants. Running the same tests, an individual who argues from (65.a) to either (65.b) or (65.b′) reasons just as impeccably:

(65) a. Maybe the butler did it using the candlestick ...

b. So, if he did it, maybe he used the candlestick.

b′. So, if he had done it, maybe he'd have used the candlestick.

Similarly, an individual who accepts either (66.a-b) or (66.a-b′) reasons just as sub-optimally:

(66) a. Anyone might buy a winning lottery ticket.

b. No-one might win the lottery, if they buy a ticket.

b′. No-one might win the lottery, if they were to buy a ticket.

Again, the observation that these inferences appear valid is not new. Gillies (forthcoming) endorses the indicative variant of And/If explicitly (and Gillies (2007) the subjunctive, implicitly). Likewise, both will constitute a reasonable inference in the framework of Stalnaker (1975). However, what has gone unnoticed is that the pairs of If/And and And/If inferences come perilously close to triggering Collapse by themselves.

3.2.3 Collapse & Quasi-Collapse.

Call the principle that $\Diamond$-indicatives and $\Diamond$-subjunctives are equivalent Quasi-Collapse:

Quasi-Collapse $\phi \dashrightarrow \Diamond\psi \dashv\vDash \phi \leadsto \Diamond\phi$

Each direction of Quasi-Collapse can be derived from one of the indicative/subjunctive variants of If/And along with the corresponding subjunctive/indicative variant of And/If.

Fact 9.

If/And and And/If imply Quasi-Collapse.

The proof of **Fact 9** involves a pair of distinctive properties of Strawson entailment, so it may be instructive to work through it in a little detail.

Take the left-to-right direction first. By indicative If/And, we know that $\Diamond\phi, \phi \dashrightarrow \Diamond\psi \models \Diamond(\phi \wedge \psi)$. By subjunctive And/If, we also know that $\Diamond(\phi \wedge \psi) \models \phi \rightsquigarrow \Diamond\psi$. Yet $\Diamond(\phi \wedge \psi)$ introduces no new presuppositions of its own.[11] Thus, by Strawson Cut, it follows that $\Diamond\phi, \phi \dashrightarrow \Diamond\psi \models \phi \rightsquigarrow \Diamond\psi$. Finally, by Indicative Licensing, $\Diamond\phi$ is a presupposition of $\phi \dashrightarrow \Diamond\psi$. Thus, by Reduction, it follows that $\phi \dashrightarrow \Diamond\psi \models \phi \rightsquigarrow \Diamond\psi$.

The two crucial steps in the above reasoning employ Strawson Cut and Reduction. The former permits us to move from the observation that $\Diamond\phi$ and $\phi \dashrightarrow \Diamond\psi$ imply $\Diamond(\phi \wedge \psi)$ and that $\Diamond(\phi \wedge \psi)$ implies $\phi \rightsquigarrow \Diamond\psi$, to the conclusion that $\Diamond\phi$ and $\phi \dashrightarrow \Diamond\psi$ imply $\phi \rightsquigarrow \Diamond\psi$. It is valid, as we noted in §3.1.1, in virtue of the fact that any context at which $\Diamond\phi$ and $\phi \dashrightarrow \Diamond\psi$ are licensed and true will be a context at which $\Diamond(\phi \wedge \psi)$ is licensed.[12]

The latter permits us to move from the observation that $\Diamond\phi$ and $\phi \dashrightarrow \Diamond\psi$ imply $\phi \rightsquigarrow \Diamond\psi$ to the conclusion that $\phi \dashrightarrow \Diamond\psi$ implies $\phi \rightsquigarrow \Diamond\psi$ by itself. It is valid, as we also noted in §3.1.1, in virtue of the fact that any context at which $\phi \dashrightarrow \Diamond\psi$ is licensed will be a context at which $\Diamond\phi$ is licensed and true.[13] Thus, the latter is redundant as a premise, since in evaluating whether the inference is Strawson valid, our attention will be restricted to contexts in which it is licensed and true already. Note that, given the definition of Strawson entailment, it is irrelevant to the validity of Reduction whether $\phi \dashrightarrow \Diamond\psi$ occurs as one of the premises or the conclusion, since in either case, we restrict our attention to cases in which it is licensed. Equivalent reasoning, *mutatis mutandis*, is sufficient to demonstrate that the right-to-left direction, $\phi \rightsquigarrow \Diamond\psi \models \phi \dashrightarrow \Diamond\psi$, follows from subjunctive If/And and indicative And/If.[14] Thus anyone who accepts both variants of the two principles (as, I've been

[11]Since, by assumption, $\delta(\Diamond(\phi \wedge \psi)) = \delta(\phi) \cup \{\phi \supset \chi | \chi \in \delta(\psi)\}$ and $\delta(\phi \dashrightarrow \Diamond\psi) \subseteq \delta(\phi) \cup \{\phi \supset \chi | \chi \in \delta(\psi)\}$)

[12]In fact, the requirements of **Strawson Cut** are weaker than this, as we saw. Nevertheless, this formulation suffices for our purposes.

[13]Since, by assumption, $\delta(\phi) \cup \{\Diamond\phi\} \subseteq \delta(\phi \dashrightarrow \psi)$ and $\delta(\Diamond\phi) = \delta(\phi)$ (see §3.1.1).

[14]NB: in the final step of the right-to-left direction of proof, Reduction will employ the fact that the conclusion has

suggesting, they should) is committed to QUASI-COLLAPSE.

COLLAPSE $\quad \phi \dashrightarrow \psi \mathrel{=}\!\mid\!\models \phi \rightsquigarrow \psi.$

QUASI-COLLAPSE and COLLAPSE (repeated above) are distinct. However, they can be shown to entail one another, given the following (pair of) principle(s), taken by many to have substantial appeal:

DUALITY	(*i.*)	$\phi \dashrightarrow \Diamond\psi$	$\mathrel{=}\!\mid\!\models$	$\neg(\phi \dashrightarrow \neg\psi)$	INDICATIVE
	(*ii.*)	$\phi \rightsquigarrow \Diamond\psi$	$\mathrel{=}\!\mid\!\models$	$\neg(\phi \rightsquigarrow \neg\psi)$	SUBJUNCTIVE

DUALITY says that a $\Diamond$-conditional is equivalent to the negation of the contrary 'bare'-conditional. §3.3 provides a survey of arguments in favor of (each variant of) the principle.

Fact 10.

QUASI-COLLAPSE and DUALITY imply COLLAPSE.

Since substitution of equivalents is not a Strawson safe inference, the proof of **Fact 10** is not immediate. Instead, it requires repeated applications of Strawson Cut.

First, consider the left-to-right direction of COLLAPSE. From (right-to-left) subjunctive DUALITY, we know that $\neg(\phi \rightsquigarrow \neg\psi) \models \phi \rightsquigarrow \Diamond\psi$. From (right-to-left) QUASI-COLLAPSE, we know that $\phi \rightsquigarrow \Diamond\psi \models \phi \dashrightarrow \Diamond\phi$. Since the presuppositions of $\neg(\phi \rightsquigarrow \neg\psi)$ are trivially entailed by the presuppositions of $\phi \rightsquigarrow \Diamond\psi$, it follows by Strawson Cut that $\neg(\phi \rightsquigarrow \neg\psi) \models \phi \dashrightarrow \Diamond\phi$.[15]

But, from (left-to-right) indicative DUALITY, we know that $\phi \dashrightarrow \Diamond\psi \models \neg(\phi \dashrightarrow \neg\psi)$. Since the presuppositions of $\neg(\phi \dashrightarrow \neg\psi)$ entail the presuppositions of $\phi \dashrightarrow \Diamond\psi$, it follows by a second application of Strawson Cut that $\neg(\phi \rightsquigarrow \neg\psi) \models \neg(\phi \dashrightarrow \neg\psi)$. Finally, by contraposition (which is a Strawson safe inference pattern) we can derive $\phi \dashrightarrow \neg\psi \models \phi \rightsquigarrow \neg\psi$, which is (equivalent to) the left-to-right direction of COLLAPSE.

Equivalent reasoning, *mutatis mutandis*, is sufficient to derive the right-to-left direction. Combining

$\Diamond\phi$ as a presupposition, rather than the premise.

[15]Note that it is important, at this step, that $\neg$ and $\Diamond$ are transparent to presuppositions (see §3.1.1).

Facts 9 and 10, anyone who accepts If/And, And/If and Duality in their indicative and subjunctive variants is committed to Collapse.

This is a striking result. Collapse looks, at first glance, to be in tension with the claim that indicatives and subjunctives differ in meaning. In §3.5, I show that this tension is, in fact, merely apparent. Collapse is compatible with orthodoxy. Indeed, I will suggest, the principle fits naturally into an appealing picture of the semantic and pragmatic differences between indicatives and subjunctives, one of which the view defended in Chapters 1-2 is an instance. The primary goal of §3.5 is to provide an account of divergent judgments about indicatives and subjunctives which is consistent with Collapse.

Before taking up that task, however, §3.3 examines in greater detail the status of the principles from which Collapse is derived. In particular, it investigates their relationship to two other principles, both of which are subjects a large body of exegetical discussion: Contraposition and CEM.

3.3 The Sources of Collapse

3.3.1 Contraposition

Consider the pair of inference patterns below:

Contraposition	i.	$\phi \dashrightarrow \psi \quad \dashv\models \quad \neg\psi \dashrightarrow \neg\phi$		Indicative
	ii.	$\phi \rightsquigarrow \psi \quad \dashv\models \quad \neg\psi \rightsquigarrow \neg\phi$		Subjunctive

Contraposition says that a conditional is equivalent to the converse of its inverse. Both indicative and subjunctive variants of the principle are related, albeit slightly differently, to If/And and And/If. §3.3.1.1 considers the subjunctive variant; §3.3.1.2 the indicative.

3.3.1.1 Subjunctive Contraposition

Contraposition is widely taken to be bad for subjunctives (Goodman (1947), Lewis (1973), Stalnaker (1968), Skyrms (1974) Slote (1978), Kratzer (2008) amongst others for classical discussion;

though cf. Urbach (1988), Hunter (1993), von Fintel (1997*a*), for dissent).

Pairs like (67.a-b) are one source of concern for the subjunctive variant of the principle.[16] A speaker who asserted (67.a) could reasonably go on to deny (67.b). Yet the latter is just the contrapositive of the former.

(67) a. If it were to rain, it wouldn't pour.

b. If it were to pour, it wouldn't rain.

The failure of the inference from (67.a) to (67.b) might, at first glance, give us reason to re-evaluate If/And, And/If, and Duality. After all, the subjunctive variants of the principles jointly entail Quasi-Contraposition (proof in footnote):[17]

Quasi-Contraposition $\Diamond\neg\psi, \phi \rightsquigarrow \psi \models \neg\psi \rightsquigarrow \neg\phi$ Subjunctive

Contraposition and Quasi-Contraposition are similar. It would be reasonable to wonder: do supposed counterexamples to the former like the one above also provide grounds for rejecting latter?

They do not. As noted by von Fintel (1997*a*), someone who accepted that it might pour could not coherently go on to assert that it wouldn't, if it were to rain. The acceptability of (67.a) is dependent upon raining-and-pouring possibilities already being ruled out.

Pairs like (68.a-b) represent a different source of concern regarding subjunctive Contraposition. Here is some backstory: Cathy is considering leaving Manhattan to take a West Coast job. Nick will go wherever Cathy goes. In contrast, whether Nick chooses to follow her is no factor in Cathy's decision making.

(68) a. If Cathy were to move to the West Coast, Nick would leave NY.

[16] Adapted from Adams (1975, 15).

[17] The proof is simple. From If/And and And/If and the commutativity of $\wedge$, we know that $\Diamond\psi, \psi \rightsquigarrow \Diamond\phi \models \Diamond(\psi \wedge \phi)$ and $\Diamond(\psi \wedge \phi) \models \phi \rightsquigarrow \Diamond\psi$. So, by Strawson Cut, $\Diamond\psi, \psi \rightsquigarrow \Diamond\phi \models \phi \rightsquigarrow \Diamond\psi$. Since, from Duality we know that $\phi \rightsquigarrow \Diamond\psi \mathrel{=\!\mid\!\models} \neg(\phi \rightsquigarrow \neg\psi)$, it follows by two applications of Strawson Cut that $\Diamond\psi, \neg(\psi \rightsquigarrow \neg\phi) \models \neg(\phi \rightsquigarrow \neg\psi)$. Finally, by two applications of $\neg$-Introduction/Elimination (which is Strawson safe), we can conclude that $\Diamond\psi, \phi \rightsquigarrow \neg\psi \models \psi \rightsquigarrow \neg\phi$ (which is equivalent to Quasi-Contraposition).

b. If Nick were to stay in NY, Cathy wouldn't move to the West Coast.

Anyone appraised of the facts should accept (68.a). Yet, it seems, they could reasonably be far more reluctant to accept (68.b).

Pairs like (68.a-b) are just as much of problem for QUASI-CONTRAPOSITION as they are for CONTRAPOSITION. The first conditional of the pair can be asserted in contexts in which it is not ruled out that Nick will stay in NY (in virtue of it not being ruled out that Cathy will not move to the West Coast). But the second remains notably marked in such contexts.

Unlike (67.a-b), the non-equivalence of contrapositives in this case appears to result from a not-at-issue implication that (the eventuality denoted by) the consequent is causally dependent on (the eventuality denoted by) the antecedent. Reversing the negated antecedent and consequent generates a different (and, in this case, false) implication.[18]

According to a popular story, this implication has something to do with the interaction of discourse structure, tense and aspect (e.g., Kamp (1979), Moens and Steedman (1988), Lascarides and Asher (1993), Kehler (2000, 2002), Asher and Lascarides (2003), Lepore and Stone (2014), Althshuler (2016)).[19] For our purposes, what is primarily important is that its effect can be controlled for by the insertion of material explicitly specifying the direction of causal dependence.

(68) b′. If Nick were to stay in NY, it would be because Cathy didn't move to the West Coast.

Any implication that Cathy moving to the West Coast causally depends upon Nick staying in NY is clearly canceled in (68.b′). Correspondingly, the inference from (68.a) to (68.b′) sounds substantially

[18] A similar implication can be seen in non-conditional environments. (§.a) does not appear interchangeable with (§.b) in all discourse contexts, for related reasons.

(§) a. Either Cathy won't move to the West Coast, or Nick will leave NY.
b. Either Nick will leave NY, or Cathy won't move to the West Coast.

[19] For contrast, it appears easier to obtain a reading of (¶.b) on which it follows from (¶.a):

(¶) a. If Cathy were to have moved to the West Coast, Nick would have left NY.
b. If Nick were to have stayed in NY, Cathy would not have moved to the West Coast.

That is, the addition of perfective aspect appears to increase the accessibility of the reading which lacks the implication of causal dependence.

better than the inference to (68.b).

Given this sensitivity to cancellable implications, I want to follow von Fintel (1997*a*, 39) and Gillies (2010, fn19) in suggesting that we should not treat pairs like (68.a-b) as counterexamples to subjunctive QUASI-CONTRAPOSITION.Rather, the examples demonstrate that in evaluating the validity of inference patterns for the conditional, we need to control for not-at-issue implications generated by the interaction of tense, aspect, and discourse structure, among other sources (for recent discussion of this issue see, in particular, (Stojnić, 2017; Stojnić, forthcoming)).

3.3.1.2 Indicative Contraposition

Contraposition is also widely taken to be bad for indicatives (Stalnaker (1968),[20] Appiah (1985), Starr (2014*a*)). A speaker who asserted (69.a) could reasonably go on to deny (69.b):

(69) a. If it rains, it won't pour.

b. If it pours, it won't rain.

However, unlike the subjunctive case, IF/AND, AND/IF and DUALITY are sufficient to establish full (indicative) CONTRAPOSITION. The proof begins in the same way as the proof of subjunctive QUASI-COLLAPSE in fn 17. However, indicatives presuppose the possibility of their antecedent. Accordingly, Reduction allows us to make a further step from $\Diamond\neg\psi, \phi \dashrightarrow \psi \models \neg\psi \dashrightarrow \neg\phi$ to indicative CONTRAPOSITION, since $\Diamond\neg\psi$ is a presupposition of the conclusion.[21] As a result, the challenge in the indicative case is harder to overcome. Merely noting that the first element of the pair is unassertable unless raining-and-pouring possibilities have already been ruled out is, unlike in the subjunctive case, insufficient to resolve the problem.

This is, however, not the final word. There is a well-known story involving shifting contexts which offers to explain the apparent non-equivalence of (69.a-b) while preserving indicative CONTRAPOSITION

[20] Stalnaker (1975) is a more complicated case, since contraposition is invalidated, but has the status of being a reasonable inference.

[21] Note that it suffices to establish one direction of contraposition, since, given $\neg\neg$-Elimination, each direction entails the other.

(its proponents include Warmbrod (1983), McCawley (1996) and Gillies (2004)).[22] Different versions of this story are possible. Here is one.

First, note that (69.a) can be asserted only if all raining-and-pouring possibilities have been ruled out. A speaker can't coherently assert that it won't pour, if it rains, while accepting that it might pour.[23] However, (69.b) can be denied only if some raining-and-pouring possibilities have not been ruled out. A speaker can't coherently assert that it will rain if it pours, while proceeding to deny that it might rain.[24] Wherever the first member of the pair can be coherently asserted, the second member of the pair will be unlicensed.

At this point, the story appeals to the ability of the context to shift in response to unsatisfied presuppositions. Where an utterance of ϕ would be licensed only if the context had some property (which it in fact lacks), we frequently observe that an utterance of ϕ brings about a change in the context so that it has that very property it needs (Lewis, 1979*b*). In line with this pattern, the argument goes, denial of (69.b) at any context at which (69.a) could be coherently asserted will trigger a shift to a new context at which some raining-and-pouring possibilities are no longer ruled out. Yet, since (69.b) is false where licensed, it can be coherently denied at the new context.

This proposal is appealing. In its favor, note that the purported counter-examples to indicative CONTRAPOSITION appears subject to precisely the kind of order effect which would be predicted. That is, an assertion that if it pours it will rain, followed by an assertion that if it rains it won't pour, appears worse, taken as a single piece of discourse, than its converse.

If the explanation succeeds, the ability of a speaker to assert (69.a) while proceeding to deny (69.b) does not constitute a counterexample to indicative CONTRAPOSITION. It is dependent upon a shift in context between the two utterances. At the prior context, (69.b) is unlicensed; at the posterior

[22]A subjunctive variant of the apologetic, originating in the apocrypha from a 1994 handout of Heim's, is developed in von Fintel (2001) and Gillies (2007). For dissent, see Moss (2012), Nichols (2017), Boylan and Schultheis (2017) and Lewis (2018).

[23]More generally, it seems we should accept: if $\psi \models \phi$, then $\phi \dashrightarrow \neg\psi \models \Box\neg\psi$.

[24]Note that the equivalence of the assertion of that it'll rain if it pours to the denial that it won't rain if it pours amounts merely to the assumption of WEAK BOETHIUS' THESIS (McCall (1966), Pizzi (1977), Pizzi and Williamson (1997)):

WEAK BOETHIUS' THESIS $\quad \phi \dashrightarrow \psi \models \neg(\phi \dashrightarrow \neg\psi)$

context, both are false. Hence, once this shift is controlled for, it turns out that there is no challenge to the Strawson equivalence of the two.

3.3.2 Duality & CEM

Duality plays an important role in the derivation of Collapse. However, it is the subject of a doctrinal schism. On one side are the proponents of Duality—most prominently, Lewis (1973) (for subjunctives) and Kratzer (1979, 1986, 2012) (for both). On the other are the proponents of Conditional Excluded Middle (hereafter, CEM)—most prominently, Stalnaker (1968, 1981, 1984) (for both). Since they are in direct opposition, evaluating the argument of §3.2 requires evaluating both principles. The aim here is not to resolve the schism, but rather to survey the arguments available to each side, and to clarify how they interact with orthodoxy regarding conditional meaning.

3.3.2.1 Duality

First, consider Duality (repeated below).

Duality	i.	$\phi \dashrightarrow \Diamond\psi$	$=\!\mid\models$	$\neg(\phi \dashrightarrow \neg\psi)$	Indicative
	ii.	$\phi \rightsquigarrow \Diamond\psi$	$=\!\mid\models$	$\neg(\phi \rightsquigarrow \neg\psi)$	Subjunctive

There are compelling arguments in favor of each variant. In the indicative case, Duality follows from (*i*) the equivalence of bare and '*must*' indicatives and (*ii*) the duality of '*must*' and '*maybe*' (proof in footnote 25).[25]

If/Must	$\phi \dashrightarrow \psi$	$=\!\mid\models$	$\phi \dashrightarrow \Box\psi$
Must/Might	$\Box\phi$	$=\!\mid\models$	$\neg\Diamond\neg\psi$

The former draws motivation from the incoherence of accepting one but not both of pairs like (70.a-b):

(70) a. If the murder occurred in the library, the vicar is innocent.

[25] For the left-to-right direction, note that by Weak Boethius' Thesis $\phi \dashrightarrow \Diamond\psi \models \neg(\phi \dashrightarrow \neg\Diamond\psi)$. Yet, by Must/Might, $\neg(\phi \dashrightarrow \neg\Diamond\psi) \models \neg(\phi \dashrightarrow \Box\neg\psi)$. So, by If/Must, $\phi \dashrightarrow \Diamond\psi \models \neg(\phi \dashrightarrow \neg\psi)$.

For the right-to-left direction, note that by If/Must, $\neg(\phi \dashrightarrow \neg\psi) \models \neg(\phi \dashrightarrow \Box\neg\psi)$. By Must/Might $\neg(\phi \dashrightarrow \Box\neg\psi) \models \neg(\phi \dashrightarrow \neg\Diamond\psi)$. So, by Weak Boethius' Thesis $\neg(\phi \dashrightarrow \neg\psi) \models \phi \dashrightarrow \Diamond\psi$.

b. If the murder occurred in the library, the vicar must be innocent.

The latter draws motivation from the incoherence of accepting one but not both of pairs like (71.a-b):

(71) a. The vicar must be innocent.

b. The vicar can't be guilty.

To the extent that denying these equivalences is taken to be heretical, the indicative variant should be taken to be orthodoxy.[26]

In the subjunctive case, note that we are reluctant to accept either conditional in pairs like (72.a-b), asserted of an unflipped fair coin (cf. Lewis (1973, 82)):[27]

(72) a. If the coin had been flipped, it'd have landed heads.

b. If the coin had been flipped, it'd have landed tails.

Our reluctance to accept either bare conditional is matched by a corresponding inclination to accept each of the contrary $\Diamond$-conditionals:

(73) a. If the coin had been flipped, maybe it'd have landed tails.

b. If the coin had been flipped, maybe it'd have landed heads.

Subjunctive Duality provides a simple explanation of this pattern of judgments. Our inclination to accept (73.a-b) explains our reluctance to accept (72.a-b), respectively, since under the principle they are pairwise inconsistent (NB: a similar abductive argument can be run in favor of the indicative variant).[28]

[26]Note that, on a view such as that of Kratzer (1979, 1986, 2012), on which all apparently bare conditionals contain a covert necessity modal in their consequent, an equivalent argument can be run for subjunctives.

[27]A closely related point can be made with pairs such as (ʃ.a-b)(attr., originally, to Quine)

(ʃ) a. If Bizet and Verdi had been compatriots, they would have been French.

b. If Bizet and Verdi had been compatriots, they would have been Italian.

[28]This argument, which takes the form of an inference to the best explanation, is hardly decisive. For discussion, see, e.g. Stalnaker (1981), DeRose (1994, 1999), Bennett (2003) Williams (2010), and Shaffer (2016).

Finally, we can test both variants by considering the equivalence of sentences in embedded environments (Goldstein (2018)). For example, consider (74)-(75):

(74) a. Sherlock doubts that if the murder occurred in the library, the vicar is guilty.

b. Sherlock thinks that if the murder occurred in the library, the vicar may be innocent.

(75) a. Sherlock doubts that if the murder had occurred in the library, the vicar would be guilty.

b. Sherlock thinks that if the murder had occurred in the library, the vicar may have been innocent.

To the extent that (74.a-b) and (75.a-b) appear equivalent, denying indicative or subjunctive Duality (respectively) will be unappealing.

3.3.2.2 CEM

Next, consider CEM:

CEM i. $\models \phi \dashrightarrow \psi \vee \phi \dashrightarrow \neg\psi$ Indicative

ii. $\models \phi \rightsquigarrow \psi \vee \phi \rightsquigarrow \neg\psi$ Subjunctive

CEM says that either a bare conditional is true, or else its contrary is. Each variant is inconsistent with the corresponding variant of Duality, on the assumption that $\Diamond$-conditionals do not entail their bare analogues.

The proof is simple. Take the indicative case. Suppose that $\phi \dashrightarrow \Diamond\psi$. Then, from Duality, it follows that $\neg(\phi \dashrightarrow \neg\psi)$. So, by disjunctive syllogism, we can derive $\phi \dashrightarrow \psi$ from CEM. Equivalent reasoning, *mutatis mutandis*, applies in the subjunctive case.

There are also compelling arguments in favor of each variant of CEM. First, note that pairs like (76.a-b) appear equivalent. Similarly, in the indicative case, for pairs like (77.a-b).

(76) a. No suspect will be found guilty if they provide an alibi.

b. Every suspect will be found innocent, if they provide an alibi.

(77) a. No suspect would be found guilty if they provided an alibi.

b. Every suspect would be found innocent, if they provided an alibi.

CEM provides a simple explanation of these judgments (von Fintel and Iatridou (2002), Higginbotham (2003), Klinedinst (2011)). According to the principle, the conditionals in the scope of the quantifiers are contradictories.[29] Hence, on the assumption that $No(\phi)(\psi)$ and $Every(\phi)(\neg\psi)$ are equivalent, each variant of CEM predicts the equivalence of the corresponding pair of quantified sentences.[30] Yet, to the extent that this argument provides motivation for adopting (one or both variants of) CEM, it provides motivation for rejecting (one or both variants of) Duality.

Second, adherents of CEM have appealed to judgments about the probabilities of conditionals such as (72.a-b) and (73.a-b) to argue against Duality (Stalnaker (1981), DeRose (1994), Edgington (2008) Schulz (2014), Dorr and Hawthorne (manuscript)). Supposing the coin in question is known to be fair, it appears as likely as not that it would have landed heads, if it had flipped. That is, it is reasonable to take (72.a) to have a probability of .5 (and likewise, by the symmetry of the case, (72.b)). However, it also appears certain that the coin might have landed tails, had it been flipped. That is, it is reasonable to take (73.a) to have a probability of 1 (likewise (73.b)). Yet, according to subjunctive Duality, this is not a coherent assignment of probabilities. Since (72.a) and (73.a) are contradictories, their probabilities must sum to 1.[31]

The judgments here appear compelling. They also generalize easily. A similar argument can be run against indicative Duality, assuming a context in which it is unknown whether the coin was flipped, and if so, how it landed.[32]

---

[29] Or, at least, they are on the assumption of Weak Boethius' Thesis.

[30] Things are not entirely simple. Schismatics on the CEM side have tended to claim that conditionals like (72.a-b) and their indicative counterparts are indeterminate. Nevertheless, as Leslie (2009, fn3) observes, we are inclined to deny both (⨍.a-b) (along with future-oriented indicative analogues):

(⨍) a. No coin would land heads if it were flipped.

b. Every coin would land tails, if it were flipped.

Yet, if the embedded conditionals are indeterminate, for each fair coin, we would predict each sentence to be indeterminate itself (cf. Klinedinst (2011)).

[31] Dorr and Hawthorne (manuscript) defend a related argument, noting that since (72.a-b) are inconsistent (by Weak Boethius' Thesis), the probability of their disjunction should equal the sum of their probabilities, i.e., in this case, 1. Yet their disjunction is simply an instance of CEM. Dorr and Hawthorne proceed to argue, convincingly, that if one accepts CEM in this case, one should be willing to accept it in any.

[32] Options are available to the proponent of Duality, particularly in the case of indicatives. For example, Rothschild

Endorsing CEM does not come without costs, however. For example, since it requires rejection of the relevant variant of DUALITY, defending indicative/subjunctive CEM commits one to the acceptability of some conjunctions with the form of (78)/(79), respectively:

(78) ?? If the butler did it, he didn't use the candlestick, but if he did it, maybe he used it.

(79) ?? If the butler had done it, he wouldn't have used the candlestick, but if he had done it, maybe he'd have used it.

To the extent that such conjunctions sound marked, CEM will evoke suspicion. Its proponents have a number of things which they can say in its defense. Some depend on the syntactically questionable stipulation that '*maybe*' takes exceptional scope over the right-hand conjunct (Stalnaker (1981), DeRose (1991, 1994, 1999)). Others, such as Mandelkern (2019), have proposed that the modal can be left *in situ*, appealing instead to constraints on licensing relative to the local context to explain the infelicity.

While Mandelkern's approach successfully predicts that (78) and (79) will be false where licensed, it does so only at the cost of generating other, less desirable, predictions. In particular, it predicts that both $\phi \rightsquigarrow \psi \vee \phi \rightsquigarrow \Box\neg\psi$ and $\phi \dashrightarrow \psi \vee \phi \dashrightarrow \Box\neg\psi$ will be Strawson validities (i.e., true wherever licensed). Yet counterinstances to the indicative variants, at least, appear easy to find.[33]

(80) Either the coin landed heads if flipped or it must have landed tails if flipped.

Where the coin was flipped and landed tails but this was unknown in the context of utterance, (80) will, intuitively, be licensed but false. The outcome of the coin toss and speaker's knowledge rule appear to rule out the left and right disjuncts, respectively, even when their local contexts are taken into account.

Finally, note that for those hoping to avoid the implications of argument in §3.2, rejecting DUALITY in favor of CEM is insufficient. First, recall that DUALITY was needed only in the derivation of

(2014) and Starr (ms) develop related accounts which would allow them to account for these judgments.

[33] The subjunctive variants are harder to evaluate, given the difficulty of embedding unambiguously epistemic necessity modals in subjunctive consequents.

Collapse. Giving up the former provides no help when it comes to avoiding Quasi-Collapse. But Quasi-Collapse seems almost as startling as Collapse.[34]

Second, note that many of the arguments in favor of CEM are directed exclusively at its subjunctive variant, remaining neutral on indicative Duality (e.g., Stalnaker (1981), DeRose (1994, 1999), Williams (2010), Schulz (2014)). And indeed, insofar as its right-to-left direction permits us to derive the falsity of a claim about what would have happened from the truth of a claim about what it is epistemic possible would have happened, subjunctive Duality does appear to warrant more skepticism than its indicative variant.[35] However, while adherents to subjunctive CEM will reject the problematic left-to-right direction of subjunctive Duality, they typically retain its right-to-left direction (which can be innocently accepted in conjunction with CEM).[36] Yet the right-to-left direction of subjunctive Duality is enough, given the left-to-right direction of indicative Duality, to derive the left-to-right (i.e., indicative-to-subjunctive) direction of Collapse from Quasi-Collapse.[37] Yet as we shall see in the next section, the left-to-right direction of Collapse is sufficient to generate the most significant philosophical consequences of the principle by itself (see, in particular, footnote 41).

3.4 Collapse Considered

Collapse says that indicatives and subjunctives are Strawson equivalent. To some, this might seem tantamount to denying a central article of conditional faith.

It might, but it shouldn't. Strawson equivalence requires that the truth values of the two conditionals

[34]Note, also, that accepting Quasi-Collapse while denying Collapse involves substantial awkwardness in many CEM-validating approaches. As noted, a standard way combining CEM with an account 'maybe'-embedded consequents is to treat modal as having wide-scope over the entire conditional (e.g., Stalnaker (1981), DeRose (1994, 1999)). This combination of views, however, will require accepting some equivalences of the form $\Diamond\phi \dashv\models \Diamond\psi$ as valid while denying the corresponding equivalences of the form $\phi \dashv\models \psi$. That is, it amounts to committing to the existence of connections between what might be the case which are not dependent on connections between what can be the case. As a result, the proponent of this brand of CEM must accept that there are contexts at which $\phi \wedge \neg\psi$ is true, but deny that there are contexts in which it is ruled out that we are not in such a context.

[35]Thanks to an anonymous referee for discussion on this point.

[36]Indeed, on the assumption that $\phi \rightsquigarrow \psi \models \phi \rightsquigarrow \Diamond\psi$, proponents of subjunctive CEM will be committed to the right-to-left direction of subjunctive Duality.

[37]Proof: Suppose that $\neg(\phi \rightsquigarrow \psi)$. By R-to-L subjunctive Duality, it follows that $\phi \rightsquigarrow \Diamond\neg\psi$. By Quasi-Collapse, it follows that $\phi \dashrightarrow \Diamond\neg\psi$. Yet, by L-to-R indicative Duality, it follows that $\neg(\phi \dashrightarrow \psi)$. Hence, by contraposition, $\phi \dashrightarrow \psi \models \phi \rightsquigarrow \psi$.

coincide at those contexts which license both. However, it allows for substantial differences in which contexts license each. As a result, it allows each conditional to be governed by a substantially different logic.[38]

The conditional articles of faith state that indicatives and subjunctives differ in meaning. This is consistent with their Strawson equivalence.[39] Indeed, one might argue that showing the conditionals to differ in their presuppositions or their logic *amounts* to showing them to differ in meaning. We do not need to defend this stronger claim here, however. What is relevant is that the adherent to Collapse is not thereby committed to the identity of indicatives' and subjunctives' meanings.

It does not, however, follow that Collapse is neutral with respect to the traditional picture of indicatives and subjunctives. While it may be compatible with attributing different meanings to the two forms, Collapse casts doubt on one of the key ways in which those meanings are typically taken to differ.

Indicative conditionals are widely held to be information sensitive (see, e.g., Gibbard (1981), Veltman (1985), Yalcin (2007, 2012) and Kolodny and MacFarlane (2010)). The truth value of an indicative in context appears to depend, in part, on what information that context makes salient. To see this, consider a case with the following structure (which is a symmetric variant of Gibbard (1981, 231)'s original 'Sly Pete' example).[40]

An individual (The GameMaster) places a ball under one of three cups (Red, Blue, Yellow). Two contestants (A, B) must guess under which cup the ball has been placed. Before they do, however, The GameMaster will privately reveal one of the empty cups to each of them. Suppose that The GameMaster places the ball under the Red cup. She reveals to Contestant A that it is not under the Blue cup, and, to Contestant B, that it is not under the Yellow cup. According to widely reported intuition, A could truthfully assert (81) (but not (82)). In contrast, B could truthfully assert (82)

[38]For example, §3.3.1 argued that the indicative, but not the subjunctive, must satisfy Contraposition.

[39]The more general point, that, for an appropriate entailment relation, equivalence between two sentences does not imply sameness of meaning is familiar in the logic of conditionals (cf., in particular, Stalnaker (1975), von Fintel (2001), Gillies (2009), Cariani and Goldstein (forthcoming) and Mandelkern (manuscript)).

[40]For extended discussion of cases of with this structure (both symmetric and asymmetric), see in particular Stalnaker (1984), Lycan (2001), Bennett (2003), Rothschild (2015), Dorr and Hawthorne (manuscript) and Goldstein (2020).

(but not (81)):

(81) If the ball is not under Red, then it is under Yellow.

(82) If the ball is not under Red, then it is under Blue.

Yet the only apparent difference between A and B's contexts of utterance is the body of information they make salient. Presumably, A's information is salient in the former, whereas B's is salient in the latter.

Subjunctives are standardly taken to be information insensitive. A common way to motivate this is to note that in normal contexts the truth of (83)-(84), unlike their indicative variants, appears to depend entirely on the dispositions of The GameMaster—it is not sensitive to what the contestants know.

(83) If the the ball hadn't been under Red, the it would have been under Yellow.

(84) If the the ball hadn't been under Red, the it would have been under Blue.

Yet according to Collapse, the truth-values of corresponding indicatives and subjunctives coincide at contexts which licenses both. Accordingly, if indicatives are information sensitive in such contexts, subjunctives must be too.[41]

We can make the same point in another, less neutral, way. An apparent symptom of the information sensitivity of indicatives is the equivalence of (85)-(86) and (81)-(82):[42]

(85) The ball must either be under Red or Yellow.

(86) The ball must either be under Red or Blue.

[41]Indeed, the left-to-right (i.e., indicative-to-subjunctive) direction of Collapse is sufficient to motivate the information sensitivity of subjunctives by itself. (83)-(84) are contrary conditionals, given the context—at most one will be true. By left-to-right Collapse, which one it is will depend on which of (81) or (82) is true. If (81) is true, then (83) will be too (and, so, (84) will be false). Similarly, if (82) is true, then (84) will be too (and (83) will be false). Yet, (81)-(82) are information sensitive. So, (83)-(84) will be information sensitive too.

[42]Closely related to this observation is Stalnaker (1975)'s proposal that the inference between an indicative and the corresponding material conditional is a reasonable inference, in his sense; if one can be felicitously asserted in a context, then the other will be accepted in the context that results from that assertion. It has gone largely unnoticed that, on Stalnaker's own proposal, the inference from $\phi \dashrightarrow \psi$ to $\neg\phi \vee \psi$ is not itself a reasonable inference.

Someone who denied (85) could not coherently accept (81). And, equally, someone who accepted (85) could not coherently deny (81). This motivates STRICTNESS (endorsed by, e.g., Warmbrod (1983), Veltman (1985), Dekker (1993), von Fintel (1999), Gillies (2004, 2009), Yalcin (2007), Starr (2014*a*, ms), and Holguin (forthcoming), amongst others):

STRICTNESS $\Box(\neg\phi \vee \psi) \dashv\models \phi \dashrightarrow \psi$

STRICTNESS says that indicative conditionals are Strawson equivalent to the epistemic necessity of the corresponding material conditional. Yet together, COLLAPSE and STRICTNESS imply EPISTEMICITY.[43]

EPISTEMICITY $\Diamond\phi \wedge (\phi \rightsquigarrow \psi) \dashv\models \Diamond\phi \wedge \Box(\neg\phi \vee \psi)$

EPISTEMICITY says that, in contexts in which its antecedent is epistemically possible, a subjunctive is equivalent to the epistemic necessity of the corresponding material conditional. Epistemic necessity claims are uncontroversially information sensitive.[44] So, given STRICTNESS, COLLAPSE implies that, in non-counterfactual contexts, subjunctives are information sensitive, too.

While this conflict with the traditional picture, it is not without precedent. Others have observed that subjunctives can sometimes permit information sensitive readings which are equivalent to their indicative counterparts (see, in particular, Edgington (2007, 211) and Yalcin (manuscript, 9)).

For example suppose Contestant A guesses that the ball is under Yellow and Contestant B that it is under Blue. After the ball is revealed to be under Red, each contestant could justify her guess along the lines of (87), *mutatis mutandis*. And, equally, a third party could rationalize each contestant along the lines of (88):

(87) Ah well—I had a 50% chance of guessing correctly: if it hadn't been under Red, it would have been under [Yellow/Blue].

[43] Proof: By the right-to-left direction of COLLAPSE, we know that $\Diamond\phi \wedge (\phi \rightsquigarrow \psi) \models \phi \dashrightarrow \psi$. By the right-to-left direction of STRICTNESS, Indicative Licensing, and Reduction we also know that $\phi \dashrightarrow \psi \models \Diamond\phi \wedge \Box(\neg\phi \vee \psi)$. Yet $\Diamond\phi \wedge (\phi \rightsquigarrow \psi) \models \bigwedge \delta(\phi \dashrightarrow \psi)$. So, by Strawson Cut, we can conclude that $\Diamond\phi \wedge (\phi \rightsquigarrow \psi) \models \Diamond\phi \wedge \Box(\neg\phi \vee \psi)$. Equivalent reasoning, with the left-to-right directions of each principle establishes the right-to-left direction of EPISTEMICITY.

[44] See, e.g., Hacking (1967), DeRose (1991), Egan et al. (2005), and von Fintel and Gillies (2007, 2010) for classic discussion of precisely what information they are sensitive too.

(88) Contestant [A/B]'s guess wasn't so bad. After all, she knew that if it hadn't been under Red, it would have been under [Yellow/Blue].

Here, both (87) and (88) ascribe past possession of the information that the contestant would have expressed with the corresponding indicative, prior to learning the location of the ball. Similarly, Khoo (2017) has recently argued for the availability an information sensitive reading of subjunctives on the basis of assumptions about the contribution of indicative and subjunctive mood.

Nevertheless, the mere availability of an information sensitive reading of subjunctives is insufficient to fully address the concerns raised by COLLAPSE. We need to explain why, in contexts which license both, subjunctives frequently permit an information insensitive reading that is not available for the indicative. And we also need to explain why, in the same contexts, the information sensitive reading of the indicative is frequently unavailable for the subjunctive. I turn to this issue in the following (and final) section.

3.5 Collapse in Context

3.5.1 Adams Pairs

In many (non-counterfactual) contexts, judgments about corresponding indicatives and subjunctives diverge. Call instance of this phenomenon 'Adams' pairs (following Adams (1970, 1975)).[45]

Here is one example: Sherlock is conducting interviews with guests who were present on the night of the murder. The butler and the vicar are by no means friends. Nevertheless, the vicar attests that he was with the butler throughout the period when the murder took place and that he (the butler) is not guilty. Given this testimony, (89.a-b) constitute an Adams pair in the context.

(89) a. If the butler did it, the vicar covered for him.

[45] The *locus classicus* here involves an observation about differences in the level of paranoia required for one to accept certain indicatives vs. subjunctives about the death of JFK. Since many readers will take Adams original pair to be counterfactuals (at least, prior to accommodation of the indicative's presuppositions), I avoid it in favor of a clearer case.

b. If the butler had done it, the vicar would have covered for him.

(89.a) appears true (as uttered by Sherlock, at least). Intuitively, it reports Sherlock's information that either the butler is innocent or the vicar lied. In contrast, (89.b) appears false (or at least uncertain). Intuitively, rather than reporting Sherlock's information, it makes a (dubious) claim about the dispositions of the vicar. Both conditionals are licensed in the context at which they are evaluated. Hence, it seems we have a counter-example to COLLAPSE.[46]

Our judgments about (89.a-b) are robust. But they are not quite conclusive. COLLAPSE requires the status of indicatives and subjunctives to coincide at any context which licenses both. At contexts which do not license both, it imposes no constraints. If, prior to evaluating one member of the pair, hearers are required to modify the common ground of the context so that it no longer licenses the other, then despite appearances, our judgments will not correspond to a counter-instance to COLLAPSE.

In fact, there is reason to think that this is precisely what occurs. As Shanon (1976) and von Fintel (2004) observe, the availability of 'Hold up/Hey, wait a minute!'-responses provides a test for the accommodation of not-at-issue material.

(90) A: The Colonel's wife was in the drawing room.

B: Hey, wait a minute! I didn't know the Colonel had a wife!

(91) A: The Colonel has a wife and she was in the drawing room.

B: ?? Hey, wait a minute! I didn't know the Colonel had a wife!

B's response in (90) is felicitous, since before evaluating A's utterance hearers must first accommodate the not-at-issue implication that the Colonel has a wife. In contrast, A's utterance in (91) requires

[46] Similar cases can be constructed in which it is even clearer that the antecedent is not ruled out. For example, I show you a coin and tell you that if I flip it and it lands heads, I'll let you know. If I don't flip or I flip it and it lands tails, I'll tell you nothing. Having heard nothing from me, it seems you should accept (ɤ.a).

(ɤ) a. If the coin was flipped, it landed tails.

b. If the coin had been flipped, it would have landed tails.

However, you definitely should not accept (ɤ.b).

no such accommodation, leading B's response to be decidedly odd.

The subjunctive members of Adams pairs pass the 'Hold up/Hey, wait a minute!'-test. In response to an utterance of (89.b) in its specified context, a hearer could reasonably object 'Hey, wait a minute! We can't rule out that the butler *did* do it yet!'.[47] In contrast, no such response is available to its indicative variant. This suggests that (89.b)—unlike (89.a)—triggers a not-at-issue implication in context that the butler must be innocent. If, prior to evaluating it, hearers accommodate this material, then the subjunctive will be assessed in a different context to the indicative. Accordingly, there will be no reason to expect that judgments about the two will coincide.

If an assertion of the subjunctive member of an Adams pair carries a not-at-issue implication that its antecedent is ruled out in context, this implication cannot take the form of a presupposition. First, such a presupposition would be incompatible with the observation that subjunctives permit non-counterfactual uses (see Anderson (1951), Stalnaker (1975) and von Fintel (1997*b*), along with §3.5.3 for discussion). Second, as we will shortly see, the implication appears defeasible—in appropriate discourse contexts, it is capable of being cancelled. Yet the presuppositions of (unembedded) sentences are standardly taken to be uncancellable (Karttunen (1971, 63), Gazdar (1979), Abbott (2006), Simons (2013), Abrusán (2016)).

Accordingly, it seems more plausible that it arises via some form of pragmatic mechanism.[48] This idea is developed below. I argue that, in fact, there is a simple and widely accepted pragmatic principle which immediately predicts the not-at-issue implications of the subjunctives.

[47]Note that the availability of this response is fragile. In particular, it is blocked in cases where the subjunctive is employed as part of a argument, via *modus tollens*, for the negation of its antecedent. This conforms to a more general rule that 'Hold up/Hey, wait a minute!'-responses are illicit in cases in which the speaker is explicitly engaged in an argument in favor of the relevant not-at-issue material

[48]Indeed, there may also be some empirical grounds for thinking that indicatives and subjunctives differ primarily in virtue of accommodation of the negation of the antecedent. Chinese lacks a dedicated irrealis mood, and indicatives/subjunctives are not morphologically distinguished. Nevertheless, recent research has argued that 'subjunctive' behavior is associated with a subclass of conditionals marked by the presence of lexical items broadly indicating their antecedent is taken to be ruled out (Yong (2013), Jiang (2019*b*,*a*)).

3.5.2 The Fluidity of Context

Differences in the presuppositions of expressions can give rise to corresponding differences in their pragmatic behavior. For instance, the determiners '*All*' and '*both*' are standardly taken to differ only at the level of their presuppositions. The latter, unlike the former, carries a presupposition that its NP complement has exactly two individuals in its denotation.

(92) a. All of the victim's children are suspects.

b. Both of the victim's children are suspects.

This difference in presuppositions is accompanied by two differences at the level of pragmatics. First, use of the former is dispreferred in contexts in which the latter is licensed. That is, if it is common ground that the victim had exactly two children then, unlike (92.b), an utterance of (92.a) will be decidedly odd. Second, and relatedly, use of the former will typically implicate that the presuppositions of the latter are not satisfied. That is, an utterance of (92.a) suggests that the victim has at least three children.[49]

While implementations differ in detail, there is broad consensus on the explanation of these observations, originating with Heim (1991, 515) and Sauerland (2003, 2008).[50] All other things being equal, it is assumed that speakers are under pragmatic pressure to employ utterances with stronger presuppositions. Or, stated a little more carefully:

Maximize Presupposition

If: $(i.)$ $\phi \Dashv\vDash \psi$;

$(ii.)$ $\delta(\phi) \subset \delta(\psi)$;

$(iii.)$ $c \vDash \bigwedge \delta(\psi)$;

then there is a preference for asserting ψ over ϕ in c.

[49]The weaker implication (that the victim does not have exactly 2 children) is blocked either by the presupposition of plurality associated with 'both' or, if 'all' is taken to lack a plurality presupposition, by the availability of the singular definite 'The victim's child is a suspect' which in this case will also have strictly stronger presuppositions than (92.a) (see Heim (1991) and Sauerland et al. (2005) for discussion).

[50]There is room for disagreement over the status of Maximize Presupposition as a pragmatic principle; see Schlenker (2012) and Lauer (2016) for discussion.

Maximize Presupposition says that if ϕ and ψ are Strawson equivalent but the presuppositions of the latter outstrip the presuppositions of the former, then as long as both are licensed, ψ should be favored over ϕ.[51]

Maximize Presupposition directly explains why use of ‘all’ is marked in contexts in which it is common ground that the victim had exactly two children. However, it also explains why, where the common ground is unopinionated about the number of children the victim has, use of ‘all’ carries a not-at-issue implication that the victim had three or more children. Assume ‘both’ and ‘all’ both carry a presupposition of plurality. The presuppositions of (92.b) are strictly stronger than the presuppositions of (92.a) (in virtue of the additional presupposition of duality associated with ‘both’). So, by Maximize Presupposition, if the speaker took the former to be licensed, she would have used it. Since she didn’t, she must assume that the speaker has at least three children.[52] Accordingly, absent objection, this information will be accommodated, leading it to be incorporated into the common ground prior to evaluating her utterance.

Crucially, the same reasoning generalizes directly to the case of conditionals. The presuppositions of subjunctives are a strict subset of the presuppositions of indicatives. Unlike the former, the latter presuppose that their antecedent is epistemically possible. As such, that a speaker uses a subjunctive can be expected to implicate that she takes its antecedent to be epistemically impossible. Absent objections, this information will be accommodated, leading it be incorporated into the common ground prior to evaluating her utterance.

While this explains the not-at-issue implication of subjunctives which the ‘Hold up/Hey, wait a minute’-test first indicated, it does not go all the way to reconciling our judgments about Adams pairs with Collapse. We have shown that the subjunctive member of a pair can be expected to trigger accommodation to a context whose common ground entails the negation of the presuppositions of the indicative. Since the two conditionals are evaluated at distinct contexts, judgments about them

[51]Since they are orthogonal to the present discussion, I set aside issues involving local accommodation, though see Percus (2006) and Singh (2011) for discussion.

[52]As with normal scalar implicatures within a neo-Gricean framework, the derivation requires the idealization that the speaker is opinionated about the presuppositions of the alternatives to her utterance. Absent this assumption, we will instead derive the implicature that the speaker is not certain that the presupposition of (92.b) is satisfied. See Sauerland (2008, §2.1) for discussion.

can diverge without threatening COLLAPSE.

However, it is not sufficient to merely explain how the pair can elicit different responses. We must also explain why, in its accommodated context, the subjunctive can receive an information insensitive reading (one which depends entirely on the vicars dispositions). In line with the discussion in §3.4, to do this we need to show that the presuppositions of the indicative will be unsatisfied in the new context. Here, the connection between the behavior of epistemic modals and what is common ground is crucial.

Let $CG(c)$ denote the common ground of c—that is, the set of claims which are mutually accepted by the participants in c (see Stalnaker (1970, 1973, 1974) for classic discussion). In the framework of the previous chapter we modeled this as an information state, which corresponding to the way the claims in the common ground collectively represent the world as being.

First, note that where it is common ground that ϕ is epistemically necessary, it can be expected to also be common ground that ϕ. That is:

$$CG(c) \models \Box\phi, \text{ then } CG(c) \models \phi.$$

Note that this merely constrains membership of the common ground. Hence, it is neutral with respect to the principle that '*must*' is weak—i.e., that $\Box\phi \not\models \phi$ (see, e.g., Karttunen (1972), Veltman (1985), Kratzer (1991), and Lassiter (2016) for discussion; cf. von Fintel and Gillies (2010, 2019) for rebuttal).

Second, ϕ cannot be epistemically possible at a context if its prejacent is incompatible with the common ground.[53] That is:

$$c \models \Diamond\phi \text{ only if } CG(c) \not\models \neg\phi.$$

Yet together, these constraints imply that in the context resulting from accommodating the not-at-

[53]Note that this appears to hold even of cases in which suppositions are added to the common ground which participants only accept hypothetically. Hence the oddity of (ʄ):

(ʄ) ?? Suppose it is raining. Nevertheless, it might not be.

issue implication of the subjunctive, the corresponding indicative will no longer be licensed. Suppose that a speaker utters $\phi \rightsquigarrow \psi$ in c. Assuming that $CG(c)$ does not entail $\Diamond\phi$, co-operative hearers can be expected to accommodate the implication that the indicative is unlicensed. This will result in a new context, c', such that $CG(c) \cup \{\neg\Diamond\phi\} \subseteq CG(c')$. Yet if $\neg\Diamond\phi \in CG(c')$, then $CG(c') \models \Box\neg\phi$ (by Must/Might, §3.3.2.1). So, from our first observation, it follows that $CG(c') \models \neg\phi$. Yet, by our second observation, it follows that $c' \not\models \Diamond\phi$. So $\phi \dashrightarrow \psi$ will not be licensed at c'.

Here is a summary of the position we have reached: where their antecedent is epistemically possible, subjunctives are equivalent to the corresponding indicatives. When evaluated in such a context, the former will receive an information sensitive reading. Indeed, according to Epistemicity, a subjunctive in a non-counterfactual context will simply express that it is epistemically impossible for its antecedent to be true but its consequent false.

However, subjunctives uttered in non-counterfactual contexts are not always evaluated at their context of utterance. Rather, due to pragmatic pressure generated by Maximize Presuppostion, they often implicate that their antecedent is epistemically impossible. Accommodating this information returns a new context. Yet, once this information is accommodated, Collapse no longer imposes an requirement that the subjunctive will receive an information sensitive reading.

3.5.3 Coda

Not all uses of subjunctives trigger context shifts of the kind just discussed. Before concluding, it is worth considering two notable categories of exception. Unlike uses of subjunctives forming Adams pairs, we should expect uses in these categories to be information sensitive, as a corollary of Collapse.

First, note that the conditionals in the subjunctive instances of If/And and And/If (i.e., (63.a-c) and (65.a-b$'$), repeated below) do not implicate that their antecedent is not epistemically possible; in neither case is a 'Hold up/Hey, wait a minute!'-response available. Accordingly, there is no reason to posit covert context shift in the arguments.

(63) a. Maybe the butler did it.

b. If he'd done it, maybe he'd have used the candlestick.

c. So, maybe the butler did it using the candlestick.

(65) a. Maybe the butler did it using the candlestick ...

b′. So, if he had done it, maybe he'd have used the candlestick.

This should be unsurprising. Not-at-issue implicatures generated by Maximize Presupposition, like other pragmatic implicatures, are widely recognized to be cancellable (see, in particular, Lauer (2016, §2.2)). In both (63.a-c) and (65.a-b′), the speaker explicitly asserts that she take the antecedent of the subjunctive to be epistemically possible. Hence, any implication that she takes the presuppositions of the corresponding indicatives to be false should be defeated.

However, explicit cancellation is not the only way in which the implicature can be cancelled. Consider the indicative and subjunctive variants of Anderson (1951)'s example:

(93) a. If Jones has taken arsenic, he's showing the symptoms he's actually showing.

b. If Jones had taken arsenic, he'd be showing the symptoms he's actually showing.

As Stalnaker (1975) and von Fintel (1997*b*) observe, (93.b) can naturally figure as part of an argument in favor of the possibility of its antecedent. In contrast, (93.a) carries a strong sense of redundancy, and cannot be expected to figure in a successful argument for anything.

Given Indicative Licensing, that (93.a) appears redundant is unsurprising. After all, it presupposes precisely what it is, intuitively, intended to establish. (93.b) has no such presupposition and, hence, can be used in an argument that Jones might have taken arsenic.

Crucially, (93.b) does not implicate that its antecedent is ruled out in context. Again, this is to be expected. The implicature of the subjunctive is generated by the need to explain why a speaker did not use the indicative. Yet, in this case, there is an independently available explanation: the indicative form presupposes what the speaker intends to establish. Accordingly, her interlocutors cannot conclude from her use of the subjunctive that she took indicative to be unlicensed—indeed, to do so would be incompatible with the intuitive point of her utterance.

3.6 Summary

COLLAPSE says that corresponding indicatives and subjunctives are Strawson equivalent; in contexts at which both are licensed, the one implies the other. COLLAPSE may be surprising, but it is not in tension with orthodoxy about indicatives/subjunctives. Since the presuppositions of indicatives and subjunctives diverge, it is compatible with their exhibiting substantially different logical properties. As the prior section demonstrated, it is also compatible with differing judgments about the members of Adams pairs.

There is a broad theoretical picture which accords nicely with this account. Conditionals (both indicative and subjunctive) involve the evaluation of their consequent at a body of information which entails their antecedent. Where their antecedent is compatible with the contextually salient information, the body of information at which the consequent is evaluated will be a subset of the information which is contextually salient. However, where it is incompatible, the contextually salient information places no constraints on the body of information at which the consequent is evaluated. Assume that epistemic modals and conditionals are evaluated with respect to same contextually salient information. Then, given Indicative Licensing, indicatives will receive an information sensitive reading where licensed—their antecedents will always be evaluated at a subset of the contextually salient information. Subjunctives will receive an information sensitive reading in contexts which are non-counterfactual. However, when evaluated in counterfactual contexts, (as, I have suggested, given their pragmatic behavior they standardly are) they will be insensitive to the contextually salient information.

The previous two chapters defended a version of this idea within a dynamic framework. On that theory, differences between indicatives and subjunctives are exhausted by differences in their presuppositions. Any other variation in their behavior is attributable to this basic difference. In this respect, the theory accords with the general approach defended by, e.g., see Stalnaker (1975), Karttunen and Peters (1979), and von Fintel (1999).

If this idea is correct, then where both are licensed, they will have the same status. Nevertheless, as long as their presuppositions do not coincide fully, the two forms of conditional may have different pragmatic effects, even in those contexts which license each.

Chapter 4

Causation & Revision

4.1 Introduction

Consider a boiler. It consists of a power switch, a water valve, and a pilot light. When the valve is open and the light is ignited, the boiler produces hot water. When either the valve is closed or the light is out, the water remains cold. When the boiler's power is off, the valve will be closed and the light out. When the boiler's power is on, the valve and light may be in any combination of states.

In the preceding chapter, we observed that corresponding indicatives and subjunctives sometimes elicit divergent judgments. Suppose that you observe that the boiler is not producing hot water. Then while you should be willing to accept an assertion of (94), you should be far less willing to accept an assertion of (95).

(94) If the valve is open, the light is out.

(95) If the valve had been open, the light would have been out.

Chapter 3 argued that this difference in judgments can be attributed to a covert shift. The subjunctive, unlike the indicative, is evaluated at a modified context—one at which the antecedent is presumed to be false. This shift in context suffices to explain why our judgments about (94) and (95) can diverge.

However, there is another, less clear-cut difference between the pair which remained unaddressed in the previous chapter: our judgments about the two forms of conditional appear to be driven by different considerations.

Our inclination to accept (94) reflects the way the available information is structured. Given the design of the boiler, observing that the water is cold puts you in a position to conclude that either the valve is closed or the light is out. However, it does not put you in a position to rule either alternative out. As a result, adding the claim that valve is open to your information would result in a state which incorporated the claim that the light is out. It is these considerations that explain why we are inclined to accept (94), given the setup (cf. Stalnaker (1975)).

Our inclination not to accept (95), in contrast, reflects the way the world is structured. We know that the temperature of the water is dependent on the state of both the valve and the light. If either the valve is closed or the light is out, the water remains cold. However, the state of the valve and the light are causally independent of one another. Changing the state of one has no effect on the state of the other. It is these considerations that appear to explain why we are reluctant to accept (95), given the setup.

Our judgments about counterfactual subjunctives appear sensitive to causal structure in a way that our judgments about indicatives are not (Bennett (2003), Edgington (2004), Schaffer (2004), Kment (2006)). According to a family of theories, this phenomenon is to be accounted for by building a sensitivity to causal structure into the semantics of subjunctives. Typically, such theories enrich the points at which subjunctives are evaluated, so that they specify not only what is the case, but also what depends on what and how it depends on what it depends on (Galles and Pearl (1998), Pearl (2000, 2009), Briggs (2012), Halpern (2013), Santorio (2019) Lassiter (2020)). As a result, points carry information about how changes in one place percolate through a system. Stated broadly, a subjunctive is accepted when making a modification to bring about its antecedent would result in bringing about its consequent.

The previous three chapters have argued that both indicatives and subjunctives are information sensitive: their status is dependent in part on the information determined by the context. Chapter 2 defended a particular form of information sensitivity. On this picture, conditionals (of both forms) can be understood as tools for exploring potential changes in information. To evaluate a subjunctive, we first revise the contextual information with its antecedent, and then evaluate its consequent at the result. Crucially, revising different bodies of information with the same claim can return different results, even when that claim is incompatible with both. Different bodies of information are structured differently. As such, even counterfactual subjunctives are predicted to exhibit information sensitivity.

These two approaches might appear starkly at odds. In evaluating a subjunctive, the former tells us to consider how changing one aspect of the world would bring about changes in other aspects. In contrast, the latter tells us to look at how accepting one piece of information would bring about changes in what other information we accept. Given this difference, it is reasonable to ask whether the kind of information sensitive account of conditionals defended above can hope accommodate the sensitivity of our judgments to causal structure.

In this concluding chapter, I demonstrate that it can. I start, in §4.2, by introducing certain basic features of structural equation models (SEMs) (Pearl (2000, 2009)). I show how these models allow us to represent causal dependence between states of affair. In §§4.3-4.4, I demonstrate how to integrate SEMs and belief revision. This provides us a way of building causal dependencies into the supposition operation at the core of our account of the conditional. Finally, in §4.5, I conclude by briefly discussing how the resulting account accommodates judgments about counterfactual subjunctives.

4.2 Modeling Causes

We could specify every potential total state of the boiler by specifying the range of cotenable states of the power switch, the water valve, and the pilot light, along with whether the boiler is producing hot water. Table 4.2 gives an exhaustive description of this form:

	S1	S2	S3	S4	S5
Power switch:	Off	On	On	On	On
Pilot light:	Out	Out	Ignited	Out	Ignited
Water valve:	Closed	Closed	Closed	Open	Open
Hot water:	No	No	No	No	Yes

Table 4.2

A description of the system of this kind leaves something out. In particular, it omits information about how different components of the system depend on one another. It is reasonable to want a description which tells us not only that either the valve is open and light ignited or the boiler is not producing hot water, but also tells us that if the boiler is not producing hot water, it is *because* either the valve is closed or the light out.

Structural Equation Models (SEMs) offer one way of describing a system with the desired level of information. Each SEM specifies not only the possible combination of states of the system, but also the way changing one state would affect the others.

4.2.1 Structural Equation Models

An SEM is a pair $\langle \mathcal{S}, \mathcal{V} \rangle$ of a set of structural equations, $\mathcal{S}$, and a set of variables, $\mathcal{V}$. We will assume that $\mathcal{V} = \{\mathsf{A}, \mathsf{B}, ..., \bot, \top\}$. That is, variables are identified with the finite set of atomic sentences of the language. Each variable in an SEM is associated with a range—the possible values it can take. As we are identifying variables with atomic sentences, we will take their range to be the set of truth values, $\{0, 1\}$. Since we hold the language fixed, we can harmlessly identify SEMs with just their set of structural equations in what follows.

A structural equation $e_i \in \mathcal{S}$ is a formula of the form:

$$\theta(\{\mathsf{A}_i, ..., \mathsf{A}_j\}) \ \Pi \ \mathsf{B}_i$$

Here, B_i is an atom, and $\{\mathsf{A}_i, ...\mathsf{A}_j\}$ a set of atoms. θ is a schematic variable over (names of) functions from a set of values of atoms to a value of atoms. Π is schematic variable over (names of) relations between values of atoms.

Intuitively, each structural equation tells us that the value of one atom depends (in part) on the values of some others. The way in which the former depends on the latter, however, depends on Π and θ. Here, we can be relatively restrained, for present purposes.

We restrict the value of θ to three functions: Max, Min, and Ident. Max and Min return the greatest and least element of their input, respectively. Ident is defined only on singletons, and returns their member.

We restrict the value of Π to four relations: $:=$, $:\neq$, $\leq$, and $\geq$. $:=$ and $:\neq$ relate values which are the same and distinct, respectively. $\leq$ and $\geq$ relate a value to values greater and smaller, respectively. Note that each relationship is asymmetric in the following sense: $\mathsf{Max}\{\mathsf{A}_1, \mathsf{A}_2\} := \mathsf{B}_1$ is a structural equation, but $\mathsf{B}_1 := \mathsf{Max}\{\mathsf{A}_1, \mathsf{A}_2\}$ is not.

Each set of structural equations determines a parent-child relationship over atoms. We define corresponding functions, $\mathsf{Parent}(\cdot)$ and $\mathsf{Child}(\cdot)$, which return the determinans and determinandum of a structural equation, respectively. $\mathsf{Parent}(\theta(\{\mathsf{A}_i, ..., \mathsf{A}_j\})\ \Pi\ \mathsf{B}_i)$ is the set $\{\mathsf{A}_i, ..., \mathsf{A}_j\}$ containing the parents of B_i in the equation. $\mathsf{Child}(\theta(\{\mathsf{A}_i, ..., \mathsf{A}_j\})\ \Pi\ \mathsf{B}_i)$ is the set $\{\mathsf{B}_i\}$ containing the child of $\{\mathsf{A}_i, ..., \mathsf{A}_j\}$ in the equation.

It will be useful to define a relationship, $<_{\mathcal{S}}$, which holds between an atom and its descendants according to the set of structural equations $\mathcal{S}$. B_i is a descendant of A_i iff A_i bears the ancestor relation (i.e., the transitive closure of the parent-child relationship) to B_i.

Definition 25.

$\mathsf{A}_i <_{\mathcal{S}} \mathsf{B}_i$ iff there is some subset $\{e_1, ..., e_n\} \subseteq \mathcal{S}$ such that:

(i) $\mathsf{A}_i \in \mathsf{Parent}(e_1)$;

(ii) $\mathsf{B}_i \in \mathsf{Child}(e_n)$;

(iii) For all k such that $1 \leq k \leq n$: $\mathsf{Child}(e_{k-1}) \subseteq \mathsf{Parent}(e_k)$.

Intuitively, $\mathsf{A}_i <_{\mathcal{S}} \mathsf{B}_i$ iff it is possible to trace a path from the former to the latter through the structural equations in $\mathcal{S}$. Since it should be clear from context, we will elide the indexation to the set of structural equations for the purposes of what follows.

Say that an SEM, $\langle \mathcal{S}, \mathcal{V} \rangle$, is recursive iff $<_{\mathcal{S}}$ is asymmetric. For the time being, we will restrict our attention to recursive SEMs. This guarantees that we are looking at SEMs in which no atom is a descendant of its own child.

Our boiler can be characterized by a recursive SEM. Let the atoms P, L, V, and W be true iff the power is on, the light is ignited, the valve is open and the water is hot, respectively. Then our model will comprise three structural equations:

$$\mathcal{S}_{\mathsf{Boiler}} = \begin{cases} e_1 : & \mathsf{Ident}(\{\mathsf{P}\}) \geq \mathsf{L}; \\ e_2 : & \mathsf{Ident}(\{\mathsf{P}\}) \geq \mathsf{V}; \\ e_3 : & \mathsf{Min}(\{\mathsf{L}, \mathsf{V}\}) := \mathsf{W}. \end{cases}$$

e_1 and e_2 say, respectively, that unless the power is on, the light is out and the valve is closed. e_3 says that hot water is produced iff the light is ignited and the valve is open.

SEMs can be partially represented using directed graphs. A directed edge ($\rightarrow$) connects one node to another iff the latter is a child of the former. As long as we restrict attention to recursive SEMs, every graph is guaranteed to be acyclic (Pearl (2000, 2009)). For example, our SEM characterizing the boiler corresponds to the directed graph in Figure 4.1.

Directed graphs offer a simple way to visualize the dependence relationships between atoms in an SEM. We can see that in our model of the boiler, L and V each depend on P, and W depends on L and V. P is the only atom which does not depend on any others. Note that the graph leaves out some important information about the model: it does not tell us how the atoms depend on one another, only whether they do. As a result, there is a many-one relationship between models and directed

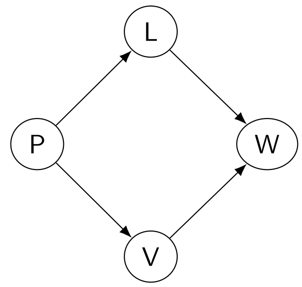

Figure 4.1: The directed graph corresponding to $\mathcal{S}_{\mathsf{Boiler}}$.

graphs.

4.2.2 Valuing SEMs.

Our SEMs carry two types of information about a system. They tell us what constraints there are on the total state of the system, by way of constraints on states of its components. And, importantly, they also tell us which components' states depend on the states of which others.

However, SEMs leave some information about a system out. In particular, they do not tell us what state the system is in fact in. $\mathcal{S}_{\mathsf{Boiler}}$ tells us that as long pilot light is out, the boiler will not produce hot water. However, it is silent on whether or not the light is out. For some purposes (such as constructing a replica of the boiler), the latter type of information will be irrelevant. For others (such as predicting what will happen if we open the valve), it will be vitally important. In order to reason about particular states of the system, we need to enrich our framework with worlds.

As above, a world is a total valuation over atoms. Structural equations are evaluated as true or false at a world, depending on whether the values it assigns to their parents/child stand in the appropriate relationship. We say that a world w verifies an SEM just in case it makes every equation in the model true.

Definition 26.

i. $w \models \theta(\{\mathsf{A}_i, ..., \mathsf{A}_j\})\ \Pi\ \mathsf{B}_i$ iff $\theta(\{w(\mathsf{A}_i), ..., w(\mathsf{A}_j)\})\ \Pi\ w(\mathsf{B}_i)$.

ii. $w \Vdash\!\models \mathcal{S}$ iff for all $e_i \in \mathcal{S} : w \models e_i$.

A world verifies an SEM iff the values it assigns to atoms conform to the constraints the SEM imposes on them. For example, consider the two worlds in table 4.2.2, representing different combinations of states of the components of the boiler.

	P	L	V	W
w	1	0	1	0
v	0	1	0	1

Table 4.2.2

According to w, the power is on, the light is out, the valve is open and the boiler is not producing hot water. w verifies our model of the boiler, since it makes every equation in the model true. According to v, in contrast, the power is off, the light is on, the valve is closed and the boiler is producing hot water. v fails to verify our model of the boiler. In particular, it makes both e_1 and e_3 false.

4.2.3 Intervening in SEMs

Two SEMs can impose the same constraints on the values of atoms, while inducing different dependence relations. To see this, consider the two SEMs below:

$$\mathcal{S}_1 = \begin{cases} \mathsf{Ident}(\{\mathsf{A}\}) := \mathsf{B}; \\ \mathsf{Ident}(\{\mathsf{B}\}) :\neq \mathsf{C}. \end{cases} \qquad \mathcal{S}_2 = \begin{cases} \mathsf{Ident}(\{\mathsf{B}\}) := \mathsf{A}; \\ \mathsf{Ident}(\{\mathsf{C}\}) :\neq \mathsf{B}. \end{cases}$$

$\mathcal{S}_1$ and $\mathcal{S}_2$ are verified by the very same worlds. A world verifies each iff either it makes A and B true and C false or it makes C true and A and B false. However, $\mathcal{S}_1$ and $\mathcal{S}_2$ characterize the dependence relationships between atoms very differently.

We can see this by looking at their corresponding directed graphs. It is easy to observe that $<_{\mathcal{S}_2}$ is the converse of $<_{\mathcal{S}_1}$.

More generally, the direction of dependence in an SEM is irrelevant to determining the truth of its constituent structural equations at a world. In order to know whether w verifies $\mathcal{S}$, we do not need

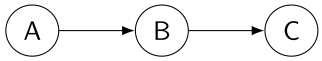

$\mathcal{S}_1$

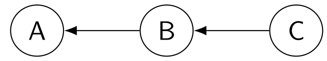

$\mathcal{S}_2$

to know anything about how atoms depend on one another according to $\mathcal{S}$. We only need to know about the global constraints it imposes on how those atoms are valued.

To differentiate SEMs which are verified by the same worlds, we need to look beyond the range of possible states a system can be in. We need to consider, in addition, the effect that changing the state of some component would have on the rest of the system.

In structural equation models, this idea is often implemented in terms of intervention (Pearl (2000), Woodward (2003)). Informally, an intervention is intended to model the result of interfering in the system to artificially fix the state of a particular component. Formally, an intervention can be thought of as an update operation on SEMs. It maps one set of structural equations to another, depending on its arguments. Intuitively, we want to be able to fix the value of an atom to be either true or false. Accordingly, our intervention operation takes three arguments: an input set of structural equations (the starting arrangement of the system), an atom (the component to be intervened on), and a truth value (the state it is to be set in).

Definition 27.

Where $n \in \{0, 1\}$:

$$\textsc{Int}(\mathcal{S}, \mathsf{A}_i, n) = \begin{cases} \{e_j \in \mathcal{S} |\ \mathsf{Child}(e_j) \neq \mathsf{A}_i\} \cup \{\mathsf{Ident}(\{\mathsf{A}_i\}) := \top\}, \text{ if } n = 1; \\ \\ \{e_j \in \mathcal{S} |\ \mathsf{Child}(e_j) \neq \mathsf{A}_i\} \cup \{\mathsf{Ident}(\{\mathsf{A}_i\}) := \bot\}, \text{ if } n = 0. \end{cases}$$

Intervening on A_i within $\mathcal{S}$ can be understood as a two-step procedure: (i) first, any equations in $\mathcal{S}$ which have A_i as a child are removed; (ii) second, a new equation is introduced, which fixes the value of A_i to coincide with either the value of the tautology (i.e., 1) or the contradiction (i.e., 0).

An intervention breaks the connection between the intervened on atom and its parents and then specifies a new value for the atom which is independent of the values of any other atoms. Despite $\mathcal{S}_1$ and $\mathcal{S}_2$ being verified by the same worlds, intervening on them produces different results. We can see this by considering the respective effect of fixing B as true:

$$\textsc{Int}(\mathcal{S}_1, \mathsf{B}, 1) = \begin{cases} \mathsf{Ident}(\{\top\}) := \mathsf{B}; \\ \\ \mathsf{Ident}(\{\mathsf{B}\}) :\neq \mathsf{C}. \end{cases} \qquad \textsc{Int}(\mathcal{S}_2, \mathsf{B}, 1) = \begin{cases} \mathsf{Ident}(\{\mathsf{B}\}) := \mathsf{A}; \\ \\ \mathsf{Ident}(\{\top\}) := \mathsf{B}. \end{cases}$$

After the intervention, the sets of structural equations are verified by different worlds.

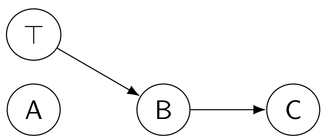

$\textsc{Int}(\mathcal{S}_1$, B, 1)

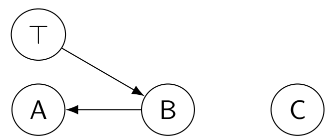

$\textsc{Int}(\mathcal{S}_2$, B, 1)

A world which makes A, B and C true will verify $\textsc{Int}(\mathcal{S}_2$, B, 1), but it will not verify $\textsc{Int}(\mathcal{S}_1$, B, 1). Conversely, a world which makes A and C false but B true will verify $\textsc{Int}(\mathcal{S}_1$, B, 1), but not $\textsc{Int}(\mathcal{S}_2$, B, 1). Note furthermore that neither world would verify either $\mathcal{S}_1$ or $\mathcal{S}_2$.

Intervention gives us a way of precisifying the idea that systems which are alike in the range of possible states they can occupy may differ in the way they respond to changes made to individual components. By investigating what happens when we intervene in a system, we can learn about how its parts depend upon one another.

For this reason, a number of authors have taken intervention to be crucial in explaining how subjunctives can serve to illuminate causal structure (Galles and Pearl (1998), Halpern (2000), Briggs (2012), Fine (2012), Zhang (2013), Lassiter (2020)). The idea is that, to evaluate a subjunctive relative to an SEM, we make an intervention corresponding to its antecedent, and evaluate the consequent at the model which results.

This idea has proven fruitful in accounting for the kinds of observations we started with. However, it has a number of shortcomings. First, unless appropriately augmented, it is insufficiently general. As a number of authors have observed, since intervention is defined only for atoms, something more needs to be said to extend the account to conditionals with complex antecedents (cf. Briggs (2012), Santorio (2019), Lassiter (2020)).

Second, it is insufficiently specific. Intervention is an operation on models alone (McDonald (2020)). Yet our judgments about subjunctives appear to be sensitive to what we know about the state of the world. That is, our judgments are guided not only by our knowledge of what depends on what, but also by our knowledge of the state the world is currently in. To accommodate this, we need to know what aspects of the world to hold fixed when it comes to evaluating the consequent of the conditional in the model resulting from intervention. We cannot hold everything fixed, since a world which satisfies an SEM need not satisfy the result of intervening on it. Yet we cannot hold nothing fixed, either. Some of our background knowledge of the state of the world is relevant to evaluating a subjunctive.

In what follows, I show how we can resolve each of these concerns while subsuming a structural equation based approach a broader, information-sensitive account of conditionals. First, however, I will briefly revisit the notion of revision.

4.3 Revisiting Revision

In chapter 2, we characterized revision as an operation which, given a pair of input bodies of information, returns a new body of information. The idea was that the operation should return the result of modifying one information state so that incorporates all of the information incorporated by the second. We considered a number of constraints on such an operation.

($*_1$) $s * s' = s \wedge s'$ *if* $s \wedge s' \neq \bot$;

($*_2$) $s * s' \leq s'$

($*_3$) $\bot < s * s'$ *if* $s' \neq \bot$;

($*_4$) $s * (s' \wedge s'') = (s * s') \wedge s''$ *if* $(s * s') \wedge s'' \neq \bot$.

It is well known that revision can be understood in terms of an ordinal ranking of worlds (Grove (1988); cf. Spohn (1988, 2012)). Let $\preccurlyeq$ be a three-place relation, which given an information state, induces a total order on worlds. Informally, we say that $w \preccurlyeq_s v$ iff v involves at least as great a departure from the way s represents things to be as w does. We assume that for any s, $\preccurlyeq s$ is well-founded.[1] We can then define revision as the minimal departure from an information state required to obtain a state which incorporates the incoming information.

$$s * s' = \{w \in s' |\ \forall v \in s' : w \preccurlyeq_s v\}$$

That is, revising s with s' just returns the 'nearest' s'-worlds to s. The ordering based approach to revision builds in $(*_2)$, $(*_3)$ and $(*_4)$. $(*_2)$ is satisfied since the nearest s'-worlds to s are guaranteed to be s'-worlds. $(*_3)$ is satisfied by the requirement that $\preccurlyeq_s$ is total and well-founded. As long as there are some s'-worlds, there are guaranteed to be some nearest s'-worlds. Finally, $(*_4)$ is satisfied since if some of the nearest s'-worlds are s''-worlds, then the nearest $s' \wedge s''$-worlds are guaranteed to be some of the nearest s'-worlds.

The remaining constraint, $(*_1)$, corresponds to the following condition on the information-relative ordering:

Minimality If $w \in s$, then $v \preccurlyeq_s w$ iff $v \in s$.

Minimality says that the worlds which come closest to the way s represents things to be are all and only those which are in s. That is, every world in s involves a minimal departure from s, and if a world is not in s, then it is further from s that any world which is. This ensures that if the incoming information is compatible with the current state, then the state resulting from revision will comprise all and only those worlds in the intersection of both.

Each information-relative ordering corresponds to a system of spheres, in the sense of Lewis (1973). This offers us a simple way to visualize revision.[2,3] For example, consider Figure 4.2, which depicts a system of spheres. The inner sphere in the system corresponds to the posterior state, s. The red line delineates the incoming information Revising it with s' returns the shaded region: the nearest

[1] If s' is an information state, there is some $w \in s'$ such that for all $v \in s'$: $w \preccurlyeq_s v$.

[2] A system of spheres is a set of information states, which is nested (i.e., if $s', s'' \in \$$, then either $s' \subseteq s''$ or $s'' \subseteq s'$), closed under union (i.e., if $X \subseteq \$$, then $\bigcup X \subseteq \$$) and intersection (i.e., if $X \subseteq \$$, then $\bigcap X \subseteq \$$.).

[3] The system of spheres induced by $\preccurlyeq_s$ is the set of information states:

$$\{\{v |\ v \preccurlyeq_s w\} |\ w \in \mathcal{W}\}.$$

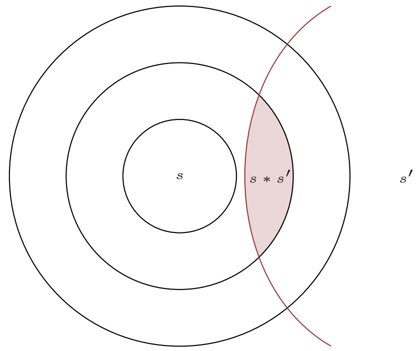

Figure 4.2: A system of spheres centered on s.

s'-worlds to s.

In the following section, I show how to construct a mapping between SEMs and information-relative orderings satisfying the appropriate constraints. As we've just seen, if we have an ordering, we have a revision operation. This revision operation is sensitive to causal structure in the following sense: revising with incompatible information returns a state in which everything causally 'upstream' is held fixed, while everything causally 'downstream' is controlled by the relevant structural equations.

4.4 Causal Revision

Here is a sketch of the plan: we first define an operation which, given a background of structural equations, returns the results of changing a world in a well-behaved way to make an arbitrary boolean sentence true. The governing idea is that for a change to be well-behaved, whatever is causally independent of what is changed is held fixed, while whatever is causally dependent on what is changed varies in accordance with the background structural equations (cf. Lewis (1979*a*, 462)).[4]

We then construct an information-relative ordering which satisfies the following condition: The nearest ϕ-worlds outside of the information state are required to be among those that result from a well-behaved change to a world in the information state which makes ϕ true. As I will show, this

[4]Lewis entertains, but rejects, an analysis of counterfactual similarity on which w is amongst the nearest worlds at which A is true iff "(1) A is true at w; (2) w is exactly like our actual world at all times before a transition period beginning shortly before t_{A} [the time A is about]; (3) w conforms to the conforms to the actual laws of nature at all times after t_{A}; and (4) during t_{A} and the preceding transition period, w differs no more from our actual world than it must to permit A to hold.". My proposal reflects the spirit of this idea in many of its broad strokes aspects. Unlike Lewis's preferred approach, my notion of counterfactual nearness makes essential appeal to causal structure.

ordering induces a revision operation which behaves the way we want when it comes to evaluating counterfactual subjunctives.

4.4.1 Aboutness

Let $\triangleright_{\mathcal{S}}(w, \phi)$ be the set of worlds which can be obtained by well-behaved change to w making ϕ true, relative to the background casual structure of $\mathcal{S}$. What worlds are included in $\triangleright_{\mathcal{S}}(w, \phi)$ will turn, in part, on what things depend on ϕ, according to $\mathcal{S}$.

Where ϕ is atomic, the structural equations immediately tell us which atoms depend on ϕ and which do not. However, where ϕ is complex, identifying the atoms which depend on it is more complicated. Accordingly, the first thing we will do is recursively define a function, η, which returns the set of atoms a complex boolean sentence is 'about'.

Definition 28.

i. $\eta(\mathsf{A}_i) = \{\mathsf{A}_i\}$.

ii. $\eta(\neg\phi) = \eta(\phi)$;

iii. $\eta(\phi \vee \psi) = \eta(\phi \wedge \psi) = \eta(\phi) \cup \eta(\phi)$

Intuitively, we can think of η as 'pulling' the atoms out of any complex formula it is given.

Recall that $\mathsf{A} <_{\mathcal{S}} \mathsf{B}$ iff it is possible to trace a path to B from A through the structural equations in $\mathcal{S}$. $\leq_{\mathcal{S}}$ is the reflexive closure of $<_{\mathcal{S}}$. It will be helpful to look at two sets of atoms in particular.

$$\mathrm{D}_{\mathcal{S}}(\phi) = \bigcup_{\mathsf{A_i} \in \eta(\phi)} \{\mathsf{B}_j \mid \mathsf{A}_i <_{\mathcal{S}} \mathsf{B}_j\}$$

$$\mathrm{I}_{\mathcal{S}}(\phi) = \bigcap_{\mathsf{A_i} \in \eta(\phi)} \{\mathsf{B}_j \mid \mathsf{A}_i \not\leq_{\mathcal{S}} \mathsf{B}_j\}$$

$\mathrm{D}_{\mathcal{S}}(\phi)$ is the set of atoms dependent upon some atom in ϕ. $\mathsf{A}_i \in \mathrm{D}_{\mathcal{S}}(\phi)$ iff it is a descendant of some atom in ϕ. $\mathrm{I}_{\mathcal{S}}(\phi)$ is the set of atoms independent from every atom in ϕ. $\mathsf{A}_i \in \mathrm{U}_{\mathcal{S}}(\phi)$ iff it is not a descendant of any atom in ϕ.

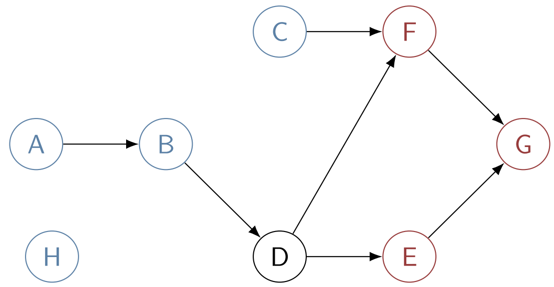

Figure 4.3: A directed graph of $\mathcal{S}^*$.

For example, in the SEM in Figure 4.3, $\mathrm{D}_{\mathcal{S}^*}(\mathsf{D} \vee \mathsf{F})$ is the set of nodes colored red. $\mathrm{I}_{\mathcal{S}^*}(\mathsf{D} \vee \mathsf{F})$ is the set of nodes colored blue.

Note that F is in $\mathrm{D}_{\mathcal{S}^*}(\mathsf{D} \vee \mathsf{F})$, since it is downstream of D. Note also that H is in $\mathrm{I}_{\mathcal{S}^*}(\mathsf{D} \vee \mathsf{F})$ despite not being upstream of either D or F, since it is not downstream of either. The two sets are guaranteed to be disjoint. What is independent of ϕ is not dependent on ϕ and what is dependent on ϕ is not independent on ϕ. However, they will not form a partition on the atoms in the SEM. In the present case, for example, D is in neither set. Note, however, that $\{\mathrm{I}_{\mathcal{S}}(\phi), \mathrm{D}_{\mathcal{S}}(\phi) \cup \eta(\phi)\}$ *is* a partition on the atoms in the SEM. We will employ this observation below.

4.4.2 Well-Behaved Change

$\rhd_{\mathcal{S}}$ is a generalization of a selection function, in the sense of Stalnaker (1968) and Stalnaker and Thomason (1970):[5] given a world and a formula, it returns a set of worlds. These are the worlds which can be obtained via a well-behaved change.

What makes a change well-behaved? Minimally, it should meet three conditions:

[5] A generalization in two senses: first, Stalnaker and Thomason require that the range of $\rhd_{\mathcal{S}}$ be the set of worlds, rather than its powerset. As Lewis (1973) observed, this is eliminable. Second, as discussed below, we impose distinct constraints on $\rhd_{\mathcal{S}}$ to Stalnaker and Thomason.

Condition 1: If $v \in \rhd_{\mathcal{S}}(w, \phi)$, then $v \in [\![\phi]\!]$.

Condition 2: If $v \in \rhd_{\mathcal{S}}(w, \phi)$ and $\mathsf{A}_i \in \mathrm{I}_{\mathcal{S}}(\phi)$, then $v(\mathsf{A}_i) = w(\mathsf{A}_i)$.

Condition 3: If $v \in \rhd_{\mathcal{S}}(w, \phi), e_i \in \mathcal{S}$ and $\mathsf{Child}(e_i) \in \mathrm{D}_{\mathcal{S}}(\phi)$, then $v \models e_i$.

Condition 1 says that well-behaved changes are successful. Any world obtained by making ϕ true had better be a ϕ-world. Condition 2 says that well-behaved changes do not make changes unnecessarily. If A_i is independent of ϕ, then any world obtained from w by making ϕ true had better agree with w on the truth of A_i. Condition 3 says that well-behaved changes do not violate causal structure unnecessarily. If A_i is dependent upon ϕ, then any world obtained from w by making ϕ true had better make the structural equations governing A_i true as well. We let $\rhd_{\mathcal{S}}$ be the weakest function satisfying Conditions 1-3; that is, if f satisfies Conditions 1-3 with respect to $\mathcal{S}$, then $f(w, \phi) \subseteq \rhd_{\mathcal{S}}$.

Condition 1 coincides with Stalnaker (1968) and Stalnaker and Thomason (1970)'s first constraint on selection functions. Stalnaker and Thomason go on to impose two additional constraints on selection functions: (i) that if ϕ is true at w, the only world that can be obtained from w by making ϕ true is w, and (ii) that if ϕ entails ψ and a ψ-world can be obtained from w by making ϕ true, then every world that can be obtained from w by making ϕ true can be obtained from w by making ψ true.

We will not impose either of these conditions on $\rhd_{\mathcal{S}}$. Indeed, it seems we should expect counterexamples to both. The former is incompatible with Condition 3: making ϕ true at a ϕ-world which does not verify $\mathcal{S}$ will not return that world itself. The latter is incompatible with Condition 2. It can be expected to fail when ϕ entails ψ, but the atoms which depend on the ϕ are not a subset of the atoms which depend on the ψ. In this case, making ϕ true can result in changes which making ψ true would not. More concretely, suppose that $\mathcal{S} = \{\mathsf{A} := \mathsf{B}\}$, $w(\mathsf{A}) = w(\mathsf{B}) = 0$, and $v(\mathsf{A}) = v(\mathsf{B}) = 1$. Then $v \in \rhd_{\mathcal{S}}(w, \mathsf{A} \wedge \mathsf{B})$ and $v \in [\![\mathsf{B}]\!]$. But $v \notin \rhd_{\mathcal{S}}(w, \mathsf{B})$.

It may be instructive to consider a concrete case of well-behaved change. Return to our boiler.

Suppose that, at w, the power is off, the valve is closed, the light is out and the water is cold.

Given the causal structure of the system, there are three well-behaved ways of making it true that either the valve is open or the light is on: we can make the first true, we can make the second true or we can make both true. These alternatives are characterized in Figure 4.4. If one disjunct is true, but the other false, then every other component of the system will remain in the same state. However, if both disjuncts are true then the water will be hot, even though the state of the water was not itself changed directly.

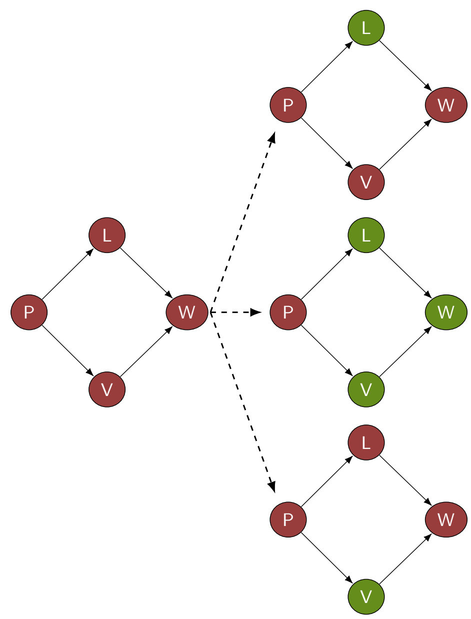

Figure 4.4: $\rhd_{\mathcal{S}_{\mathsf{Boiler}}} (w, \mathsf{L} \vee \mathsf{V})$.

Observe that every world that results from changing w to make $\mathsf{L} \vee \mathsf{V}$ true fails to verify $\mathcal{S}_{\mathsf{Boiler}}$, since the power remains off while either the light is ignited or the valve is open. As Lewis (1986) (and, later, Kment (2006)) observes, in making counterfactual changes to a world, we are faced with two competing considerations. On one hand, we want to preserve actual matters of fact. On the other, we want to preserve fidelity to actual laws. We cannot do both.

Lewis, wanting to keep causal talk out of counterfactual talk, suggests a 'system of weights'. We privilege avoiding major violations of the laws over exact coincidence in matters of fact, and we privilege exact coincidence in matters of fact over minor violations of the laws. Unencumbered by concerns about keeping counterfactual theorizing free from causal talk, I suggest that a simpler solution is available. We privilege fidelity to the laws when it comes to what is downstream of the counterfactual changes we make, and coincidence with matters of fact when it comes to everything else. Relative to a choice of causal structure, this is precisely what $\rhd$ encodes.

4.4.3 Revision

The selection function, $\rhd_{\mathcal{S}}$, gives us a way to explore the causal structure characterized by the equations in $\mathcal{S}$. It tells us, for any change we could make to a system in a particular state (i.e., a particular world), how that change could percolate through the rest of the system.

To turn this into a revision operation, we impose the following simple condition on an information-relative ordering over worlds.

Definition 29.

$\preccurlyeq_s$ is well-behaved (relative to $\mathcal{S}$) iff it satisfies the following condition:

For any $w \in \bigcup_{z \in s} \rhd_{\mathcal{S}}(z, \phi)$ and any $v \in [\![\phi]\!]$: $v \preccurlyeq_s w$ only if either $v \in s$ or $v \in \bigcup_{z \in s} \rhd_{\mathcal{S}}(z, \phi)$.

An ordering $\preccurlyeq_s$ is well-behaved iff it ranks all ϕ-worlds which can be obtained by changing some world s to make ϕ true as strictly less remote than any other ϕ-worlds outside of s. Well-behaved orderings which satisfy minimality can be thought of as dividing the ϕ-worlds into three tiers, relative to an information state: first are those ϕ-worlds included in the information state itself; second are those which can be obtained by making ϕ true at a world in the information state; third are any remaining ϕ-worlds.

Next, we translate our condition on well-behaved orderings into one on revision operations, by way of the equivalence laid out in §4.3.

Definition 30.

Where $s * s' = \{w \in s' |\, \forall v \in s' : w \preccurlyeq_s v\}$: $*$ is well-behaved relative to $\mathcal{S}$ iff $\preccurlyeq_s$ is well-behaved relative to $\mathcal{S}$.

Well-behaved revision operations care about causal structure, in the following sense. Revising an information state with incompatible information simply amounts to choosing some ways that worlds in the information state could be changed to make ϕ true.

Fact 11.

For any well-behaved revision operation $*$, if $s \cap [\![\phi]\!] = \emptyset$ and $\bigcup_{w \in s} \rhd_{\mathcal{S}}(w, \phi) \neq \emptyset$, then $s * [\![\phi]\!] \subseteq \bigcup_{w \in s} \rhd_{\mathcal{S}}(w, \phi)$

The proof is immediate. We know that if $*$ is well-behaved, then for any s and ϕ, $s * [\![\phi]\!]$ comprises all and only the $\preccurlyeq_s$-minimal ϕ-worlds, for some well-behaved order $\preccurlyeq_s$. Since $s \cap [\![\phi]\!] = \emptyset$, there are no ϕ-worlds in s. So, given that $\preccurlyeq_s$ is well-behaved, it follows that a ϕ-world is $\preccurlyeq_s$-minimal only if it is a member of $\rhd_{\mathcal{S}}(w, \phi)$, for some $w \in s$. So, $s * [\![\phi]\!]$ is a subset of $\bigcup_{w \in s} \rhd_{\mathcal{S}}(w, \phi)$.

We have seen that well-behaved revision operations exhibit the properties we want: they respect causal structure in the appropriate way when incorporating incompatible information. However, we are not yet done. Afterall, it remains to be established that there are any well-behaved revision operations. In fact, we want to show something stronger: that for any set of structural equations, there is a revision operation which is well-behaved relative to those equations and which satisfies the other conditions which we have considered for revision.

Recall that $*$ is a quasi-revision operation iff satisfies $(*_2)$-$(*_4)$. Every total well-founded ordering induces a quasi-revision operation. $*$ is full revision operation iff it is a quasi-revision operation which additionally satisfies $(*_1)$. A total well-founded ordering induces a full revision operation iff it satisfies Minimality.

Fact 12.

For any consistent set of structural equations $\mathcal{S}$, there is a well-behaved full revision operation (relative to $\mathcal{S}$).

To prove Fact 12, we start by introducing a notion of groupwise variance between worlds. Where Δ

is a set of atoms, v is a Δ-variant of w iff v differs from w on at most the values of the members of Δ and v makes any equations governing descendants of those members true.[6]

$$v \overset{\Delta}{\sim} w \text{ iff:} \quad \text{(i)} \quad w(\mathsf{A}_j) = v(\mathsf{A}_j), \text{ if } \mathsf{A}_j \notin \Delta; \text{ and}$$
$$\text{(ii)} \quad w \Vert\!\!\models \bigcup_{\mathsf{A}_i \in \Delta} \{e_j \in \mathcal{S} |\ \mathsf{A}_i < \mathsf{Child}(e_j)\}.$$

Informally, we can think of the Δ-variants of w as those worlds which can be reached from w by changing the value of, at most, the elements of Δ, with the further requirement that they must conform to that subset of the structural equations in the model which govern atoms 'downstream' of Δ.

Groupwise variance is non-reflexive: if w does not verify $\mathcal{S}$, it may not be an Δ-variant of itself, for some Δ. Nor is it symmetric: v may be an Δ-variant of w without w being a Δ-variant of v. Finally, it is not monotonic: where $\Delta \subseteq \Delta'$, v may be a Δ-variant of w without being a Δ'-variant.

Next, we recursively define a world-relative ranking function, κ_w, for each w. This function assigns a rank to each world on the basis of the size of its variance from the 'base' world.

$$\kappa_w(v) = \begin{cases} i & \textit{if } Min\{|\Delta| : v \overset{\Delta}{\sim} w\} = i; \\ \infty & \textit{otherwise.} \end{cases}$$

The rank of v relative to w is the cardinality of the smallest Δ such that v is a Δ-variant of w. Where there is no such set, the rank of v is ∞.

Relative to w, only w receives rank zero. Only w is a $\emptyset$-variant of itself. Every other world differs from w on the value of at least one atom. Worlds which are not rank zero, but can be reached from w by changing the value of one atom and verify the equations governing the descendants of that atom receive rank one. Worlds which are not ranks one or zero, but can be reached from w by changing the values of two atoms and verify the equations governing their descendants receive rank two. And so on. Worlds which cannot be reached by groupwise variation from w receive rank ∞.

[6]Note that the relation of groupwise variance is always relative to a particular SEM. For reasons of readability, we will suppress indexation to an SEM in the characterization of notions in what follows. The reader will not have difficulty identifying how indexes should be reintroduced.

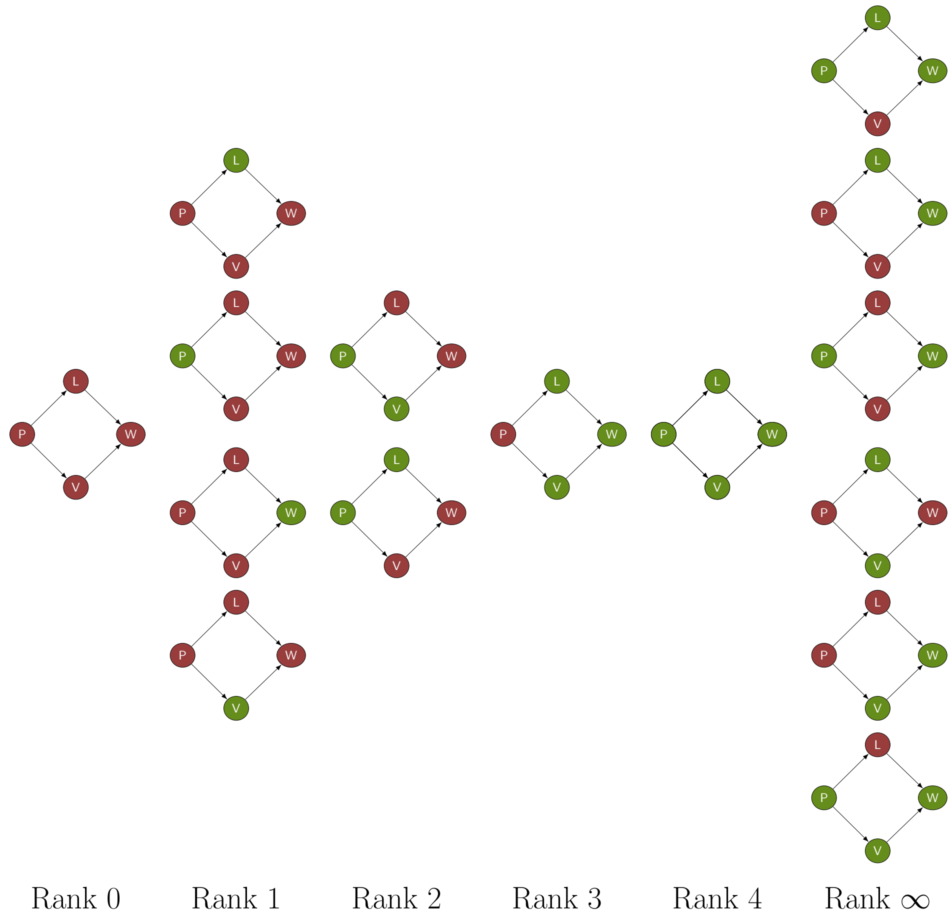

Figure 4.5: A series of ranked worlds.

It may be instructive to consider a concrete case. Figure 4.5 shows the ranking of various worlds according to $\kappa_{\mathcal{S}}^{w}$, where w is a world at which the power is off, the light out, the valve closed and the water cold. Worlds are ranked according to the extent of their variation from w, with the constraint that any world which falsifies a structural equation governing atoms downstream of an atom which is changed receives rank ∞.

We derive an information-relative ordering in two steps. The rank of a world relative to an information state is its minimal rank relative to some world in the state. That is:

$$\kappa_{\mathcal{S}}^{s}(w) = min\{\kappa_{\mathcal{S}}^{w}(v)|v \in s\}$$

The ordering over worlds relative to an information state is then simply determined by their rank. A world is no more remote from s than all and only those worlds with a rank at least as high (relative to s) That is:

$$w \preccurlyeq_s v \text{ iff } \kappa^s_{\mathcal{S}}(w) \leq \kappa^s_{\mathcal{S}}(v).$$

For every $w \in s$, $\kappa^s_{\mathcal{S}}(w) = 0$. Thus, $\preccurlyeq_s$ is guaranteed to satisfy Minimality. Since $\preccurlyeq_s$ is both total and well-founded, it follows that it induces a full revision operation.

Finally, we need to show that $\preccurlyeq_s$ is well-behaved. That is, if w can be reached by a well-behaved change to a member of s making ϕ true, then it is nearer to s than any ϕ-world outside of s which cannot be reached via a well-behaved change to s making ϕ true. That is, we will prove that for any $w \in \bigcup_{z \in s} \rhd_{\mathcal{S}}(z, \phi)$: if $v \in [\![\phi]\!]$ but $v \notin \bigcup_{z \in s} \rhd_{\mathcal{S}}(z, \phi)$ and $v \notin s$, then $w \prec_s v$.

So, suppose that $w \in \bigcup_{z \in s} \rhd_{\mathcal{S}}(z, \phi)$. We know that there is some $z \in s$ such that $w \in \rhd_{\mathcal{S}}(z, \phi)$. By Condition 1, w is a ϕ-world. By Condition 2, $w(\mathsf{A}_i) = z(\mathsf{A}_i)$, if $\mathsf{A}_i \notin \mathrm{I}_{\mathcal{S}}(\phi)$. But $\mathsf{A}_i \notin \mathrm{I}_{\mathcal{S}}(\phi)$ iff $\mathsf{A}_i \in \mathrm{D}_{\mathcal{S}}(\phi) \cup \eta(\phi)$. So w agrees with z on any atoms which are neither in $\eta(\phi)$ nor in the descendants of $\eta(\phi)$. By Condition 3, if $\mathsf{A}_i \in \eta(\phi)$: and $\mathsf{A}_i < \mathsf{Child}(e_i)$, then $z \models e_i$. That is, z verifies the equations which govern the descendants of $\eta(\phi)$. But the descendants of the descendants of $\eta(\phi)$ are a (proper) subset of the descendants of $\eta(\phi)$. So z verifies the equations which govern the descendants of any atom in $\mathsf{A}_i \in \mathrm{D}_{\mathcal{S}}(\phi) \cup \eta(\phi)$. As a result, it follows that w is a $\mathrm{D}_{\mathcal{S}}(\phi) \cup \eta(\phi)$-variant of z. Accordingly, we know that $\kappa^s_{\mathcal{S}}(w) \leq |\mathrm{D}_{\mathcal{S}}(\phi) \cup \eta(\phi)|$. Since our language has a finite supply of atoms, the rank of z relative to s and $\mathcal{S}$ is finite. So now all we need to show is that the rank of z is strictly lower than the rank of any ϕ-world outside of s which cannot be reached by a well-behaved change from s.

Suppose that v is a ϕ-world, but $v \notin s$ and there is no $z \in s$ such that $v \in \rhd_{\mathcal{S}}(z, \phi)$. Then it follows that for any $z \in s$, either v differs from z on one of the atoms in $\mathrm{I}_{\mathcal{S}}(\phi)$ or v falsifies some e_i governing a descendant of $\eta(\phi)$. Thus there is no $z \in s$ such that v is a $\mathrm{D}_{\mathcal{S}}(\phi) \cup \eta(\phi)$-variant of z. Since v is a ϕ-world, it follows that $\kappa^z_{\mathcal{S}}(v) > |\mathrm{D}_{\mathcal{S}}(\phi) \cup \eta(\phi)|$, for all $z \in s$. So, $\kappa^s_{\mathcal{S}}(v) > |\mathrm{D}_{\mathcal{S}}(\phi) \cup \eta(\phi)|$ and hence, $w \prec_s v$.

Summarizing, we know that the order induced by a ranking function $\kappa^s_{\mathcal{S}}$ will be well-behaved. And, correspondingly, the revision operation it returns will respect causal structure: in revising with incompatible information, it privileges changing matters causally downstream, in ways which are governed by the relevant structural equations.

4.5 Summary

Returning to our starting point, well-behaved full revision operations give us a way of capturing the idea that how we incorporate counterfactual information is sensitive to causal structure. When we modify an information state with some erstwhile incompatible claim, we care more about ensuring that what is causally upstream is left unchanged and what is causally downstream is law-abiding, than we do ensuring that the change is a 'minimal' departure from our old information (in the sense of retaining as much information as possible).

Nevertheless, well-behaved full revision operations still characterize a weak notion of minimal change. Any such operation will satisfy the AGM postulates $(*_1)$-$(*_4)$ of Alchourrón et al. (1985) (though see Rott (2000) for arguments that this notion is too weak to capture the intuitive notion of minimal change).

Returning to conditionals, well-behaved full revision operations give us the kind of variation in behavior we want across indicatives and subjunctives. Suppose that we are looking at the boiler. We can see that the power is on and the water is cold but we cannot see the state of either the pilot light or the valve. Accordingly, our information state will be compatible with three possible states of the world: one in which the light is ignited and the tap closed; one in which the tap is open and the light out; and one in which both the light is out and the tap closed.

As noted, our judgments about (94) and (95) diverge.

(94) If the valve is open, the light is out.

(95) If the valve had been open, the light would have been out.

Since the light is out in every state compatible with our information in which the valve is open, (94) is judged true. Not so for (95). Due to its competition with the indicative, it first triggers accommodation of the claim that the valve is closed. In the resulting context, however, it is not judged true. Indeed, there appears to be not inconsiderable pressure to judge it false.

A well-behaved proper revision operation predicts such behavior. Since the context at which (94) is evaluated is compatible with its antecedent, by Minimality, revising with the claim that the valve is open returns that subset of the information state at which it is true. Yet this subset incorporates the information that the light is out.

The context at which (95) is evaluated incorporates the information that the valve is closed. It is compatible with the light being either ignited or out. The state of the light is causally independent of the state of the valve. Accordingly, a well-behaved revision of the latter information state with the information that the valve is open need not result in a state which incorporates the information that the light is out. Indeed, if we adopt the particular revision operation discussed in the previous section, the conditional will be judged false, since the resulting information state will be both compatible with the light being ignited and compatible with it being out.

In order to evaluate whether a revision operation is well-behaved, we need to select a particular structural equation model against which to assess it. There is an obvious outstanding question: how is this choice of model motivated? Without offering a concrete answer, we can point to at some general considerations guiding SEM choice.

Any conversational context can be expected to come with a subject matter: a set of issues which are of significance to the interlocutors (Lewis (1988*a*,*b*)).[7] The subject matter of a conversational context is informed by the contributions qualify as relevant.

In the present setting, a subject matter is represented by our choice of variables in an SEM (i.e., the atoms of the language). Since worlds are simply ways of valuing atoms, they distinguish only between those states of affairs which can be expressed by some complex formula. If no atom (or

[7] See Yablo (2014), Hoek (2018) and Hawke (2018) for more recent discussion related to subject matter.

complex of atoms) corresponds to a given state of affairs, then whether that state of affairs obtains will play no role in individuating worlds for the purposes of our model. As such, we can think of the choice of variables in an SEM (i.e., the choice of atomic sentences in the language) as driven by the subject matter of the conversation (cf. Santorio (2019)).

What about the set of structural equations? Plausibly, these should be determined by interlocutors opinions about the causal relationships between the relevant issues. So far, we have been assuming that each revision operation corresponds to a single SEM. A conversational context is modeled by a pair of an information state and revision operation. In informal terms, this amounts to making the assumption that interlocutors are collectively opinionated on what causal relationships hold between the various issues which are of significance to them. However, this is clearly an idealization. Sometimes, the causal structure of the world is itself at issue.

It is relatively easy to see how the present framework could be extended to account for indecision or ignorance about causal structure. Rather than modeling a context with a pair of an information state and revision operation, we could instead model it as a set of such pairs. Each revision operation would correspond to a different set of structural equations, representing a different candidate for the causal structure of the world. Versions of this idea, for various forms of epistemic vocabulary, have been proposed by, e.g., **?**, Yalcin (2012), Willer (2013), and Goldstein (2020, forthcoming). In such a setting, the result of uttering a conditional is to return the set of pairs in the prior context which are fixed points of the test it denotes.

As has been widely observed, this allows us to model the way an utterance of a conditional can be informative. The context that results from uttering a conditional need not coincide with the prior context. However, in the present setting it does more than this. It allows us to model the way an utterance of a subjunctive, in particular, can tell us something about the what depends on what. Where distinct pairs agree on their information, whether they survive in the posterior context will depend entirely on their revision operation. As a result, uttering a conditional can convey something about the causal structure of the world.

Chapter 5

Conclusion

I have argued that we need to think about conditionals in a new way. Indicatives and subjunctives have the same at-issue meaning. Any differences in their behaviour arise purely from differences in their presuppositions.

What is that shared at-issue meaning? Chapters 1 and 2 proposed an answer involving central appeal to supposition. This picture is not new. It combines aspects of a variety of different classic approaches to conditionals. Before concluding, I will draw out some of these connections.

The conditional defended here is most closely related to the AGM conditional (Gärdenfors (1988)). Gärdenfors entertains the idea that, to evaluate a conditional, we evaluate its consequent at the result of revising our contextual information with its antecedent. Given that supposition is understood in terms of revision, this proposal can be understood as an instance of that idea.

It differs in two respects. First, unlike the AGM conditional, a conditional cannot itself be evaluated at a body of information. This avoids issues with Persistence from which Gärdenfors (1988) derives impossibility results (cf. Bradley (2000)). Second, unlike the AGM conditional, the environment at which its consequent is evaluated is not merely the result of revising the contextual information with the antecedent. It also includes information on how to perform future revisions in a way that reflects the information one has gained. That is, it changes the operation which is relevant for processing its consequent. Since this operation need not satisfy the full range of AGM axioms (cf. §2.2), the logics of the two conditionals will differ.

The conditional defended above is also closely related to the static variably strict conditional proposed by Stalnaker (1968), Stalnaker and Thomason (1970) and Lewis (1973). Stalnaker, Thomason and Lewis take their conditionals to be evaluated at points which are maximally informative, settling every matter in which we could be interested. The conditional defended here more general, being evaluated at points which need not be so heavy on detail. Our information states differ from worlds in leaving some matters unsettled.

However, putting this difference aside, the underlying idea is very similar. Both require the context to supply some measure of distance. For Stalnaker & Thomason, this is a selection function; for us, a revision operation. It is easy to establish a mapping between the two in terms of an ordering over worlds (§4.3). Given that measure, on both theories a conditional is to be accepted iff consequent is accepted at the 'nearest' point at which its antecedent is.

Like Stalnaker (1975), I have argued that the only difference between indicatives and subjunctives is in their not-at-issue properties. Like Lewis, I have argued that the indicative is equivalent to the material conditional.[1] The most significant deviation is in the predictions for conditionals in embedded environments and dynamic settings. The static variably strict conditional has an uninteresting logic for embedded conditionals. The distance measure at one world imposes no constraints on the distance measure at another. By employing a modification of the revision operation of the context to evaluated embedded material, the approach in Chapter 2 is able to generate different predictions. For example, whereas the static variably strict conditional validates *Modus Ponens* and gives up Import/Export, it validates Import/Export and gives up *ModusPonens.* At the same time, the conditional defended her is also closely related to the dynamic strict conditional, as proposed by Dekker (1993), Gillies (2004, 2009) and Starr (2014*a*), amongst others. As we saw in Chapter 2, it

[1] Lewis is, additionally, in favor of its identity with the material conditional.

can be given a decompositional analysis in terms of the strict conditional and supposition. Indeed, the indicative is strawson equivalent to the dynamic strict conditional.

Unlike the dynamic strict conditional, the present proposal does not validate ex falso inferences. Whereas the former can be inferred from the negation of its antecedent regardless of its consequent, the latter allows for non-trivial counterfactual uses. Crucially, though, this difference shows up only when we consider subjunctives, since counterfactual use of indicatives is undefined.

Finally, the subjunctive has a close relationship to accounts of the conditional which make use of causal models (Halpern (2000), Pearl (2000, 2009), Briggs (2012), Lassiter (2020)). As we saw, we can define a revision ordering which incorporates the information carried by a causal model and ensures that the revision operation is sensitive to the direction of causal dependence. The present proposal has a couple of advantages, however. Unlike many existing causal model theories, it allows for significant embedding of conditionals, along with complex antecedents. It also provides a solution to the problem of what facts about the world to hold fixed when making counterfactual changes (McDonald (2020)).

While they differ substantially in their details, it is possible to identify a common core shared by all of these theories. On each, conditionals are characterized by a rule for moving between points of information, whether those be worlds, information states or structural equation models. At this level of abstraction, what unifies the different theories is a particular view about the function of conditionals. They offer us a means of reasoning hypothetically and a tool for exploring information.

Bibliography

Abbott, B. (2006), Where have some of the presuppositions gone, *in* B. J. Birner and G. Ward, eds, 'Drawing the Boundaries of Meaning: Neo-Gricean studies in pragmatics and semantics in honor of Laurence R. Horn', pp. 1–20.

Abrusán, M. (2016), 'Presupposition cancellation: Explaining the 'hard' trigger distinction', *Natural Language Semantics* **24**(2), 165–202.

Abusch, D. (2010), 'Presupposition triggering from alternatives', *Journal of Semantics - J SEMANT* **27**, 37–80.

Adams, E. (1965), 'The logic of conditionals', *Inquiry: An Interdisciplinary Journal of Philosophy* **8**(1-4), 166–197.

Adams, E. W. (1970), 'Subjunctive and indicative conditionals', *Foundations of Language* **6**(1), 89–94.

Adams, E. W. (1975), *The Logic of Conditionals: An Application of Probability to Deductive Logic*, D. Reidel Pub. Co., ht, Holland.

Alchourrón, C. E., Gärdenfors, P. and Makinson, D. (1985), 'On the logic of theory change: Partial meet contraction and revision functions', *Journal of Symbolic Logic* **50**(2), 510–530.

Althshuler, D. (2016), *Events, States and Times. An essay on narrative discourse*, de Gruyter.

Anderson, A. R. (1951), 'A note on subjunctive and counterfactual conditionals', *Analysis* **12**(2), 35–38.

Appiah, A. (1985), *Assertion and Conditionals*, Cambridge University Press.

Arregui, A. (2007), 'When aspect matters: The case of would-conditionals', *Natural Language Semantics* **15**(3), 221–264.

Arregui, A. (2009), 'On similarity in counterfactuals', *Linguistics and Philosophy* **32**(3), 245–278.

Asher, N. and Lascarides, A. (2003), *Logics of Conversation*, CUP.

Asher, N. and McCready, E. (2007), 'Were, would, might and a compositional account of counterfactuals', *Journal of Semantics* **24**(2), 93–129.

Austin, J. L. (1962), *How to Do Things with Words*, Clarendon Press.

Barker, S. J. (1995), 'Towards a pragmatic theory of 'if'', *Philosophical Studies* **79**(2), 185–211.

Barnett, D. (2006), 'Zif is if', *Mind* **115**(459), 519–566.

Beaver, D. I. (1993), 'The kinematics of presupposition'.

Beaver, D. I. (2001), *Presupposition and Assertion in Dynamic Semantics*, CSLI Publications.

Bennett, J. (2003), *A Philosophical Guide to Conditionals*, Oxford University Press.

Boylan, D. and Schultheis, G. (2017), Strengthening principles and counterfactual semantics, *in* T. van Gessel & Floris Roelofsen Alexandre Cremers, ed., 'Proceedings of the 21st Amsterdam

Colloquium', pp. 155–164.

Boylan, D. and Schultheis, G. (2019), 'Attitudes, conditionals and the qualitative thesis', *manuscript* .

Bradley, R. (2000), 'A preservation condition for conditionals', *Analysis* **60**(3), 219–222.

Briggs, R. (2012), 'Interventionist counterfactuals', *Philosophical Studies* **160**(1), 139–166.

Cariani, F. and Goldstein, S. (forthcoming), 'Conditional heresies', *Philosophy and Phenomenological Research* .

Carter, S. (ms), 'Degrees of assertability'.

Chemla, E. and Schlenker, P. (2012), 'Incremental vs. symmetric accounts of presupposition projection: an experimental approach', *Natural Language Semantics* **20**(2), 177–226.

Ciardelli, I. (2020), 'Indicative conditionals and graded information', *Journal of Philosophical Logic* **49**(3), 509–549.

Darwiche, A. and Pearl, J. (1997), 'On the logic of iterated belief revision', *Artificial Intelligence* **89**, 1–29.

Davis, W. A. (1979), 'Indicative and subjunctive conditionals', *Philosophical Review* **88**(4), 544–564.

Dekker, P. (1993), Transsentential Meditations: Ups and Downs in Dynamic Logic, PhD thesis, University of Amsterdam: Institute for Logic, Language and Computation.

DeRose, K. (1991), 'Epistemic possibilities', *Philosophical Review* **100**(4), 581–605.

DeRose, K. (1994), 'Lewis on 'might' and 'would' counterfactual conditionals', *Canadian Journal of Philosophy* **24**(3), 413–418.

DeRose, K. (1996), 'Knowledge, assertion and lotteries', *Australasian Journal of Philosophy* **74**(4), 568–580.

DeRose, K. (1999), 'Can it be that it would have been even though it might not have been?', *Philosophical Perspectives* **13**(s13), 385–413.

DeRose, K. (2010), 'The conditionals of deliberation', *Mind* **119**(473), 1–42.

DeRose, K. and Grandy, R. E. (1999), 'Conditional assertions and "biscuit" conditionals', *Noûs* **33**(3), 405–420.

Dorr, C. and Hawthorne, J. (2013), 'Embedding epistemic modals', *Mind* **122**(488), 867–914.

Dorr, C. and Hawthorne, J. (manuscript), 'If:... a theory of conditionals'.

Dowty, D. (1979), *Word Meaning and Montague Grammar: The Semantics of Verbs and Times in Generative Semantics and in Montague's PTQ*, Dordrecht: Reidel.
URL: *https://books.google.com/books?id=SxhtCQAAQBAJ*

Dudman, V. H. (1992), 'Probability and assertion', *Analysis* **52**(4), 204–211.

Edgington, D. (1986), 'Do conditionals have truth-conditions', *Cr'itica* **18**(52), 3–30.

Edgington, D. (1995), 'On conditionals', *Mind* **104**(414), 235–329.

Edgington, D. (2004), Counterfactuals and the benefit of hindsight, *in* P. Dowe and P. Noordhof, eds, 'Cause and Chance: Causation in an Indeterministic World', Routledge.

Edgington, D. (2007), On conditionals, *in* D. Gabbay and F. Guenthner, eds, 'Handbook of Philosopical Logic', Vol. 14, pp. 127 – 221.

Edgington, D. (2008), 'Counterfactuals', *Proceedings of the Aristotelian Society* **108**(1pt1), 1–21.

Egan, A., Hawthorne, J. and Weatherson, B. (2005), Epistemic modals in context, *in* G. Preyer and

G. Peter, eds, ‘Contextualism in Philosophy’, Oxford University Press, pp. 131–170.

Filip, H. (2012), Lexical aspect, Oxford University Press.

Fine, K. (2012), ‘Counterfactuals without possible worlds’, *Journal of Philosophy* **109**(3), 221–246.

Fitelson, B. (2015), ‘The strongest possible lewisian triviality result’, *Thought: A Journal of Philosophy* **4**(2), 69–74.

Galles, D. and Pearl, J. (1998), ‘An axiomatic characterization of causal counterfactuals’, *Foundations of Science* **3**(1), 151–182.

Gärdenfors, P. (1988), ‘Knowledge in flux. modeling the dynamics of epistemic states’, *Studia Logica* **49**(3), 421–424.

Gazdar, G. (1979), *Pragmatics: Implicature, Presupposition and Logical Form*, Academic Press.

Geurts, B. (1999), *Presuppositions and Pronouns*, Elsevier.

Gibbard, A. (1981), Two recent theories of conditionals, *in* W. Harper, R. C. Stalnaker and G. Pearce, eds, ‘Ifs’, Reidel, pp. 211–247.

Gillies, A. S. (2001), ‘A new solution to moore’s paradox’, *Philosophical Studies* **105**(3), 237–250.

Gillies, A. S. (2004), ‘Epistemic conditionals and conditional epistemics’, *Noûs* **38**(4), 585–616.

Gillies, A. S. (2007), ‘Counterfactual scorekeeping’, *Linguistics and Philosophy* **30**(3), 329–360.

Gillies, A. S. (2009), ‘On truth-conditions for if (but not quite only if)’, *Philosophical Review* **118**(3), 325–349.

Gillies, A. S. (2010), ‘Iffiness’, *Semantics and Pragmatics* **3**(4), 1–42.

Gillies, A. S. (2020), ‘Updating data semantics’, *Mind* **129**(513), 1–41.

Gillies, A. S. (forthcoming), ‘Updating data semantics’, *Mind* p. fzy008.

Goldstein, S. (2018), ‘Conditional excluded middle and might conditionals’.

Goldstein, S. (2020), ‘Hypothetical conditionals’, *Synthese* .

Goldstein, S. (forthcoming), ‘Hypothetical assertion’, *Synthese* .

Goodman, J. and Salow, B. (2018), ‘Taking a chance on kk’, *Philosophical Studies* **175**(1), 183–196.

Goodman, N. (1947), ‘The problem of counterfactual conditionals’, *Journal of Philosophy* **44**(5), 113–128.

Green, M. S. (2000), ‘The status of supposition’, *Noûs* **34**(3), 376–399.

Grice, H. P. (1967), Logic and conversation, *in* P. Grice, ed., ‘Studies in the Way of Words’, Harvard University Press, pp. 41–58.

Groenendijk, J. and Stokhof, M. (1991), Two theories of dynamic semantics, *in* J. van Eijck, ed., ‘Logics in AI’, Springer Berlin Heidelberg, Berlin, Heidelberg, pp. 55–64.

Grove, A. (1988), ‘Two modellings for theory change’, *Journal of Philosophical Logic* **17**(2), 157–170.

Hacking, I. (1967), ‘Possibility’, *Philosophical Review* **76**(2), 143–168.

Hájek, A. (2011), ‘Triviality pursuit’, *Topoi* **30**(1), 3–15.

Halpern, J. Y. (2000), Axiomatizing causal reasoning, *in* ‘Proceedings of the Fourteenth Conference on Uncertainty in Artificial Intelligence’, UAI’98, Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, p. 202–210.

Halpern, J. Y. (2013), ‘From causal models to counterfactual structures’, *Review of Symbolic Logic* **6**(2), 305–322.

Han, C.-h. (1998), The Structure and Interpretation of Imperatives: Mood and Force in Universal Grammar, PhD thesis.

Hawke, P. (2018), 'Theories of aboutness', *Australasian Journal of Philosophy* **96**(4), 697–723.

Hawthorne, J. (2003), *Knowledge and Lotteries*, Oxford University Press.

Heim, I. (1983), On the projection problem for presuppositions, *in* P. Portner and B. H. Partee, eds, 'Formal Semantics - the Essential Readings', Blackwell, pp. 249–260.

Heim, I. (1990), Presupposition projection, *in* R. van der Sandt, ed., 'Reader for the Nijmegen Workshop on Presupposition, Lexical Meaning, and Discourse Processes', University of Nijmegen.

Heim, I. (1991), Artikel und definitheit, *in* A. von Stechow and D. Wunderlich, eds, 'Handbuch der Semantik', de Gruyter, pp. 487–535.

Higginbotham, J. (2003), 'Conditionals and compositionality', *Philosophical Perspectives* **17**(1), 181–194.

Hobbs, J. (1985), On the coherence and structure of discourse, Technical Report.

Hoek, D. (2018), 'Conversational exculpature', *Philosophical Review* **127**(2), 151–196.

Holguin, B. (forthcoming), 'Ramsey's thesis and introspection principles for belief and knowledge', *Noûs* .

Hunter, G. (1993), 'The meaning of 'if' in conditional propositions', *Philosophical Quarterly* **44**(172), 279–297.

Iatridou, S. (2000), 'The grammatical ingredients of counterfactuality', *Linguistic Inquiry* **31**(2), 231–270.

Ippolito, M. (2003), 'Presuppositions and implicatures in counterfactuals', *Natural Language Semantics* **11**(2), 145–186.

Jiang, Y. (2019*a*), *Chinese and Counterfactual Reasoning*.

Jiang, Y. (2019*b*), 'Ways of expressing counterfactual conditionals in mandarin chinese', *inguistics Vanguard* **5**.

Kamp, H. (1979), Events, instants and temporal reference, *in* R. Bäuerle, U. Egli and A. von Stechow, eds, 'Semantics From Different Points of View', Springer Verlag, pp. 376–418.

Kaplan, D. (1989), Demonstratives, *in* J. Almog, J. Perry and H. Wettstein, eds, 'Themes From Kaplan', Oxford University Press, pp. 481–563.

Karttunen, L. (1971), 'Some observations on factivity', *Research on Language & Social Interaction* **4**, 55–69.

Karttunen, L. (1972), Possible and must, Academic Press.

Karttunen, L. (1973), 'Presuppositions of compound sentences', *Linguistic Inquiry* **4**(2), 169–193.

Karttunen, L. (1974), 'Presupposition and linguistic context', *Theoretical Linguistics* **1**, 181–194.

Karttunen, L. and Peters, S. (1979), Conventional implicatures, *in* C.-k. Oh and D. Dineen, eds, 'Syntax and Semantics', Vol. 11 of *1*, Academic Press, pp. 1–56.

Kehler, A. (2000), Resolving temporal relations using tense meaning and discourse interpretation, *in* S. K. In Martina Faller and M. Pauly, eds, 'Formalizing the Dynamics of Information', CSLI, pp. 189–205.

Kehler, A. (2002), *Coherence, Reference and the Theory of Grammar*, CSLI.

Khoo, J. (2017), 'Backtracking counterfactuals revisited', *Mind* **126**(503), 841–910.

Klinedinst, N. (2011), 'Quantified conditionals and conditional excluded middle', *Journal of Semantics* **28**(1), 149–170.

Kment, B. (2006), 'Counterfactuals and explanation', *Mind* **115**(458), 261–310.

Kolbel, M. (2000), 'Edgington on compounds of conditionals', *Mind* **109**(433), 97–108.

Kolodny, N. and MacFarlane, J. (2010), 'Ifs and oughts', *Journal of Philosophy* **107**(3), 115–143.

Kratzer, A. (1977), 'What 'must' and 'can' must and can mean', *Linguistics and Philosophy* **1**(3), 337–355.

Kratzer, A. (1979), Conditional necessity and possibility, *in* R. Bäuerle, U. Egli and A. von Stechow, eds, 'Semantics From Different Points of View', Springer Verlag, pp. 117–147.

Kratzer, A. (1986), 'Conditionals', *Chicago Linguistics Society* **22**(2), 1–15.

Kratzer, A. (1991), Modality, *in* von Stechow and Wunderlich, eds, 'Semantics: An International Handbook of Natural Language Meaning', De Gruyter, pp. 639 – 650.

Kratzer, A. (2008), *The Notional Category of Modality*, Wiley-Blackwell, chapter 12, pp. 289–323. **URL:** *https://onlinelibrary.wiley.com/doi/abs/10.1002/9780470758335.ch12*

Kratzer, A. (2012), *Modals and Conditionals*, OUP.

Krifka, M. (2001), 'Quantifying into question acts', *Natural Language Semantics* **9**(1), 1–40.

Krifka, M. (2004), 'Semantics below and above speech acts'.

Kyburg, H. E. (1961), *Probability and the Logic of Rational Belief*, Wesleyan University Press.

Lakoff, G. (1970), 'Linguistics and natural logic', *Synthese* **22**(1-2), 151–271.

Lascarides, A. and Asher, N. (1993), 'Temporal interpretation, discourse relations and commonsense entailment', *Linguistics and Philosophy* **16**(5), 437–493.

Lassiter, D. (2016), 'Must, knowledge, and directness', *Natural Language Semantics* **24**(2), 117–163.

Lassiter, D. (2020), 'Causation and probability in indicative and counterfactual conditionals'.

Lauer, S. (2016), 'On the status of 'maximize presupposition'', *Semantics and Linguistic Theory* **26**, 980.

Lepore, E. and Stone, M. (2014), *Imagination and Convention: Distinguishing Grammar and Inference in Language*, Oxford University Press.

Leslie, S. (2009), 'if', 'unless', and quantification, *in* R. Stainton and C. Viger, eds, 'Compositionality, Context, and Semantic Values'.

Levi, I. (1977), 'Subjunctives, dispositions and chances', *Synthese* **34**(4), 423–455.

Lewis, D. (1976), 'Probabilities of conditionals and conditional probabilities', *Philosophical Review* **85**(3), 297–315.

Lewis, D. (1979*a*), 'Counterfactual dependence and time's arrow', *Noûs* **13**(4), 455–476.

Lewis, D. (1979*b*), 'Scorekeeping in a language game', *Journal of Philosophical Logic* **8**(1), 339–359.

Lewis, D. (1986), 'Counterfactual dependence and time's arrow', *Philosophical Papers* **2**.

Lewis, D. K. (1973), *Counterfactuals*, Blackwell.

Lewis, D. K. (1988*a*), 'Relevant implication', *Theoria* **54**(3), 161–174.

Lewis, D. K. (1988*b*), 'Statements partly about observation', *Philosophical Papers* **17**(1), 1–31.

Lewis, K. S. (2018), 'Counterfactual discourse in context', *Noûs* **52**(3), 481–507.

Lycan, W. G. (2001), *Real Conditionals*, Oxford University Press UK.

Mackie, J. L. (1972), *Truth, Probability and Paradox: Studies in Philosophical Logic*, Oxford, Clarendon Press.

Makinson, D. C. (1965), 'The paradox of the preface', *Analysis* **25**(6), 205.

Mandelkern, M. (2019), 'Talking about worlds', *Philosophical Perspectives* pp. 1–34.

Mandelkern, M. (manuscript), If p, then p!

Mandelkern, M., Romoli, J., Schwarz, F. and Zehr, J. (2017), 'Asymmetry in presupposition projection: The case of conjunction', *Proceedings of Semantics and Linguistic Theory 27* pp. 504–524.

McCall, S. (1966), 'Connexive implication', *Journal of Symbolic Logic* **31**(3), 415–433.

McCawley, J. (1996), Conversational scorekeeping and the interpretation of conditional sentences, *in* M. Shibatani and S. Thompson, eds, 'Grammatical Constructions', Clarendon Press, pp. 77–101.

McDonald, J. R. (2020), 'The importantion problem for structural semantics'.

Moens, M. and Steedman, M. (1988), 'Temporal ontology and temporal reference', *Comput. Linguist.* **14**(2), 15–28.
URL: *http://dl.acm.org/citation.cfm?id=55056.55058*

Moss, S. (2012), 'On the pragmatics of counterfactuals', *Noûs* **46**(3), 561–586.

Murray, S. E. and Starr, W. B. (2018), Force and conversational states, *in* D. Fogal, D. Harris and M. Moss, eds, 'New Work on Speech Acts', Oxford University Press, New York, pp. 202–236.

Murray, S. E. and Starr, W. B. (Forthcoming), The structure of communicative acts.

Nichols, C. (2017), 'Strict conditional accounts of counterfactuals', *Linguistics and Philosophy* **40**(6), 621–645.

Pearl, J. (2000), *Causality: Models, Reasoning, and Inference*, Cambridge University Press.

Pearl, J. (2009), *Causality*, Cambridge University Press.

Percus, O. (2006), Anti-presuppositions, *in* A. Ueyama, ed., 'Theoretical and empirical studies of reference and anaphora: Toward the establishment of generative grammar as an empirical science. Report of the Grant-in-Aid for Scientific Research (B), Project No. 15320052', Japan Society for the Promotion of Science, pp. 52–73.

Pizzi, C. (1977), 'Boethius' thesis and conditional logic', *Journal of Philosophical Logic* **6**(1), 283–302.

Pizzi, C. and Williamson, T. (1997), 'Strong boethius' thesis and consequential implication', *Journal of Philosophical Logic* **26**(5), 569–588.

Popper, K. (1959), *The Logic of Scientific Discovery*, Routledge.

Portner, P. (1992), Situation theory and the semantics of propositional expressions, PhD thesis, UMass-Amherst.

Portner, P. (1997), 'The semantics of mood, complementation, and conversational force', *Natural Language Semantics* **5**(2), 167–212.

Ramsey, F. P. (1931), *The Foundations of Mathematics and Other Logical Essays*, Paterson, N.J., Littlefield, Adams.

Rényi, A. (1955), 'On a new axiomatic theory of probability', *Acta Mathematica Academiae Scientiarum Hungarica* **6**, 285–335.

Roberts, C. (1989), 'Modal subordination and pronominal anaphora in discourse', *Linguistics and Philosophy* **12**(6), 683–721.

Roberts, C. (1996), Information structure in discourse: Towards an integrated formal theory of pragmatics, *in* J. Yoon and A. Kathol, eds, 'OSUWPL Volume 49: Papers in Semantics'.

Roberts, C. (2012), 'Information structure in discourse: Towards an integrated formal theory of pragmatics', *Semantics and Pragmatics* **5**, 1–69.

Rothschild, D. (2008), 'Presupposition projection and logical equivalence', *Noûs* **42**(1), 473–497.

Rothschild, D. (2014), 'Capturing the relationship between conditionals and conditional probability with a trivalent semantics', *Journal of Applied Non-Classical Logics* **24**(1-2), 144–152.

Rothschild, D. (2015), 'Conditionals and propositions in semantics', *Journal of Philosophical Logic* **44**(6), 781–791.

Rott, H. (2000), 'Two dogmas of belief revision', *Journal of Philosophy* **97**(9), 503–522.

Russell, J. S. and Hawthorne, J. (2016), 'General dynamic triviality theorems', *Philosophical Review* **125**(3), 307–339.

Sadock, J. M. and Zwicky, A. M. (1985), Speech acts distinctions in syntax, *in* T. Shopen, ed., 'Language Typology and Syntactic Description', Cambridge University Press, Cambridge, pp. 155–196.

Sampanis, K. (2012), A semantic features analysis for the subjunctive mood.

Santorio, P. (2019), 'Interventions in premise semantics', *Philosophers' Imprint* **19**.

Sauerland, U. (2003), 'Implicated presuppositions'.

Sauerland, U. (2008), Implicated presuppositions, *in* A. Staube, ed., 'Sentence and Context', de Gruyter.

Sauerland, U., Andersen, J. and Yatsushiro, K. (2005), *The Plural is Semantically Unmarked*.

Schaffer, J. (2004), 'Counterfactuals, causal independence and conceptual circularity', *Analysis* **64**(4), 299–308.

Schlenker, P. (2007), 'Anti-dynamics: Presupposition projection without dynamic semantics', *Journal of Logic, Language and Information* **16**(3), 325–356.

Schlenker, P. (2008), 'Be articulate: A pragmatic theory of presupposition projection', *Theoretical Linguistics* **34**(3), 157–212.

Schlenker, P. (2009), 'Local contexts', *Semantics and Pragmatics* **2**(3), 1–78.

Schlenker, P. (2012), 'Maximize presupposition and gricean reasoning', *Natural Language Semantics* **20**(4), 391–429.

Schulz, K. (2007), Minimal models in semantics and pragmatics: free choice, exhaustivity, and conditionals, PhD thesis, University of Amsterdam: Institute for Logic, Language and Computation.

Schulz, M. (2014), 'Counterfactuals and arbitrariness', *Mind* **123**(492), 1021–1055.

Searle, J. R. and Vanderveken, D. (1985), *Foundations of Illocutionary Logic*, Cambridge University Press.

Shaffer, M. J. (2016), 'What if bizet and verdi had been compatriots?', *Logos and Episteme* **7**(1), 55–73.

Shanon, B. (1976), 'On the two kinds of presuppositions in natural language', *Foundations of Language* **14**(2), 247–249.

Simons, M. (2013), On the conversational basis of some presuppositions, *in* 'Perspectives on Linguistic Pragmatics', pp. 329–348.

Singh, R. (2011), 'Maximize presupposition! and local contexts', *Natural Language Semantics* **19**(2), 149–168.

Skyrms, B. (1974), 'Contraposition of the conditional', *Philosophical Studies* **26**(2), 145–147.

Slote, M. A. (1978), 'Time in counterfactuals', *Philosophical Review* **87**(1), 3–27.

Smiley, T. J. (1967), 'Mr. strawson on the traditional logic', *Mind* **76**(301), 118–120.

Spohn, W. (1988), Ordinal conditional functions. a dynamic theory of epistemic states, *in* W. L. Harper and B. Skyrms, eds, 'Causation in Decision, Belief Change, and Statistics, vol. II', Kluwer Academic Publishers.

Spohn, W. (2012), *The Laws of Belief: Ranking Theory and its Philosophical Applications*, Oxford University Press.

Stalnaker, R. (1973), 'Presuppositions', *Journal of Philosophical Logic* **2**(4), 447–457.

Stalnaker, R. (1974), Pragmatic presuppositions, *in* R. Stalnaker, ed., 'Context and Content', Oxford University Press, pp. 47–62.

Stalnaker, R. (1975), 'Indicative conditionals', *Philosophia* **5**(3), 269–286.

Stalnaker, R. (1978), 'Assertion', *Syntax and Semantics (New York Academic Press)* **9**, 315–332.

Stalnaker, R. (1984), *Inquiry*, Cambridge University Press.

Stalnaker, R. (2002), 'Common ground', *Linguistics and Philosophy* **25**(5-6), 701–721.

Stalnaker, R. (2009*a*), 'Iterated belief revision', *Erkenntnis* **70**(2), 189–209.

Stalnaker, R. (2014), *Context*, Oxford University Press.

Stalnaker, R. C. (1968), 'A theory of conditionals', *Americal Philosophical Quarterly* pp. 98–112.

Stalnaker, R. C. (1970), 'Pragmatics', *Synthese* **22**(1-2), 272–289.

Stalnaker, R. C. (1981), A defense of conditional excluded middle, *in* W. Harper, R. C. Stalnaker and G. Pearce, eds, 'Ifs', Reidel, pp. 87–104.

Stalnaker, R. C. (2009*b*), 'On hawthorne and magidor on assertion, context, and epistemic accessibility', *Mind* **118**(470), 399–409.

Stalnaker, R. C. and Thomason, R. H. (1970), 'A semantic analysis of conditional logic', *Theoria* **36**(1), 23–42.

Starr, W. (ms), 'Indicative conditionals, strictly'.

Starr, W. B. (2014*a*), 'Force, mood and truth', *ProtoSociology* **31**, 160–181.

Starr, W. B. (2014*b*), 'A uniform theory of conditionals', *Journal of Philosophical Logic* **43**(6), 1019–1064.

Starr, W. B. (2014*c*), 'What 'if'?', *Philosophers' Imprint* **14**.

Stenius, E. (1967), 'Mood and language-game', *Synthese* **17**(1), 254–274.

Stojnić, U. (2017), 'One's modus ponens: Modality, coherence and logic', *Philosophy and Phenomenological Research* **95**(1), 167–214.

Stojnić, U. (2019), 'Content in a dynamic context', *Noûs* **53**(2), 394–432.

Stojnić, U. (forthcoming), 'Content in a dynamic context', *Noûs* .

Strawson, P. F. (1952), *Introduction to Logical Theory*, New York: Wiley.

Strawson, P. F. (1986), If and "", *in* R. E. Grandy and R. Warner, eds, 'Philosophical Grounds of Rationality: Intentions, Categories, Ends', Oxford University Press, pp. 363–380.

Urbach, P. (1988), 'What is a law of nature? a humean answer', *British Journal for the Philosophy of Science* **39**(2), 193–209.

van Benthem, J. (1996), *Exploring Logical Dynamics*, Center for the Study of Language and Information, USA.

Veltman, F. (1985), *Logics for Conditionals*, Jurriaans.

Veltman, F. (1996), 'Defaults in update semantics', *Journal of Philosophical Logic* **25**(3), 221–261.

von Fintel, K. (1997*a*), 'Bare plurals, bare conditionals, and only', *Journal of Semantics* **14**(1), 1–56.

von Fintel, K. (1997*b*), The presupposition of subjunctive conditionals, MIT Press.

von Fintel, K. (1999), 'Npi licensing, strawson entailment, and context dependency', *Journal of Semantics* **16**(2), 97–148.

von Fintel, K. (2001), 'Counterfactuals in a dynamic context', *CSLS* **36**, 123–152.

von Fintel, K. (2004), Would you believe it? the king of france is back! presuppositions and truth-value intuitions, *in* A. Bezuidenhout and M. Reimer, eds, 'Descriptions and Beyond', Oxford University Press.

von Fintel, K. and Gillies, A. S. (2007), An opinionated guide to epistemic modality, *in* T. Gendler and J. Hawthorne, eds, 'Oxford Studies in Epistemology: Volume 2', Oxford University Press, pp. 32–62.

von Fintel, K. and Gillies, A. S. (2010), 'Must . . . stay . . . strong!', *Natural Language Semantics* **18**(4), 351–383.

von Fintel, K. and Gillies, A. S. (2019), 'Still going strong'.

von Fintel, K. and Iatridou, S. (2002), 'If and when if-clauses can restrict quantifiers', http://web.mit.edu/fintel/fintel-iatridou-2002-ifwhen.pdf.

Warmbrod, K. (1983), 'Epistemic conditionals', *Pacific Philosophical Quarterly* **64**(3), 249–265.

Willer, M. (2013), 'Dynamics of epistemic modality', *Philosophical Review* **122**(1), 45–92.

Williams, J. R. G. (2010), 'Defending conditional excluded middle', *Noûs* **44**(4), 650–668.

Williamson, T. (2000), *Knowledge and its Limits*, Oxford University Press.

Woodward, J. (2003), *Making Things Happen: A Theory of Causal Explanation*, Oxford University Press.

Yablo, S. (2014), *Aboutness*, Princeton University Press.

Yalcin, S. (2007), 'Epistemic modals', *Mind* **116**(464), 983–1026.

Yalcin, S. (2011), Nonfactualism about epistemic modality, *in* A. Egan and B. Weatherson, eds, 'Epistemic Modality', Oxford University Press.

Yalcin, S. (2012), 'A counterexample to modus tollens', *Journal of Philosophical Logic* **41**(6), 1001–1024.

Yalcin, S. (manuscript), 'Notes on iffy knowledge'.

Yong, Q. (2013), Typological stage of counterfactuals in mandarin, *in* 'PACLIC'.

Zhang, J. (2013), 'A lewisian logic of causal counterfactuals', *Minds and Machines* **23**(1), 77–93.