INFORMATIVE DYNAMIC SEMANTICS

BY SIMON GOLDSTEIN

A dissertation submitted to the

Graduate School—New Brunswick

Rutgers, The State University of New Jersey

in partial fulfillment of the requirements

for the degree of

Doctor of Philosophy

Graduate Program in Philosophy

Written under the direction of

Anthony S. Gillies & Ernie Lepore

and approved by

New Brunswick, New Jersey

October, 2017

ABSTRACT OF THE DISSERTATION

Informative dynamic semantics

by Simon Goldstein

Dissertation Directors: Anthony S. Gillies & Ernie Lepore

In dynamic semantics, the meaning of a sentence is modeled as a rule for how a body of information grows when the sentence is accepted. Recent work in dynamic semantics has analyzed sentences involving modals and conditionals as tests. Tests are a special type of dynamic meaning. When an agent learns a test, her information is guaranteed to either stay the same, or become absurd.

The aim of this dissertation is to separate dynamic semantics from the notion of a test. The dissertation offers new analyses of modals, conditionals, and disjunction on which they are not tests. Nonetheless, these analyses are still genuinely dynamic. The meaning of sentences involving these expressions cannot be represented as a rule instructing the agent to update her information with the claim that she inhabits one of a fixed set of possibilities—that a certain proposition is true. Rather, the rule for how she is to update her information must make essential reference to her current information, providing different instructions depending on what that information is like.

Acknowledgements

This dissertation owes the most to my dissertation advisors. Ernie Lepore has been a constant support, mentor, raconteur, teacher, and therapist throughout my time at Rutgers. Ernie has given more of his time than I can ever repay, and has been a rich source of data for many of my papers. Above all, Ernie has taught me by his own example how to respond to objections with an open mind.

Thony Gillies is the inspiration for this dissertation. He has been actively involved in every step of the process, and at every level of resolution. He taught me every tool that I apply in what follows. His own work has been my model for combining formal semantics and epistemology.

My other committee members have also provided crucial help along the way. Since arriving at Rutgers, I've attended more of Jeff King's seminars than anyone else's. His careful attention to teaching provided a foundation for my study, and set the agenda for my own research. As Andy Egan's research assistant, I've learned a tremendous amount, working through paper after paper, with countless hours of happy conversation. Finally, Seth Yalcin's constant encouragement has been an invaluable help. He's gone above and beyond his duties as an external member of the committee.

The linguistics department at Rutgers is extremely welcoming to philosophers, and provided most of the training in semantics that I've relied on in this dissertation. Seminars with Veneeta Dayal, Maria Bittner, Roger Schwarzschild, and Simon Charlow have guided me through every year of my time at Rutgers. Other professional mentors inside and outside Rutgers include Fabrizio Cariani, Branden Fitelson, John Hawthorne, Martin Lin, Brian McLaughlin, Daniel Rothschild, Larry Temkin, and many others.

Since arriving at Rutgers, my fellow graduate students have taught me so much. David Black gave me many hours of patient math tutoring early in graduate school. Sam Carter's wild and boundless creativity has been an inspiration to me. Nico Kirk-Giannini's earnest criticism kept me honest, as much as possible. Besides the above, more thanks are due to Alex Anthony, Nick Beckstead, Ben Bronner, Marco Dees, Peter van Elswyk, Verónica Gómez, Jeremy Goodman, Ben Levinstein, Carlotta Pavese, and Una Stojnić. Without their support and conversation, my life and my work would not have been the same.

I've learned the most at Rutgers from two people. First, Bob Beddor has been my constant friend and interlocutor since I arrived at Rutgers. Bob taught me how to speak, and how to listen, when doing philosophy. Then he taught me how to write. Coauthoring with Bob has been one of the most fulfilling intellectual experiences of my life. I hope we may do so again. Above all, his emphasis on clarity and argument has given me an ideal to aspire to in the future.

Simon Charlow showed up at Rutgers just as Bob was leaving. He's been a worthy replacement. Simon has taught me more about semantics than anyone else. His constant refrain to think like a functional programmer, his encyclopedic knowledge of the field, and above all else his overwhelming ambition, have been an inspiration to never be content with my own work.

None of my work over the past six years would have been possible without the constant support of my parents and step-parents, who gave me the love of learning that has made my time at Rutgers such a joy. Finally, without Kathryn's intellectual and emotional support, no form of this dissertation would exist. She's been my constant companion for the last dozen years. Without her, I don't know who I would be.

Table of Contents

Chapter 1

Introduction

In dynamic semantics, meanings are richer than truth conditions. Rather than thinking about the meaning of a sentence as the conditions under which it is true, dynamic semantics conceives of the meaning of a sentence as a rule for how to learn that the sentence is true. This rule takes the form of a recipe for taking any agent's information state (or context), and updating that context with the sentence.

In most of what follows, information states are interpreted as sets of possible worlds.

Definition 1.1. A world w is a function from atomic sentences to truth values. An information state (or context) s is a set of worlds. A context change potential is a function from contexts to contexts. An interpretation function $[\cdot]$ assigns each sentence of the language a context change potential. s supports φ ($s \models \varphi$) just in case $s[\varphi] = s$.

Dynamic semantics is not the first theory of how a context should be updated when a sentence is learned. Stalnaker 1978 famously proposed that when a sentence is asserted, the context is updated by intersection with a fixed set of worlds, the proposition expressed by the sentence. In particular, each sentence in the language is assigned a set of worlds as its meaning, using some function $[\![\cdot]\!]$. Then updating with $[\varphi]$ is equivalent to intersecting s with $[\![\varphi]\!]$.

Definition 1.2. φ is intersective iff $\exists p \forall s : s[\varphi] = s \cap p$.

In dynamic semantics, a variety of meanings cannot be interpreted in this way. In these cases, updating a context cannot be interpreted in terms of learning that a certain proposition is true.

We can also understand nonintersective updating in another way: van Benthem [1986, 1989] proved that intersectivity is equivalent to the possession of two further properties. First, that updating a context never introduces new possibilities. Second, that updating a set of worlds is equivalent to updating each world in the set individually, and unioning the results.

Definition 1.3.

1. φ is eliminative iff $\forall s : s[\varphi] \subseteq s$.

2. φ is distributive iff $\forall s : s[\varphi] = \bigcup \big\{ \{w\}[\varphi] \mid w \in s \big\}$.

Throughout this dissertation, we will hold fixed the eliminativity of updating. This means that every failure of intersectivity that we consider will also be a failure of distributivity.

If $[\varphi]$ is not reducible to an intersective update, what can $[\varphi]$ be? Previous work in dynamic semantics has largely offered a single kind of answer to this question. Intersective updating has been replaced with a different kind of meaning: the test.[1] The job of a test isn't to tell the agent what world she inhabits. Rather, the job of a test is to explore structural properties of one's current information state. If the information has the relevant property, it stays the same; otherwise, it becomes absurd.

Definition 1.4. φ is a test iff $\forall s : s[\varphi] = s$ or $s[\varphi] = \emptyset$.

Tests are not in the business of providing new information about the world. They only provide any information when they provide too much information, narrowing the state down to $\emptyset$.

We've now introduced two kinds of context change potentials. The first type just amounted to intersecting the context with the truth conditions of the sentence. The second type was the test. These notions are disjoint, except in the limit.

Fact 1.1. φ is a test and intersective iff $\forall s : s[\varphi] = s$ or $\forall s : s[\varphi] = \emptyset$.

No interesting test is intersective. But there are failures of intersectivity that are not tests.

The project of this dissertation is to investigate what sorts of phenomena require a nonintersective kind of meaning, and what sorts of phenomena require treating meanings as tests. Again and again, we will find data that rules out intersectivity, but does not require treating any sentence as a test. This gives us an opportunity to develop a variety of new dynamic operators. These dynamic meanings do not amount to intersecting with a fixed proposition. But they can also tell us quite a lot about which worlds we inhabit, because they are not tests. These operators are dynamic, yet also informative.

---

[1] See for example Veltman 1996 and Harel et al. 2000.

In the chapters that follow, we examine three different kinds of expressions. In Chapter 2, we focus on possibility and necessity modals. We zoom in on a classic data point that has motivated a test semantics for modals: the inconsistency of epistemic contradictions, sentences of the form *p and might not p*.

EPISTEMIC CONTRADICTIONS $\varphi \wedge \Diamond\neg\varphi \models \bot$

To begin, we develop a general framework for theorizing about modal operators within dynamic semantics. The signature tool of truth conditional analyses of modals is an accessibility relation, which relates a world w to all worlds v possible from w. This tool can be generalized to the case of dynamic semantics by enriching the type of accessibility relations, letting them be information sensitive. The resulting theory allows a characterization of epistemic contradictions in terms of accessibility. Their inconsistency requires a nonintersective theory of modals. But this does not require that modals are tests. Ultimately, we defend a theory of modals that restricts an underlying accessibility relation to the worlds in the context. This theory does not make modals tests. But it predicts the inconsistency of epistemic contradictions, in addition to the validity of several other inferences. In particular, in this framework the strength of necessity modals and the positive introspection principle (that $\Box\varphi$ implies $\Box\Box\varphi$) turn out to follow from the inconsistency of epistemic contradictions.

Chapter 3 enriches the above discussion of modals with a second operator: disjunction. We turn to a surprising interaction effect between these two operators, Free Choice, where the possibility of a disjunction appears equivalent to a conjunction of possibilities.

FREE CHOICE $\Diamond(\varphi \vee \psi) \dashv\models \Diamond\varphi \wedge \Diamond\psi$

We begin by showing that within the framework of dynamic semantics, Free Choice requires a semantics for disjunction that is not distributive, and hence is not intersective. Then we characterize the family of Free Choice validating operators more precisely, finding the weakest and strongest such operators. Ultimately, this allows us to derive the traditional truth conditions for disjunction from the validity of Free Choice. These truth conditions are the strongest ones available that validate Free Choice, when combined with a new bit of dynamic meaning for disjunction: the requirement that each disjunct is possible.

Chapter 4 turns from modals and disjunction to conditionals, and also marks a change in methodology. Previous chapters focus on the constraints that the validity of various inferences impose on the choice of meaning within dynamic semantics. By contrast, in Chapter 4 we consider directly a variety of plausible constraints on updating. In particular, we explore the relationship between the following three constraints on learning a conditional. First, that any agent who has learned a conditional and that conditional's antecedent is committed to the consequent.

MODUS PONENS $s[\varphi \to \psi][\varphi] \models \psi$

Second, that the negation of a conditional's antecedent screens off the conditional itself.

SCREENING OFF $s[\varphi \to \psi][\neg\varphi] = s[\neg\varphi]$

Third, that any agent believes a conditional just in case she would believe the consequent if she learned the antecedent:

RAMSEY TEST $s \models \varphi \to \psi$ iff $s[\varphi] \models \psi$

We prove a pair of impossibility results. First, no test can satisfy the first two conditions. Second, there is a unique intersective meaning for the conditional that satisfies all three conditions: the material conditional. In the face of these results, we develop a variety of new dynamic meanings for the conditional that satisfy all three conditions, without collapsing to either a test or the material conditional. Finally, Chapter 5 provides proofs for the main facts in the previous chapters.

There is a large space of meanings that are nonintersective, but are also not tests. One of the projects of this dissertation is to narrrow down this space of meanings as much as possible. To do so, several general constraints on updating are quite useful. First, one natural idea is to be as conservative as possible in the failure of intersectivity. To do so, we will consider a variety of ways in which an update can be decomposed into two separate updates, one of which is intersective and one of which is not.

In Chapter 2, this decomposition is achieved by feeding a nonintersective component into a classical, intersective meaning for modals. On a classical theory, modals predicate possibility or necessity of a proposition, a set of worlds. On our nonintersective variant of this theory, the relevant proposition of which possibility or necessity is predicated is derived dynamically, as a function of both the input to the modal operator, and the context. Where $*$ denotes some function from a context and a context change potential to a proposition, our factorized modals behave as follows:

$$s[\Diamond\varphi] = s \cap [\![\Diamond]\!](*(s, [\varphi]))$$

$$s[\Box\varphi] = s \cap [\![\Box]\!](*(s, [\varphi]))$$

In our investigation of Free Choice, decomposition takes a different form. There, we let an update with the disjunction $\varphi \vee \psi$ be factorized into the intersection of two different updates, where one is intersective and the other is not. First, the context is updated intersectively with a propositional operation $\bowtie$ determined by $[\varphi]$ and $[\psi]$. Second, the context is updated nonintersectively with a test operation $*$ determined by $[\varphi]$ and $[\psi]$. Then the results of these two updates are combined together to produce a nonintersective update that is also not a test.

$$s[\varphi \vee \psi] = (s \cap [\![\varphi \bowtie \psi]\!]) \cap *(s, [\varphi], [\psi])$$

In our investigation of conditionals, decomposition takes yet another form. There, we let an update with the conditional be factorized into the union of two updates, again where one is intersective and one is not. Here, the intersective update is simply learning that one inhabits a world where the antecedent is false. The nonintersective update instead selects some combination of worlds from the context where the antecedent and consequent are both true. Crucially, the function $*$ that selects the relevant combination of worlds can be nonintersective, varying its choice of worlds with the choice of input context.

$$s[\varphi \to \psi] = (s \cap [\![\neg\varphi]\!]) \cup *(s, [\varphi], [\psi])$$

In addition to decomposing our updates into an intersective and nonintersective part, we can also narrow down the space of nonintersective meanings in other ways. One constraint that consistently leads to interesting results is the requirement that φ be supported in any state after that state has been updated with φ:

Definition 1.5. φ is idempotent iff $\forall s : s[\varphi] \models \varphi$.

Idempotence is a plausible requirement on any kind of learning. $s[\varphi]$ is a body of information committed to many things. But what could possibly follow from $s[\varphi]$, if not φ itself?

Idempotence interacts in interesting ways with the various forms of factorizability described above. In the case of possibility and necessity modals, it turns out that any factorizable modal is idempotent if and only if it satisfies the requirement of shift reflexivity, that $\Box(\Box\varphi \supset \varphi)$ is valid. In the case of disjunction, it turns out that there is a unique strongest factorizable theory of disjunction that is idempotent and satisfies Free Choice. In the case of conditionals, it turns out that our preferred class of factorizable conditionals satisfying the constraints above are all idempotent.

Idempotence is also interestingly related to intersectivity. Rothschild and Yalcin 2015b show that intersectivity is more or less equivalent to the conjunction of idempotence and commutativity:

Definition 1.6. φ and ψ are commutative iff for any s, $s[\varphi][\psi] = s[\psi][\varphi]$.

Much of this dissertation holds fixed both the failure of intersectivity and the satisfaction of idempotence, along with eliminativity. This means that the semantics below is a study in the failure of not only distributivity, but also commutativity: what requires these failures, and what else they imply about meaning.

Chapter 2

Generalized update semantics

2.1 Introduction

In a traditional possible worlds semantics for epistemic modals, *might* ($\Diamond$) and *must* ($\Box$) quantify over accessible possible worlds. $\Diamond\varphi$ is true at w iff φ is true at some world v accessible from w. $\Box\varphi$ is true at w iff φ is true at every world v accessible from w.[1]

Within this framework, various logical principles about modals correspond to constraints on accessibility. For example, the principle that *must* is strong (so that $\Box\varphi$ implies φ) is valid just in case every world is accessible from itself. The positive introspection principle (that $\Box\varphi$ implies $\Box\Box\varphi$) corresponds to the frame condition that whenever v is accessible from w and u is accessible from v, u is accessible from w. The B axiom (that truth implies necessary possibility) corresponds to the symmetry of accessibility.

One advantage of this framework is that it allows us to investigate the relationship between various logical principles. For example, we can convince ourselves that positive introspection and the B axiom imply that *must* is strong by observing that transitivity and symmetry imply reflexivity.

While this framework has many advantages, it also faces a well known problem predicting the inconsistency of 'epistemic contradictions' like the following:

(1) #It's raining but it might not be raining.

In the theory above, the inconsistency of epistemic contradictions (sentences of the form $\varphi \wedge \Diamond\neg\varphi$) corresponds to the trivializing result that every world is possible only from itself.

[1] Kripke 1963; Kratzer 2012.

The inconsistency of epistemic contradictions is better explained by a rival semantic framework: update semantics.[2] To predict this inconsistency, update semantics claims that modal sentences are *tests*. Briefly, this means that when we add the information carried by a modal claim to an agent's information, that agent's information either remains unchanged or becomes absurd.

The test semantics for modals does not merely predict the inconsistency of epistemic contradictions. It also validates all of the logical principles discussed above. These predictions are controversial. For example, there is a rich debate in the literature about whether *must* is strong.[3] And in recent work, Moss [2015] has challenged the thesis that iterated modals collapse.

Whether these principles are valid or not, we should still ask for more from update semantics. In the classical framework above, each logical principle under discussion corresponds to some model theoretic property in the semantics. We can use these model theoretic properties to better understand the logical principles we care about, and their relationship. Yet update semantics loses the correspondence between logical principles and constraints on accessibility.

This leads to one especially frustrating problem. Update semantics does not give us the tools to investigate which other logical principles are implied by the inconsistency of epistemic contradictions. This inconsistency is hard to predict, and so it is natural to wonder whether it implies, say, that iterated modals collapse or that *must* is strong. Update semantics does not provide any straightforward answer to this question.

In this chapter, I develop a new dynamic semantics for epistemic modals, Generalized Update Semantics (GUS), within which we can investigate the logical properties of epistemic modals. The major idea in GUS is that accessibility relations have a place within a dynamic theory of modals. But in order to earn their keep, accessibility relations must be be sensitive to the state or context in which a modal claim is interpreted. GUS is general enough to include update semantics and traditional possible worlds semantics as special cases. It allows us to characterize exactly what it takes for epistemic contradictions to be inconsistent, along with a variety of logical principles.

---

[2]For some defenses of update semantics, see Veltman 1996, Groenendijk et al. 1996, Beaver 2001, Gillies 2004, and Willer 2013. For some nearby views, see Yalcin 2007, Swanson 2011, Swanson 2012, Moss 2015, and Ninan 2016.

[3]See for example: Karttunen 1972; Kratzer 1991; von Fintel and Gillies 2010; and Lassiter 2016.

First, we will see that if epistemic contradictions are inconsistent, then the positive introspection principle (that $\Box\varphi \models \Box\Box\varphi$) is also valid. Second, we will see that if epistemic contradictions are inconsistent in a well behaved way, then *must* is strong. This second result connects the strength of *must* to the concept of *idempotence*: that asserting a sentence twice has the same effect as asserting it once. Surprisingly, once epistemic contradictions are inconsistent the strength of *must* is equivalent to the idempotence of *might* claims. This provides a powerful new argument for strength. Ultimately, then, a major claim of this chapter is that the inconsistency of epistemic contradictions places significant, unappreciated constraints on the logic of epistemic modality.

Finally, in addition to uncovering the logic of epistemic modals, the theory below also explicates the relationship between the traditional theory we started with and its update-theoretic rival. In what follows, we will develop a procedure for transforming any given classical semantics for modality into a dynamic semantics. The procedure is guaranteed to produce a semantics that predicts the inconsistency of epistemic contradictions, but where the old theory is still recognizable. This procedure is a tool that any theorist of epistemic modality may add to her existing theory in order to rule out epistemic contradictions.

2.2 Background

Let's start by reviewing two popular theories of modality: a classical, possible worlds semantics and a dynamic alternative.

In a traditional possible worlds semantics for epistemic modals, *might* ($\Diamond$) and *must* ($\Box$) quantify over accessible possible worlds. $\Diamond$ is an existential quantifier, while $\Box$ is a universal quantifier. To make this idea precise, we must first clarify the notion of accessibility. An accessibility relation allows what is possible to vary from world to world.

Definition 2.1. An accessibility relation R ($\subseteq W \times W$) relates a world w to the set of worlds v that are possible relative to w.

It is helpful to use a visual metaphor here: when wRv, we can think of v as visible from or seen by w. Possibility and necessity are then interpreted in terms of accessibility:

Definition 2.2.

1. $[\![\Diamond\varphi]\!] = \{w \mid \exists v : wRv \;\&\; v \in [\![\varphi]\!]\}$

2. $[\![\Box\varphi]\!] = \{w \mid \forall v : wRv \supset v \in [\![\varphi]\!]\}$

With this definition in place, we can say that some premises classically entail a conclusion just in case every world where the premises are true is a world where the conclusion is true.

Definition 2.3. Γ classically entail δ ($\Gamma \models_{\text{CL}} \delta$) iff $\bigcap_{\gamma \in \Gamma} [\![\gamma]\!] \subseteq [\![\delta]\!]$.

Within this framework, various logical principles about modals correspond to constraints on accessibility.[4]

Definition 2.4.

1. R is reflexive iff for any world w, wRw.

2. R is transitive iff for any worlds w, v, u, if wRv and vRu then wRu.

3. R is euclidean iff for any worlds w, v, u, if wRv and wRu then vRu.

4. R is symmetric iff for any worlds w, v, if wRv then vRw.

Observation 2.1.

1. R is reflexive iff $\Box\varphi \models_{\text{CL}} \varphi$.

2. R is transitive iff $\Box\varphi \models_{\text{CL}} \Box\Box\varphi$.

3. R is euclidean iff $\Diamond\varphi \models_{\text{CL}} \Box\Diamond\varphi$.

4. R is symmetric iff $\varphi \models_{\text{CL}} \Box\Diamond\varphi$.

While this theory is extremely fruitful, it struggles to predict the inconsistency of epistemic contradictions.

Definition 2.5. R is isolated iff for any worlds w, v, if wRv then $w = v$.[5]

[4] Strictly speaking, the entailment relation here is relativized to frames which agree on a choice of accessibility relation. For a precise definition of entailment, see the appendix.

[5] See Kaufmann and Kaufmann 2015.

Observation 2.2. R is isolated iff $\varphi \wedge \Diamond\neg\varphi \models_{CL} \bot$.

So in this framework, epistemic contradictions are inconsistent only if modal claims are trivial.

The inconsistency of epistemic contradictions is much easier to explain within update semantics. Update semantics is a type of dynamic semantics.[6] According to dynamic semantics, the meaning of a sentence is not its truth conditions. Rather, the meaning of a sentence is its ability to change the context in which it is said—its *context change potential*:

> You know the meaning of a sentence if you know the change it brings about in the information state of anyone who accepts the news conveyed by it.[7]

To give an update semantics, we need two things: a definition of information states (or contexts), and an interpretation function which assigns a context change potential to each sentence in our language. Veltman 1996 models an information state as a set of possible worlds. Then an interpretation function assigns every sentence a function from sets of worlds to sets of worlds.

Definition 2.6. A possible world w assigns every atomic sentence α a truth value. W is the set of all possible worlds. An information state s is a set of possible worlds. A context change potential is a function from one information state s to a new one. An interpretation function $[\cdot]$ assigns every sentence a context change potential. $s[\varphi]$ is the result of inputting s into $[\varphi]$. For any non-modal sentence φ, the classical content of φ ($[\![\varphi]\!]$) is the set of worlds that are unchanged by updating with φ ($\{w \mid \{w\}[\varphi] = \{w\}\}$).

Once we have a representation of information states, we can then recursively define our interpretation function, $[\cdot]$. Atomic sentences narrow down an information state to the worlds where they are true. A negation $\neg\varphi$ removes from a state any world that would survive updating with φ. A conjunction $\varphi \wedge \psi$ first updates a state with φ, and then updates the result with ψ.

Definition 2.7.

1. $s[\alpha] = \{w \in s \mid w(\alpha) = 1\}$

---

[6] See Stalnaker 1973; Karttunen 1974; Heim 1982; Heim 1983; Veltman 1985; Groenendijk and Stokhof 1990; Groenendijk and Stokhof 1991a; and many others.

[7] See Veltman 1996, p. 221.

2. $s[\neg\varphi] = s - s[\varphi]$

3. $s[\varphi \wedge \psi] = s[\varphi][\psi]$

Modal sentences don't give us new information about what world we are in. Instead, they explore properties of the current state. $\Diamond\varphi$ explores whether s can be consistently updated with φ; $\Box\varphi$ explores whether updating s with φ has no effect. In either case, states that satisfy the relevant property are unchanged by updating. States that fail to satisfy the relevant property result in the absurd state $\emptyset$.[8]

Definition 2.8.

1. $s[\Diamond\varphi] = \{w \in s \mid s[\varphi] \neq \emptyset\}$

2. $s[\Box\varphi] = \{w \in s \mid s[\varphi] = s\}$

This last definition implies that modal sentences are tests. They always either leave the information state the same, or produce the empty state.

Observation 2.3 (Tests). For any information state s, $s[\Diamond\varphi] = s$ or $s[\Diamond\varphi] = \emptyset$.

This feature is surprising from the perspective of a Stalnakerian theory of communication, on which the goal of assertion is to narrow down the common ground.[9] Modal sentences would then be guaranteed to never satisfy the goal of assertion.

Finally, to get predictions about the inconsistency of epistemic contradictions, we need a definition of entailment. Say that a state supports a sentence φ just in case updating the state with φ has no effect. Then an argument is dynamically valid just in case any state that supports the premises also supports the conclusion.[10]

Definition 2.9.

---

[8] See Veltman 1996. Officially, definitions 2.7 and 2.8 are one recursive definition.

[9] For example: "A speaker should not assert what he presupposes to be true, or what he presupposes to be false...To assert something incompatible with what is presupposed is self-defeating: one wants to reduce the context set, but not eliminate it altogether. And to assert something which is already presupposed is to attempt to do something that is already done." (Stalnaker 1978 88-9.)

[10] See Veltman 1996 and van Benthem 1996 139-41 for a discussion of the different entailment relations available within update semantics. Much of our discussion below could also be conducted using other entailment relations, for example update-to-test entailment.

1. s supports φ ($s \models \varphi$) iff $s[\varphi] = s$.

2. Γ entails δ ($\Gamma \models_{\text{US}} \delta$) iff for any information state s, if $s \models \gamma$ for every $\gamma \in \Gamma$, then $s \models \delta$.

To get a feel for this semantics, note that it successfully predicts the inconsistency of epistemic contradictions.

Observation 2.4 (Epistemic Contradictions). $\varphi \wedge \Diamond\neg\varphi \models_{\text{US}} \bot$

To see why, consider Figure 2.1. To update an information state with $\varphi \wedge \Diamond\neg\varphi$, one first updates with the first conjunct, zooming in to the worlds where φ is true. Then one updates the resulting state with $\Diamond\neg\varphi$. But, crucially, the $\Diamond\neg\varphi$ update is applied not to the original information state, but to the result of learning φ in that state. Once we have zoomed in to the φ worlds, the $\Diamond\neg\varphi$ test is guaranteed to fail. So the only state that supports $\varphi \wedge \Diamond\neg\varphi$ is absurd.

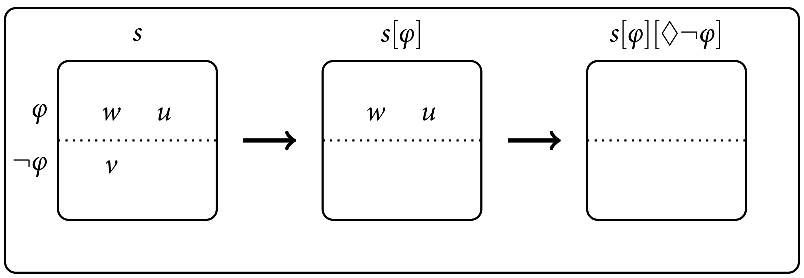

Figure 2.1: Epistemic contradictions

In addition to predicting the inconsistency of epistemic contradictions, update semantics has several surprising properties. First, since modals are tests, update semantics implies that iterated strings of modals are equivalent to their innermost member. In particular, it validates both of the following:

Observation 2.5 (Positive Introspection). $\Box\varphi \models_{\text{US}} \Box\Box\varphi$

Observation 2.6 (Negative Introspection). $\Diamond\varphi \models_{\text{US}} \Box\Diamond\varphi$

To see why, note that a state s is either changed by φ, or unchanged by φ. If the former, then $s[\Box\varphi]$ is $\emptyset$, and so $s[\Box\Box\varphi]$ is also $\emptyset$. If the latter, then $s[\Box\varphi] = s$, and so $s[\Box\Box\varphi] = s$ as well. So we know that $[\Box\varphi] = [\Box\Box\varphi]$. Similar reasoning applies to $\Diamond\varphi$.

From a big picture philosophical perspective, the collapse of iterated modality is somewhat surprising. Epistemic modals have something to do with knowledge. But knowledge doesn't obey positive and negative introspection.[11] Moreover, Moss 2015 argues on empirical grounds that the above are invalid. To do so, she develops cases where iterated modals have a different effect than their uniterated counterparts. For example, consider a scenario in which several candidates are being assessed for a job. Several factors are relevant (resume strength, recommendation letters, etc). Suppose that we know the status of some candidates with respect to a factor, but do not know how other candidates fare. Then sentences like the following become coherent:

(2) Alice is a likely hire, and Bob might be a likely hire.

(3) Alice is a possible hire, and Bob is probably also a possible hire.

(4) Alice is definitely a possible hire, and Bob might be a possible hire too.

The exact interpretation of such iterated modals is controversial.[12] But whether or not positive and negative introspection are actually valid for epistemic modals, we benefit from understanding what exactly they require of a theory.

Besides the collapse of iterated modals, update semantics also predicts that *must* is strong. That is, the T axiom is valid:

Observation 2.7 (Strength). $\Box\varphi \models_{\text{US}} \varphi$

After all, $\Box\varphi$ is supported by a state just in case φ is supported. But this prediction is quite controversial. Many have argued against the strength of *must* claims, citing the felt indirectness of claiming that φ must be the case, instead of simply asserting φ.[13]

We have now seen that update semantics explains the inconsistency of epistemic contradictions, but that it also forces us to make several more commitments about the logical properties of modals. Again, whether these predictions are good ones or not, our ultimate aim is to answer

[11] See for example Williamson 2000.

[12] For some options consistent with the collapse of iterated modals to their innermost modal, see Anand and Hacquard 2013 and Kratzer 2013.

[13] See again Karttunen 1972; Kratzer 1991; von Fintel and Gillies 2010; and Lassiter 2016.

a more abstract question. We will be investigating whether the inconsistency of epistemic contradictions is related to the validity of these various other logical principles, or whether they are conceptually independent. To do so, we will now develop a new semantics for epistemic modals within which the logical properties above are not mere consequences, but are instead characterizable in terms of some parameter in the theory. This will give us the tools to critically assess what exactly is required of a semantics in order to make epistemic contradictions inconsistent, to make *must* strong, or to collapse iterated modals.

2.3 Generalized update semantics

In *generalized update semantics* (GUS), we will integrate update semantics with the signature tool from the classic possible worlds semantics above: accessibility relations. To do so, we need one new idea. Notice how in that traditional framework, we can think of each accessibility relation as generating an information state for each possible world, from the set of worlds accessible from the base world. In GUS, we instead say that a world only determines a body of information relative to another, initial information state. In other words, we will let accessibility relations themselves be information sensitive. A traditional accessibility relation relates two possible worlds w and v. By contrast, an *information sensitive* accessibility relation instead relates a world w and information state s to any world v possible from the perspective of w and s.

Definition 2.10. An information sensitive accessibility relation $R_{(\cdot)}$ relates a world w and information state s to the set of worlds v that are possible relative to w, given information state s. When v is possible with respect to w and s, we say wR_sv.

Armed with information sensitive accessibility relations, we can now give a new dynamic meaning to $\Diamond$ and $\Box$. To update a state s with $\Diamond\varphi$, one narrows s down to a new set of worlds w. Which ones? Those where the information at w, relative to s, is consistent with φ. In particular, for every world w in s, we find the set of worlds accessible from w relative to s (that is, $\{v \mid wR_sv\}$). This set of worlds is itself an information state. We then see whether this information state can be updated with φ without becoming absurd. If so, w survives update with $\Diamond\varphi$.

Similarly, to update s with $\Box\varphi$, we again find $\{v \mid wR_sv\}$, the set of worlds accessible from w given s. For w to survive update with $\Box\varphi$, we require that this derivative information state support

φ.[14]

Definition 2.11.

1. $s[\Diamond\varphi] = \{w \in s \mid \{v \mid wR_s v\}[\varphi] \neq \emptyset\}$

2. $s[\Box\varphi] = \{w \in s \mid \{v \mid wR_s v\} \models \varphi\}$

Crucially, this theory of modals allows the information at each possible world to vary. This means that modal updates are not automatically tests. A modal update can eliminate some but not all worlds from the input state.

To get a feel for this semantics, let's work through an example (Figure 2.2). Imagine a state s with three worlds w, v, and u, where relative to s w and v can see each other, but u can see only itself. Suppose w and u are φ worlds, while v is not. Updating with $\Diamond\neg\varphi$ will retain w and v, while removing u. Conversely, updating s with $\Box\varphi$ would produce a new state that eliminates w and v, but retains u. While w and u are both φ worlds, only world u is isolated to see only φ worlds relative to s. Finally, we can consider which worlds are possible relative to each other in our new state $s[\Diamond\neg\varphi] = \{w, v\}$. In this model, we do not have $uR_{\{w,v\}}u$; so removing u from s also changes which worlds are possible relative to each other.

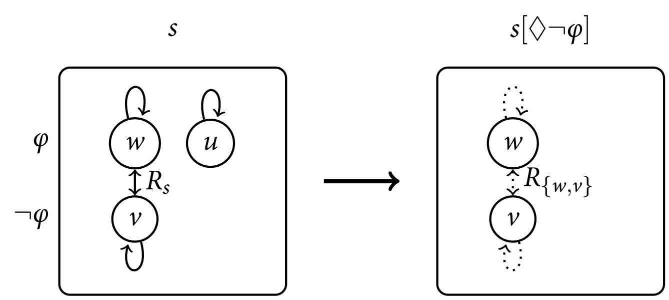

Figure 2.2: An example

So far we have considered one abstract example of our semantics at work. To get a better understanding of the theory, let's turn to a more concrete example of how information sensitive accessibility might function.

[14]Officially, Definitions 2.7 and 2.11 form one recursive definition. In addition, we will retain the definitions of support and entailment from Definition 2.9, but will refer to GUS entailment with $\models$.

2.4 Factorizability

In this section we will develop a procedure for deriving an information sensitive accessibility relation from an underlying classical accessibility relation. To do so, let's say that an information sensitive accessibility relation is *factorizable* into an underlying classical accessibility relation just in case the possibilities relative to w and s are simply the classical possibilities relative to w that are members of s.[15]

When an information sensitive accessibility relation can be constructed in this way, let's call it *factorizable* into a truth conditional and dynamic component. Let R be a classical accessibility relation.[16] Then:

Definition 2.12.

1. The dynamic lift of R is $R^{\uparrow}_{(\cdot)}$, where:

   $\forall s \forall w \forall v\ [wR^{\uparrow}_{s} v \text{ iff } (wRv\ \&\ v \in s)]$.

2. $R'_{(\cdot)}$ is factorizable into R iff $R'_{(\cdot)} = R^{\uparrow}_{(\cdot)}$.

3. $R_{(\cdot)}$ is factorizable iff there is some R that $R_{(\cdot)}$ is factorizable into.

When a relation is factorizable, we will automatically have that $\{v \mid wR_s v\}$ is equal to $\{v \mid wRv\} \cap s$, for some choice of R. That is, factorizable relations restrict the second argument of an underlying accessibility relation down to worlds in the relevant information state. What does this mean? We can illuminate this requirement in several ways.

First, we can think of factorizable accessibility relations in terms of syntactic and semantic scope. The inconsistency of epistemic contradictions can be thought of as following from a more general exceptional scope principle. We can say that our epistemic modals exhibit exceptional scope whenever they semantically interpret content outside their syntactic scope as if it was inside their scope. More precisely, consider the following equivalence:

EXCEPTIONAL SCOPE $\varphi \wedge \Diamond\psi \dashv\vDash \varphi \wedge \Diamond(\varphi \wedge \psi)$

---

[15] Thanks to Jeremy Goodman for help here.

[16] As we will see in detail later, any classical accessibility relation corresponds to a special kind of information sensitive accessibility relation.

The inconsistency of epistemic contradictions follows from the Exceptional Scope principle as a special case.[17] When ψ is of the form $\neg\varphi$, the Exceptional Scope principle implies that any consistent epistemic contradiction would require the consistency of $\Diamond(\varphi \wedge \neg\varphi)$.

Factorizable relations are tailor-made to explain exceptional scope. When a relation is factorizable, we can interpret $\Diamond$ as predicating possibility of the combination of its prejacent with the current information state. To illustrate, suppose for simplicity that φ is a non-modal, so that its possible worlds propositional content $[\![\varphi]\!]$ is defined. Then updating s with $\Diamond\varphi$ is identical to narrowing down s to the worlds where the proposition $s \cap [\![\varphi]\!]$ is classically possible—to the worlds where some world in $s \cap [\![\varphi]\!]$ is classically accessible. Similarly, updating with $\Box\varphi$ is a matter of zooming in to worlds where the material conditional if s then φ (that is, $(W - s) \cup [\![\varphi]\!]$) is classically necessarily—to the worlds where every classically accessible world is in $(W - s) \cup [\![\varphi]\!]$. To state this fact precisely, let's introduce for any traditional accessibility relation R the classical modals $\Diamond_R$ and $\Box_R$. Further, let's assign these modals traditional meanings as sentential operators. So $[\![\Diamond_R]\!]$ is a function from some proposition p to the set of worlds where p is possible (to the worlds w where some p-world v is R-accessible). Similarly, $[\![\Box_R]\!]$ is a function from p to the set of worlds where p is necessary. Then:

Fact 2.1. When (i) φ is non-modal, and (ii) $R_{(\cdot)}$ is factorizable into R:

1. $s[\Diamond\varphi] = s \cap [\![\Diamond_R]\!](s \cap [\![\varphi]\!])$
2. $s[\Box\varphi] = s \cap [\![\Box_R]\!]((W - s) \cup [\![\varphi]\!])$

(See Chapter 5 for proofs of the main facts in what follows.)

---

[17] This principle is a modal counterpart of a property of dynamic theories of anaphora, such as DPL (Groenendijk and Stokhof 1991a). In those theories, we reach equivalences like the following, for arbitrary indefinite quantifiers Q (such as 'some' and 'exactly one'):

ANAPHORIC EXCEPTIONAL SCOPE $Qx(Fx) \wedge Gx \dashv\models Qx(Fx) \wedge Qx(Fx \wedge Gx)$.

For example, consider the discourse:

(5) Exactly one lawyer came in. He sat down.

This sentence implies not only that exactly one lawyer came in and sat down, but also that there is exactly one lawyer (Charlow 2017).

Traditionally, modal operators were thought to predicate possibility and necessity of the proposition expressed by their input. Now, we can instead think of modal operators as predicating possibility or necessity of a proposition found by combining the current information state with the proposition expressed by their input. In the case of $\Diamond\varphi$, $\Diamond$ predicates possibility of the intersection of the current information with the proposition that φ. Similarly, $\Box\varphi$ allows $\Box$ to predicate necessity of the material conditional that if the current information is true, then so is the proposition that φ. Dynamic semantics certainly exploits propositions; but only after they have been enriched with the current information.

Here's a last way to think about factorizable modals. In general, any accessibility relation can be derived from a set of propositions. More precisely, following Kratzer 2012 we can let a modal base be a function from worlds to sets of propositions. Then the accessible worlds at w are the worlds consistent with every proposition in the modal base at w. With this background, we can think of information sensitive accessibility relations as allowing the modal base to be information sensitive: we now have a function from a world and an information state to a set of propositions. Now, factorizability corresponds to the constraint that the information state is always a member of the modal base at a world, relative to that information state. This is a very natural constraint. Epistemic modals care about a body of relevant information. Factorizable modals require that the current information state is always relevant.

Now that we have developed GUS and considered one instance of it, let's put the framework to work. First, I will show that the semantics above is conservative, in that both possible worlds semantics and update semantics are special cases of generalized update semantics. Second, I will show that generalized update semantics is fruitful. It turns out that a rich family of logical principles governing modals are characterizable in terms of constraints on accessibility relations.

2.5 Conservativity

A natural question to ask in this framework is whether we can express previous theories of modals within it. It turns out that both traditional possible worlds semantics and update semantics correspond to particular constraints on information sensitive accessibility.

To generate the traditional possible worlds semantics we started with, we can allow information sensitive accessibility to ignore the information state. On this view, v is possible with respect to w and s just in case v is possible with respect to w and s', for any arbitrary s'. That is, the choice of information state is irrelevant to what is possible.

Definition 2.13.

1. $R_{(\cdot)}$ is information insensitive just in case for any w, s, s', $\{v \mid wR_s v\} = \{v \mid wR_{s'} v\}$.

2. When $R_{(\cdot)}$ is information insensitive, say that $wR_{\downarrow}v$ iff for every s, $wR_s v$.

When a relation $R_{(\cdot)}$ is informative insensitive, we can use it to construct a traditional accessibility relation, using $\downarrow$. The result will relate w and v just in case the original relation connects w and v relative to any s. Then updating s with $\Diamond\varphi$ narrows down s to the possible worlds where φ is classically possible via $R_{\downarrow}$:

Fact 2.2. If $R_{(\cdot)}$ is information insensitive, $s[\Diamond\varphi] = s \cap \{w \mid \exists v : wR_{\downarrow}v \;\&\; v \in [\![\varphi]\!]\}$.

This same point can be understood another way. Say that a formula φ is *distributive* when updating any state with φ is equivalent to updating each world in the state with φ and unioning the results. Failures of distributivity are often thought of as essential to dynamic semantics.[18] This is because in the framework above, a sentence is distributive just in case updating any state with that sentence is tantamount to intersecting the state with some fixed set of possible worlds.

Definition 2.14. φ is distributive iff for any state s, $s[\varphi] = \bigcup_{w \in s} \{w\}[\varphi]$.

It turns out that the distributivity of modal claims is more or less equivalent to the information insensitivity of $R_{(\cdot)}$. That is, information sensitivity is equivalent, in our current framework, to genuine dynamics.

To see why, note first that our information sensitive accessibility relations are a bit richer than they need to be. These relations connect w and v via s. But these relations do not require that the initial world w be a member of s. Nonetheless, inspection of our semantics above reveals that we can safely ignore any cases where $w \notin s$. Any such w is irrelevant to the context change potential

[18] See van Benthem 1996; Veltman 1996; and Rothschild and Yalcin 2015b.

of any sentence in our language, since $\Diamond\varphi$ and $\Box\varphi$ narrow s down to the worlds that are related to the relevant φ worlds. So we can call any two relations semantically equivalent when they agree on accessibility for any choice of w in s:

Definition 2.15. $R_{(\cdot)}$ and $R'_{(\cdot)}$ are semantically equivalent iff for any w, v, and s, if $w \in s$ and $w \in s'$, then $wR_s v$ iff $wR'_{s'} v$.

Now we can state precisely the relationship between information insensitivity and distributivity. First, if accessibility is information insensitive then all modal sentences are distributive. Second, whenever our modals are distributive, they induce an accessibility relation that is semantically equivalent to an information insensitive relation.

Fact 2.3.

1. If $R_{(\cdot)}$ is information insensitive, then $\Diamond\varphi$ is distributive for every φ.
2. If $\Diamond\varphi$ is distributive for every φ, then $R_{(\cdot)}$ is semantically equivalent to an information insensitive $R'_{(\cdot)}$.

We have now seen that a traditional possible worlds semantics results whenever $R_{(\cdot)}$ is insensitive to its informational argument. Similarly, the traditional update semantics for modals results when $R_{(\cdot)}$ ignores its world argument. Here, we need to take some care. We can first say that a relation is world insensitive just in case $wR_s v$ iff $w'R_s v$. But this alone does not give us the test semantics from before. For it is consistent with this requirement that, relative to s, every world is related to some set of worlds s' distinct from s. So we can instead say that a relation is *strongly* world insensitive just in case relative to s, w sees exactly the worlds in s. The test semantics above turns out to be characterizable in terms of exactly this relation:

Definition 2.16. $R_{(\cdot)}$ is strongly world insensitive just in case for any w, s $\{v \mid wR_s v\} = s$.

Fact 2.4. $R_{(\cdot)}$ is strongly world insensitive iff for every state s, $s[\Diamond\varphi] = \{w \in s \mid s[\varphi] \neq \emptyset\}$.

We can also understand the test semantics in terms of factorizability. Recall that we earlier introduced a lift from any traditional accessibility relation R into $R^{\uparrow}_{(\cdot)}$, an information sensitive accessibility relation. This lift generates a new relation where v is possible from w and s just in case v

is classically possible from w and v is in s. It turns out that the test semantics corresponds to one special instance of this lift, where we start from the universal accessibility relation ($W \times W$).

Fact 2.5. $(W \times W)^{\uparrow}_{(\cdot)}$ is the unique strongly world insensitive relation.

So the test semantics from earlier is produced by starting with the purest classical conception of possibility (where every world is possible from every other), and lifting it into a dynamic meaning.

Summing up, we have seen that each theory we started with occupies an endpoint on a spectrum of GUS meanings. GUS possibility is sensitive to the world and the information state. When the world is utterly ignored, we reach the test semantics. When the information state is ignored, we reach the classical semantics. But in between these two endpoints is a wide space of meaning. Now we can go on to see what new work we can do with our more complex theory of modality.

2.6 Epistemic contradictions

Now that accessibility relations are information sensitive, we can find a new set of constraints that characterize the logical principles above. Here, we will zoom in to cases where φ is non-modal, so that $[\![\varphi]\!]$ denotes the set of worlds where φ is true ($\{w \mid \{w\} \models \varphi\}$).

Our first goal is to characterize the inconsistency of epistemic contradictions. This principle turns out to be equivalent to the requirement that only worlds consistent with a body of information are possible relative to that information. More precisely: that whenever w is in s, the set of worlds visible from w and s is a subset of s.

Definition 2.17. $R_{(\cdot)}$ is eliminative iff $\forall s \forall v \forall w \in s\ [wR_s v \supset v \in s]$.

Fact 2.6. $R_{(\cdot)}$ is eliminative iff $\varphi \wedge \Diamond\neg\varphi \models \bot$.

(For proofs of the main facts in what follows, see the appendix.) To get a feeling for why the above holds, consider the model in Figure 2.3, where eliminativity fails. Here we have a world w in state s that can see a world outside s. Now let φ be some formula true at w (and any other world in s) and false at the world outside s. Updating with φ has no effect on s. Similarly, updating with $\Diamond\neg\varphi$ keeps the state the same, since w can see a $\neg\varphi$ world. So the epistemic contradiction $\varphi \wedge \Diamond\neg\varphi$ is supported by this state.

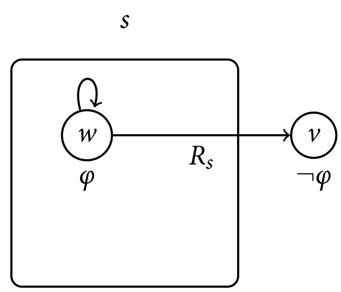

Figure 2.3: A failure of eliminativity

Fact 2.6 has several consequences. First, it is clear that there are many eliminative relations that are not strongly world insensitive. So the inconsistency of epistemic contradictions does not actually require the test semantics for modals from Veltman 1996. Epistemic contradictions may be inconsistent, even though possibility claims tell us more about which world we are in.

Nonetheless, the test operator occupies a privileged place among eliminative possibility modals. It turns out that the test semantics for $\Diamond$ is the weakest eliminative possibility modal. Let $\Diamond$ be a possibility modal with accessibility relation $R_{(\cdot)}$. Let $\Diamond_t$ be the test operator (so that $s[\Diamond_t\varphi] = \{w \in s \mid s[\varphi] \neq \emptyset\}$). Then:

Fact 2.7. If $R_{(\cdot)}$ is eliminative, then $\Diamond\varphi \models \Diamond_t\varphi$.

The test semantics for $\Diamond$ is the weakest theory of possibility modals which predicts that epistemic contradictions are inconsistent.

While the inconsistency of epistemic contradictions does not require a test semantics for modals, it does require a genuinely dynamic semantics for epistemic modals. We saw above that our semantics is equivalent to possible world semantics just in case our information sensitive accessibility relation is information insensitive. But whenever this occurs, the inconsistency of epistemic contradictions is trivializing:

Fact 2.8. If $R_{(\cdot)}$ is information insensitive and $\varphi \wedge \Diamond\neg\varphi \models \bot$ for all descriptive φ, then $wR_{\downarrow}v$ is isolated.

While classical accounts of modals do not predict the inconsistency of epistemic contradictions, factorizable theories do.

Fact 2.9. If $R_{(\cdot)}$ is factorizable, then $R_{(\cdot)}$ is eliminative.

Fact 2.9 is one of the main results of this chapter, because it offers a recipe for converting any classical theory of modals into a dynamic explanation of epistemic contradictions. Factorizability shows us how to take any static theory of epistemic modals, encode it into an R and then convert it into $R^{\uparrow}_{(\cdot)}$, to explain the inconsistency of epistemic contradictions. For example, start with a contextualist theory of modals, where $\Diamond\varphi$ says that the speaker's knowledge is consistent with φ. Say that wCv iff the speaker's knowledge at w is consistent with v. Then $C^{\uparrow}_{(\cdot)}$ gives the dynamic lift of this contextualist theory, which predicts that epistemic contradictions are inconsistent.

This last point allows a division of labor in semantic theorizing. Over the last few decades, there has been a flurry of truth conditional investigation into epistemic modals. There have been a variety of debates about whose information matters for the evaluation of epistemic modal claims (the speaker, the group, some exocentric point of view?) and what kind of information is important (belief, knowledge, certainty?). These debates can all be thought of as debates about an underlying accessibility relation. But these debates are not concerned with epistemic contradictions. Then dynamic theories of modality came along that predicted the inconsistency of epistemic contradictions, but also took a stand on these old debates—for example requiring that *must* be strong. What we have now seen is that the question of epistemic contradictions is somewhat independent of other debates that were characterizable in terms of an underlying accessibility relation. Once any of those debates have been settled, and an appropriate accessibility relation is determined, that relation can be lifted with $\uparrow$ to generate a factorizable information sensitive accessibility relation that predicts the inconsistency of epistemic contradictions. So we have a clear division of labor between static and dynamic approaches to modality.

2.7 Strength

We've now determined what the inconsistency of epistemic contradictions amounts to in GUS, and seen that it requires dynamic semantics, but does not require a test semantics. Now let's turn to some other principles. First, let's consider the thesis that *must* is strong. The T axiom that $\Box\varphi \models \varphi$ is no longer equivalent to the constraint that every world is accessible from itself; that constraint ignores information sensitivity. Rather, the T axiom is now equivalent to the requirement that for any state s and any w in s, there is some world in s that can see w. This constraint is essentially *global*, in that it involves properties of s not reducible to individual worlds. Instead of requiring that every world sees itself, global reflexivity requires that every part of a state is seen by some part of that state.

Definition 2.18. $R_{(\cdot)}$ is globally reflexive iff $\forall s \forall w \in s\, \exists w' \in s : w' R_s w$.

Fact 2.10. $R_{(\cdot)}$ is globally reflexive iff $\Box\varphi \models \varphi$.

For a sense of why this holds, consider Figure 2.4. Here we have a world w in state s that is not seen by any world in s relative to s. When φ is true at every visible world in s but false at w, s will support $\Box\varphi$ but not φ.

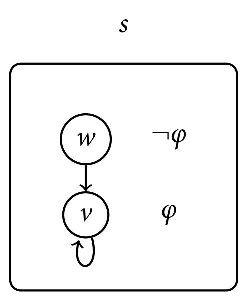

Figure 2.4: A failure of global reflexivity

One interesting feature of the present semantics is that entailment is not automatically guaranteed to contrapose. For example, in update semantics $\varphi \models_{\text{US}} \Box\varphi$, but we do not have the contrapositive $\Diamond\varphi \models_{\text{US}} \varphi$. In the present framework, then, it is worth considering the contrapositive of T, that $\varphi \models \Diamond\varphi$. This requirement is equivalent to an inverse of global reflexivity—that for any state s and any w in s, w can see some world in s.[19]

Definition 2.19. $R_{(\cdot)}$ is globally inverse reflexive iff $\forall s \forall w \in s\, \exists w' \in s : w R_s w'$.

Fact 2.11. $R_{(\cdot)}$ is globally inverse reflexive iff $\varphi \models \Diamond\varphi$.

When a relation is information insensitive, global reflexivity and global inverse reflexivity coincide, and are equivalent to the normal reflexivity condition (consider $s = \{w\}$). Once we give

[19]Thanks to Fabrizio Cariani for help here.

up information insensitivity, these two properties are two different generalizations of ordinary reflexivity.

In our framework, global reflexivity is not immediately implied by eliminativity. This means that our semantics in principle allows for theories on which epistemic contradictions are inconsistenct, but *must* is not strong. Nonetheless, we can use factorizability to give a powerful new argument for strength. First, it turns out that for factorizable relations, global reflexivity and reflexivity are equivalent:

Fact 2.12. Suppose $R_{(\cdot)}$ is factorizable into R. Then $R_{(\cdot)}$ is globally reflexive iff R is reflexive.

Given factorizability, we can also understand global reflexivity in another way, which will give us our new argument for strength. Our argument requires one premise: that modal sentences be idempotent.

Definition 2.20. φ is idempotent iff for any s, $s[\varphi] \models \varphi$.

Idempotence is the requirement that learning be successful. This is a natural constraint, satisfied by modal sentences in Veltman 1996.[20]

We will now see that within factorizable theories, idempotence comes close to implying that *must* is strong. We will consider the case of both $\Box\varphi$ and $\Diamond\varphi$. First, it turns out that it is much easier to make $\Box\varphi$ idempotent than to make $\Diamond\varphi$ idempotent.

Fact 2.13. If $R_{(\cdot)}$ is factorizable and φ is descriptive, then $\Box\varphi$ is idempotent.

While factorizability is sufficient for the idempotence of necessity claims, it is not sufficient for the idempotence of possibility claims. Given factorizability, however, the idempotence of possibility claims is equivalent to the requirement that accessibility be shift-reflexive, so that any world visible at all is visible to itself:

Definition 2.21. R is shift-reflexive iff for $\forall w \forall v[wRv \supset vRv]$.

Fact 2.14. Suppose $R_{(\cdot)}$ is factorizable into R. Then $\Diamond\varphi$ is idempotent for any descriptive φ iff R is shift-reflexive.

---

[20] Here, a modal sentence is any sentence with $\Diamond$ and $\Box$ at outermost scope. The test semantics does predict failures of idempotence for sentences of the form $\Diamond\varphi \wedge \neg\varphi$; but it does not predict failures of idempotence for sentences with a modal operator at widest scope.

Shift reflexivity is equivalent to the validity of $\Box(\Box\varphi \supset \varphi)$, which is extremely close to the T axiom that $\Box\varphi \models \varphi$. It would be bizarre to accept shift reflexivity (any visible world can see itself) while rejecting reflexivity (every world sees itself). To do so requires that some worlds are not seen by any other possible world. For this reason, the leading opponents of reflexivity also deny shift-reflexivity.[21]

Summing up, Fact 2.14 constitutes a powerful new argument for the thesis that *must* is strong. Current discussions of strength focus on particular examples where speakers use *must* sentences to convey some sort of indirect information.[22] These cases are difficult to explain, because it is unclear whether the relevant phenomenon should be analyzed as an ordinary entailment or as some sort of presuppositional meaning. Our argument for *strength* is more abstract. We have sidestepped any direct judgments about the use of *must p* versus *p*, and shown that a pair of elegant structural properties connect epistemic contradictions and strength. Our only assumption was the theory-neutral, universally attractive property of idempotence.

2.8 Positive introspection

Now let's turn to a second consequence of epistemic contradictions: the positive introspection principle. In GUS, the principle that $\Box\varphi \models \Box\Box\varphi$ is equivalent to a generalization of transitivity. Global transitivity requires that whenever v is visible from w relative to s, and u is visible from v relative to the information at w in s, then u can be seen relative to s by some world in s. This generalizes transitivity in two ways. First, it is *global*: w does not need to see u; rather, some world in s must see u. Second, it is *shifty*: we first consider what is accessible relative to s, and then consider what is accessible relative to a different state. (In general, any principle governing nested modals will involve such a shift in the relevant state.)

Definition 2.22. $R_{(\cdot)}$ is globally transitive iff $\forall s \forall v \forall u \forall w \in s : [wR_s v \;\&\; vR_{\{v|wR_s v\}} u \supset \exists w' \in s : w' R_s u]$.

Fact 2.15. $R_{(\cdot)}$ is globally transitive iff $\Box\varphi \models \Box\Box\varphi$.

---

[21] See for example Kratzer 1991.

[22] See again: Karttunen 1972; Kratzer 1991; von Fintel and Gillies 2010; and Lassiter 2016.

For an illustration of this last fact, consider Figure 2.5. Here we have a case where global transitivity fails. u is not visible from any point in s relative to s. Nonetheless, u is visible from v relative to the information at w and s, and v is visible from w relative to s. To see why positive introspection fails, let φ be true at w and v but false at u. $\Box\varphi$ is supported in s, since every world visible from a world in s relative to s is a φ world. Nonetheless, $\Box\Box\varphi$ eliminates w from s. This is because relative to the information at w and s, u is possible from v.

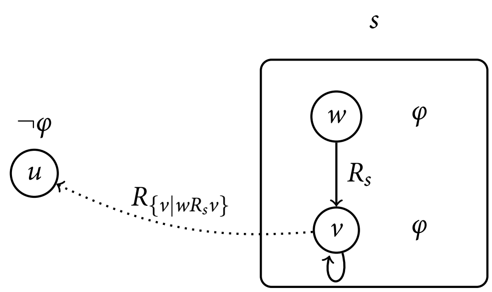

Figure 2.5: A failure of global transitivity

Again, one of the main advantages of the current framework is that we can determine exactly what follows from the inconsistency of epistemic contradictions. Here is a case in point: the inconsistency of Epistemic Contradictions implies that Positive Introspection is valid. This is somewhat surprising, as we have only assumed that epistemic contradictions are inconsistent for non-modal φ, and yet this has repercussions for formulas like $\Box\Box\varphi$, which involve multiple modal operators.

Fact 2.16. If $R_{(\cdot)}$ is eliminative, then $R_{(\cdot)}$ is globally transitive.

We can also understand this fact from a more abstract perspective. For suppose that $\Box$ is upwards monotonic, so that:

UPWARD MONOTONICITY If $\varphi \models \psi$, then $\Box\varphi \models \Box\psi$.

The inconsistency of epistemic contradictions quickly implies that $\varphi \models \Box\varphi$. When we instantiate this inference in upward monotonicity, we reach positive introspection.

Fact 2.17. Suppose Upward Monotonicity is valid and Epistemic Contradictions are inconsistent. Then Positive Introspection is valid.

Putting together the last few sections, we have now seen that the inconsistency of epistemic contradictions places heavy demands on the logic of epistemic modals. In particular, we have seen that this logic must be at least as strong as the combination of the inconsistency of epistemic contradictions with the modal logic S4 (characterized by positive introspection and the T axiom).

2.9 Avoiding collapse

A natural question at this point is whether the inconsistency of epistemic contradictions implies that iterated modals always collapse to the innermost modal. But this is not so. Eliminativity does not imply that negative introspection is valid. For the negative introspection principle that $\Diamond\varphi \models \Box\Diamond\varphi$ is equivalent to the requirement that whenever w in s can see some world v relative to s, there is some world in s where every world it can see is visible to v relative to the worlds w can see.

Definition 2.23. $R_{(\cdot)}$ is globally euclidean iff $\forall s \forall v \forall w \in s : [wR_s v \supset \exists w' \in s : \forall v'[w'R_s v' \supset vR_{\{v|wR_s v\}} v']]$.

Fact 2.18. $R_{(\cdot)}$ is global euclidean iff $\Diamond\varphi \models \Box\Diamond\varphi$.

So one way we can avoid the collapse of iterated modals is to accept that $R_{(\cdot)}$ is eliminative, but reject the claim that it is globally euclidean. While the resulting modals obey positive introspection, they do not satisfy negative introspection.

In a similar vein, we can separate the inconsistency of epistemic contradictions from the validity of the B axiom, on which truth implies necessary possibility (that $\varphi \models \Box\Diamond\varphi$). In our framework, this principle corresponds to a generalization of the symmetry requirement on accessibility. Global symmetry requires that whenever w can see v relative to s, there is some world in s that can be seen by v relative to whatever w could see relative to s.

Definition 2.24. $R_{(\cdot)}$ is globally symmetric iff $\forall s \forall v \forall w \in s : [wR_s v \supset \exists w' \in s : vR_{\{v|wR_s v\}} w']$.

Fact 2.19. R satisfies global symmetry iff $\varphi \models \Box\Diamond\varphi$.

In the previous section, we saw that epistemic contradictions require a modal logic at least as strong as S4. Here, we have seen that epistemic contradictions do not require a modal logic as strong as S5, where iterated modals collapse completely.

2.10 Conclusion

This chapter has introduced a simple new tool: information sensitive accessibility. We implanted this tool within the meaning of modal operators to uncover a wide family of new meanings. A few results stand out. First, we developed tools for transforming any classical modal operator into a dynamic operator that predicts the inconsistency of epistemic contradictions. Second, we uncovered the relationship between the inconsistency of epistemic contradictions and a variety of other logical principles. In particular, we showed that epistemic contradictions require that *must* is strong, and that positive introspection is valid.

There are at least two applications of GUS that require further research. First, a natural question is how to extend the semantics for modality above to the theory of conditionals. For example, dynamic theories of conditionals have had great success validating the Import-Export principle.[23] But what exactly does the validity of this principle require of a dynamic semantics for conditionals? Second, a further challenge remains to integrate the dynamic semantics for modals above with leading dynamic theories of anaphora.[24] For example, will such an integration continue to validate the exceptional scope principles discussed above? In future work, I hope to explore both of these avenues of inquiry.

In an appendix, I present GUS with greater precision.

[23]See Gibbard 1981, McGee 1985, Gillies 2004, and Gillies 2009.

[24]For some attempts, see Groenendijk et al. 1996 and Büring 1998.

Appendix

Definition 2.25.

1. Let a *GUS* frame $\mathcal{F}$ be a tuple $\langle W, I, R_{(\cdot)} \rangle$, where:
   (a) W is a set of possible worlds.
   (b) The set of information states I is $\mathcal{P}(W)$.
   (c) An information sensitive accessibility relation $R_{(\cdot)}$ relates a world w in W and an s in I to a world v in W.
2. Let a *GUS* model $\mathcal{M}$ be a pair $\langle \mathcal{F}, V \rangle$ of a *GUS* frame and a valuation function V, where V assigns a truth value in $\{0, 1\}$ to every atomic sentence.

Definition 2.26. Let $\mathcal{M} = \langle F, V \rangle$ be any GUS model. An interpretation function $[\cdot]_{\mathcal{M}}$ assigns each sentence in $\mathcal{L}$ a context change potential, as follows.

1. $s[\alpha] = \{w \in s \mid w \in V(\alpha)\}$
2. $s[\neg\varphi] = s - s[\varphi]$
3. $s[\varphi \wedge \psi] = s[\varphi][\psi]$
4. $s[\Diamond\varphi] = \big\{w \in s \mid \{v \mid wR_s v\}[\varphi] \neq \emptyset\big\}$
5. $s[\Box\varphi] = \big\{w \in s \mid \{v \mid wR_s v\}[\varphi] = \{v \mid wR_s v\}\big\}$

Definition 2.27.

1. $s \models_{\mathcal{M}} \varphi$ iff $s \in I_{\mathcal{M}}$ and $s[\varphi]_{\mathcal{M}} = s$.
2. $\Gamma \models_{\mathcal{F}} \delta$ iff for any *GUS* model $\mathcal{M}$ containing $\mathcal{F}$: for any $s \in I_{\mathcal{M}}$, if $s \models_{\mathcal{M}} \gamma$ for every $\gamma \in \Gamma$, then $s \models_{\mathcal{M}} \delta$.

Definition 2.28.

1. $R_{(\cdot)}$ is eliminative iff $\forall s \forall v \forall w \in s\ [wR_s v \supset v \in s]$.

2. $\mathcal{F} = \langle W, I, R_{(\cdot)} \rangle$ is eliminative iff $R_{(\cdot)}$ is eliminative.

Definition 2.29.

1. $R_{(\cdot)}$ is globally reflexive iff $\forall s \forall w \in s\ \exists w' \in s : w' R_s w$.

2. $\mathcal{F} = \langle W, I, R_{(\cdot)} \rangle$ is globally reflexive iff $R_{(\cdot)}$ is globally reflexive.

Definition 2.30.

1. $R_{(\cdot)}$ is globally inverse reflexive iff $\forall s \forall w \in s\ \exists w' \in s : w R_s w'$.

2. $\mathcal{F} = \langle W, I, R_{(\cdot)} \rangle$ is globally inverse reflexive iff $R_{(\cdot)}$ is globally inverse reflexive.

Definition 2.31.

1. $R_{(\cdot)}$ is globally transitive iff $\forall s \forall v \forall u \forall w \in s : [w R_s v \ \&\ v R_{\{v \mid w R_s v\}} u \supset \exists w' \in s : w' R_s u]$.

2. $\mathcal{F} = \langle W, I, R_{(\cdot)} \rangle$ is globally transitive iff $R_{(\cdot)}$ is globally transitive.

Definition 2.32.

1. $R_{(\cdot)}$ is globally euclidean iff $\forall s \forall v \forall w \in s : [w R_s v \supset \exists w' \in s : \forall v' [w' R_s v' \supset v R_{\{v \mid w R_s v\}} v']]$.

2. $\mathcal{F} = \langle W, I, R_{(\cdot)} \rangle$ is globally euclidean iff $R_{(\cdot)}$ is globally euclidean.

Definition 2.33.

1. $R_{(\cdot)}$ is globally symmetric iff $\forall s \forall v \forall w \in s : [w R_s v \supset \exists w' \in s : v R_{\{v \mid w R_s v\}} w']$.

2. $\mathcal{F} = \langle W, I, R_{(\cdot)} \rangle$ is globally symmetric iff $R_{(\cdot)}$ is globally symmetric.

Chapter 3

Free choice impossibility results

3.1 Introduction

When disjunctions interact with possibility modals, they famously give rise to a Free Choice inference.[1] For example, (1-a) appears equivalent to (1-b):

(1) a. Mary might be in New York or Los Angeles.
 b. Mary might be in New York and Mary might be in Los Angeles.

In (1-a), the possibility modal *might* ($\Diamond$) takes scope over the disjunction *Mary is in New York or Los Angeles* ($\mathsf{nyc} \vee \mathsf{la}$). The possibility of a disjunction appears equivalent to a conjunction of possibility claims.

Free Choice is a surprisingly robust phenomenon. First, it occurs with many different flavors of modality. Above, we looked at the epistemic case. Here's a deontic variant:

(2) a. Mary may go to New York or Los Angeles.
 b. Mary may go to New York and Mary may go to Los Angeles.

Willer 2015 observes that Free Choice also occurs in the consequent of conditionals:

(3) a. If Mary doesn't go to Pisa, she might go to Lisbon or Rome.
 b. If Mary doesn't go to Pisa, she might go to Lisbon, and if Mary doesn't go to Pisa, she might go to Rome.

Again, this is independent of modal flavor: it also occurs in subjunctive conditionals.

[1] See von Wright 1968 and Kamp 1974.

(4) a. If Mary had not gone to Pisa, she might have gone to Lisbon or Rome.

b. If Mary had not gone to Pisa, she might have gone to Lisbon, and if Mary had not gone to Pisa, she might have gone to Rome.

Free Choice also persists in other embedded environments. For example, Miller 2016 observes the phenomenon with attitude verbs:

(5) a. Susan thinks Mary might be in New York or Los Angeles.

b. Susan thinks Mary might be in New York and Susan thinks Mary might be in Los Angeles.

While systematic, Free Choice is surprising. A standard possible worlds semantics for possibility and disjunction predicts that Free Choice is invalid. On that theory, (1-a) requires that there is a possible world where one of nyc or la is true. (1-b) instead requires that there is one possible world where nyc is true, and another world where la is true. These two conditions are logically independent.

Free Choice isn't just surprising; it is also paradoxical. Kamp 1974 showed that there is serious tension between the validity of Free Choice, Disjunction Introduction, and the upwards monotonicity of $\Diamond$:

DISJUNCTION INTRODUCTION $\varphi \models \varphi \vee \psi$

UPWARDS MONOTONICITY If $\varphi \models \psi$, then $\Diamond\varphi \models \Diamond\psi$

With these assumptions, Free Choice implies the equivalence of any two possibility claims:

EXPLOSION $\Diamond\varphi \models \Diamond\psi$

Observation 3.1 (Kamp). Assume Transitivity and Upwards Monotonicity of $\Diamond$. Then Free Choice and Disjunction Introduction imply Explosion.

After all, $\Diamond\varphi$ implies $\Diamond(\varphi \vee \psi)$ by Disjunction Introduction and Upwards Monotonicity, which by Free Choice implies $\Diamond\psi$.

Nonetheless, a variety of recent work in philosophy and linguistics has attempted to validate Free Choice. Semantic theories of Free Choice have, one way or another, proposed a radical departure from a traditional intensional semantics for disjunction. In that tradition, $\vee$ represents a propositional operator, which takes two sets of possible worlds as input and returns a new set of possible worlds (say, their union). Proponents of Free Choice have departed from a propositional theory of disjunction. For example, Kratzer and Shimoyama 2002, Simons 2005 and Aloni 2007 propose an alternative semantics for disjunction in the tradition of Hamblin 1973, on which the meaning of $\varphi \vee \psi$ is not a set of possible worlds, but is instead a set of propositions: the set containing the propositions φ and ψ. Fusco 2015 proposes a two-dimensional semantics for disjunction, on which disjunctions denote functions from worlds to sets of propositions (the disjuncts true at that world). Charlow 2015, Willer 2015, and Starr 2016 have all offered some form of dynamic semantics for disjunction. On these proposals, the meaning of $\varphi \vee \psi$ is not a set of possible worlds, but is instead a *context change potential*, a rule for modifying bodies of information. Finally, Zimmermann 2000 proposes that sentences with the English expression *or* are actually, at the level of LF, disguised conjunctions of possibility claims.[2] Here the expression $\varphi \vee \psi$ is not even an LF, let alone semantically interpreted as a set of possible worlds.[3]

All of these proposals for validating Free Choice require a radical departure from an orthodox theory of disjunction. It is natural to wonder whether validating Free Choice requires such a departure.[4]

In this chapter, I argue that validating Free Choice really does require this sort of radical revision. In particular, I prove three main results. First, I consider the prospects for validating Free Choice within a traditional possible worlds semantics. Say that disjunction is propositional just in case the meaning of $\vee$ is a function from two sets of worlds to a new set of worlds. Now hold fixed a standard propositional theory of possibility modals. On this view, possibility modals are existential quantifiers over a set of worlds, saying that their prejacent holds at one of those worlds. The first main result is that given a propositional theory of possibility modals, there is no way to

---

[2] See Geurts 2005 for discussion.

[3] Strictly speaking, Zimmermann 2000 focused on a close cousin of what I call Free Choice above, where disjunction takes wide scope to the possibility modal.

[4] Of course, a semantic account of Free Choice may not be necessary in the first place. See Aloni 2007 and Fox 2009 for well developed pragmatic accounts of Free Choice.

validate Free Choice with a propositional theory of disjunction.

Given this first result, there are two natural strategies one might employ to validate Free Choice. We could give up a traditional theory of possibility modals, or give up a traditional theory of disjunction. Our second main result shows that the first option alone is not so promising. In particular, we consider the prospects for validating Free Choice within a dynamic theory of meaning. On this theory, possibility modals are *tests*, which explore whether an agent's information state can consistently be updated with their complement. Then say that disjunction is propositional in this framework just in case updating an agent's information with that disjunction is simply a matter of learning that some fixed proposition is true. The second main result is that validating Free Choice within this dynamic framework also requires a nonpropositional theory of disjunction.

With these negative results in hand, I turn to a positive theory of Free Choice. I show that there are a variety of meanings for disjunction that validate Free Choice within a simple dynamic semantics. This family of meanings is well behaved. In particular, it has a weakest and strongest member. Ultimately, I argue that the meaning of a disjunction can be factorized into a dynamic and a truth conditional component. The dynamic component requires that each disjunct is possible, while the truth conditional component requires that one of the disjuncts is true. This leads to our third main result: that this meaning for disjunction is actually the strongest possible meaning that can validate Free Choice within a simple dynamic semantics. This result explicates the connection between the validity of free choice and the truth conditions for disjunction. The traditional truth conditions for disjunction are the strongest truth conditions that can be assigned to *or*, consistent with the validity of Free Choice.

3.2 Assumptions

Our main negative result is that Free Choice is inconsistent with a propositional theory of disjunction. To prove this, we first need to state Free Choice more precisely. The most natural statement of free choice would be to say that $\Diamond(\varphi \vee \psi)$ is equivalent to $\Diamond\varphi \wedge \Diamond\psi$. However, our goal below is to determine the constraints that Free Choice imposes on the meaning of disjunction. To do so, it helps simplify the statement of Free Choice as much as possible. For this reason, we can omit any object language use of conjunction. Rather, we can let Free Choice consist in the validity of the following principles:

FC I $\Diamond(\varphi \vee \psi) \models \Diamond\varphi$ and $\Diamond(\varphi \vee \psi) \models \Diamond\psi$

FC II $\Diamond\varphi; \Diamond\psi \models \Diamond(\varphi \vee \psi)$

The principles above only contain two operators: $\Diamond$ and $\vee$. So our strategy in the rest of this chapter will be to consider various meanings for possibility modals, and show how these different meanings combine with Free Choice to constrain the meaning of *or*.

Since our principle of interest only contains possibility modals and disjunction, we restrict our attention below to the interpretation of a sparse formal language, consisting of a stock of atomic sentences supplemented by a disjunction operator and a family of possibility modals. As we saw above, Free Choice appears valid for a variety of modal flavors. So here we introduce not one but several possibility modals. Yet we have no need for either negation or conjunction operators.

Definition 3.1. Let $\mathscr{L}$ be a language consisting of a set of atomic formulae $\alpha, \beta, \ldots$ closed under $\vee$ and a family of possibility operators $\Diamond_1, \ldots, \Diamond_n$.

Our first result concerns a simple intensional semantics, in which the meaning of any sentence is a set of possible worlds.

Definition 3.2. Let a possible world w be a function from atomic sentences to truth values. Let W be the set of possible worlds. Let $A, B, \ldots$ be variables for subsets of W. Let $[\![\cdot]\!]$ be a function from $\mathscr{L}$ to $\mathcal{P}(W)$.

Our result focuses on the space of possible meanings for disjunction. So we hold fixed a traditional theory of all the expressions in our language other than $\vee$. First, we assume that atomic sentences are semantically associated via $[\![\cdot]\!]$ with the set of worlds at which they are true.

Definition 3.3. $[\![\alpha]\!] = \{w \in W \mid w(\alpha) = 1\}$

Then we hold fixed a standard possible worlds semantics for our possibility operators $\Diamond_1, \ldots, \Diamond_n$. Each possibility operator $\Diamond_i$ gets its distinctive flavor by being associated with its own accessibility relation R_i. For any i, $\Diamond_i\varphi$ existentially quantifies over the worlds R_i-accessible to w, requiring that there be some such accessible world where φ is true.

Definition 3.4. $[\![\Diamond_i \varphi]\!] = \{w \mid \exists w' : wR_i w' \;\&\; w' \in [\![\varphi]\!]\}$

For example, an epistemic possibility claim $\Diamond_e \varphi$ is true at w just in case the relevant agent's knowledge at w is consistent with φ. But a deontic claim $\Diamond_d \varphi$ is true at w just in case there is a φ world v consistent with the normative ideals (just in case $wR_d v$).

The last thing we need is a definition of validity. Here, let's suppose that validity is simply preservation of truth. So a set of premises entail a conclusion just in case every world at which the premises are true is a world at which the conclusion is true:

Definition 3.5. $\Gamma \models \delta$ iff $\bigcap_{\gamma \in \Gamma} [\![\gamma]\!] \subseteq [\![\delta]\!]$.

Everything we've said so far leaves unsettled the meaning of disjunction. Usually, the disjunction $\varphi \vee \psi$ is interpreted as unioning the meaning of φ and ψ.

$$[\![\varphi \vee \psi]\!] = [\![\varphi]\!] \cup [\![\psi]\!]$$

But the possible worlds framework we are using is not committed to this particular meaning for disjunction. Rather, the only commitment so far is that $\vee$ is propositional, so that it denotes *some* function $*$ from $[\![\varphi]\!]$ and $[\![\psi]\!]$ to the set of worlds where $\varphi \vee \psi$ is true:

Definition 3.6. Disjunction is propositional iff there exists some function $*$ such that (i) $[\![\varphi \vee \psi]\!] = [\![\varphi]\!] * [\![\psi]\!]$; and (ii) $[\![\varphi \vee \psi]\!] \subseteq W$.

We've now clarified what Free Choice is, what it takes for disjunction to be propositional, and how possibility modals work. Let's proceed to our first result.

3.3 Free choice and propositional disjunction

We will now see that if Free Choice is valid, then disjunction is not propositional.

To see why, let's think about what it would take for Free Choice to be valid, given the framework so far. So suppose that $\vee$ satisfies propositionality, with some function $*$ as the witness. This means that $[\![\varphi \vee \psi]\!]$ is always equal to $[\![\varphi]\!] * [\![\psi]\!]$. Then the validity of Free Choice for a possibility modal $\Diamond_i$ amounts to the requirement that whenever the set of R_i accessible worlds is consistent with $A * B$, for some propositions A and B, that set of accessible worlds is also consistent with each of A and B. That is:

Observation 3.2. $\Diamond_i$ satisfies Free Choice just in case for any world w:

$\exists w' : wR_iw' \;\&\; w' \in A * B$ iff $\exists w' : wR_iw' \;\&\; w' \in A$ and $\exists w'' : wR_iw'' \;\&\; w'' \in B$.

In other words, Free Choice requires that there be a special operator, $*$, where $A * B$ is consistent with a body of information just in case A and B are both consistent with that body of information. We will now see that this condition is hard to satisfy.

In particular, we will see that when accessibility relations have a certain structure, the equation above is unsatisfiable for any choice of $*$. The structure we are looking for is not hard to find. Sometimes, the information available at one possible world can be split up into two mutually exclusive and jointly exhaustive bodies of information that are each accessible at some other world. Each body of information in turn can be generated by some accessibility relation and world, and each such accessibility relation can be associated with some possibility operator. Finally, we are interested in cases where each body of information is nameable, so that it is the semantic value of some sentence in our language. Summing up, we are interested in families of possibility operators where the information relevant to one operator is a partition of the information relevant to the other ones. When all this occurs, let's say that the resulting set of possibility operators is *partitional*.

Definition 3.7. A set of possibility operators $\Diamond_1, \ldots, \Diamond_n$ with corresponding accessibility relations $R_1, \ldots, R_n$ is partitional just in case there exist accessibility relations R, R', R'' in $\{R_1, \ldots, R_n\}$ and worlds w, w', w'' such that:

1. $\{v \mid wRv\}, \{v' \mid w'R'v'\}$, and $\{v'' \mid w''R''v''\}$ are not empty.
2. $\{v \mid wRv\} = \{v' \mid w'R'v'\} \cup \{v'' \mid w''R''v''\}$
3. $\{v' \mid w'R'v'\} \cap \{v'' \mid w''R''v''\} = \emptyset$
4. $\exists \varphi, \psi \in \mathscr{L} : [\![\varphi]\!] = \{v' \mid w'R'v'\} \;\&\; [\![\psi]\!] = \{v'' \mid w''R''v''\}$.

Our first result is that no such set of possibility operators can all satisfy Free Choice when disjunction is propositional.

Fact 3.1. If a set of possibility operators is partitional and all satisfy Free Choice, then disjunction is not propositional.

To get a sense of what Fact 3.1 says, and why it's true, let's walk through some examples. First, it is commonplace to think that epistemic modals make claims about the information of some relevant group of agents. So in particular let's imagine that we have three agents: Alex, Billie, and Carrie. We can then imagine three different possibility modals, $\Diamond_{\mathsf{al}}$, $\Diamond_{\mathsf{bi}}$, and $\Diamond_{\mathsf{ca}}$, whose accessibility relations are indexed to each agent's information. These operators model the truth conditions of an utterance of a *might* sentence in each agent's context of utterance. Again let **ny** and **la** represent the propositions that Mary is in New York and that Mary is in Los Angeles.

Now let's imagine a world w where Alex's information is partitioned by Billie's information at v and Carrie's at u (see Figure 1). In particular, suppose that at w Mary could be in New York or Los Angeles for all Alex knows. For simplicity, let's then assume that the set of world's consistent with what Alex knows is just $\mathbf{ny} \cup \mathbf{la}$. But Billie at v knows that Mary is in New York (and in particular Billie's epistemic possibilities are simply **ny**). By contrast, Carrie in world u knows that Mary is in Los Angeles (so that Billie's epistemic possibilities are **la**). We now have a family of partitional modals, since Alex's information at w is partitioned by Billie's at v and Carrie's at u.

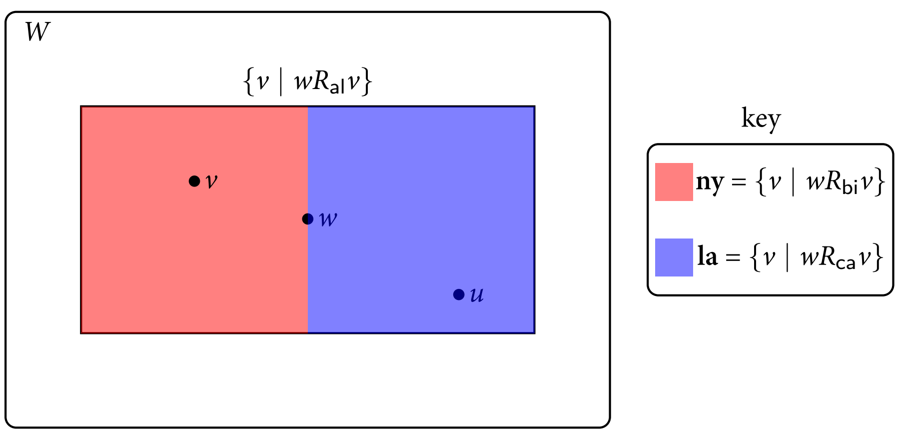

Figure 3.1: Partitional modals

To understand Fact 3.1, now suppose that each modal validates Free Choice. First, we can infer that $\Diamond_{\mathsf{al}}(\mathsf{nyc} \vee \mathsf{la})$ is true at w. After all, by Free Choice this is equivalent to the requirement that $\Diamond_{\mathsf{al}}\mathsf{nyc}$ and that $\Diamond_{\mathsf{al}}\mathsf{la}$. Since Alex's epistemic possibilities at w are just $\mathbf{nyc} \cup \mathbf{la}$, these latter requirements are satisfied. But if $\Diamond_{\mathsf{al}}(\mathsf{nyc} \vee \mathsf{la})$ is true at w, we can infer that one of Alex's epistemic possibilities at w is in $\mathbf{nyc} * \mathbf{la}$. Now consider whether $\Diamond_{\mathsf{bi}}(\mathsf{nyc} \vee \mathsf{la})$ is true at v. Since Billie knows that Mary is in New York, we can infer that $\Diamond_{\mathsf{bi}}\mathsf{la}$ is false at v. So by Free Choice again

$\Diamond_{\mathsf{bi}}(\mathsf{nyc} \vee \mathsf{la})$ is false at v. But this means that none of Billie's epistemic possibilities at v are in **nyc** $*$ **la**. By symmetric reasoning, we can infer that none of Carrie's epistemic possibilities at u are in **nyc** $*$ **la**. But this leads to a contradiction. Alex's epistemic possibilities are the union of Billie's and Carrie's. So **nyc** $*$ **la**, whatever proposition it is, cannot be consistent with Alex's epistemic possibilities while being inconsistent with both Billie's and Carrie's.

Summing up, the rough idea of the proof goes as follows. We consider three bodies of information A, B, and C, where B and C partition A. Then we find two sentences φ and ψ true at B and C uniquely. Applying Free Choice to the modal associated with A, we know that $\varphi \vee \psi$ is consistent with A, since φ is consistent with B and ψ is consistent with C. But, applying Free Choice to the modal associated with B, $\varphi \vee \psi$ is inconsistent with B since ψ is inconsistent with B. Similarly, $\varphi \vee \psi$ is inconsistent with C since φ is inconsistent with B. But this all is contradictory: $\varphi \vee \psi$, whatever proposition it is, can't be consistent with A while being inconsistent with both B and C, because B and C partition A.

We relied above on there being three different epistemic modals: one for each agent's information. However, this assumption was unnecessary. We can also generate a family of partitional modals using different flavors of modality. For example, it's natural to think that there is a family of deontic modals corresponding to different sorts of normative requirements. So we can think about what is legally permitted ($\Diamond_l$), morally permitted ($\Diamond_m$), and what is prudentially permitted ($\Diamond_{pr}$). It is possible for there to be a single world w where only φ worlds are morally permissible, only $\neg\varphi$ worlds are prudentially permissible, and both kinds of worlds are legally permissible. Yet if Free Choice holds for all three types of deontic modality, then disjunction can't be propositional.

Next, it's worth considering what constraints Fact 3.1 places on the validity of Free Choice for any individual modal. A special case of partitionality occurs when a single modal with accessibility relation R features three worlds w, v, and u where the R-information at w is partitioned by the R-information at v and at u. Then that single modal itself is partitional. So Fact 3.1 implies that if this single modal validates Free Choice, then disjunction is not propositional.

It is unclear whether any single modal should actually be partitional. For example, at least in the case of epistemic modality it is common to assume that accessibility is an equivalence relation, so that the modals obey an S5 logic, where iterated strings of modals collapse to their innermost member. In this case, no particular modal will be partitional. Nonetheless, Fact 3.1 still has relevance here, since three modals could each obey an S5 logic, while the information for one is partitioned into the information relevant for the other two.

Finally, those skeptical of partitionality should also consider the conditional form of Free Choice, as in (3)-(4). This inference's validity is inconsistent with propositional disjunction, independently of partitionality. In particular, suppose we model *might* conditionals using a two-place operator $\Diamond_i(\cdot)(\cdot)$.[5] The first argument of $\Diamond_i(\cdot)(\cdot)$ restricts the set of accessible worlds with which the second argument must be consistent:

Definition 3.8. $[\![\Diamond_i(\varphi)(\psi)]\!] = \{w \mid \exists w' : wR_iw' \;\&\; w' \in [\![\varphi]\!] \;\&\; w' \in [\![\psi]\!]\}$

So (3) has a logical form like $\Diamond_e(\neg\mathsf{Pisa})(\mathsf{Lisbon} \vee \mathsf{Rome})$. It is true at *w* just in case *w*'s epistemic possibilities contain a world where Mary does not go to Pisa, but where $\mathsf{Lisbon} \vee \mathsf{Rome}$ is true.

The validity of Conditional Free Choice amounts to the following principle:

CFC I $\Diamond(\chi)(\varphi \vee \psi) \models \Diamond(\chi)(\varphi)$ and $\Diamond(\chi)(\varphi \vee \psi) \models \Diamond(\chi)(\psi)$

CFC II $\Diamond(\chi)(\varphi); \Diamond(\chi)(\psi) \models \Diamond(\chi)(\varphi \vee \psi)$

But it turns out that if any conditional modal $\Diamond_i(\cdot)(\cdot)$ validates this inference, then disjunction is not propositional.

Fact 3.2. If any conditional possibility operator $\Diamond_i(\cdot)(\cdot)$ satisfies Conditional Free Choice, then $\vee$ is not propositional.

Essentially, the antecedent of the conditional can mimic the quantification over bodies of information from the partitionality constraint.

To bring home the main ideas of this section, it's worth thinking about what we've established from the perspective of rational belief formation. After all, we saw above that Free Choice holds for attitude verbs:

(6) a. Susan thinks Mary might be in New York or Los Angeles.

[5] See Kratzer 1986.

b. Susan thinks Mary might be in New York and Susan thinks Mary might be in Los Angeles.

Let's introduce a special operator $\blacklozenge_S$ to model what agent S believes might be the case.[6] Then the equivalence of the above amounts to the validity of the inference from $\blacklozenge_S(\varphi \vee \psi)$ to $\blacklozenge_S\varphi$ and $\blacklozenge_S\psi$, and vice versa.[7] To see whether this is consistent with a propositional theory of disjunction, we'd need an account of what it is to believe a possibility claim. According to one popular recent tradition, an agent believes $\Diamond\varphi$ just in case her beliefs are consistent with $[\![\varphi]\!]$.[8] To implement this idea, we can treat $\blacklozenge_S$ using the semantics from Definition 3.4, and simply associate each agent S with an accessibility relation R_S that models S's doxastic possibilities at any world. Given this assumption, we can make good intuitive sense of what Free Choice requires. If disjunction is propositional, then there must be some operation $*$ where any rational agent's beliefs are consistent with $A * B$ just in case they are consistent with A and they are consistent with B. But another consequence of Fact 3.1 is that no such operation can exist. Here, it's a bit easier to see the problem, since for just about any set of worlds we can imagine an agent who believes she is among them. This makes it easy to generate families of partitional accessibility relations, by taking three agents where the first agent's beliefs are partitioned by the others'.

So for example consider an agent S who is maximally ignorant about the world she inhabits, and consider two propositions A and $\overline{A}$ which partition logical space. S believes that A might be true and that $\overline{A}$ might be true, since her doxastic possibilities include both kinds of worlds. So by Free Choice S believes that A or $\overline{A}$ might be true. However, now consider two more agents S' and S''. S' has come to believe A, while S'' has come to believe $\overline{A}$. S' does not believe $\overline{A}$, since her doxastic possibilities contain only A worlds. So by Free Choice S' does not believe that A or $\overline{A}$ might be true. Similarly, S'' has come to believe $\overline{A}$ and so she does not believe that A might be true. So by Free Choice S'' does not believe that A or $\overline{A}$ might be true. But this all is inconsistent. For this implies that S's belief state is consistent with $A * \overline{A}$ while S' and S'' are not, even though S' and S'''s belief states partition S's. In short, there can't be a propositional operator $*$ where any

[6] See Yalcin 2007.

[7] Of course, any account of believing-might should ultimately be decomposed into a semantics for *believes* and a semantics for *might*. But we can suppress this complexity for now.

[8] See Stephenson 2007, Yalcin 2007, 2012, Swanson 2011, 2012, Rothschild 2012, and Moss 2015.

agent's beliefs are consistent with $A * B$ just in case they are consistent with each of A and B.

These last points about rational belief can also be thought about in terms of rational credence. So let's grant ourselves a set of probability functions defined on a set of worlds:

Definition 3.9. Let W be the set of possible worlds. Let $A, B, \ldots$ be variables for subsets of W. Let a probability function Pr be a function on some Boolean algebra over W to $[0, 1]$ satisfying:

1. $Pr(A) \geq 0$ NON-NEGATIVITY
2. $Pr(W) = 1$ UPPER BOUND
3. If $A \cup B = \emptyset$, then $Pr(A \cup B) = Pr(A) + Pr(B)$ ADDITIVITY

Suppose that an agent believes that might φ just in case she assigns some credence to φ.[9] Then the validity of free choice for ♦ corresponds to the principle that there be some propositional operator $*$ where any rational agent assigns a positive credence to $A * B$ iff she assigns a positive credence to A and to B. That is, Free Choice corresponds to the following thesis:

Definition 3.10 (Credal Free Choice)**.** There is some operator $*$ such that for any rational Pr, A, B:

$Pr(A * B) > 0$ iff $Pr(A) > 0$ and $Pr(B) > 0$.

To see why this thesis is absurd, we can imagine an agent first who assigns some credence to A and to $W - A$. Then she will assign some positive credence to $A * (W - A)$. Now if our agent conditionalizes on A, she should still count as rational. Yet now she will assign no credence to $W - A$, and hence will assign no credence to $A * (W - A)$. Similarly, if she conditionalized on $(W - A)$ she would assign no credence to A, and hence no credence to A. But by the law of total probability her current credence in $A * (W - A)$ is a weighted sum of her posterior credence in $A * (W - A)$ after conditionalizing on each of A and $W - A$. So if each posterior credence is 0, her current credence is 0 after all. That is, the existence of such an operator $*$ implies maximal opinionation.

Fact 3.3. If Credal Free Choice holds, then $Pr(A) = 0$ or $Pr(A) = 1$ for every rational Pr, A.

---

[9] Hájek 2013 has developed a variety of counterexamples to the left-to-right direction of this principle, involving infinite partitions. For example, imagine throwing an infinitely small dart at the real number line between 0 and 1. What is the probability it will land at $\frac{1}{2}$? Seemingly zero. But one can still believe that the dart might land at $\frac{1}{2}$. Here, we can ignore these worries, since our result will still make trouble when zooming in on finite partitions.

In the context of rational credence, then, our principle about possibility leads to a triviality result in the style of Lewis 1976.

In this section, we have shown that if we hold fixed a possible worlds semantics for possibility modals, Free Choice is inconsistent with a propositional theory of disjunction. This leads to a natural question. Perhaps if we give up on a standard possible world semantics for possibility modals, we can validate Free Choice without giving up a propositional theory of disjunction. In the next section, we will see that this intuitive idea is mistaken. If we hold fixed the leading dynamic semantics for possibility modals, the validity of Free Choice still requires a revision to a traditional understanding of disjunction.

3.4 A dynamic impossibility result

So far, we have seen that holding fixed a propositional theory of possibility modals, there is no way to validate Free Choice for suitably rich operators and retain a propositional theory of disjunction. In this section, we will see that the previous result was not an artifact of the particular theory of possibility that we started with. An analogous result can be proved within a dynamic framework. If we hold fixed a standard dynamic meaning for possibility operators, we will see that there is no way to validate Free Choice without positing an irreducibly dynamic meaning for disjunction.

In particular, we will look at one version of dynamic semantics: update semantics.[10] According to dynamic semantics, the meaning of a sentence is not its truth conditions. Rather, the meaning of a sentence is its ability to change the context in which it is said—its *context change potential*:

> You know the meaning of a sentence if you know the change it brings about in the information state of anyone who accepts the news conveyed by it.[11]

To give an update semantics, we need two things: a definition of information states (or contexts), and an interpretation function which assigns a context change potential to each sentence in our language. Veltman 1996 models an information state as a set of possible worlds. Then an interpretation function assigns every sentence a function from sets of worlds to sets of worlds.

---

[10] See Stalnaker 1973; Karttunen 1974; Heim 1982; Heim 1983; Veltman 1985; Groenendijk and Stokhof 1990; Groenendijk and Stokhof 1991a; and many others.

[11] See Veltman 1996, p. 221.

Definition 3.11. A possible world w assigns every atomic sentence α a truth value. W is the set of all possible worlds. An information state s is a set of possible worlds. A context change potential is a function from one information state s to a new one. An interpretation function $[\cdot]$ assigns every sentence a context change potential. $s[\varphi]$ is the result of inputting s into $[\varphi]$.

Once we have a representation of information states, we can then recursively define our interpretation function, $[\cdot]$. Again, we are working with a very simple language. So we only need to give the meanings of atomic sentences and possibility modals before turning to disjunction. Here, we will assume that atomic sentences simply narrow down an information state to the worlds where they are true.

Definition 3.12. $s[\alpha] = \{w \in s \mid w(\alpha) = 1\}$

Updating a state with an atomic sentence communicates new information about which world we inhabit. We can move from the initial state to an interesting, non-empty subset of it. By contrast, modal sentences don't give us new information about what world we are in. Instead, they are tests, exploring properties of the current state. $\Diamond\varphi$ explores whether s can be consistently updated with φ; If so, the initial state is unchanged by updating. Otherwise, the absurd state $\emptyset$ results.[12]

Definition 3.13. $s[\Diamond\varphi] = \{w \in s \mid s[\varphi] \neq \emptyset\}$

Before turning to disjunction, we also need a definition of entailment. Say that a state supports a sentence φ just in case updating the state with φ has no effect. Then an argument is dynamically valid just in case any state that supports the premises also supports the conclusion.[13]

Definition 3.14.

1. s supports φ ($s \models \varphi$) iff $s[\varphi] = s$.

2. Γ entails δ ($\Gamma \models \delta$) iff for any information state s, if $s \models \gamma$ for every $\gamma \in \Gamma$, then $s \models \delta$.

Within this framework, we can now investigate the space of possible meanings for disjunction. Let's start by saying, within dynamic semantics, what it takes for a sentence to have an ordinary,

[12] Definitions 3.12 and 3.13 are one recursive definition.

[13] See Veltman 1996 and van Benthem 1996 139-41 for an exploration of how to define entailment in dynamic semantics. Our results below could also be established using update-to-test entailment.

static meaning. A sentence φ has a static meaning when it updates every context s distributively, so that updating s with φ is equivalent to updating each world in s with φ, and unioning the result:

Definition 3.15. φ is distributive iff for all s, $s[\varphi] = \bigcup_{w \in s} \{w\}[\varphi]$

When an operator is distributive, updating a state is simply a matter of exploring the properties of each world in the state. In this case, the operator isn't really sensitive to its input context. No matter what context it operates on, it always has the same effect. So let's say that disjunction has a static meaning only if for any atomic φ and ψ, $\varphi \vee \psi$ is distributive.

With this background, we can now show that if Free Choice is valid, disjunction is not distributive. To do so, let's first think about what it takes for Free Choice to be valid in this framework. In this framework, possibility modals test that their complement's update produces a non-empty result. So the validity of Free Choice requires that for any state, updating with $\varphi \vee \psi$ is consistent just in case updating with both φ and ψ are.

Observation 3.3. Free Choice is valid just in case for any state s:
$s[\varphi \vee \psi] \neq \emptyset$ iff $s[\varphi] \neq \emptyset$ and $s[\psi] \neq \emptyset$.

We can now show that this requirement is inconsistent with the distributivity of $\varphi \vee \psi$. This equation is only satisfiable if disjunctions can access global properties of an input state, rather than updating the state one world at a time.

Fact 3.4. If Free Choice is valid, then $\varphi \vee \psi$ is not distributive for any atomic φ and ψ.

The explanation of Fact 3.4 is similar to the one before—in particular, to our discussion of rational belief above. We can simply imagine three information states s, s', and s'' where the first is partitioned by the other two. Then we can let φ and ψ be non-modal (distributive) claims true at exactly the worlds in s' and s''. This implies that updating s with either of φ and ψ will not result in the absurd state, and so by the validity of Free Choice s can be updated with $\varphi \vee \psi$ without reaching the absurd state. But if $\varphi \vee \psi$ is distributive, then this requires that there is some world in s where $\varphi \vee \psi$ is true. But then this world must be present in either s' in s''. This in turn implies that either s' supports $\Diamond(\varphi \vee \psi)$ or s'' does. But in this case, one of those states must be consistently updatable with both φ and ψ. This is a contradiction.

There was been a flurry of recent work investigating what really distinguishes static and dynamic semantics.[14] In the light of Fact 3.4, we can think of Free Choice as a test case. Free Choice is the sort of phenomenon that can only be explained by abandoning a static, propositional theory of disjunction.

3.5 Dynamic disjunction

Our results so far show that we can't validate Free Choice while thinking of disjunction propositionally. Now we will consider whether a genuinely dynamic meaning for disjunction can validate Free Choice. I characterize a family of dynamic semantics for which Free Choice is valid, and find a strongest and weakest member of this family.

To validate Free Choice, we need one simple idea: disjunction requires each disjunct to be possible.[15] That is, we need a semantics on which the following inference is valid:

DISJUNCTION-POSSIBILITY LINK $\varphi \vee \psi \models \Diamond\varphi$ and $\varphi \vee \psi \models \Diamond\psi$

This isn't an arbitrary choice. Rather, given a few background assumptions it is equivalent to the validity of Free Choice. To see why, note that the dynamic semantics for possibility modals we have endorsed above validates the T axiom, which say that anything true is possible:

T $\varphi \models \Diamond\varphi$

Given the duality of *might* and *must*, the T axiom is equivalent to the requirement that *must* is strong, so that $\Box\varphi$ implies φ.[16]

Disjunction-Possibility Link follows from Free Choice, holding fixed the T axiom and the transitivity of entailment.

TRANSITIVITY $\varphi \models \psi \;\&\; \psi \models \chi \Longrightarrow \varphi \models \chi$

---

[14] See for example Groenendijk and Stokhof 1991b, Veltman 1996, van Benthem 1996, Rothschild and Yalcin 2015a, and Russell and Hawthorne 2016.

[15] See Zimmermann 2000 for a static implementation of this idea.

[16] For discussion of whether this principle is valid for epistemic modals, see Karttunen 1972; Kratzer 1991; von Fintel and Gillies 2010; and Lassiter 2016.

After all, the T axiom implies that $\varphi \vee \psi \models \Diamond(\varphi \vee \psi)$. By Free Choice, this immediately implies $\Diamond\varphi$ and $\Diamond\psi$. Since the dynamics semantics for possibility modals above validates the T axiom, we must find a semantics for disjunction that validates Disjunction-Possibility Link, if we want to validate Free Choice.

Furthermore, Disjunction-Possibility Link itself comes close to implying Free Choice. For suppose we accept the 4 axiom, that anything possibly possible is itself possible:

4 $\Diamond\Diamond\varphi \models \Diamond\varphi$

If we combine the 4 axiom with the upwards monotonicity of $\Diamond$, Disjunction-Possibility Link implies FC I.[17] By Disjunction-Possibility Link, $\varphi \vee \psi$ implies $\Diamond\psi$. So by the upwards monotonicity of $\Diamond$, $\Diamond(\varphi \vee \psi)$ implies $\Diamond\Diamond\psi$, which by the 4 axiom implies $\Diamond\psi$. The 4 axiom and the T axiom together generate the modal logic S4. So, summing up, if our possibility operators satisfy S4 then we can move back and forth between FC I and Disjunction-Possibility Link.

Fact 3.5. Suppose that entailment is transitive, that $\Diamond$ is upwards monotonic, that the T axiom is valid, and that the 4 axiom is valid. Then FC I is valid iff Disjunction-Possibility Link is valid.

There are a variety of dynamic meanings for disjunction that validate Free Choice and Disjunction-Possibility Link. On the simplest and weakest such meaning, disjunctions merely state that each disjunct is possible. That is, we could let $[\varphi \vee \psi] = [\Diamond\varphi][\Diamond\psi]$, where $\Diamond$ is the dynamic possibility operator discussed above. This meaning for disjunction is a test, so I will represent it with the symbol $\vee_t$.

Definition 3.16. $s[\varphi \vee_t \psi] = s[\Diamond\varphi][\Diamond\psi]$

Unpacking the meaning of $\Diamond$, $\varphi \vee_t \psi$ tests a state s to require that updating with each of φ and ψ produces a nonempty state. That is: $s[\varphi \vee_t \psi] = \{w \in s \mid s[\varphi] \neq \emptyset \;\&\; s[\psi] \neq \emptyset\}$.

This particular dynamic meaning is of interest to us because it is the weakest meaning for disjunction that can validate Free Choice. First, we can see that this operation does satisfy Free Chioce.

Observation 3.4. (i) $\Diamond(\varphi \vee_t \psi) \models \Diamond\varphi$; (ii) $\Diamond(\varphi \vee_t \psi) \models \Diamond\psi$; (iii) $\Diamond\varphi; \Diamond\psi \models \Diamond(\varphi \vee_t \psi)$

[17] See Zimmermann 2000 for a similar observation.

The reason this semantics validates Free Choice is that the dynamic semantics for possibility modals above collapses iterated modals, so that $\Diamond\Diamond\varphi$ implies $\Diamond\varphi$.

Next, we can see that $\vee_t$ is the weakest operator that can validate Free Choice.

Definition 3.17. $*$ is at least as strong as $\bowtie$ iff for any φ, ψ: $\varphi * \psi \models \varphi \bowtie \psi$.

Fact 3.6. If $*$ validates Free Choice, then $*$ is at least as strong as $\vee_t$.

This is really just a semantic analogue of Fact 3.5. By the T axiom, we know that $\varphi * \psi$ entails $\Diamond(\varphi * \psi)$. Since $*$ satisfies Free Choice, this later claim entails each of $\Diamond\varphi$ and $\Diamond\psi$, which in turn entail $\varphi \vee_t \psi$. So $\varphi * \psi$ entails $\varphi \vee_t \psi$.

This semantics for disjunction can validate Free Choice because it is non-distributive. After all, consider a state $s = \{w, v\}$, where α is true at w and false at v. Suppose further that α and β partition W. Then s supports $\alpha \vee_t \beta$, but neither $\{w\}$ nor $\{v\}$ do. That is: $\alpha \vee_t \beta$ is true at state s without being true at each part of s. Its truth is an irreducibly global affair.

So far, we've found the weakest possible dynamic semantics for disjunction that validates Free Choice. But this semantics is unsatisfying as a theory of disjunction. It is simply too weak. First, note that on this semantics, $\varphi \vee_t \psi$, $\Diamond(\varphi \vee_t \psi)$, and $\Box(\varphi \vee_t \psi)$ are all equivalent. The problem is that $\varphi \vee_t \psi$ is a test. So embedding this test within another test, like $\Diamond$ or $\Box$, will have no further effect. $\varphi \vee_t \psi$ goes wrong as a theory of disjunction because it does not allow the assertion of a disjunction to contribute any information about what world one inhabits. Rather, we would like a theory of disjunction where in addition to requiring that each disjunct is possible, the whole disjunction also contributes new information about what the world is like.

To find more meanings for disjunction that can validate Free Choice, we can build on our test meaning, $\vee_t$. We can construct a family of well behaved meanings for disjunction that combine the requirement that each disjunct is possible with another operation that narrows down the worlds in a state. That is, asserting a disjunction does two things: (i) it narrows down the input state to the worlds where the disjunction is true; and (ii) it requires that both disjuncts are possible.

We can make more precise predictions by constraining these meanings so that the updating component of disjunction is well-behaved. In particular, let's assume that this first component of the meaning of a disjunction is itself a static operation. That is, updating with $\varphi \vee \psi$ will involve first updating with $\varphi \bowtie \psi$, where $\varphi \bowtie \psi$ is distributive whenever φ and ψ are. Intuitively, the idea is this: there is some fixed set of worlds where $\varphi \vee \psi$ is true. Updating a state with $\varphi \vee \psi$ involves both zooming in to this set of worlds, and requiring that both φ and ψ are possible. When a meaning for disjunction has this sort of structure, let's say that it is factorizable into a static and dynamic component.

Definition 3.18. $*$ is factorizable iff there is some distributive operator $\bowtie$ such that $[\varphi * \psi] = s[\varphi \bowtie \psi][\varphi \vee_t \psi]$.

This is an intuitive constraint. The underlying idea is that we can factorize the meaning of a disjunction into a dynamic and a static component. The dynamic component guarantees Free Choice, while the static component implements the other aspects of a disjunction's meaning.

Now we can generate a family of dynamic operators to serve as the meaning of disjunction, by combining $[\vee_t]$ with some other function. One natural candidate stands out. Let $\vee_u$ denote the usual dynamic meaning of disjunction, on which $\varphi \vee \psi$ updates a context with φ, and takes the union of this with the result of updating with ψ:

Definition 3.19. $s[\varphi \vee_u \psi] = s[\varphi] \cup s[\psi]$

This meaning for disjunction is distributive whenever φ and ψ is. In particular, updating with $\varphi \vee_u \psi$ amounts to zooming in to the worlds where one of φ or ψ are true. That is, $\vee_u$ is a dynamic restatement of the traditional possible worlds semantics for disjunction.

We can build a new, dynamic meaning for disjunction by combining $\vee_u$ and $\vee_t$.

Definition 3.20. $s[\varphi \vee_{ut} \psi] = s[\varphi \vee_u \psi][\varphi \vee_t \psi]$.

Unpacking our previous definitions, $s[\varphi \vee_{ut} \psi] = (s[\varphi] \cup s[\psi])[\Diamond\varphi][\Diamond\psi] = \{w \in s[\varphi] \cup s[\psi] \mid s[\varphi] \neq \emptyset \;\&\; s[\psi] \neq \emptyset\}$. On this proposal, the meaning of a disjunction is factorized into two components. First, $[\varphi \vee_u \psi]$ performs an update, allowing a disjunction to contribute some information about the world. Next, $[\varphi \vee_t \psi]$ tests that each of φ and ψ are possible in the resulting state.

$\varphi \vee_u \psi$ is a distributive theory of disjunction. Whenever φ and ψ are non-modal, $\varphi \vee_u \psi$ is true at a world just in case one of φ or ψ is. $\varphi \vee_{ut} \psi$ inherits some of these truth conditional properties. Whenever $\varphi \vee_{ut} \psi$ is uttered in a state where φ and ψ are possible, the state will be narrowed down to the union of the φ and ψ worlds in that context.

Now I will vindicate the idea that $[\vee_{ut}]$ is a particularly natural meaning for disjunction. In particular, we will see that $\vee_{ut}$ is the strongest factorizable meaning for disjunction that validates Free Choice. First, this semantics validates Free Choice:

Observation 3.5. (i) $\Diamond(\varphi \vee_{ut} \psi) \models \Diamond\varphi$; (ii) $\Diamond(\varphi \vee_{ut} \psi) \models \Diamond\psi$; (iii) $\Diamond\varphi; \Diamond\psi \models \Diamond(\varphi \vee_{ut} \psi)$

Second, $\vee_{ut}$ is the strongest factorizable meaning that validates Free Choice, whenever its inputs are not ordered by strength.

Fact 3.7. if $*$ is factorizable and satisfies Free Choice, and if $\varphi \not\models \psi$ and $\psi \not\models \varphi$, then $\varphi \vee_{ut} \psi \models \varphi * \psi$.

Here, the key property of $\varphi \vee_{ut} \psi$ is that its first component, $\varphi \vee_u \psi$, never eliminates a world where φ or ψ is true. Suppose that $*$ is factorizable into $\bowtie$, validates Free Choice, and does eliminate such a world w (where φ but not ψ is true). Then consider a state where w is the only φ world, but there are also some ψ worlds. That state will support $\Diamond\varphi$ and $\Diamond\psi$. But it will not support $\Diamond(\varphi * \psi)$, since $\varphi \bowtie \psi$ eliminates w, at which point $\varphi \vee_t \psi$ produces the empty state.[18] [19]

There is a deeper significance to the results above. Fact 3.7 shows that the truth conditional behavior of disjunction is ultimately derivable from the validity of Free Choice. The truth functional theory of disjunction—$[\vee_u]$—is the strongest well-behaved operation that can compose with $[\vee_t]$ to create an operation that satisfies Free Choice. So the truth functional properties of disjunction flow naturally from the attempt to communicate as much information as possible with a certain expression, while our use of this expression satisfies the Free Choice inference.

---

[18]One arbitrary feature of factorizability was that it involved first updating with $\varphi \bowtie \psi$, and second updating with $\varphi \vee_t \psi$. This order effect is somewhat independent of the underlying facts, however. We could instead defined factorizability so that $s[\varphi * \psi] = [\varphi \vee_t \psi][\varphi \bowtie \psi]$. Or, we could have instead defined it so that $s[\varphi * \psi] = s[\varphi \bowtie \psi] \cap s[\varphi \vee_t \psi]$. In each case we could have proved an analogous result, so that the strongest $\bowtie$ that can validate Free Choice was $\vee_u$. However, if we used either of these definitions, we would have needed one more assumption to complete our proof: that $\varphi * \psi$ be idempotent, so that $s[\varphi * \psi]$ always supports $\varphi * \psi$. Idempotent would require the test imposed by $\varphi \vee_t \psi$ to be satisfied after updating with $\varphi \bowtie \psi$, which was the key step relevant to proving Fact 3.7. For arguments that idempotence is a natural constraint to impose on a semantic theory, see Yalcin 2015.

[19]Another caveat to the fact above is that it only concerns disjuncts that are not ordered by strength. Free Choice can be validated by operators that are stronger than $\vee_{ut}$ when φ and ψ are ordered by strength. In general, the strongest such operator will intersect $s[\varphi]$ and $s[\psi]$ whenever φ and ψ are ordered by strength, and otherwise will take their union. While this operator does validate Free Choice, it may be possible to rule out such a meaning on other grounds. For example, this meaning intuitively involves a failure of continuity. Imagine varying the inputs to such an operator ever so slightly, by starting with a φ and ψ that are ordered by strength, and swapping out φ for φ', which is not ordered by strength with ψ. Even a small change from φ to φ' can produce a dramatically different result. Thanks to Jiji Zhang for help here.

3.6 Dynamic disjunction; static possibility

So far, our discussion of dynamic semantics has focused on epistemic modality. In addition, our discussion has held fixed a dynamic semantics for epistemic possibility modals. In this section, I will show how we can generalize our results so far to other flavors of modality, without assuming that these modals are essentially dynamic.

Above, we appealed to possibility operators in two separate ways. First, Free Choice involved embedding disjunction under some sort of possibility modal. Second, our semantics had disjunctions themselves require that each disjunct be possible. In this section, we will see that these two appeals to possibility modals can be treated separately. In particular, we will see that we can integrate a static (distributive) semantics for deontic modals with our dynamic semantics for disjunction in a way that will validate Free Choice.

In particular, we can incorporate our earlier possible worlds semantics for modals into our dynamic framework in a straightforward way. For any possibility modal $\Diamond_i$, we can introduce an update rule for $\Diamond_i\varphi$ that simply zooms in to the worlds where φ is classically possible:

Definition 3.21. $s[\Diamond_i\varphi] = \{w \in s \mid \{v \mid wR_iv\}[\varphi] \neq \emptyset\}$

This semantics for possibility modals is distributive. When φ is distributive, updating s with $\Diamond_i\varphi$ narrows down s to the worlds that are i-related to a world where φ is true. When φ is not distributive, $\Diamond_i\varphi$ still operates on each world in s separately. In these cases, φ can still have an interesting effect, because it operates globally on $\{v \mid wR_iv\}$. However, since this set is itself controlled by w and not s, the resulting update is still distributive.

This semantics for modals allows us to model arbitrary flavors of modality. For example, we can model Free Choice for deontic modals by appeal to a deontic accessibility relation R_d, which relates w to the worlds that are consistent with the normative rules at w.

Both of our dynamic semantics above for disjunction validate Free Choice for our static possibility modal. That is, while Free Choice requires a dynamic semantics for disjunction, it does not require a dynamic semantics for possibility modals.

Observation 3.6.

1. $\Diamond_i(\varphi \vee_t \psi) \models \Diamond_i\varphi$; (ii) $\Diamond_i(\varphi \vee_t \psi) \models \Diamond_i\psi$; (iii) $\Diamond_i\varphi; \Diamond_i\psi \models \Diamond_i(\varphi \vee_t \psi)$

2. $\Diamond_i(\varphi \vee_{ut} \psi) \models \Diamond_i\varphi$; (ii) $\Diamond_i(\varphi \vee_{ut} \psi) \models \Diamond_i\psi$; (iii) $\Diamond_i\varphi; \Diamond_i\psi \models \Diamond_i(\varphi \vee_{ut} \psi)$

To see why Free Choice is valid, consider the case of our weakest dynamic meaning for disjunction, $\vee_t$, on which $\varphi \vee_t \psi$ simply tests that each of φ and ψ are possible. $\Diamond_i(\varphi \vee_t \psi)$ then involves two layers of possibility modals: the inner layer is $\Diamond\varphi$ and $\Diamond\psi$, where this is a dynamic possibility modal. Now $\Diamond_i(\varphi \vee_t \psi)$ is supported by s just in case for every w in s, the i-possibilities at w ($\{v \mid wR_iv\}$) can be consistently updated with $\varphi \vee_t \psi$. Updating $\{v \mid wR_iv\}$ with $\varphi \vee_t \psi$ in turn involves considering $\{v \mid wR_iv\}[\Diamond\varphi][\Diamond\psi]$. This is non-empty just in case w is R_i-related to a φ world and to a ψ world. But this is exactly what is required by w in order to survive update with $\Diamond_i\varphi$ and $\Diamond_i\psi$. To summarize, the crucial property needed here is that the following principle is valid:

OUTER NEGATIVE INTROSPECTION $\Diamond_i\Diamond\varphi \models \Diamond_i\varphi$.

The reason that we are able to validate this principle regardless of the structure of R_i is that $\Diamond$ is a dynamic operator. $\Diamond\varphi$ predicates possibility of a body of information. But because $\Diamond\varphi$ is dynamic, it can in principle predicate possibility of different bodies of information, depending on the environment in which it occurs. So if $\Diamond\varphi$ applies to the whole context, s, $\Diamond\varphi$ will predicate epistemic possibility of φ—consistency with what is common knowledge, say. By contrast, when $\Diamond\varphi$ is embedded under an information shifting operator, like $\Diamond_i$, it has a different effect. Under $\Diamond_i$, $\Diamond\varphi$ says that φ is consistent with $\{v \mid wR_iv\}$. That is, $\Diamond$ says the same thing under the scope of $\Diamond_i$ as $\Diamond_i$ says when unembedded. For example, $\Diamond_d\Diamond\varphi$ says that φ is permissible. This sensitivity to local context is a hallmark of dynamic semantics.

Furthermore, nothing in the explanation above appealed to particular features of the underlying accessibility relation. Rather, any distributive possibility modal will validate Free Choice when combined with our dynamic semantics for disjunction. Our theory thus goes a long way towards explaining why Free Choice appears valid for arbitrary flavors of modality. Summing up, we have seen that validating Free Choice does not actually require a dynamic semantics for possibility modals. Rather, disjunction itself contributes a dynamic possibility test that need not be contributed by possibility modals themselves. We can take any classical semantics for a possibility claim and convert it into a simple update that zooms in to the worlds where that claim is true. When the resulting static possibility operators embed our dynamic disjunction operator, Free Choice is valid.

3.7 Wide free choice

So far, we've focused our attention on cases where possibility modals take wide scope to disjunction. In this section we extend our results to the opposite scope relation. As many have observed, a similar conjunctive effect also occurs when disjunctions take wide scope over possibility modals.[20] Call such an effect Wide Free Choice.

WFC 1 $\Diamond\varphi \vee \Diamond\psi \models \Diamond\varphi$ and $\Diamond\varphi \vee \Diamond\psi \models \Diamond\psi$

WFC 2 $\Diamond\varphi; \Diamond\psi \models \Diamond\varphi \vee \Diamond\psi$

In this section, we will see that similar points to those above can be made for Wide Free Choice. First, we will develop an analogue of our earlier propositional impossibility results. In particular, we will see that Wide Free Choice is inconsistent with almost any propositional semantics for $\vee$. The only available propositional semantics is one on which $\vee$ is almost identical to intersection. Second, we will turn to our dynamic framework. There, we will see that our earlier dynamic semantics for disjunction allow us to validate Wide Free Choice. But in order to do so, we must somewhat constrain our semantics for possibility modals. Thus Wide Free Choice requires somewhat less of disjunction than Free Choice, but requires somewhat more of possibility modals.

Our first result is that in a propositional framework, Wide Free Choice requires disjunction to behave almost exactly like conjunction. More precisely, we can show that Wide Free Choice requires disjunction to intersect any two propositions that can be expressed using possibility modals. To state our result precisely, we need to clarify this last notion.

So suppose that $\Diamond$ is an existential quantifier over accessible worlds. Then we can say that a proposition A is expressible using $\Diamond$ just in case there is some other proposition B where $A = [\![\Diamond]\!](B)$. That is, A is expressible using $\Diamond$ just in case A is equivalent to the claim that some proposition B is possible via $\Diamond$.

[20] See for example Zimmermann 2000.

We can say all this more precisely using accessibility relations. Suppose that $\Diamond$ is modeled using the accessibility relation R. Then A is expressible using $\Diamond$ just in case A is true at exactly the worlds at which some B world is R-accessible.[21]

Definition 3.22. A is expressible via R iff $\exists B : A = \{w \mid \exists v : wRv \ \& \ v \in B\}$.

In other words, A must be the union of a variety of $\{w \mid wRv\}$, for any v in some fixed B.

Accessibility relations differ with respect to how many propositions they can express. For example, suppose that R is an equivalence relation, providing a partition of logical space. In that case, A is expressible just in case there is some B where A is a union of cells of the partition that overlap B. In fact, given any union of cells of the partition, we can find some B that overlaps all and only those cells. This means that for equivalence relations, A is expressible just in case A is a union of cells of the partition.

Armed with this notion of expressibility, we can define a family of operators that are conjunctive whenever B is expressible via R.

Definition 3.23. $*$ is quasi-conjunctive relative to R iff for any A and B, if A and B are expressible via R, then $A * B = A \cap B$.

Now that we have this notion of a quasi-conjunctive operation, we can return to Free Choice. It turns out that in a propositional framework, Wide Free Choice requires that disjunction is quasi-conjunctive.

Fact 3.8. If Wide Free Choice is valid for $\Diamond_R$ and disjunction is propositional, then disjunction is quasi-conjunctive relative to R.

The main idea here is straightforward. Wide Free Choice says that a disjunction of two claims is equivalent to their conjunction. This can only hold if the operation expressed by disjunction is guaranteed to intersect the two claims. This means that disjunction must behave conjunctively whenever its inputs satisfy the schema in Wide Free Choice. But only possibility claims are relevant to that schema. So Wide Free Choice requires that any propositions that are equivalent to a possibility claim behave conjunctively when embedded in a disjunction.

[21] Thanks to Andrew Bacon and Jeremy Goodman for critical help here.

Since disjunction intuitively should express set intersection, this all suggests that validating Wide Free Choice should be done without a propositional theory of disjunction. Fortunately, we need not look far to find a semantics for disjunction that is up to the task. In earlier sections, we developed several semantics for disjunction that can validate Free Choice. It turns out that each of these semantics also validates Wide Free Choice, when combined with the test semantics for possibility modals.

Observation 3.7.

1. $\Diamond\varphi \vee_t \Diamond\psi \models \Diamond\varphi$; (ii) $\Diamond\varphi \vee_t \Diamond\psi \models \Diamond\psi$; (iii) $\Diamond\varphi; \Diamond\psi \models \Diamond\varphi \vee_t \Diamond\psi$
2. $\Diamond\varphi \vee_{ut} \Diamond\psi \models \Diamond\varphi$; (ii) $\Diamond\varphi \vee_{ut} \Diamond\psi \models \Diamond\psi$; (iii) $\Diamond\varphi; \Diamond\psi \models \Diamond\varphi \vee_{ut} \Diamond\psi$

For simplicity, let's focus on our simple test semantics for disjunction, $\vee_t$. Then s supports $\Diamond\varphi \vee_t \Diamond\psi$ just in case s supports each of $\Diamond\Diamond\varphi$ and $\Diamond\Diamond\psi$. Here, the crucial detail is that $\Diamond$ satisfies the 5 axiom, so that $\Diamond\Diamond\varphi$ implies $\Diamond\varphi$. This is because $\Diamond$ is a test. So $s \models \Diamond\Diamond\varphi$ just in case $s[\Diamond\varphi]$ is non-empty, which holds just in case $s \models \Diamond\varphi$.

The next natural question is whether this explanation essentially relies on a dynamic semantics for possibility modals. It turns out that the answer is no. To investigate this question, we can consider the validity of the inference from $\Diamond_i\varphi \vee_t \Diamond_i\psi$ to $\Diamond_i\varphi$, using our distributive semantics for $\Diamond_i$ from the previous section (where $s[\Diamond_i\varphi] = \{w \in s \mid \{v \mid wRv\}[\varphi] \neq \emptyset\}$). This inference can be valid, provided that we constraint accessibility significantly. In particular, we must require that accessibility is universal, so that whenever v is visible from some world, it is visible from any world.[22]

Fact 3.9. Wide Free Choice is valid for $\Diamond_i$ and $\vee_t$ iff for any w, v, u: if wR_iv then uR_iv.

To see why, note that s supports $\Diamond_i\varphi \vee_t \Diamond_i\psi$ just in case s supports $\Diamond\Diamond_i\varphi$ and $\Diamond\Diamond_i\psi$. Thus, to validate Wide Free Choice, we must guarantee the validity of the following principle:

INNER NEGATIVE INTROSPECTION $\Diamond\Diamond_i\varphi \models \Diamond_i\varphi$.

[22] Here our result parallels some of the ideas in Zimmermann 2000, who relies on various introspection principles to validate wide scope versions of Free Choice.

This requires that whenever s can be consistently updated with $\Diamond_i\varphi$, s supports $\Diamond_i\varphi$. That is, whenever the information at any world in s is consistent with φ, the information at every world in s is consistent with φ. This immediately implies that R_i is universal, so that whenever v is accessible from any w, v is accessible from every w. For suppose v is accessible from w but not u. Now let $s = \{w, u\}$, and let φ be a claim true at v uniquely. s supports $\Diamond\Diamond_i\varphi$, since $s[\Diamond_i\varphi] = \{w\} \neq \emptyset$. But s does not support $\Diamond_i\varphi$, since s contains u, which cannot see v.

Summing up, Wide Free Choice can be validated with a dynamic semantics for disjunction. But its validity is only guaranteed if we also constrain our semantics for possibility modals. This is a difference between Free Choice and Wide Free Choice. The former inference could be validated independently of any constraint on accessibility, while the latter cannot.

3.8 Negation

We have now explored a range of dynamic meanings for disjunction that can validate Free Choice. Now we will turn to the most serious problem for the semantics of Free Choice: negation. It is well known that the free choice inference disappears under negation.[23] In particular, the following Dual Prohibition inferences appears valid:

DP I $\neg\Diamond(\varphi \vee \psi) \models \neg\Diamond\varphi$ and $\neg\Diamond(\varphi \vee \psi) \models \neg\Diamond\psi$

DP II $\neg\Diamond\varphi; \neg\Diamond\psi \models \neg\Diamond(\varphi \vee \psi)$

Unfortunately, it is very difficult to validate both this inference and Free Choice. To see why, first note that the validity of both inferences is inconsistent with the contraposition and transitivity of entailment.

CONTRAPOSITION $\varphi \models \psi \Longrightarrow \neg\psi \models \neg\varphi$

TRANSITIVITY $\varphi \models \psi \;\&\; \psi \models \chi \Longrightarrow \varphi \models \chi$

Together, these assumptions lead to the absurd result that any possibility claim implies any other:

EXPLOSION $\Diamond\varphi \models \Diamond\psi$

[23] See Alonso-Ovalle 2006; Aloni 2007; Barker 2010; Willer 2015; and Starr 2016.

Fact 3.10. Free Choice, Dual Prohibition, Contraposition, and Transitivity imply Explosion.

As a matter of fact, Contraposition is invalid in the sorts of dynamic systems we've looked at above. In particular, these theories usually predict that while $\varphi \models \Box\varphi$, we do not have that $\neg\Box\varphi \models \neg\varphi$.

This should give us some hope that we could validate both Free Choice and Dual Prohibition in our dynamic semantics. Unfortunately, there is no straightforward way to do so. Let's hold fixed the standard dynamic meaning for negation, where updating with $\neg\varphi$ removes any worlds that survive update with φ.

Definition 3.24. $s[\neg\varphi] = s - s[\varphi]$

Given this assumption, Dual Prohibition then amounts to the requirement that updating with $\varphi \vee \psi$ produces the absurd state iff updating with φ and ψ each produce the absurd state.

Observation 3.8. Dual Prohibition is valid just in case for any state s:
$s[\varphi \vee \psi] = \emptyset$ iff $s[\varphi] = \emptyset$ and $s[\psi] = \emptyset$.

When we combine Observation 3.8 with our earlier observation about what it takes to validate Free Choice (Observation 3.3), we reach a serious problem. Free Choice is valid just in case for any state s, updating with $\varphi \vee \psi$ is consistent iff updating with each of φ and ψ is consistent. Dual Prohibition is valid just in case updating with $\varphi \vee \psi$ is inconsistent iff updating with each of φ and ψ is inconsistent. But these conditions together crowd out the space of meanings for disjunction. For we know that either updating with $\varphi \vee \psi$ is consistent, or it is inconsistent. But if the left hand sides of our two observations are exhaustive, then their right hand sides must be exhaustive also. But this means that either updating with each of φ and ψ is consistent, or updating with each of φ and ψ is inconsistent. But this is trivializing: it rules out a case where φ can be learned consistently, but ψ cannot. That is, we validate Explosion.

Fact 3.11. If Free Choice and Dual Prohibition are valid, then Explosion is valid.

We saw earlier in this chapter that there is a small range of dynamic meanings for disjunction that can validate Free Choice, holding fixed a usual dynamic meaning for possibility modals. Unfortunately, we've now seen that if we hold fixed the usual dynamic meaning for negation, there is no nontrivial semantics for disjunction that can validate both Free Choice and Dual Prohibition.

There are a few ways to respond to this problem. First, we might pursue a pragmatic account of Dual Prohibition, as in Barker 2010. Second, we might enrich our meanings further. Rather than simply letting our meanings be a function from sets of worlds to sets of worlds, we might let them contain more information. Willer 2015 and Starr 2016 have both pursued this strategy to explain Dual Prohibition. Third, we might investigate new dynamic meanings for negation, and find one that allows some natural meanings for disjunction where Free Choice and Dual Prohibition are valid. It is clear that there are at least some such meanings. As a toy example, if we interpret negation as the identity function ($\lambda f.f$) then Free Choice and Dual Prohibition will be validated by exactly the same theories of disjunction. Of course, this interpretation is not acceptable. So a task remains to formulate some natural constraints on the meaning of negation within dynamic semantics, relative to which we could investigate this question productively.

3.9 Conclusion: a context sensitive alternative

Our main investigation is now concluded. In this chapter, we first saw that there is no way to validate Free Choice without giving up a propositional theory of disjunction. In response, we introduced a family of dynamic meanings for disjunction.

But perhaps we have interpreted the validity of Free Choice in a way that is overly strong. So far, we have required that there be a single function $*$ such that any body of information is consistent with $A * B$ just in case it is consistent with both A and with B. But another option would be to let the choice of $*$ vary with the body of information. On this proposal, disjunctions themselves would be information-sensitive, relativized to an accessibility relation, in a way that can covary with the modal under which they are embedded. For each modal $\Diamond_i$ associated with accessibility relation R_i there is some disjunction operation $\vee_i$ that is associated with R_i.[24] Perhaps a propositional semantics for $\vee_i$ can be constructed that could then validate Free Choice for each possibility modal:

SHIFTY FREE CHOICE 1 $\Diamond_i(\varphi \vee_i \psi) \models \Diamond_i\varphi$ and $\Diamond_i(\varphi \vee_i \psi) \models \Diamond_i\psi$

SHIFTY FREE CHOICE 2 $\Diamond_i\varphi; \Diamond_i\psi \models \Diamond_i(\varphi \vee_i \psi)$

[24] This strategy parallels the context sensitive response to Lewis 1976's triviality results for conditionals. For a recent version of this context sensitive strategy, see Bacon 2015.

To conclude this chapter, let's develop such a shifty response in detail. The basic idea will be to again incorporate the insights of Zimmermann 2000, but this time in a propositional theory of disjunction. The proposition expressed by $\varphi \vee_i \psi$ says that one of φ and ψ are actually true, and that both are true at some i-accessible possibility. Since disjunction is relativized to an accessibility relation, we could imagine different flavors of disjunction, depending on the relevant accessibility relation. So we will actually let our operator $\vee_i$ itself be relativized to R_i:

Definition 3.25. $[\![\varphi \vee_i \psi]\!] = ([\![\varphi]\!] \cup [\![\psi]\!]) \cap [\![\Diamond_i \varphi]\!] \cap [\![\Diamond_i \psi]\!]$

For what follows, let's also assume that R_i is an equivalence relation: reflexive, symmetric, and transitive.

Now that we have set up our assumptions, let's see how the resulting theory engages with our earlier impossibility results. First, since R_i is an equivalence relation it is not partitional. So we can avoid the first challenge to propositionality.

Second, let's consider the case of partitional families of possibility operators. This is where our shifty relativization will help. Suppose we have three possibility operators $\Diamond_1$, $\Diamond_2$ and $\Diamond_3$ that are together a partitional set. Then we can build three corresponding disjunction operators $\vee_1$, $\vee_2$, and $\vee_3$ that are sensitive to the corresponding accessibility relation. We can then validate Shifty Free Choice for each pair of $\Diamond_i$ and $\vee_i$. But, crucially, no one disjunction operator satisfies Free Choice for all three possibility operators. So we can avoid the challenge to propositionality above. Summing up, then, the semantics above can validate Shifty Free Choice, once the accessibility relation is assumed to be an equivalence relation.

Observation 3.9. Shifty Free Choice is valid.

To see why, consider that Shifty Free Choice is valid iff the following biconditional holds: for any w, there is some R_i accessible v in $[\![\varphi \vee_i \psi]\!]$ just in case there is some R_i accessible v in $[\![\varphi]\!]$ and some R_i accessible u in $[\![\psi]\!]$. There is some R_i accessible v in $[\![\varphi \vee_i \psi]\!]$ just in case there is some u where vR_iu such that $u \in [\![\varphi]\!]$ and there is some u' where vR_iu' such that $u' \in [\![\psi]\!]$. Since R_i is an equivalence relation, this holds just in case u and u' are both accessible to w.

We've now developed a contextualist alternative to the dynamic semantics we explored earlier. On this theory, we gave up the hope of representing disjunction with a single propositional operator, and instead introduced a family of such operators. This could to some extent salvage a traditional, propositional theory of disjunction.

How can we adjudicate between these two strategies? There are two related concerns for the contextualist strategy above. First, it is inelegant, because it requires an enrichment of logical form to include indices on every occurrence of a disjunction. Second, it overgenerates. On this strategy, we have a family of possibility modals, and a family of disjunction operators. This should allow mixed sentences, where we embed for example a deontic disjunction under an epistemic possibility modal. This should predict that we could find English sentences that violate Free Choice all the time, since the relevant LFs should be available.

Chapter 4

A theory of conditional assertion

4.1 Introduction

According to one tradition, uttering an indicative conditional involves a special kind of speech act: a conditional assertion.

> An affirmation of the form 'if p then q' is [not] an affirmation of a conditional [but] a conditional affirmation of the consequent. If, after we have made such an affirmation, the antecedent turns out true, then we consider ourselves committed to the consequent…[If] the antecedent turns out to have been false, our conditional affirmation is as if it had never been made.[1]

The slogan above is intuitive. Quine [1959] took it to be the usual, "everyday" attitude. Yet one way or another, conditional assertion is not one of the dominant theories of conditionals today. This is for many reasons, perhaps best summed up in DeRose 1999:

> We believe these are mostly problems with understanding the view in the first place—understanding just what conditional assertions are, how they operate, and how they interact with the common machinery of current philosophy of language —machinery that was not set up with such a speech act in view…It can also seem problematic what kind of logic conditionals could have if they were devices of conditional assertion.[2]

The crucial task in giving an account of conditional assertion is to model the intuitive ideas above using the tools of contemporary semantics and pragmatics. Several attempts to do so have been judged "immediately implausible" or incoherent.[3] For example, several leading accounts of conditional assertion imply that when the antecedent of a conditional is false, the conditional fails

---

[1] Quine 1959 p. 12. See also von Wright 1957 p. 131; and Mackie 1973 and discussion. For developments of the view, see Jeffrey 1963; Belnap 1970; Belnap 1973; Manor 1974; Barker 1995; Milne 1997; and DeRose 1999 among others.

[2] See DeRose 1999.

[3] See for example Dummett 1973 and Lycan 2006.

to express a proposition.[4] Here, conditional assertion is analogous to handing someone a sealed envelope containing a record of the consequent, with instructions to open only if the antecedent is true.[5] Yet there are many cases where this seems wrong. Consider:

(1) If you don't finish mowing the yard this afternoon, you can't go to the mall after dinner.

Lycan 2006 observes that the addressee learns something immediately upon hearing this conditional. They need not wait to learn whether the antecedent is true, to determine whether anything has been expressed. Moreover, they cannot control whether the speaker fails to express a proposition simply by mowing the lawn.

It is unclear, however, whether the problems above are essential to the theory of conditional assertion. In fact, a wide range of theorists have adopted the slogan of conditional assertion in some way or other. Jeffrey 1963, Belnap 1970, Manor 1974, and McDermott 1996 all offer implementations of the ideas above that involve a failure of truth value in the case when the antecedent is false. But Edgington 1995 and Barker 1995 instead implement the crucial ideas behind conditional assertion using an NTV theory. Stalnaker 2011 argues that the material conditional itself satisfies most of the goals of conditional assertion, and then argues that his own theory of the indicative does too. Starr Forthcoming shows how dynamic conditionals can accomplish much of the same.

Summing up: conditional assertion is the usual, everyday account of conditionals; and it is immediately implausible. It has been largely ignored by the literature; and is also a component of every leading theory of conditionals. With this much disagreement on what conditional assertion even is, it's worth taking a step back and searching for a more precise definition.

This chapter offers a new theory of conditional assertion. We first develop a formal framework that can model speech acts like conditional assertion. This framework is ecumenical, usable even by those who are skeptical of a semantic treatment of speech acts. All we assume is that we have a set of contexts, that agents are committed to certain claims in context, and that utterances change

[4] Belnap 1970 develops a version of this theory in possible worlds semantics. Jeffrey 1963, Manor 1974, Stalnaker 1975 (fn. 2) and McDermott 1996, Huitink 2009a, and Huitink 2009b all explore a view where the conditional is truth-valueless when the antecedent is false.

[5] See Jeffrey 1963.

the context. We can then use this model to provide a precise characterization of what conditional assertion is.

Our characterization of conditional assertion is weak enough to be consistent with a broad range of theories. It does not have any immediately paradoxical conclusions, such as denying that conditionals with false antecedents express propositions. Nonetheless, our theory is still strong enough to be interesting. Our first result is that relying on just a few background assumptions, we can prove that any theory of conditional assertion in our sense implies the validity of several inferences involving conditionals. First, conditional assertion requires the validity of Modus Ponens. Second, conditional assertion requires the validity of one of the paradoxes of material implication: that the negation of the antecedent implies the conditional.

With these logical results under our belt, we turn to semantic implementations of conditional assertion, in three increasingly rich theoretical frameworks. First, we consider whether the conditional assertion theory can be implemented within a truth conditional theory of conditionals. We will see that in this setting, conditional assertion requires that the conditional is almost truth functional. The conditional is true whenever the antecedent is false, and false whenever the antecedent is true and the consequent is false. But the conditional's truth is unsettled when the antecedent and consequent are both true.

After considering truth conditional implementations, we turn to dynamic semantics.[6] We identify a new family of dynamic conditionals that satisfy the requirements of conditional assertion. On all these theories, the conditional tests whether learning the antecedent suffices to accept the consequent. However, these theories differ in what happens when this test fails. Extant dynamic conditionals (for example, that in Gillies 2004) predict that in cases of failure, we reach an absurd state. By relaxing this assumption, we generate a family of conditional assertion theories. We let failure of the test create a state containing all worlds from the context in which the antecedent is false, and then some. The result is a new kind of dynamic meaning.

The theories we examine up to this point all predict that conditional assertions affect which worlds in the context are live possibilities. In the final section of the chapter, we draw on work from Veltman 1985 to develop an enriched model of context which relaxes this assumption. In addition

[6] For some defenders of dynamic semantics, see Veltman 1996, Groenendijk et al. 1996, Beaver 2001, Gillies 2004, and Willer 2013. For some nearby views, see Yalcin 2007, Swanson 2011, Swanson 2012, and Moss 2015.

to the live possibilities of the context, the context also contains a record of which extensions of the live possibilities are admissible. We let conditional assertion update this record, without changing which possibilities are currently live.

4.2 Conditional assertion

We model a formal language $\mathscr{L}$ consisting of an ordinary propositional language $\mathscr{L}^0$, supplemented by an indicative conditional operator $\rightarrow$. We focus on conditionals that do not themselves contain another conditional.[7]

Definition 4.1. Let $\mathscr{A}$ be the set of atomic sentences α. Let $\mathscr{L}^0$ be a language containing $\mathscr{A}$ and closed under negation $\neg$ and conjunction $\wedge$. Let $\mathscr{L}$ be the smallest set which contains $\mathscr{L}^0$ and which, for any sentences φ and ψ in $\mathscr{L}^0$, contains the indicative conditional $\varphi \rightarrow \psi$ and its negation $\neg(\varphi \rightarrow \psi)$. Let $\vee$ and $\supset$ abbreviate $\neg(\neg\varphi \wedge \neg\psi)$ and $\neg(\varphi \wedge \neg\psi)$.

We assume that utterances take place in context, which we represent with the variable s. We let $[\cdot]$ be a function modeling the effect of performing a speech act in context. So an utterance of φ in s generates a new context $s[\varphi]$. If φ is a declarative sentence (in $\mathscr{L}^0$), $s[\varphi]$ models the effect of asserting φ in s. By contrast, $s[\varphi \rightarrow \psi]$ models the effect of conditionally asserting ψ given φ. Finally, we assume that within any context the relevant agents are committed to some claims. We use $\models$ to express this relation, so that $s \models \varphi$ just in case the participants of context s are committed to φ. Finally, we call a context absurd just in case it is committed to every claim.[8]

Definition 4.2. Let s be a context. Let C be the set of contexts. Let $[\cdot]$ be a function from sentences in $\mathscr{L}$ to functions from C to C. Let $\models$ be a relation between contexts s and sentences φ in $\mathscr{L}$. s is absurd just in case $s \models \varphi$ for every $\varphi \in \mathscr{L}$.

To start with, we make no assumptions about the exact structure of contexts. For now, our goal is to take the intuitive idea behind conditional assertion, and see how far this idea alone takes us towards a precise theory of conditionals.

---

[7] The language below allows embeddings of conditionals under negation. This may seem worrisome for some defenders of conditional assertion, since it requires negating a conditional speech act. Fortunately, none of our results rely especially on this feature of the language. In particular, in Fact 4.1 the inference $\varphi; \neg\psi \models \neg(\varphi \rightarrow \psi)$ could be replaced with $\varphi; \varphi \rightarrow \psi \models \psi$.

[8] Here we rely on many of the formal tools introduced within update semantics. For some representative samples, see Veltman 1996, Beaver 2001, and Gillies 2004.

Given the framework we have adopted, the theory of conditional assertion can be expressed as two constraints on $[\varphi \rightarrow \psi]$. Suppose that $\varphi \rightarrow \psi$ is uttered in context s. This moves us to state $s[\varphi \rightarrow \psi]$. The theory of conditional assertion tells us what happens if the antecedent φ turns out true, and also what happens if the antecedent turns out false. We model these two scenarios using $[\varphi]$ and $[\neg\varphi]$. So the conditional assertion theory tells us something about both $s[\varphi \rightarrow \psi][\varphi]$ and $s[\varphi \rightarrow \psi][\neg\varphi]$.

In the former case, the theory says that we are committed to the consequent. Since we are using $\models$ to model commitment, this means that $s[\varphi \rightarrow \psi][\varphi] \models \psi$. In other words: given any body of information, if the conditional and its antecedent are added to the information, then the new information contains the consequent.

In the second case, where the antecedent turns out false, our original conditional assertion is 'as if it had never been made'. We model this last idea through the identity $s[\varphi \rightarrow \psi][\neg\varphi] = s[\neg\varphi]$. In other words: when a state is updated with the conditional and then the negation of its antecedent, the state might as well just have been updated with the antecedent's negation immediately. We can also think about this last requirement more generally as a thesis about what it is to learn a conditional. No matter what information an agent has, if they first learn the conditional and then learn that the antecedent is false, they arrive in the same state as if they had just learned that the antecedent were false from the beginning. Learning $\neg\varphi$ screens off learning $\varphi \rightarrow \psi$. Putting these two ideas together, we reach the following theory:

Definition 4.3. $\rightarrow$ is a conditional assertion operator iff:

1. $s[\varphi \rightarrow \psi][\varphi] \models \psi$
2. $s[\varphi \rightarrow \psi][\neg\varphi] = s[\neg\varphi]$

In the rest of this chapter, we explore the implications of the definition above for the theory of conditionals. But before we do this, it's worth flagging a few points where we could have made a different choice.

First, when the antecedent turns out true, we said that we are committed to the consequent. But here we might replace the notion of commitment with that of actually uttering the consequent.

Then we would say that uttering a conditional has the same effect as uttering the consequent whenever the antecedent turns out true. That gives us the equation $s[\varphi \to \psi][\varphi] = s[\varphi][\psi]$. This condition is significantly stronger than our own. In particular, we show later that our own condition is compatible with the failure of $\varphi \wedge \psi$ to imply $\varphi \to \psi$. This inference's validity is guaranteed by the stronger condition here.

Second, in both clauses above we analyzed the notion of the antecedent turning out true or false in a particular way: in turns of updating the context with the antecedent. However, one might distinguish the bare truth or falsity of the antecedent from its being learned. In that case, one might say that $s[\varphi \to \psi] \models \psi$ whenever φ is true, and $s[\varphi \to \psi] = s$ whenever φ is false. Something like this idea is suggested by Belnap 1970, who argued that the proper interpretation of our second condition requires the conditional to fail to express a proposition:

> If the conditional asserted something according to semantics, we could not on the pragmatic level treat its utterance as if it had never happened.[9]

We are now in a position to see why this is too strong. The conditional can assert something when the antecedent is false, as long as whatever it asserts is screened off by the assertion of the antecedent's negation. In this case, the conditional assertion might as well never have happened.

One concern with the gappy theory above, compared with our own, is that it requires a conditional's effect on context to depend on something other than the context itself: the antecedent's actual truth value. This idea is hard to make sense of in the model above. In addition, it is hard to know how conversational participants could implement this idea when s leaves the antecedent unsettled. That is, this idea violates Stalnaker [1978]'s uniformity principle, on which a sentence should have the same effect on context at every world consistent with the information of the context.

Our definition of conditional assertion is quite different from some of the traditional accounts, especially those involving truth value gaps. Ultimately, the results below serve as our argument for thinking of conditional assertion this way, rather than as it has been interpreted before. Whether or not this is what others have had in mind when talking about conditional assertion, we show

[9] See Belnap 1973, p. 49.

that our definition is fruitful, characterizing a new and well behaved region of semantic and logical space.

Before jumping in to the details of the theory, it's worth flagging a few benefits of the account right from the start. First, the account avoids the 'initial implausibility' objection to earlier theories of conditional assertion, such as Jeffrey 1963 and Belnap 1970. Our definition in no way requires that the conditional is meaningless, truth valueless, or fails to express a proposition when the antecedent is false. Rather, we understand the falsity condition in a dramatically different way, as a screening off condition. That learning $\neg\varphi$ screens off the information in $\varphi \rightarrow \psi$ is perfectly consistent with the fact that we can use conditionals to communicate information even when as a matter of fact the antecedent is false.

Another initial benefit of the theory is that it is in several senses ecumenical. First, several proponents of conditional assertion have accepted an NTV view of conditionals, denying that they express propositions that can be assigned a truth value. While the theory above is consistent with conditionals expressing propositions, it does not require this. So our theory provides a neutral framework within which propositional and NTV proponents alike can investigate the consequences of conditional assertion.

The theory is also ecumenical in another way. Conditional assertion has primarily been thought about within the theory of speech acts. But our set of contexts above could be anything. In particular, we can interpret contexts simply as the possible beliefs that a rational agent might possess. Then $[\cdot]$ is the rule by which an agent updates her beliefs over time, and $s \models \varphi$ says that an agent in state s believes φ. Under this interpretation, the theory above defines an epistemic concept of conditional learning, rather than a pragmatic concept of conditional assertion. The results below double as a characterization of this special kind of learning.

Finally, the theory is ecumenical within the theory of pragmatics. One natural goal with conditional assertion is to provide a general account of the conditional that connects conditional assertion with other speech acts like conditional promising, questioning, and commanding.[10] Again, because our definition of a context is completely neutral, our framework can also be used to model

[10] See Barker 1995 and Lycan 2006 for discussion, among others.

these other conditional speech acts. For example, we might represent a context as a stack of questions under discussion, as a to-do list of commands, or with all of the above.[11] In each case, our definition of conditional assertion above generalizes straightforwardly to other conditional speech acts. For example, on the definition above a conditional command would add the consequent's command to the to-do list of a state when the antecedent is learned, but would have no effect on the to-do list of a context if the antecedent's negation is learned.

We now have a precise statement of what conditional assertion is, along with a rough sense of how it compares to a few alternatives. In the rest of this chapter, we explore the implications of this model for the logic, truth conditions, and conversational dynamics of conditionals.

4.3 The logic of conditional assertion

Our first goal is to investigate the logic of conditional assertion operators. In particular, we introduce a few minimal assumptions, and use them to show that any conditional assertion operator must validate several inferences, including Modus Ponens and the False Antecedent inference.

Our model contains a rule for updating contexts with an utterance, and also a representation of the commitments that each context contains. These two features of our model allow us to formulate a definition of entailment. In particular, let's say that an argument is valid just in case we are committed to its conclusion whenever we utter the premises.[12]

Definition 4.4. $\varphi_1, \ldots, \varphi_n \models \psi$ iff $\forall s\; s[\varphi_1] \ldots [\varphi_n] \models \psi$

The model above is ecumenical. It need not be interpreted as a semantic theory of conditionals. We could instead think of it merely as a formal pragmatics, describing how various speech acts affect the information of speakers. In that case, we can think of entailment as a formal model of how information changes in systematic ways when speech acts are performed. That being said, we also show in later sections that the model can be interpreted truth conditionally. In that setting, this definition coincides with preservation of truth.

---

[11]See Roberts 2012, Portner 2007, and Murray and Starr Forthcoming for examples.

[12]See van Benthem 1996 for discussion. Our results below also hold for another popular definition of entailment: that an argument is valid just in case whenever we are committed to all the premises, we are committed to the conclusion.

To prove our results, we need to add a few more assumptions. Basically, we assume that the fragment of our language not containing conditionals is conservative, obeying a variety of structural rules that bring it close to behaving truth conditionally. But we make no assumptions about the behavior of conditionals. First, our entailment relation allows us to make a few well behavedness assumptions about negation. Without accepting any specific semantics for negation, we can commit to the law of noncontradiction. In our framework, this requires that any state is committed to anything after being updated with both of φ and $\neg\varphi$. In addition, we assume that any coherent context which is not committed to φ will remain coherent after an utterance of $\neg\varphi$.

Definition 4.5. $\neg$ is well behaved only if:

1. $\varphi; \neg\varphi \models \psi$
2. If $s \not\models \varphi$, then $s[\neg\varphi]$ is not absurd.

We also need several assumptions about updating. First, we rely on a strong principle linking commitment and update. We suppose that the participants of a context are committed to something just in case uttering it leaves their state unchanged. Second, we suppose that updating a context always gives us more information. More precisely, we suppose that we can never get back to s by uttering two sentences φ and ψ, at least one of which changes the state.[13]

Definition 4.6. $\models$ and $[\cdot]$ are well-behaved only if:

1. $s[\varphi] = s$ iff $s \models \varphi$.
2. If $s[\varphi] \neq s$ or $s[\varphi][\psi] \neq s[\varphi]$, then $s[\varphi][\psi] \neq s$.

Finally, we focus on sentences that are well behaved in a few ways. First, we only consider antecedents and consequents that preserve commitment under update.[14] Second, we focus on antecedents that are idempotent, so that any context is committed to them after updating with them.

Definition 4.7.

[13] See Groenendijk and Stokhof 1991b and Rothschild and Yalcin 2015a for discussion.

[14] See Veltman 1996 for discussion. We do not require that the conditional or its negation is persistent.

1. φ is persistent iff for any s, ψ, whenever $s \models \varphi$, $s[\psi] \models \varphi$.

2. φ is idempotent iff for any s, $s[\varphi] \models \varphi$.

With all of these assumptions in place, we now show that the conditional assertion theory imposes a significant amount of truth functionality on the conditional. To see why, note that we can use our definition of entailment above to express a version of truth functionality.[15] The conditional is truth functional just in case a commitment to whether φ and ψ each hold is sufficient to generate a commitment on whether $\varphi \to \psi$:

Definition 4.8. $\to$ is truth functional iff:

1. $\varphi; \psi \models \varphi \to \psi$ or $\varphi; \psi \models \neg(\varphi \to \psi)$, and

2. $\neg\varphi; \psi \models \varphi \to \psi$ or $\neg\varphi; \psi \models \neg(\varphi \to \psi)$, and

3. $\varphi; \neg\psi \models \varphi \to \psi$ or $\varphi; \neg\psi \models \neg(\varphi \to \psi)$, and

4. $\neg\varphi; \neg\psi \models \varphi \to \psi$ or $\neg\varphi; \neg\psi \models \neg(\varphi \to \psi)$.

This notion of truth functionality does not assume that the conditional has truth conditions. Rather, this is a commitment-theoretic notion of truth functionality, requiring that our commitment to the conditional supervenes on our commitments to its parts. Again, we show later that all the key notions here can be interpreted truth conditionally, in which case the definition above coincides with the usual one.

We now show that any conditional assertion operator is at least three quarters truth functional. A commitment to $\neg\varphi$ implies a commitment to $\varphi \to \psi$, and a commitment to φ and to $\neg\psi$ implies a commitment to $\neg(\varphi \to \psi)$. Nonetheless, this leaves open the possibility that a state could be committed to φ and to ψ while remaining agnostic on whether $\varphi \to \psi$.

Fact 4.1. Suppose $\models$, $[\cdot]$, and $\neg$ are well-behaved, φ and $\neg\psi$ are persistent, and $\neg\varphi$ is idempotent. Suppose $\to$ is a conditional assertion operator. Then:

1. $\varphi; \neg\psi \models \neg(\varphi \to \psi)$

---

[15]See Fine 2016 for discussion.

2. $\neg\varphi \models \varphi \to \psi$.

(For proofs of the main facts in what follows, see the appendix.)

The first condition in the above is basically equivalent to the validity of Modus Ponens. While we can also show that Modus Ponens must be valid for any conditional assertion operator, we stick with the formulation above because of its connection with truth functionality.

Fact 4.1 shows that conditional assertion places surprising demands on the logic of conditionals. After all, the False Antecedent inference from $\neg\varphi$ to $\varphi \to \psi$ is controversial. Consider:

(2) a. The butler did it.

b. ??So: if the butler didn't do it, then the gardener did.

This inference is one of the signature properties of the material conditional. Nonetheless, conditional assertion does not imply collapse to the material conditional. While the False Antecedent inference is required of any conditional assertion operator, the True Consequent inference is not. In particular, the truth of the antecedent and consequent need not imply the truth of the corresponding conditional. This And to If inference is controversial. For example, following McDermott 2007 imagine that a coin is tossed twice, and Sally bets that it will land heads both times. Suppose the coin does land heads twice. In this scenario (3) seems false:

(3) If at least one heads came up, Sally won.

After all, in any scenario in which one coin landed heads and the other landed tails, Sally loses the bet. This is just one of many counterexamples to And to If in the literature. Others have focused on cases where φ and ψ are irrelevant to one another, or where ψ is extremely unlikely in the presence of φ.[16] Moreover, Lycan 2006 objects to the theories of conditional assertion developed in Jeffrey 1963 and Belnap 1970 precisely because they validate And to If. We've now seen that this objection does not extend to all theories of conditional assertion.

The theory of conditional assertion imposes heavy demands on the logic of conditionals. But what about the semantics? In the rest of this chapter, we consider which semantic theories of

---

[16]For discussion see among others Lewis 1973, p. 27, Bennett 1974, Fine 1975, Penczek 1997, McDermott 2007, McGlynn 2012, and He 2016.

conditionals are consistent with conditional assertion. To do so, we assign progressively richer meanings to the conditional. We start within a truth conditional framework. Later, we let the conditional have a dynamic meaning. Finally, we develop a richer model of meaning in which a context contains not only a set of live possibilities, but also a record of how the context may evolve in the future.

4.4 The truth conditions of conditional assertion

We now explore the space of conditional assertion theories for truth conditional operators. To do so, we now interpret contexts as sets of possible worlds.[17]

Definition 4.9. Let a world w assign a truth value to every atomic sentence. Let a context s be a set of worlds.

Once contexts are sets of worlds, we can model assertion and conditional assertion in terms of narrowing down this set. Since we are assuming the conditional is truth conditional, we can assign every sentence a set of possible worlds where it is true:

Definition 4.10. Let $[\![\cdot]\!]$ assign every sentence in $\mathscr{L}$ a set of possible worlds.

Then we can introduce a truth conditional update function $[\cdot]_1$, which lets an utterance of any sentence narrow down a context to the worlds where it is true.[18]

Definition 4.11. $s[\varphi]_1 = s \cap [\![\varphi]\!]$

Say that $\rightarrow$ is propositional just in case $[\![\varphi \rightarrow \psi]\!]$ is defined for any φ, ψ in $\mathscr{L}^0$ and updating with $\varphi \rightarrow \psi$ amounts to intersecting with $[\![\varphi \rightarrow \psi]\!]$ (that is, just in case $[\varphi \rightarrow \psi] = [\varphi \rightarrow \psi]_1$). Our task now is to determine what constraints are imposed on $[\![\cdot]\!]$ if we require $\rightarrow$ to be a conditional assertion operator.

Our earlier result, Fact 4.1, has a corollary in the propositional domain. We can show the conditional is a propositional conditional assertion operator just in case (i) $\varphi \rightarrow \psi$ is true at any world where φ is false, and (ii) $\varphi \rightarrow \psi$ is false at any world where φ is true and ψ is false.

[17] See Stalnaker 1978.

[18] Throughout, we omit subscripts on our update functions whenever possible.

Fact 4.2. If $\to$ is propositional and $\neg$ is well behaved, then $\to$ is a conditional assertion operator iff:

1. $[\![\varphi]\!] \cap [\![\neg\psi]\!] \cap [\![\varphi \to \psi]\!] = \emptyset$
2. $[\![\neg\varphi]\!] \subseteq [\![\varphi \to \psi]\!]$.

Stalnaker 2011 observes that updating with the material conditional achieves the goals of the conditional assertion theory. We have now seen that the material conditional is not alone in doing so. Rather, there is a family of truth conditional theories with this property.

In this section, we figured out the truth conditions of any conditional assertion operator, assuming that it has truth conditions. But perhaps we should not model conditional assertion using truth conditions. Fortunately, this does not require that we give up on semantically interpreting the conditional. Rather, in the next section we uncover a rich family of conditional assertion operators within the framework of dynamic semantics. On these accounts, conditionals do not express propositions. Updating with them is not a matter of simply learning which world one inhabits.

4.5 The dynamics of conditional assertion

To explore nonpropositional theories of conditionals, we retain some of the assumptions from the previous section. We again assume that contexts are sets of worlds. But now we only semantically associate declarative sentences with sets of worlds.

Definition 4.12. Let $[\![\cdot]\!]$ assign every sentence in $\mathscr{L}^0$ a set of possible worlds.

Then we introduce an update function $[\cdot]_2$ which behaves conservatively for ordinary assertions, but not for conditional assertions. $[\cdot]_2$ lets ordinary assertions update s intersectively.

Definition 4.13. If $\varphi \in \mathscr{L}^0$, then $s[\varphi]_2 = s \cap [\![\varphi]\!]$.

Crucially, we do not require that $[\varphi \to \psi]_2$ itself update the context in a traditional, truth conditional manner. That is, we do not assume the existence of some fixed proposition p where for any state s, $s[\varphi \to \psi]_2 = s \cap p$. Rather, we allow $[\varphi \to \psi]_2$ to be any function whatsoever from sets of worlds to sets of worlds. Now we explore what constraints are imposed on $[\varphi \to \psi]_2$ by conditional assertion.

With Fact 4.1, we saw that the validity of Modus Ponens and False Antecedent are necessary for conditional assertion. Then with Fact 4.2 we saw that their joint validity is also sufficient for conditional assertion in a truth conditional framework. This sufficiency claim fails when the conditional is genuinely dynamic, not amounting to intersection with a fixed set of worlds. Consider the dynamic conditional from Gillies 2004, where $\varphi \to \psi$ tests s to see whether $s[\varphi]$ is committed to ψ.

$$s[\varphi \to \psi] = \{w \in s \mid s[\varphi] \models \psi\}$$

This semantics validates Modus Ponens and False Antecedent. However, it rejects the screening off component of conditional assertion. For a wide range of contexts, $s[\varphi \to \psi][\neg\varphi] \neq s[\neg\varphi]$. For example, suppose s contains a world w where φ and $\neg\psi$, and a world v where $\neg\varphi$. $s[\neg\varphi]$ is not empty, since it contains v. But $s[\varphi \to \psi][\neg\varphi]$ is empty, since $s[\varphi \to \psi]$ is empty. Summing up, the test conditional behaves like a conditional assertion operator whenever the test is passed; but when the test fails, it departs from the requirements of conditional assertion.

We can say more about the class of meanings that are not conditional assertion operators. The dynamic conditional above is an example of a test. Tests are the paradigmatic example of a dynamic meaning. They have been proposed as the meaning not only of conditionals, but also of possibility and necessity modals (Veltman 1996). To be more precise, a context change potential is a test just in case it always returns either the input, or the absurd state.

Definition 4.14.

1. $[\varphi]$ is a test iff for every state s: $s[\varphi] = s$ or $s[\varphi] = \emptyset$.
2. $\to$ is a test operator iff for any sentences φ and ψ: $[\varphi \to \psi]$ is a test.

The dynamic conditional above is a test operator, but is not a conditional assertion operator. This is no coincidence: no operator is both a test and a conditional assertion operator.

Fact 4.3. If $\to$ is a test operator, then $\to$ is not a conditional assertion operator.

Here is a general diagnosis of why tests are not conditional assertion operators. When a test fails, the absurd state results. This means that any test semantics for the conditional has the potential to remove worlds from the state where the antecedent is false. But in this case, $\neg\varphi$ does not screen off $\varphi \to \psi$. So the test is not a conditional assertion operator. If we want to find an interestingly dynamic conditional assertion operator, we need a new type of dynamic meaning.[19]

Tests are not conditional assertion operators because they can eliminate worlds where the antecedent is false. This points us in the direction of characterizing the class of conditional assertion operators. The first requirement of this class is simply that $s[\neg\varphi] \subseteq s[\varphi \to \psi]$. This guarantees that $s[\varphi \to \psi][\neg\varphi]$ and $s[\neg\varphi]$ are identical.

To produce a conditional assertion operator, we must meet one more requirement: that $s[\varphi \to \psi][\varphi] \models \psi$. To achieve this goal, let's consider $s[\varphi]$ rather than $s[\neg\varphi]$. Here, we must guarantee that after updating with $[\varphi \to \psi]$, s has entirely excluded $s[\varphi][\neg\psi]$. To achieve this result, we let $s[\varphi \to \psi]$ be the union of two updates: $s[\neg\varphi]$, and some privileged portion of $s[\varphi][\psi]$. Any component of $s[\varphi][\psi]$ will do, as long as $s[\varphi][\neg\psi]$ is eliminated.

To implement these ideas, we need one more tool. Say that f is a generalized selection function just in case f takes s, $[\varphi]$, and $[\psi]$ as input, and returns some subset of $s[\varphi][\psi]$. These generalized selection functions take a context and two propositions as input, and return some selected worlds in the context where both propositions are true. Generalized selection functions give us a precise way to think about the privileged portion of $s[\varphi][\psi]$. Then we can represent updating with the conditional in terms of the union of $s[\neg\varphi]$ with $f(s, [\varphi], [\psi])$.[20]

Definition 4.15.

1. A generalized selection function f is a function from a context and two *ccp*s to a new context, where $f(s, [\varphi], [\psi]) \subseteq s[\varphi][\psi]$.

2. $\to$ is selective iff $s[\varphi \to \psi] = s[\neg\varphi] \cup f(s, [\varphi], [\psi])$.

Selective conditionals split up the job of conditional assertion into two steps. To guarantee that $s[\varphi \to \psi][\varphi]$ is committed to ψ, the conditional only includes a region of $s[\varphi]$ that survives update

[19] Other extant dynamic conditionals in the literature also fail to be conditional assertion operators, for example the conditional in Russell and Hawthorne 2016, which is just like the test conditional above except that it behaves like the material conditional in cases where $s[\varphi]$ does not support $\neg\psi$.

[20] This formal tool allows us to require some sort of connection between φ and ψ. This bears some resemblance to previous work in Chisholm 1946, Goodman 1947, Mackie 1962, Rescher 1964, and especially Gabbay 1972, who explicitly introduces an operator that takes both φ and ψ as input, and delivers a set of worlds relevant to the evaluation of the conditional. Our own theory departs significantly from all of the above, since it allows that the conditional is true whenever the antecedent is false.

with $[\psi]$, selected by f. To guarantee that the conditional is screened off by the negated antecedent, the conditional also includes all of the context that survives update with the negated antecedent. It turns out that $\rightarrow$ is a conditional assertion operator just in case it is selective in this way.

Fact 4.4. $\rightarrow$ is a conditional assertion operator iff $\rightarrow$ is selective.

We now have a characterization of conditional assertion in dynamic semantics. This result is in a sense less general than Fact 4.1, our logical characterization of conditional assertion, since it assumes that contexts are sets of worlds. For example, in a later section we develop a different model of contexts, and introduce a conditional assertion operator in that setting. The resulting operator is not selective, but still satisfies Modus Ponens and False Antecedent. But while our latest result is weaker than Fact 4.1 in one sense, in another sense it is stronger. It provides necessary and sufficient conditions for conditional assertion, while Fact 4.1 provided only necessary conditions. The result is also more general than Fact 4.2, which governed the truth conditional case. In particular, each conditional assertion operator in the truth conditional setting is also a selective operator, corresponding to a particular choice of selector. To guarantee truth conditionality, all that is required is that the selector be a function of $s \cap [\![\varphi]\!]$ and $s \cap [\![\psi]\!]$.

Importantly, however, there are also many nonpropositional selective conditionals. To see how they work, let's return to the test semantics above. There, we saw that the test semantics fails the screening off component of conditional assertion whenever the test for $s[\varphi] \models \psi$ is failed. A natural thought is then to change the test semantics for this failure condition. In the semantics above, failure of the test results in $\emptyset$, which eliminates $\neg\varphi$ worlds from s. We can modify the meaning above by letting failure of the test result in a different state. In particular, we can let it result in $s[\neg\varphi] \cup f(s, [\varphi], [\psi])$, our selective output from before. When the conditional has this form, guaranteeing that the context is committed to the consequent after updating with the antecedent, but otherwise updating with $s[\neg\varphi] \cup f(s, [\varphi], [\psi])$, let's call it testlike.

In fact, once a conditional is selective, we can derive this last operation by constraining f in a particular way. To achieve this result, call a selector testlike when it returns all of $s[\varphi][\psi]$ whenever $s[\varphi]$ is committed to ψ. When $s[\varphi]$ is committed to ψ, s is identical to $s[\neg\varphi] \cup s[\varphi][\psi]$. This implies that testlike selectors allow the conditional to return the entire state whenever $s[\varphi] \models \psi$.

Definition 4.16.

1. f is testlike iff $f(s, [\varphi], [\psi]) = s[\varphi][\psi]$ whenever $s[\varphi] \models \psi$.

2. $\rightarrow$ is testlike iff $s[\varphi \rightarrow \psi] = \begin{cases} s & \text{if } s[\varphi] \models \psi \\ s[\neg\varphi] \cup f(s, [\varphi], [\psi]) & \text{otherwise.} \end{cases}$

Fact 4.5. $\rightarrow$ is testlike just in case $\rightarrow$ is selective and f is testlike.

Testlike operators retain much of the test semantics above. In particular, whenever $s[\varphi] \models \psi$, testlike operators act like tests, returning the input state. But they are also selective, which implies that they are conditional assertion operators. Testlike operators are thus a new type of dynamic meaning, designed to satisfy the demands of conditional assertion.

To better understand the family of testlike operators, let's explore a few representative examples. Testlike operators are a special case of selective operators. In a similar way, the material conditional is a special case of a testlike operator. To see why, note that as f returns more worlds, the conditional becomes weaker. The weakest testlike operator makes f the identity function. This operator never removes any worlds from $s[\varphi][\psi]$.

$$s[\varphi \rightarrow \psi] = \begin{cases} s & \text{if } s[\varphi] \models \psi \\ s[\neg\varphi] \cup s[\varphi][\psi] & \text{otherwise.} \end{cases}$$

This last operator is simply the material conditional in disguise, assuming $s[\varphi \supset \psi] = s[\neg\varphi] \cup s[\varphi][\psi]$. After all, either $s[\varphi]$ supports ψ or not. In the former case, $s[\varphi \rightarrow \psi] = s[\varphi \supset \psi] = s$. In the latter case, $s[\varphi \rightarrow \psi]$ simply eliminate all $\varphi \wedge \neg\psi$ possibilities from s. But this is exactly what $[\varphi \supset \psi]$ does in that case.

At the other extreme, the strongest testlike operator always removes every world in $s[\varphi][\psi]$. This conditional looks more like the test semantics from earlier:

$$s[\varphi \rightarrow \psi] = \begin{cases} s & \text{if } s[\varphi] \models \psi \\ s[\neg\varphi] & \text{otherwise.} \end{cases}$$

This gives us a conditional assertion operator; but it is a strange one. Suppose s contains some $\varphi \wedge \psi$ worlds and some $\varphi \wedge \neg\psi$ worlds. Then $s[\varphi \rightarrow \psi]$ removes all the $\varphi \wedge \psi$ worlds. This result seems too strong. So it seems that we want some choice of testlike operator in between these two extremes.

Not every testlike operator is nonpropositional, since the material conditional is testlike. Nonetheless, we can show that any choice of testlike f other than the identity function will result in a nonpropositional theory of $\rightarrow$. For these operators, updating is not just intersecting with a fixed set of possible worlds.

Fact 4.6. If $\rightarrow$ is testlike and $[\rightarrow] \neq [\supset]$, then $\rightarrow$ is not propositional.

For any such f, there will be some witnessing s, φ, and ψ where $f(s, [\varphi], [\psi]) \neq s[\varphi][\psi]$. This means that there will be some world w removed from $s[\varphi][\psi]$ by f. This w is a φ and ψ world. However, now consider the context $\{w\}$. $\{w\} \models [\varphi \rightarrow \psi]$, since $\{w\}$ passes the test imposed by $\varphi \rightarrow \psi$. This means that whether $[\varphi \rightarrow \psi]$ removes w from a context depends on the structure of that entire context. We cannot model the effect of $[\varphi \rightarrow \psi]$ in terms of intersecting s with a fixed set of possible worlds.

While testlike conditionals are essentially dynamic, they are only dynamic in a limited way. It turns out that any testlike conditional can be factorized into a static and a dynamic component. First, note that a standard semantics for disjunction within update semantics consists of updating the state with each disjunct, and unioning the result, so that $s[\varphi \vee \psi] = s[\varphi] \cup s[\psi]$. This semantics for disjunction is not essentially dynamic. In particular, whenever each of φ and ψ are propositional, so is $\varphi \vee \psi$. Since testlike conditionals are all selective, we know that we can express their meaning with the identity $s[\varphi \rightarrow \psi] = s[\neg\varphi] \cup f(s, [\varphi], [\psi])$. Given the semantics for disjunction above, that means that we can define any conditional assertion operator in terms of disjunction. In particular, we could introduce into our object language a new connective $\overline{\wedge}$ which takes two sentences as input and returns the result of selecting some worlds in s where both claims are true. For example, this operator might be the meaning of *because*, which entails its inputs and expresses some connection between them. Then we can define the conditional in terms of $\vee$ and $\overline{\wedge}$:

Definition 4.17.

1. $s[\varphi \vee \psi] = s[\varphi] \cup s[\psi]$
2. $s[\varphi \mathbin{\overline{\wedge}} \psi] = f(s, [\varphi], [\psi])$

Fact 4.7. If $\rightarrow$ is selective, then $[\varphi \rightarrow \psi] = [\neg\varphi \vee (\varphi \mathbin{\overline{\wedge}} \psi)]$.

Whenever $\rightarrow$ is selective, and hence whenever $\rightarrow$ is a conditional assertion operator, it can be expressed as a disjunction of $\neg\varphi$ and $\varphi \mathbin{\overline{\wedge}} \psi$. This means that for testlike operators, all of the nonpropositionality can be factorized into $\overline{\wedge}$. This in turn means that all of the essentially dynamic component of the conditional's meaning is contributed by the selector.

We have now identified a subclass of selective conditional operators. We know that these are all conditional assertion operators, and that with the exception of the material conditional they are all essentially dynamic. Above, we said that these testlike operators retain something from the test semantics: an account of what happens when $s[\varphi]$ is committed to ψ. This feature from the test semantics is important, because it corresponds to the Ramsey Test. Here, consider Ramsey 1990's famous idea connecting conditional belief and learning:

> 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.

We can represent this idea in our own framework as yet another requirement on $[\varphi \rightarrow \psi]$. In particular, let's say that an agent believes $\varphi \rightarrow \psi$ just in case she would be committed to ψ if she learned φ. Now let's interpret belief in terms of our commitment relation $\models$, and let's interpret learning using $[\cdot]$. Putting these together, we reach the following requirement:

Definition 4.18. $\rightarrow$ is Ramseyan just in case $s \models \varphi \rightarrow \psi$ iff $s[\varphi] \models \psi$.

According to this version of the Ramsey test, a state supports $\varphi \rightarrow \psi$ just in case it would support ψ if it was updated with φ.

In a truth conditional setting, the requirement above is quite strong. In fact, in a truth conditional setting the material conditional is the unique Ramseyan operator.[21] To see why, note that the above formulation of the Ramsey Test is equivalent to the Deduction Theorem, given our definition of entailment. It guarantees that ψ follows from any premise combined with φ just in case that premise itself implies $\varphi \rightarrow \psi$.

By contrast, in a dynamic setting there are a variety of Ramseyan operators. For example, consider again the test operator from Gillies 2004. This operator is Ramseyan, because it does not allow an agent to believe $\varphi \rightarrow \psi$ if her current information does not support ψ when updated

[21] See Gillies 2009 for discussion.

with φ. However, the test operator says more than just this. It also says that whenever $s[\varphi]$ does not support ψ, s is actually committed to the negation of $\varphi \to \psi$.

Now consider our testlike operators. It turns out that when we focus on selective operations, the testlike operators precisely characterize the Ramsey Test.

Fact 4.8. If $\to$ is selective, then $\to$ is testlike iff $\to$ is Ramseyan.

We saw earlier that the selective operators are just the conditional assertion operators. This means that once one accepts the theory of conditional assertion, satisfaction of the Ramsey Test is equivalent to accepting a testlike theory of the conditional. On the other hand, we already saw before that the Ramsey Test does not guarantee conditional assertion operatorhood, since the test conditional in Gillies 2004 satisfies the former but not the latter. Likewise, there are many conditional assertion operators that are not Ramseyan, since there are selective operators that are not testlike.[22]

We've now introduced a preferred family of conditional assertion operators, and characterized it in terms of the Ramsey Test. It turns out that testlike operators have another independently plausible structural property: they are idempotent, so that $s[\varphi \to \psi]$ always supports $\varphi \to \psi$. Some dynamic meanings lack this property, while others possess it.[23]

Fact 4.9. If $\to$ is testlike, then $\to$ is idempotent.

Conditional assertion operators are precisely those that are selective. Within this family of operators, we have now located two desirable properties: satisfaction of the Ramsey Test, and idempotence. Possession of the former property is equivalent to being testlike. Each of these properties in turn guarantees idempotence.

We started this section with a negative result: the quintessential example of a dynamic meaning, the test, cannot be a conditional assertion operator. To model conditional assertion in dynamic semantics, we then developed a new kind of dynamic meaning. To start, we introduced the notion of a selective operator, proving this was equivalent to conditional assertion. Then we

[22] The discussion above is not the last word on selective conditionals. There are many more interesting constraints on selection to consider. For example, one constraint worth considering is to make f a function of $s[\varphi][\psi]$, rather than of s, $[\varphi]$, and $[\psi]$ individually.

[23] Say $\to$ is idempotent iff $\varphi \to \psi$ is idempotent for any φ and ψ.

developed the notion of a testlike operator. This turned out to be a special case of a selective operator, and was genuinely nonpropositional. We saw that testlike operators therefore occupy the intersection of two plausible constraints: conditional assertion and the Ramsey Test. In this way, the testlike operators retain one of the signature benefits of the test semantics for conditionals.

The dynamic and the truth conditional meanings we explored so far have a common consequence. They agree that performing a conditional assertion of $\varphi \rightarrow \psi$ has the potential to change which worlds we regard as possible. Otherwise, we would have that for every s, $s[\varphi \rightarrow \psi] = s$. But in this case, we would have no model at all of any effect of conditional assertion on context.

One might think, however, that the conditional assertion theory should guarantee that by itself, uttering a conditional does not change which worlds are possible in the context. Rather, updating the live possibilities of a state should also require a judgment about the antecedent. To have any chance of implementing such a theory, we need to let contexts contain more information than simply a set of possible worlds. Otherwise, conditional assertion would have no effect whatsoever on context. In the next section, we extend our representation of context to model this idea.

4.6 The future of conditional assertion

The crucial idea we need is that any context contains a set of possible extensions of the current live possibilities. Even though our context currently tells us we are in one of some set of possible worlds, not every subset of that state is a candidate for future information we might accept. Only some subsets are admissible.

To model this idea, we enrich our representation of context. We let a context c contain not only a set of worlds s, but also (building on Veltman 1985) a set σ of subsets of s. Intuitively, we can interpret σ as the set of possible extensions of the live possibilities s present in the context.

We need only one constraint on σ. We must suppose that whenever s' and s'' are possible extensions of s, so is their union. Intuitively: if we could learn one of two things, then we could also learn their disjunction. This constraint guarantees that for any descriptive claim φ, there is a unique weakest state that would result from learning φ in context c.

Definition 4.19. Let a context c be a pair $\langle s, \sigma \rangle$, where s is a set of possible worlds, and σ is a set of subsets of s closed under $\cup$. Let s_c and σ_c denote the s and σ components of c. Let a context c

be absurd ($\bot$) iff either $s = \emptyset$ or $\sigma = \emptyset$.

These higher order contexts can now figure in the meaning of a sentence. In particular, we can introduce a third update function $[\cdot]_3$, which assigns each sentence a function from $\langle s, \sigma \rangle$ pairs to $\langle s, \sigma \rangle$ pairs.

Now that we have extra structure in our contexts, we can model the intuitive idea from above, that conditional assertion does not directly change the live possibilities of a state. We can now say that $\rightarrow$ is a strong conditional assertion operator just in case it is a conditional assertion operator that does not affect s_c:

Definition 4.20. $\rightarrow$ is a strong conditional assertion operator iff $\rightarrow$ is a conditional assertion operator and for every context c: $s_{c[\varphi \rightarrow \psi]} = s_c$.

Now our task is to incorporate a theory of ordinary and conditional assertion into this model of context. Even in the case of ordinary assertion, we need a new theory. Now that contexts contain two components, s and σ, we need our updates to be sensitive to both. Ordinary update cannot simply narrow down the worlds in s: in that case, σ would play no role in constraining the future evolution of the live possibilities.

Instead, we let an assertion of φ update the live possibilities by selecting from σ the weakest possible evolution of the context that entails φ.[24] To be well behaved, we also need an ordinary assertion of φ to update σ, by filtering out any evolution of the context that does not imply φ. So suppose again that $[\![\cdot]\!]$ assigns every sentence in $\mathscr{L}^0$ a set of possible worlds. Then:

Definition 4.21. Where $\varphi \in \mathscr{L}^0$:

1. $\sigma \sqcap \varphi = \{s \in \sigma \mid s \subseteq [\![\varphi]\!]\}$

2. $c[\varphi] = \langle \bigcup \sigma_{c[\varphi]}, \sigma_c \sqcap \varphi \rangle$

In this theory, updating with φ can fail even when the live possibilities contain a φ world. If no evolution of the state (no $s \in \sigma$) implies φ, then updating with φ is still absurd.

Now we can implement a theory of strong conditional assertion. We can let the conditional assertion of $\varphi \rightarrow \psi$ pave the way for a later assertion of φ. Conditional assertion guarantees that if

[24] This is the major difference between the theory below and that of Veltman 1985, which we would more naturally interpret as requiring an update with φ to eliminate any φ worlds from the live possibilities.

the antecedent is asserted later, the consequent will then be supported. More precisely, conditional assertions operate on σ. They remove any $s \in \sigma$ which implies φ, but does not imply ψ. That is, they eliminate any extension of the context where we are committed to the antecedent, but are not committed to the consequent.

Definition 4.22.

1. $\sigma \sqcap \varphi \to \psi = (\sigma - (\sigma \sqcap \varphi)) \cup (\sigma \sqcap \psi)$

2. $c[\varphi \to \psi] = \langle s_c, \sigma_c \sqcap \varphi \to \psi \rangle$

There are three types of extensions of c ($s \in \sigma_c$) within $\sigma_{c[\varphi \to \psi]}$. Some are committed to $\neg\varphi$, some are committed to φ, and some are agnostic with respect to φ. The extensions of c that are either committed to $\neg\varphi$ or agnostic with respect to φ are preserved in $\sigma_{c[\varphi \to \psi]}$. Of the extensions committed to φ, only those remain from σ_c that are also committed to ψ. Interestingly, this all allows that that σ_c contains some s with $\varphi \wedge \neg\psi$ worlds. Later, we will consider another strong conditional assertion operator that eliminates such extensions (relying on $\sigma \sqcap \neg\varphi$ where the above uses $\sigma - (\sigma \sqcap \varphi)$). But it's worth starting with this weaker operation, since even this will give us a strong conditional assertion operator.

Now that we have our definitions in place, let's prove that we have created a strong conditional assertion operator. First, we know that $c[\varphi \to \psi][\varphi]$ is always committed to ψ. This is tantamount to the validity of Modus Ponens. To see why it holds, first consider what $c[\varphi \to \psi]$ looks like. We know $c[\varphi \to \psi]$ affects σ_c, by removing any s from it where φ holds throughout, and $\neg\psi$ does not. Next, $[\varphi]$ narrows down $s_{c[\varphi \to \psi]}(= s_c)$ to the weakest s in $\sigma_{c[\varphi \to \psi]}$ which implies φ, and narrows $\sigma_{c[\varphi \to \psi]}$ down to φ implying states. All such states imply ψ. So $s_{c[\varphi \to \psi]}$ and $\sigma_{c[\varphi \to \psi]}$ are unchanged by ψ.

Now let's turn to the screening off condition, that $c[\varphi \to \psi][\neg\varphi]$ is the same state as $c[\neg\varphi]$. To see why this holds, consider that $\sigma_{c[\neg\varphi]}$ contains only s that imply $\neg\varphi$. But $[\varphi \to \psi]$ has no effect on such states in c. Rather, $\sigma_{c[\varphi \to \psi]}$ is just like c_σ except that it eliminates some s that imply φ and do not imply ψ. So now consider $c[\varphi \to \psi][\neg\varphi]$. This updates c_s by moving to the weakest s in $c[\varphi \to \psi]_\sigma$ that implies $\neg\varphi$. This s will be the same as the weakest s in c_σ that implies $\neg\varphi$, since $[\varphi \to \psi]$ does not update $\neg\varphi$-implying subsets of σ_c. For this reason, $c[\varphi \to \psi][\neg\varphi] = c[\neg\varphi]$.

Finally, it is clear that $\to$ is a strong conditional assertion operator, so that $s_{c[\varphi\to\psi]}$ is always identical to s_c. We constructed $[\varphi \to \psi]$ so that it only updates σ_c, and never updates the live possibilities.

We have now developed a strong conditional assertion operator. The crucial idea was to distinguish updating the live possibilities from updating the way in which those possibilities can evolve. We now show that this sort of separation leads to surprising downstream effects on the logic of conditionals.

Since conditional assertion only affects the possible evolutions of the context, it is possible for $c[\varphi \to \psi]$ to eliminate from σ_c any possible evolution containing some world (say, a $\varphi \wedge \neg\psi$ world), even though that world remains in s_c. Now consider a context c that contains some $\neg\psi$ worlds, and contains only states that are opinionated about φ. Imagine updating with both $\varphi \to \psi$ and $\neg\varphi \to \psi$. The resulting σ will contain two types of evolutions: those that entail φ and ψ, and those that entail $\neg\varphi$ and ψ. So every evolution of the state implies ψ. Nonetheless, the live possibilities of the resulting context still contain $\neg\psi$ worlds. So updating the state further with ψ will have an effect. (It will not affect σ, but will affect s). This leads to a failure of the following principle:

$$\varphi \to \psi; \neg\varphi \to \psi \not\models \psi$$

That is: even when ψ is a sure thing to be learned, updating with ψ can have an effect.

Here's another way to see the same thing. Let $\top$ be a claim true at every possible world, and now consider conditionals like $\top \to \varphi$. These conditionals, like any other, update σ without updating s. They remove from σ any evolution that does not imply φ. We can think of these as one way to introduce a *must* operator $\Box$ into our language, using the definition $(\top \to \varphi) \equiv \Box\varphi$.

Definition 4.23. $c[\Box\varphi] = \langle s_c, \sigma_c \sqcap \varphi \rangle$

But take a context where the live possibilities originally include some φ worlds. Updating with $\Box\varphi$ narrows down σ to those states that imply φ. But since this update leaves the live possibilities unchanged, a further update with φ still has an effect. So c is not committed to φ. This leads to a failure of the strength of *must*:

$$\Box\varphi \not\models \varphi$$

Indeed, Veltman 1985 develops a similar theory in order to achieve exactly this result. Our theory can thus explain the felt weakness or indirectness of epistemic necessity claims. For illustration, suppose that we are sitting in a room and a friend walks in with a wet umbrella. We can say:

(4) It must be raining.

But imagine instead that we are looking out the window and see the rain. Then we would instead say:

(5) It's raining.

So it appears that uttering φ requires more evidence than simply uttering $\Box\varphi$.[25]

Interestingly, however, the failure of $\top \to \varphi$ to imply φ does not lead to a failure of Modus Ponens. While $\top \to \varphi \not\models \varphi$, we do have that $\top \to \varphi; \top \models \varphi$. Once we have learned $\top \to \varphi$, updating with the tautology has an effect. It eliminates any live possibility that is not present in any extension of the state. Moreover, recall that our set of evolutions σ is closed under unions. This means that whenever we have updated c with $\Box\varphi$, there is a unique weakest state in σ that supports $\top$: the intersection of c_s and $[\![\varphi]\!]$.

This last result leads to further logical consequences. For again let $\top$ be a claim true at every possible world. We are no longer guaranteed that:

$$\varphi \models \top$$

After all, uttering a conditional can lead to a state where even uttering the tautology can have an effect. While $\top$ itself can sometimes change the context, our language also contains $\Box\top$, which is guaranteed to leave every state unchanged. This expression is implied by any claim whatsoever, and so perhaps better fulfills the functional role of tautologies. Alternatively, we might introduce a new notion of entailment, where $\Gamma \models \delta$ just in case Γ implies $\Box\delta$ on our earlier definition.

[25] For other perspectives on the same phenomenon, see: Karttunen 1972; Kratzer 1991; von Fintel and Gillies 2010; and Lassiter 2016.

In what follows, it will help to introduce a kind of assertion operator, !, that encodes the minimal effect of $\top$. This assertion operator removes from the live possibilities any world that has been eliminated from every extension of the context.

Definition 4.24. $c! = c[\top] = \langle \bigcup \sigma_c, \sigma_c \rangle$

! and $\Box$ allow us to move back and forth between modal and non-modal updates. In particular, we are guaranteed that whenever $[\varphi]$ is an ordinary assertion, $c[\Box\varphi]! = c[\varphi]$.

So far, we've introduced a theory of assertion and conditional assertion, and explored a few of its consequences. But we have not given a compositional semantics that can underwrite the definitions above. Since our definition of context is somewhat unfamiliar, it's worth pausing to actually provide a semantics that delivers the definitions above. This is worth doing, because it will further explicate the relationship between the theory above, and more traditional forms of update semantics.

For atomic sentences, we can keep basically the same clause for ordinary assertion defined above. Similarly, we can retain the usual dynamic semantics for conjunction, in terms of sequential update. But for negation, we need something a bit more sophisticated. We need to answer two questions. First, how does $\neg\varphi$ update σ_c? Second, how does $\neg\varphi$ update s_c? In each case, we will see that there is a weak notion of negation as denial that behaves well with modal claims, and a strong notion of negation that behaves well with non-modal claims. Ultimately, we will craft a single semantics for negation that incorporates both the weak and strong operators as components.

Let's start with the first question, about how $\neg\varphi$ affects σ_c. Here, there is a natural distinction to draw between weak and strong negation. First, consider a weak negation operator where updating c with $\neg\varphi$ removed from σ_c any extension that would survive updating with φ. This weak negation would leave behind extensions that are agnostic with respect to φ. For example, suppose $s = \{w, v\}$, where φ is true at w and false at v. This state would be eliminated by φ, and hence would be retained by updating with a weak $\neg\varphi$. By contrast, s would be removed by a strong negative update. This update would remove from σ_c both the extensions that are committed to φ, and the extensions that are agnostic about φ. To implement this view, we could construct, for any extension s in σ_c, a new context where the live possibilities are s, and σ is every extension of s from σ_c. Then $\neg\varphi$ could eliminate s from σ_c just in case this hypothetical context would be destroyed by updating with φ.

These two negations are analogous to similar operators in a three valued setting, where a strong negation operator says that its input is false, while a weak negation operator says that its input is false or indeterminate. These notions of falsehood and indeterminacy can be modeled in update semantics through the distinction between cases where updating with s makes the context absurd, and cases where it changes the context. Summarizing:

Definition 4.25.

1. $\sigma_{c[\neg^w\varphi]} = \sigma_c - \sigma_{c[\varphi]}$

2. $\sigma_{c[\neg^s\varphi]} = \{s \in \sigma_c \mid \langle s, \{s' \in \sigma_c \mid s' \subseteq s\}\rangle[\varphi] = \bot\}$

Strong negation is a good model of how negations behave over ordinary assertions. Negating an ordinary assertion should narrow down the context to those worlds where the original assertion is false. On the other hand, weak negation is a better model of how negations behave over modal assertions. Consider $\neg\Box\varphi$. It is natural to think that updating with $\neg\Box\varphi$ should remove exactly the extensions that consist entirely of φ worlds. But strong negation does more, eliminating extensions that contain some φ worlds and some $\neg\varphi$ worlds, so that $[\neg^s\Box\varphi] = [\Box\neg^s\varphi]$. Such s can have subsets in σ_c with only φ worlds, and so $\langle s, \{s' \in \sigma_c \mid s' \subseteq s\}\rangle[\Box\varphi]$ need not be empty.

Now that we've explored two options for modeling the effect of $\neg\varphi$ on σ_c, let's turn to s_c. Again we can imagine a weak and strong version of negation. The first choice, $\neg_A$, likens uttering $\neg\varphi$ to asserting any other sentence, by letting $s_{c[\neg\varphi]}$ be the weakest evolution in $\sigma_{c[\neg\varphi]}$. The second option, $\neg_{CA}$, instead says that $\neg\varphi$ has no effect on s_c.

Definition 4.26.

1. $s_{c[\neg_A\varphi]} = \bigcup \sigma_{c[\neg_A\varphi]}$

2. $s_{c[\neg_{CA}\varphi]} = s_c$

Given the other operators we have already introduced, these two negations are interdefinable: $[\neg_{CA}\varphi] = [\Box\neg_A\varphi]$, and $c[\neg_A\varphi] = c[\neg_{CA}\varphi]!$, for any c.

Above, we said that ordinary assertions affect s_c, while modal and conditional assertions only σ_c. Correspondingly, $\neg_A$ is a good model of negations that outscope ordinary assertions, while

$\neg_{CA}$ is a good model of negations that outscope modals and conditionals. For example, suppose that α and β are atomic sentences that partition logical space. Then $[\neg_A \alpha] = [\beta]$. By contrast, $\neg_{CA}\alpha$ is modal, identical in meaning to $\Box\beta$. Similarly, we get weird results from embedding $\Box$ under $\neg_A$. For example, we do not have that $[\neg_A \neg_A \Box\varphi] = [\Box\varphi]$. Rather, double negation forces its input to update the live possibilities, guaranteeing that $[\neg_A \neg_A \Box\varphi] = [\Box\varphi]!$. This leads to downstream trouble in defining possibility modals in terms of $\neg_A$ and $\Box$.

Above, we've introduced two dimensions of variation for negation. In each case, the stronger kind of negation was right for ordinary assertion, while the weaker negation was right for modal assertion. On first glance, this seems to suggest that we need two different negation operators in our language. Fortunately, however, this is actually not required at all. Our meanings above are rich enough to allow us to semantically distinguish any conditional or modal assertion from any ordinary assertion. Without looking at the syntactic structure of a sentence, we can tell from its meaning alone whether it is an assertion or a conditional assertion. The former always narrow down the live possibilities of some states, while the latter do not.

Definition 4.27. $[\varphi]$ is an assertion iff $\exists c : s_c \neq s_{c[\varphi]}$.

Since being an assertion rather than a conditional assertion is a property of meanings rather than of syntactic forms, operators like negation can access this property of their inputs. So we can let negation operate differently depending on whether its input is an assertion or a conditional assertion. We can define $\neg\varphi$ piecewise, so that it behaves like $\neg^s_A\varphi$ when φ is an assertion, and like $\neg^w_{CA}\varphi$ when φ is a conditional assertion.

With atoms, negation, and conjunction out of the way, let's turn to conditionals and modals. First, it's worth restating our semantics for conditionals and necessity modals to be a bit more general. We can replace the use of $\sqcap$ above with direct appeal to $[\cdot]$. Second, we can introduce a slight modification of our semantics for the conditional above, for the sake of further exploring the landscape of strong conditional assertion operators. Above, conditional assertions preserved extensions of the context that were agnostic about the antecedent, even when such contexts contained worlds where the antecedent was true and the consequent was false. This update is definable in terms of weak negation: $\sigma_{c[\varphi\to\psi]} = \sigma_{c[\neg^w\varphi]} \cup \sigma_{c[\psi]}$. But it will be of interest in what follows to instead remove any evolutions that contain worlds where the antecedent is true and the consequent is false. To do so, we can replace the weak negation above with a strong negation: $\sigma_{c[\varphi\to\psi]} = \sigma_{c[\neg^s\varphi]} \cup \sigma_{c[\psi]}$. This update contains two types of extensions: those that consist of only worlds where φ is false, and those that consist of only ψ worlds. This means that $c[\varphi \to \psi]!$ is guaranteed to contain no $\varphi \wedge \neg\psi$ possibilities. While somewhat stronger than our earlier semantics for the conditional, this is still a strong conditional assertion operator.

Next, it's worth introducing a possibility modal into our framework. One natural choice is to let $\Diamond\varphi$ remove from σ any evolution that rules out every φ world. Equivalently, $\Diamond\varphi$ removes from σ exactly the extensions that survive update with $\neg\varphi$. Also equivalently, $[\Diamond\varphi] = [\Box\neg^w_A\neg\varphi]$ since $\sigma_{c[\neg\varphi]}$ contains the evolutions that rule out every φ world, and hence $\neg^w_A\neg\varphi$ removes exactly those evolutions from σ, while $\Box$ quarantines this effect from the live possibilities. Also equivalently, $\Diamond$ is the dual of $\Box$, because $[\neg\Box\neg\varphi] = [\Diamond\varphi]$ and $[\neg\Diamond\neg\varphi] = [\Box\varphi]$. This last and most important equivalence would have failed if we relied only on $\neg^s_A$ rather than our pairwise defined $\neg$. Interestingly, $[\neg^s_A\Diamond\neg^s_A\varphi]$ is identical to not $[\Box\varphi]$, but rather $[\Box\varphi]!$, and hence $[\varphi]$. That definition would thus distinguish (4) and (6).

(6) It can't not be raining.

Summing up, we have now defended the following semantics for $[\cdot]_3$:

Definition 4.28.

1. $c[\alpha] = \langle \bigcup \sigma_{c[\alpha]}, \{s \in \sigma_c \mid \forall w \in s : w(\alpha) = 1\}\rangle$

2. $c[\neg\varphi] = \begin{cases} c[\neg^s_A\varphi] & \text{if } [\varphi] \text{ is an assertion} \\ c[\neg^w_{CA}\varphi] & \text{otherwise} \end{cases}$

3. $c[\varphi \wedge \psi] = c[\varphi][\psi]$

4. $c[\varphi \to \psi] = \langle s_c, \sigma_{c[\neg\varphi]} \cup \sigma_{c[\psi]}\rangle$

5. $c[\Box\varphi] = \langle s_c, \sigma_{c[\varphi]}\rangle$

6. $c[\Diamond\varphi] = \langle s_c, \sigma_c - \sigma_{c[\neg\varphi]}\rangle$

The semantics above required making several important decisions, especially involving negation. There are several reasons why we made the particular decisions that we did. First, this semantics produces results that coincide almost exactly with our earlier definitions of assertion and conditional assertion. When φ does not contain $\rightarrow$, $\Box$, or $\Diamond$, updating c with φ narrows σ down to exactly the extensions where φ is classically true, and narrows down the live possibilities to the weakest such extension. When φ is a conditional or modal, only σ is affected by updating; the live possibilities are unchanged. The only difference concerns the conditional, where we have now explored a stronger conditional assertion operator, which eliminates any $\varphi \wedge \neg\psi$ world from any extension of the context.

We focused on this stronger conditional assertion operator for a reason. This operator places the semantics above in a rich structural correspondence with a more traditional form of dynamic semantics. Consider the following update function on sets of worlds, from Veltman 1996 and Gillies 2004:

Definition 4.29.

1. $s[\alpha]_2 = \{w \in s \mid w(\alpha) = 1\}$

2. $s[\neg\varphi]_2 = s - s[\varphi]_2$

3. $s[\varphi \wedge \psi]_2 = s[\varphi]_2[\psi]_2$

4. $s[\varphi \rightarrow \psi]_2 = \{w \in s \mid s[\varphi]_2 \models \psi\}$

5. $s[\Box\varphi]_2 = \{w \in s \mid s \models \varphi\}$

6. $s[\Diamond\varphi]_2 = \{w \in s \mid s[\varphi]_2 \neq \emptyset\}$

This update function includes the test conditional from above, along with a treatment of $\Box\varphi$ and $\Diamond\varphi$ as tests that the context supports or is consistent with φ.

It turns out that there is a structural relationship between this traditional update function, and the one developed in this section. Any update $[\varphi]_2$ on sets of worlds can be transformed into an update $[\varphi]_2^{\uparrow}$ on our current contexts, that filters out any extensions of the context that do not support the original update.[26]

---

[26]See Willer 2013 for a simpler application of this filter, which represents contexts as merely containing the structure of σ rather than $\langle s, \sigma \rangle$.

Definition 4.30. $c[\varphi]_2^{\uparrow} = \langle s_c, \{s \in \sigma_c \mid s[\varphi]_2 = s\}\rangle$

We can use $\uparrow$ and our earlier assertion operator ! to state the simple relationship between the two semantics above. First, for any conditional or modal assertion, our current update function is simply the lift of the traditional update function. Second, for any ordinary assertion, our current update function is the lift of the traditional update function, enriched with our above assertion operator.

Fact 4.10.

1. If φ does not contain $\rightarrow$, $\Box$, or $\Diamond$, then for every c, $c[\varphi]_3 = c[\varphi]_2^{\uparrow}!$.

2. If φ contains $\rightarrow$, $\Box$, or $\Diamond$, then $[\varphi]_3 = [\varphi]_2^{\uparrow}$.

One consequence of this last result is that $[\varphi \rightarrow \psi]_3 = [\varphi \rightarrow \psi]_2^{\uparrow}$. This is remarkable, because we proved earlier that the test semantics for $\rightarrow$ is not a conditional assertion operator. Indeed, no test can be one. Nonetheless, our current semantics for the conditional is a conditional assertion operator. This means that our lift operation $[\cdot]_2^{\uparrow}$ is a procedure for taking a traditional dynamic meaning, and transforming it into a conditional assertion operator.

4.7 Conclusion

In this chapter we explored the consequences of an intuitive picture of what it is to utter a conditional. Along the way, we've offered a new methodology for investigating the meaning of the conditional. Most work on conditionals proceeds by starting with data concerning the truth conditions or valid inferences of the conditional, and goes on to design a semantics that predicts this data. By contrast, we started with an intuitive idea about what it is to learn a conditional: crucially, that learning the antecedent is false screens off learning the conditional. We went on to characterize several families of meanings for the conditional that are consistent with that theory of learning.

Chapter 5

Results

This chapter provides proofs for the main facts in the previous chapters.

5.1 Generalized update semantics

Definition 2.25.

1. Let a *GUS* frame $\mathcal{F}$ be a tuple $\langle W, I, R_{(\cdot)} \rangle$, where:
   (a) W is a set of possible worlds.
   (b) The set of information states I is $\mathcal{P}(W)$.
   (c) An information sensitive accessibility relation $R_{(\cdot)}$ relates a world w in W and an s in I to a world v in W.
2. Let a *GUS* model $\mathcal{M}$ be a pair $\langle \mathcal{F}, V \rangle$ of a *GUS* frame and a valuation function V, where V assigns a truth value in $\{0, 1\}$ to every atomic sentence.

Definition 2.26. Let $\mathcal{M} = \langle F, V \rangle$ be any GUS model. An interpretation function $[\cdot]_{\mathcal{M}}$ assigns each sentence in $\mathcal{L}$ a context change potential, as follows.

1. $s[\alpha] = \{w \in s \mid w \in V(\alpha)\}$
2. $s[\neg\varphi] = s - s[\varphi]$
3. $s[\varphi \wedge \psi] = s[\varphi][\psi]$
4. $s[\Diamond\varphi] = \big\{w \in s \mid \{v \mid wR_s v\}[\varphi] \neq \emptyset\big\}$
5. $s[\Box\varphi] = \big\{w \in s \mid \{v \mid wR_s v\}[\varphi] = \{v \mid wR_s v\}\big\}$

(In what follows I will omit relativization to $\mathcal{M}$ when possible.)

Definition 2.27.

1. $s \models_{\mathcal{M}} \varphi$ iff $s \in I_{\mathcal{M}}$ and $s[\varphi]_{\mathcal{M}} = s$.
2. $\Gamma \models_{\mathcal{F}} \delta$ iff for any *GUS* model $\mathcal{M}$ containing $\mathcal{F}$: for any $s \in I_{\mathcal{M}}$, if $s \models_{\mathcal{M}} \gamma$ for every $\gamma \in \Gamma$, then $s \models_{\mathcal{M}} \delta$.

Definition 2.12.

1. The dynamic lift of R is $R^{\uparrow}_{(\cdot)}$, where:

   $\forall s \forall w \forall v \, [wR^{\uparrow}_{s} v \text{ iff } (wRv \,\&\, v \in s)]$.
2. $R'_{(\cdot)}$ is factorizable into R iff $R'_{(\cdot)} = R^{\uparrow}_{(\cdot)}$.
3. $R_{(\cdot)}$ is factorizable iff there is some R that $R_{(\cdot)}$ is factorizable into.

Where $[\![\Diamond_R]\!] = \lambda p.\{w \mid \exists v : wRv \,\&\, v \in p\}$, and $[\![\Box_R]\!] = \lambda p.\{w \mid \forall v : wRv \supset v \in p\}$:

Fact 2.1. When (i) φ is non-modal, and (ii) $R_{(\cdot)}$ is factorizable into R:

1. $s[\Diamond\varphi] = s \cap [\![\Diamond_R]\!](s \cap [\![\varphi]\!])$
2. $s[\Box\varphi] = s \cap [\![\Box_R]\!]((W - s) \cup [\![\varphi]\!])$

Proof. $s[\Diamond\varphi] = \{w \in s \mid \{v \mid wR_s v\} \cap [\![\varphi]\!] \neq \emptyset\} = \{w \in s \mid (\{v \mid wRv\} \cap s) \cap [\![\varphi]\!] \neq \emptyset\} = \{w \in s \mid \exists v : wRv \,\&\, v \in s \,\&\, v \in [\![\varphi]\!]\} = s \cap \{w \mid \exists v : wRv \,\&\, v \in (s \cap [\![\varphi]\!])\} = s \cap [\![\Diamond_R]\!](s \cap [\![\varphi]\!])$. $s[\Box\varphi] = \{w \in s \mid \{v \mid wR_s v\} \subseteq [\![\varphi]\!]\} = \{w \in s \mid (\{v \mid wRv\} \cap s) \subseteq [\![\varphi]\!]\} = \{w \in s \mid \forall v : wRv \,\&\, v \in s \supset v \in [\![\varphi]\!]\} = s \cap [\![\Box_R]\!](s \supset [\![\varphi]\!])$. □

Definition 2.13.

1. $R_{(\cdot)}$ is information insensitive just in case for any w, s, s', $\{v \mid wR_s v\} = \{v \mid wR_{s'} v\}$.
2. When $R_{(\cdot)}$ is information insensitive, say that $wR_{\downarrow}v$ iff for every s, $wR_s v$.

Fact 2.2. If $R_{(\cdot)}$ is information insensitive, $s[\Diamond\varphi] = s \cap \{w \mid \exists v : wR_{\downarrow}v \,\&\, v \in [\![\varphi]\!]\}$.

Proof. $s[\Diamond\varphi] = \{w \in s \mid \{v \mid wR_s v\}[\varphi] \neq \emptyset\} = \{w \in s \mid \{v \mid wR_{\downarrow} v\}[\varphi] \neq \emptyset = s \cap \{w \mid \exists v : wR_{\downarrow} v \;\&\; v \in [\![\varphi]\!]\}$. □

Definition 2.14. φ is distributive iff for any state s, $s[\varphi] = \bigcup_{w \in s} \{w\}[\varphi]$.

Definition 2.15. $R_{(\cdot)}$ and $R'_{(\cdot)}$ are semantically equivalent iff for any w, v, and s, if $w \in s$ and $w \in s'$, then $wR_s v$ iff $wR'_{s'} v$.

Fact 2.3.

1. If $R_{(\cdot)}$ is information insensitive, then $\Diamond\varphi$ is distributive for every φ.

2. If $\Diamond\varphi$ is distributive for every φ, then $R_{(\cdot)}$ is semantically equivalent to an information insensitive $R'_{(\cdot)}$.

Proof. For claim 1, suppose $R_{(\cdot)}$ is information insensitive. Then $\Diamond\varphi$ is distributive since $s[\Diamond\varphi]$ simply narrows down s to the worlds w where $\exists v : wR_{\downarrow} v \;\&\; v \in [\![\varphi]\!]$. For claim 2, Suppose that $\Diamond\varphi$ is distributive for any φ. Now suppose that $w \in s$ and $w \in s'$. We will show that $wR_s v$ iff $wR_{s'} v$. This will suffice to show that $R_{(\cdot)}$ is semantically equivalent to an information insensitive relation, since this guarantees that $R_{(\cdot)}$ is information insensitive for any world w inside any information state s. To show this, suppose that $wR_s v$, and let φ be a claim true at v uniquely. By the semantics for $\Diamond$, we know that $w \in s[\Diamond\varphi]$. By distributivity, this implies that $\{w\} \models \Diamond\varphi$. By a second application of distributivity, this implies that $w \in s'[\Diamond\varphi]$, and so $wR_{s'} v$. □

Definition 2.16. $R_{(\cdot)}$ is strongly world insensitive just in case for any w, s $\{v \mid wR_s v\} = s$.

Fact 2.4. $R_{(\cdot)}$ is strongly world insensitive iff for every state s, $s[\Diamond\varphi] = \{w \in s \mid s[\varphi] \neq \emptyset\}$.

Proof. When $R_{(\cdot)}$ is strongly world insensitive, $\{v \mid wR_s v\} = s$. So $s[\Diamond\varphi] = \{w \in s \mid \{v \mid wR_s v\}[\varphi] \neq \emptyset\} = \{w \in s \mid s[\varphi] \neq \emptyset\}$. Now suppose $R_{(\cdot)}$ is not strongly world insensitive. Then there is some w, s where $\{v \mid wR_s v\} \neq s$. In this case, there is some φ where $s[\Diamond\varphi] \neq \{w \in s \mid s[\varphi] \neq \emptyset\}$, since φ could hold at a world in s without holding at $\{v \mid wR_s v\}$, or vice versa. □

Fact 2.5. $(W \times W)^{\uparrow}_{(\cdot)}$ is the unique strongly world insensitive relation.

Proof. $\{v \mid \langle w, v\rangle \in (W \times W)^{\uparrow}_s\} = \{v \mid v \in s\} = s$. □

In what follows, let φ be descriptive, containing no occurrences of $\Diamond$ or $\Box$. When φ is descriptive, let $[\![\varphi]\!]_{\mathcal{M}} = \{w \mid \{w\} \models_{\mathcal{M}} \varphi\}$.

Definition 2.28.

1. $R_{(\cdot)}$ is eliminative iff $\forall s \forall v \forall w \in s\,[wR_s v \supset v \in s]$.

2. $\mathcal{F} = \langle W, I, R_{(\cdot)} \rangle$ is eliminative iff $R_{(\cdot)}$ is eliminative.

Fact 2.6. $\mathcal{F}$ is eliminative iff $\varphi \models_{\mathcal{F}} \Box\varphi$.

Proof. For the left to right direction, consider any eliminative frame $\mathcal{F} = \langle W, I, R_{(\cdot)} \rangle$, and any model $\mathcal{M}$ based on it. To show that $\varphi \models_{\mathcal{F}} \Box\varphi$, we must take an arbitrary $s \in I$, suppose that $s \models \varphi$, and prove that $s \models \Box\varphi$, so that $s[\Box\varphi] = s$. To do so, we can take an arbitrary $w \in s$, and show that $w \in s[\Box\varphi]$. This requires showing that $\{v \mid wR_s v\} \models \varphi$. Since $\mathcal{F}$ is eliminative, we know that $\{v \mid wR_s v\} \subseteq s$. Since $s \models \varphi$, we know that $s \subseteq [\![\varphi]\!]$. So $\{v \mid wR_s v\} \subseteq [\![\varphi]\!]$, and so $\{v \mid wR_s v\} \models \varphi$. Since w was arbitrary in s, this implies that $s \models \Box\varphi$.

For the right to left direction, suppose for contraposition that we have some $\mathcal{F} = \langle W, I, R_{(\cdot)} \rangle$ that is not eliminative. Then there is some $s \in I$ and some $w \in s$ where $\exists v : wR_s v \;\&\; v \notin s$. We will construct a valuation $\mathcal{V}$ where for some α, we have that $s \models_{\mathcal{M}} \alpha$ but $s \not\models_{\mathcal{M}} \Box\alpha$. So let $\mathcal{V}(\alpha) = s$, and let $\mathcal{M} = \langle \mathcal{F}, V \rangle$ (see Figure 1). $\alpha \models_{\mathcal{F}} \Box\alpha$ only if whenever $s \models_{\mathcal{M}} \alpha$, we have that $s \models_{\mathcal{M}} \Box\alpha$. By hypothesis, we have that $s \models \alpha$, since $V(\alpha) = s$. Nonetheless, we don't have that $s \models \Box\alpha$. For consider our w above that can see a world v not in s. w is in s. But $\{v \mid wR_s v\} \not\models \alpha$, since $v \notin V(\alpha)$. So $w \notin s[\Box\alpha]$, and $s \neq s[\Box\alpha]$. $\Box$

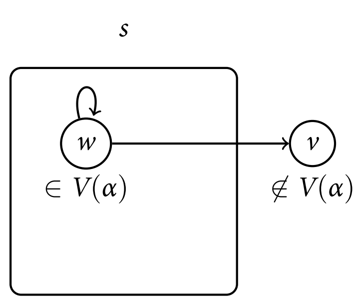

Figure 5.1: A failure of eliminativity

Fact 2.7. If $R_{(\cdot)}$ is eliminative, then $\Diamond\varphi \models \Diamond_t\varphi$.

Proof. Suppose $R_{(\cdot)}$ is eliminative, and $s \models \Diamond\varphi$. This implies that for any world $w \in s$, $\exists v$: $wR_s v$ & $v \in [\![\varphi]\!]$. Since $R_{(\cdot)}$ is eliminative, we know that every such v is in s. So $s[\varphi] \neq \emptyset$. So $s \models \Diamond_t\varphi$. □

Fact 2.8. If $R_{\cdot}$ is information insensitive and $\varphi \wedge \Diamond\neg\varphi \models \bot$ for all descriptive φ, then $wR_{\downarrow}v$ is isolated.

Proof. Suppose $R_{(\cdot)}$ is information insensitive and yet $R_{\downarrow}$ is not isolated. Then there is some $v \neq w$ where $wR_{\downarrow}v$. Now let φ be a claim true at w but false at v, and let $s = \{w\}$. $s \models \varphi \wedge \Diamond\neg\varphi$, since $s[\varphi] = s$, and $s[\Diamond\neg\varphi] = s$ (since $wR_{\{w\}}v$). □

Fact 2.9. If $R_{(\cdot)}$ is factorizable, then $R_{(\cdot)}$ is eliminative.

Definition 2.29.

1. $R_{(\cdot)}$ is globally reflexive iff $\forall s \forall w \in s\ \exists w' \in s : w'R_s w$.

2. $\mathcal{F} = \langle W, I, R_{(\cdot)}\rangle$ is globally reflexive iff $R_{(\cdot)}$ is globally reflexive.

Fact 2.10. $\mathcal{F}$ is globally reflexive iff $\Box\varphi \models_{\mathcal{F}} \varphi$.

Proof. For the left to right direction, consider any globally reflexive frame $\mathcal{F} = \langle W, I, R_{(\cdot)}\rangle$, and any model $\mathcal{M}$ based on it. To show that $\Box\varphi \models_{\mathcal{F}} \varphi$, we must take an arbitrary $s \in I$, suppose that $s \models \Box\varphi$, and prove that $s \models \varphi$, so that $s[\varphi] = s$. To do so, we can take an arbitrary $w \in s$, and show that $w \in [\![\varphi]\!]$. First, applying global reflexivity, we know that there is some w' in s where $w'R_s w$. Now since $s \models \Box\varphi$, we know that $s[\Box\varphi] = s$. This implies that $w' \in s[\Box\varphi]$, and so $\{v \mid w'R_s v\} \models \varphi$. Since $w'R_s w$, this means that $w \in [\![\varphi]\!]$.

For the right to left direction, suppose for contraposition that we have some $\mathcal{F} = \langle W, I, R_{(\cdot)}\rangle$ that is not reflexive. Then there is some $s \in I_{\mathcal{F}}$ and some $w \in s$ where $\neg\exists w' \in s : w'R_s^{\mathcal{F}} w$. We will construct a valuation $\mathcal{V}$ where for some α, we have that $s \models_{\mathcal{F}} \Box\alpha$ but $s \not\models_{\mathcal{F}} \alpha$. So let $\mathcal{V}(\alpha) = \bigcup_{w \in s}\{v \mid wR_s v\}$, and let $\mathcal{M} = \langle \mathcal{F}, V\rangle$ (see Figure 2). $\Box\alpha \models_{\mathcal{F}} \alpha$ only if whenever $s \models_{\mathcal{M}} \Box\alpha$, we have that $s \models_{\mathcal{M}} \alpha$. By hypothesis, we have that $s \models \Box\alpha$. For let w^* be an arbitrary world in s. We must show that $w^* \in s[\Box\alpha]$, which in turn requires that $\{v \mid w^* R_s v\} \models \alpha$. This last requires that $\{v \mid w^* R_s v\} \subseteq V(\alpha)$, which is guaranteed. Nonetheless, we don't have that $s \models \alpha$. For $w \notin s[\alpha]$. After all, $w \notin V(\alpha)$, since $\neg \exists w' \in s : w' R_s w$. $\square$

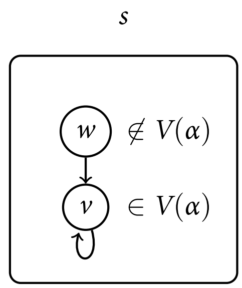

Figure 5.2: A failure of global reflexivity

Definition 2.30.

1. $R_{(\cdot)}$ is globally inverse reflexive iff $\forall s \forall w \in s\, \exists w' \in s : w R_s w'$.

2. $\mathcal{F} = \langle W, I, R_{(\cdot)} \rangle$ is globally inverse reflexive iff $R_{(\cdot)}$ is globally inverse reflexive.

Fact 2.11. $\mathcal{F}$ is globally inverse reflexive iff $\varphi \models_{\mathcal{F}} \Diamond \varphi$.

Proof. For the left to right direction, consider any globally inverse reflexive frame $\mathcal{F} = \langle W, I, R_{(\cdot)} \rangle$, and any model $\mathcal{M}$ based on it. To show that $\varphi \models_{\mathcal{F}} \Diamond \varphi$, we must take an arbitrary $s \in I$, suppose that $s \models \varphi$, and prove that $s \models \Diamond \varphi$, so that $s[\Diamond \varphi] = s$. To do so, we can take an arbitrary $w \in s$, and show that $\exists v : wRv \,\&\, v \in [\![\varphi]\!]$. Applying global inverse reflexivity, we know that there is some w' in s where $w R_s w'$. Now since $s \models \varphi$, we know that $s[\varphi] = s$. This implies that $w' \in s[\varphi]$, and so $w' \in [\![\varphi]\!]$.

For the right to left direction, suppose for contraposition that we have some $\mathcal{F} = \langle W, I, R_{(\cdot)} \rangle$ that is not globally inverse reflexive. Then there is some $s \in I_{\mathcal{F}}$ and some $w \in s$ where $\neg \exists w' \in s : w R_s^{\mathcal{F}} w'$. We will construct a valuation $\mathcal{V}$ where for some α, we have that $s \models_{\mathcal{F}} \alpha$ but $s \not\models_{\mathcal{F}} \Diamond \alpha$. So let $\mathcal{V}(\alpha) = s$, and let $\mathcal{M} = \langle \mathcal{F}, V \rangle$. $\alpha \models_{\mathcal{F}} \Diamond \alpha$ only if whenever $s \models_{\mathcal{M}} \alpha$, we have that $s \models_{\mathcal{M}} \Diamond \alpha$. By hypothesis, we have that $s \models \alpha$, since α is true at exactly the worlds in s. We must show that $w \notin s[\Diamond \alpha]$. Since there is no v where $w R_s v$ and $v \in s$, we know that this holds: the information at w relative to s excludes α. Since $w \in s$, this means that $s[\Diamond \alpha] \neq s$. $\square$

Fact 2.12. Suppose $R_{(\cdot)}$ is factorizable into R. Then $R_{(\cdot)}$ is globally reflexive iff R is reflexive.

Proof. Suppose $R_{(\cdot)}$ is factorizable into R. For the right to left direction, suppose that $w \in s$ and R is reflexive. It follows that $wR_s w$, since wRw and $w \in s$. So $\exists w' \in s : w'R_s w$. For the left to right direction, suppose that $R_{(\cdot)}$ is globally reflexive. Consider some arbitrary world w. We know that $\exists w' \in \{w\} : w'R_{\{w\}}w$; so $wR_{\{w\}}w$. So wRw. So R is reflexive. □

Definition 2.20. φ is idempotent iff for any s, $s[\varphi] \models \varphi$.

Fact 2.13. If $R_{(\cdot)}$ is factorizable and φ is descriptive, then $\Box\varphi$ is idempotent.

Proof. Suppose $R_{(\cdot)}$ is factorizable. Consider some arbitrary state s. To show that $s[\Box\varphi] \models \Box\varphi$, let's take some arbitrary $w \in s[\Box\varphi]$, and show that $\{v \mid wR_{s[\Box\varphi]}v\} \models \varphi$. This in turn requires that $(\{v \mid wRv\} \cap s[\Box\varphi]) \subseteq [\![\varphi]\!]$. Since $w \in s[\Box\varphi]$, we know that $(\{v \mid wRv\} \cap s) \subseteq [\![\varphi]\!]$. Since $s[\Box\varphi] \subseteq s$, it follows that $(\{v \mid wRv\} \cap s[\Box\varphi]) \subseteq [\![\varphi]\!]$. □

Definition 2.21. R is shift-reflexive iff for $\forall w \forall v[wRv \supset vRv]$.

Fact 2.14. Suppose $R_{(\cdot)}$ is factorizable into R. Then $\Diamond\varphi$ is idempotent for any descriptive φ and any model M built on $R_{(\cdot)}$ iff R is shift-reflexive.

Proof. Suppose $R_{(\cdot)}$ is factorizable into R. For the left to right direction, suppose that R is not shift-reflexive. Then there are some worlds w and w' where $w'Rw$ but it is not the case that wRw. We must show there is some state s and valuation $\mathcal{V}$ where $s[\Diamond\alpha] \not\models \Diamond\alpha$. So let $s = \{w, w'\}$, and let $\mathcal{V}(\alpha) = \{w\}$. $s[\Diamond\alpha] = \{w'\}$: first, $\{v \mid w'R_s v\}[\alpha] \neq \emptyset$, since $w'Rw$, $w \in s$, and $w \in \mathcal{V}(\alpha)$. Second, $\{v \mid wR_s v\}[\alpha] = \emptyset$, since w can't see any α world in s (w can't see itself). However, $s[\Diamond\alpha][\Diamond\alpha] = \emptyset \neq \{w'\} = s[\Diamond\alpha]$. For $w' \notin s[\Diamond\alpha][\Diamond\alpha]$, since there is no α world in $s[\Diamond\alpha]$.

For the right to left direction, suppose $R_{(\cdot)}$ is shift-reflexive. Now take some arbitrary s. If $s[\Diamond\varphi]$ is empty then the proof is trivial, so we will suppose that it is non-empty. To show that $s[\Diamond\varphi] \models \Diamond\varphi$, we can take an arbitrary $w \in s[\Diamond\varphi]$ and show that $\{v \mid wR_{s[\Diamond\varphi]}v\}[\varphi] \neq \emptyset$. Since $w \in s[\Diamond\varphi]$, we know that there is some v where wRv, $v \in s$, and $v \in [\![\varphi]\!]$. By shift-reflexivity, we know that vRv. So v is also in $s[\Diamond\varphi]$. But this means that $\{v \mid wRv\} \cap s[\Diamond\varphi] \cap [\![\varphi]\!] \neq \emptyset$. □

Definition 2.31.

1. $R_{(\cdot)}$ is globally transitive iff $\forall s \forall v \forall u \forall w \in s : [wR_s v \ \& \ vR_{\{v|wR_s v\}} u \supset \exists w' \in s : w'R_s u]$.

2. $\mathcal{F} = \langle W, I, R_{(\cdot)} \rangle$ is globally transitive iff $R_{(\cdot)}$ is globally transitive.

Fact 2.15. $\mathcal{F}$ is globally transitive iff $\Box\varphi \models_{\mathcal{F}} \Box\Box\varphi$.

Proof. For the left to right direction, consider any globally transitive frame $\mathcal{F} = \langle W, I, R_{(\cdot)} \rangle$, and any model $\mathcal{M}$ based on it. To show that $\Box\varphi \models_{\mathcal{F}} \Box\Box\varphi$, we must take an arbitrary $s \in I$, suppose that $s \models \Box\varphi$, and prove that $s \models \Box\Box\varphi$, so that $s[\Box\Box\varphi] = s$. To do so, we can take an arbitrary $w \in s$, and show that $w \in s[\Box\Box\varphi]$. This requires showing that $\{v \mid wR_s v\} \models \Box\varphi$. This in turn requires taking an arbitrary such v, and showing that $\{u \mid vR_{\{v|wR_s v\}} u\} \models \varphi$. To show this, we can take an arbitrary u in this last set, and show that $u \in [\![\varphi]\!]$. Since $wR_s v$ and $vR_{\{v|wR_s v\}} u$, we know that $\exists w' \in s : w'R_s u$. Since $s \models \Box\varphi$, we know that $\{v' \mid w'R_s v'\} \subseteq [\![\varphi]\!]$, and so since $w'R_s u$, we know that $u \in [\![\varphi]\!]$.

For the right to left direction, suppose for contraposition that we have some $\mathcal{F} = \langle W, I, R_{(\cdot)} \rangle$ that is not globally transitive. Then there is some $s \in I$ and some $w \in s$ where $wR_s v$ and $vR_{\{v|wR_s v\}} u$, but there is no $w' \in s$ where $w'R_s u$. We will construct a valuation $\mathcal{V}$ where for some α, we have that $s \models_{\mathcal{M}} \Box\alpha$ but $s \not\models_{\mathcal{M}} \Box\Box\alpha$. So let $\mathcal{V}(\alpha) = \bigcup_{w \in s} \{v \mid wR_s v\}$, and let $\mathcal{M} = \langle \mathcal{F}, V \rangle$ (see Figure 3). $\Box\alpha \models_{\mathcal{F}} \Box\Box\alpha$ only if whenever $s \models_{\mathcal{M}} \Box\alpha$, we have that $s \models_{\mathcal{M}} \Box\Box\alpha$. By hypothesis, we have that $s \models \Box\alpha$. For consider an arbitrary $w^* \in s$. We can show that $w^* \in s[\Box\alpha]$, by showing that $\{v \mid w^* R_s v\} \models \alpha$. Nonetheless, we don't have that $s \models \Box\Box\alpha$. For consider our w above that can see a world v relative to s that can in turn see a world u unseen by any w' in s. Crucially, $u \notin V(\alpha)$, since $V(\alpha)$ is the set of worlds seen by some w' in s. Now we will see that $w \notin s[\Box\Box\alpha]$. Here, the key is that $\{v \mid wR_s v\} \not\models \Box\alpha$, since $vR_{\{v|wR_s v\}} u$ but $u \notin V(\alpha)$. So $w \notin s[\Box\Box\alpha]$, and $s \neq s[\Box\Box\alpha]$. $\Box$

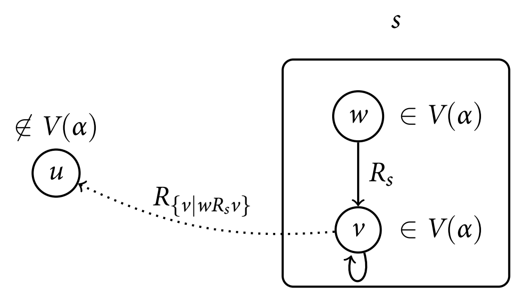

Figure 5.3: A failure of global transitivity

Fact 2.16. If $R_{(\cdot)}$ is eliminative, then $R_{(\cdot)}$ is globally transitive.

Fact 2.17. Suppose Upward Monotonicity is valid and Epistemic Contradictions are inconsistent. Then Positive Introspection is valid.

Definition 2.32.

1. $R_{(\cdot)}$ is globally euclidean iff $\forall s \forall v \forall w \in s : [wR_s v \supset \exists w' \in s : \forall v'[w'R_s v' \supset vR_{\{v|wR_s v\}} v']]$.

2. $\mathcal{F} = \langle W, I, R_{(\cdot)} \rangle$ is globally euclidean iff $R_{(\cdot)}$ is globally euclidean.

Fact 2.18. $\mathcal{F}$ is globally euclidean iff $\Diamond\varphi \models_{\mathcal{F}} \Box\Diamond\varphi$.

Proof. For the left to right direction, consider any globally euclidean frame $\mathcal{F} = \langle W, I, R_{(\cdot)} \rangle$, and any model $\mathcal{M}$ based on it. To show that $\Diamond\varphi \models_{\mathcal{F}} \Box\Diamond\varphi$, we must take an arbitrary $s \in I$, suppose that $s \models \Diamond\varphi$, and prove that $s \models \Box\Diamond\varphi$, so that $s[\Box\Diamond\varphi] = s$. To establish that $s \models \Box\Diamond\varphi$, we can take an arbitrary $w \in s$, and show that $\{v \mid wR_s v\} \models \Diamond\varphi$, which in turn requires taking an arbitrary v where $wR_s v$ and showing that there is some φ world u where $vR_{\{v|wR_s v\}} u$. By our assumption, we know that there is some $w' \in s$ where whenever $w'R_s v'$, we have that $vR_{\{v|wR_s v\}} v'$. Since $s \models \Diamond\varphi$, we know that there is some v' where $w'R_s v'$ and $v' \in [\![\varphi]\!]$. Therefore $vR_{\{v|wR_s v\}} v'$, and hence $\{v \mid wR_s v\} \models \Diamond\varphi$.

For the right to left direction, suppose for contraposition that we have some $\mathcal{F} = \langle W, I, R_{(\cdot)} \rangle$ that is not globally euclidean. Then there is some $s \in I$ and some $w \in s$ and v where $wR_s v$ but there is no $w' \in s$ where $\{v' \mid w'R_s v'\} \subseteq \{u \mid vR_{\{v|wR_s v\}} u\}$. We will construct a valuation $\mathcal{V}$ where for some α, we have that $s \models_{\mathcal{M}} \Diamond\alpha$ but $s \not\models_{\mathcal{M}} \Box\Diamond\alpha$. To construct $\mathcal{V}(\alpha)$, we will need the axiom of choice. In particular, for each $w \in s$, we must find $\{v' \mid wR_s v' \; \& \; \neg(vR_{\{v|wR_sv\}}v')\}$—the set of worlds accessible from w, but not from v. Now we must use a choice function to select one world from each of these sets. In particular, where f is a choice function, let $\mathcal{V}(\alpha) = \bigcup_{w \in s} f(\{v' \mid wR_s v' \; \& \; \neg(vR_{\{v|wR_sv\}}v')\})$. Then let $\mathcal{M} = \langle \mathcal{F}, V \rangle$ (see Figure 4). $\Diamond\alpha \models\!\!\models_{\mathcal{F}} \Box\Diamond\alpha$ only if whenever $s \models\!\!\models_{\mathcal{M}} \Diamond\alpha$, we have that $s \models\!\!\models_{\mathcal{M}} \Box\Diamond\alpha$. By hypothesis, we have that $s \models \Diamond\alpha$. For consider some arbitrary $w \in s$. We can show that $w \in s[\Diamond\alpha]_{\mathcal{M}}$, by showing that $\{v \mid wR_s v\}[\alpha]_{\mathcal{M}} \neq \emptyset$. After all, $\mathcal{V}(\alpha)$ contains a world from $\{v \mid wR_s v\}$: namely, $f(\{v' \mid wR_s v' \; \& \; vR_{\{v|wR_sv\}}v'\})$. Nonetheless, we don't have that $s \models\!\!\models_{\mathcal{M}} \Box\Diamond\alpha$. For consider our w and v from above. $w \notin s[\Box\Diamond\alpha]_{\mathcal{M}}$, since $\{v \mid wR_s v\} \not\models \Diamond\alpha$. To see why $\{v \mid wR_s v\} \not\models \Diamond\alpha$, consider an arbitrary u where $vR_{\{v|wR_sv\}}u$. $u \notin \mathcal{V}(\alpha)$, since α is a set of worlds accessible from some $w \in s$, but not accessible from v via $R_{\{v|wR_sv\}}$. So $w \notin s[\Box\Diamond\alpha]$, and $s \neq s[\Box\Diamond\alpha]$. $\Box$

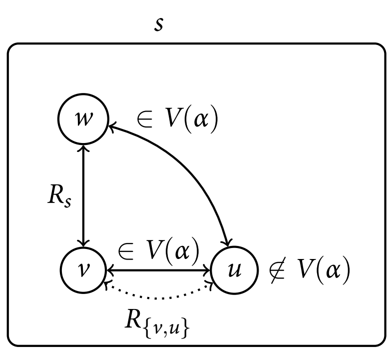

Figure 5.4: A failure of global euclideanness

Definition 2.33.

1. $R_{(\cdot)}$ is globally symmetric iff $\forall s \forall v \forall w \in s : [wR_s v \supset \exists w' \in s : vR_{\{v|wR_sv\}}w']$.

2. $\mathcal{F} = \langle W, I, R_{(\cdot)} \rangle$ is globally symmetric iff $R_{(\cdot)}$ is globally symmetric.

Fact 2.19. $\mathcal{F}$ is globally symmetric iff $\varphi \models\!\!\models_{\mathcal{F}} \Box\Diamond\varphi$.

Proof. For the left to right direction, consider any globally symmetric frame $\mathcal{F} = \langle W, I, R_{(\cdot)} \rangle$, and any model $\mathcal{M}$ based on it. To show that $\varphi \models\!\!\models_{\mathcal{F}} \Box\Diamond\varphi$, we must take an arbitrary $s \in I$, suppose that $s \models \varphi$, and prove that $s \models \Box\Diamond\varphi$, so that $s[\varphi] = s$. To do so, we can take an arbitrary $w \in s$, and show that $\{v \mid wR_s v\} \models \Diamond\varphi$. To show this, we can take an arbitrary $v \in \{v \mid wR_s v\}$ and show that there is some u where $vR_{\{v|wR_s v\}}u$ and $u \in [\![\varphi]\!]$. Applying global symmetry, we know that since $wR_s v$, there is some w' in s where $vR_{\{v|wR_s v\}}w'$. Since $s \models \varphi$, this means that $w' \in [\![\varphi]\!]$. So w' is our witnessing u. That is, $\{v \mid wR_s v\} \models \Diamond\varphi$ since every such v sees some world in s, where φ holds. So $s \models \Box\Diamond\varphi$.

For the right to left direction, suppose for contraposition that we have some $\mathcal{F} = \langle W, I, R_{(\cdot)}\rangle$ that is not globally symmetric. Then there is some $s \in I_{\mathcal{F}}$ and some $w \in s$ where $wR^{\mathcal{F}}_s v$ and yet there is no $w' \in s$ where $vR^{\mathcal{F}}_{\{v|wR_s v\}}w'$. We will construct a valuation $\mathcal{V}$ where for some α, we have that $s \models_{\mathcal{F}} \alpha$ but $s \not\models_{\mathcal{F}} \Box\Diamond\alpha$. So let $\mathcal{V}(\alpha) = s$, and let $\mathcal{M} = \langle \mathcal{F}, V\rangle$. $\alpha \models_{\mathcal{F}} \Box\Diamond\alpha$ only if whenever $s \models_{\mathcal{M}} \alpha$, we have that $s \models_{\mathcal{M}} \Box\Diamond\alpha$. By hypothesis, we have that $s \models \alpha$, since α is true at exactly the worlds in s. But we do not have $s \models \Box\Diamond\alpha$. w is in s but $w \notin s[\Box\Diamond\alpha]$. For $w \in s[\Box\Diamond\alpha]$ only if $\{v \mid wR_s v\} \models \Diamond\alpha$. But this last holds just in case every such v' can see an α world. But our witnessing v from above counterexamples this claim, since $wR_s v$ and yet there is no $w' \in s$ where $vR^{\mathcal{F}}_{\{v|wR_s v\}}w'$. Since v sees no world in s, v sees no world where α is true. So $\{v \mid wR_s v\} \not\models \Diamond\alpha$, and hence $s \not\models \Box\Diamond\alpha$. □

5.2 Free choice impossibility results

Definition 3.7. A set of possibility operators $\Diamond_1, \ldots, \Diamond_n$ with corresponding accessibility relations $R_1, \ldots, R_n$ is partitional just in case there exist accessibility relations R, R', R'' in $\{R_1, \ldots, R_n\}$ and worlds w, w', w'' such that:

1. $\{v \mid wRv\}, \{v' \mid w'R'v'\}$, and $\{v'' \mid w''R''v''\}$ are not empty.
2. $\{v \mid wRv\} = \{v' \mid w'R'v'\} \cup \{v'' \mid w''R''v''\}$
3. $\{v' \mid w'R'v'\} \cap \{v'' \mid w''R''v''\} = \emptyset$
4. $\exists\varphi, \psi \in \mathscr{L} : [\![\varphi]\!] = \{v' \mid w'R'v'\}\ \&\ [\![\psi]\!] = \{v'' \mid w''R''v''\}$.

Fact 3.1. If a set of possibility operators is partitional and all satisfy Free Choice, then disjunction is not propositional.

Proof. Let $\Diamond_A$, $\Diamond_B$, and $\Diamond_C$ be the witnessing operators. Let B and C denote $\{v' \mid w'R_B v'\}$, and $\{v'' \mid w''R_C v''\}$ respectively, where $A = B \cup C$. Suppose disjunction is propositional, with $*$ as the witnessing operator. Let φ and ψ be true at exactly B and C respectively. Then by the validity of Free Choice for $\Diamond_A$, $\Diamond_A(\varphi \vee \psi)$ is true at w just in case $\Diamond_A \varphi$ and $\Diamond_A \psi$ are true at w. $[\![\varphi \vee \psi]\!] = B * C$. So $B * C$ is consistent with $B \cup C$ just in case B is consistent with $B \cup C$ and C is consistent with $B \cup C$. However, we also know that $B * C$ is inconsistent with C if B is inconsistent with C, since $\Diamond_C(\varphi \vee \psi)$ is true at w'' only if $\Diamond_C \varphi$ is true at w''. Since B is inconsistent with C, it follows that $B * C$ is inconsistent with C. Similarly, we know that $B * C$ is inconsistent with B if C is inconsistent with B, since $\Diamond_B(\varphi \vee \psi)$ is true at w' only if $\Diamond_B \psi$ is true at w'. Since C is inconsistent with B, it follows that $B * C$ is inconsistent with B. But if $B * C$ is inconsistent with both B and with C, then $B * C$ is inconsistent with $B \cup C$. Contradiction. □

Fact 3.2. If any conditional possibility operator $\Diamond_i(\cdot)(\cdot)$ satisfies Conditional Free Choice, then $\vee$ is not propositional.

Proof. Let $\Diamond(\cdot)(\cdot)$ be the witnessing operator with accessibility relation R. Now take some arbitrary world w such that $\Diamond(\cdot)(\cdot)$ has a non-trivial accessibility relation at that world, such that there are at least two worlds v and u where wRv and wRu. Now I will assume that there are two propositions A and B that partition $\{v \mid wRv\}$, and are the semantic value of some sentences φ and ψ. Now let $\top$ be a claim true at every possible world. By Conditional Free Choice, $w \in [\![\Diamond(\top)(\varphi \vee \psi)]\!]$ just in case $w \in [\![\Diamond(\top)(\varphi)]\!]$ and $w \in [\![\Diamond(\top)(\psi)]\!]$. In this semantic framework, the last condition in turn requires that $\exists v : wRv \ \& \ v \in A * B$ iff $\exists v : wRv \ \& \ v \in A$ and $\exists v : wRv \ \& \ v \in B$. By the construction of R, we know that this condition is satisfied. But now consider the conditionals $\Diamond(\varphi)(\varphi \vee \psi)$ and $\Diamond(\psi)(\varphi \vee \psi)$. If $\Diamond$ satisfies Conditional Free Choice, then we know that $w \in [\![\Diamond(\varphi)(\varphi \vee \psi)]\!]$ only if $w \in [\![\Diamond(\varphi)(\psi)]\!]$. This in turn implies that $\exists v : wRv \ \& \ v \in A \ \& \ v \in A * B$ only if $\exists v : wRv \ \& \ v \in A \ \& \ v \in B$. But there is no such v in both A and B, since these propositions are inconsistent. So there is no accessible v where both A and $A * B$ hold. Likewise, $w \in [\![\Diamond(\psi)(\varphi \vee \psi)]\!]$ only if $w \in [\![\Diamond(\psi)(\varphi)]\!]$; which in turn implies that $\exists v : wRv \ \& \ v \in B \ \& \ v \in A * B$ only if $\exists v : wRv \ \& \ v \in A \ \& \ v \in B$. Since this last condition cannot obtain, we know that there is no accessible v where both B and $A * B$. But every accessible v world is either an A or a B world. So $A * B$ can't hold at any accessible world in the first place. □

Definition 3.10 (Credal Free Choice)**.** There is some operator $*$ such that for any rational Pr, A, B: $Pr(A * B) > 0$ iff $Pr(A) > 0$ and $Pr(B) > 0$.

Fact 3.3. If Credal Free Choice holds, then $Pr(A) = 0$ or $Pr(A) = 1$ for every rational Pr, A.

Proof. $Pr(A*(W-A)) = Pr(A*(W-A)) \times Pr(A*(W-A) \mid A) + Pr(W-A) \times Pr(A*(W-A) \mid W-A)$. $Pr(A \mid A) = 1$, which means $Pr(W-A \mid A) = 0$, which implies $Pr(A*(W-A) \mid A) = 0$ by Credal Free Choice. By similar reasoning, $Pr(A*(W-A) \mid W-A) = 0$. So $Pr(A*(W-A)) = 0$. So by Credal Free Choice again, $Pr(A) = 0$ or $Pr(W-A) = 0$. So $Pr(A)$ is 0 or 1. □

Fact 3.4. If Free Choice is valid, then $\varphi \vee \psi$ is not distributive for arbitrary atomic φ and ψ.

Proof. Consider some arbitrary atomic α and β. Let s be some context which contains some α worlds, some β worlds, and no worlds where both α and β are false. Since $s[\alpha] \neq \emptyset$ and $s[\beta] \neq \emptyset$, we know by Dynamic Free Choice that $s[\alpha \vee \beta] \neq \emptyset$. Now let $s_{\neg\alpha}$ be $\{w \in s \mid w(\alpha) = 0\}$, and let $s_{\neg\beta}$ be $\{w \in s \mid w(\beta) = 0\}$. By Dynamic Free Choice, $s_{\neg\alpha}[\alpha \vee \beta] = \emptyset$, since $s_{\neg\alpha}[\alpha] = \emptyset$. Similarly, $s_{\neg\beta}[\alpha \vee \beta] = \emptyset$. Now suppose $\alpha \vee \beta$ is distributive. Then for any $w \in s_{\neg\alpha}$, $\{w\}[\alpha \vee \beta] = \emptyset$. Similarly, for any $w \in s_{\neg\beta}$, $\{w\}[\alpha \vee \beta] = \emptyset$. But $s = s_{\neg\alpha} \cup s_{\neg\beta}$. So for any $w \in s$, $\{w\}[\alpha \vee \beta] = \emptyset$. So, supposing $\alpha \vee \beta$ is distributive, $s[\alpha \vee \beta] = \emptyset$. This is inconsistent; so $\alpha \vee \beta$ is not distributive. □

Fact 3.5. Suppose that entailment is transitive, that $\Diamond$ is upwards monotonic, that the T axiom is valid, and that the 4 axiom is valid. Then FC I is valid iff Disjunction-Possibility Link is valid.

Proof. By the T axiom, $\varphi \vee \psi$ implies $\Diamond(\varphi \vee \psi)$. By Free Choice, this last implies $\Diamond\varphi \wedge \Diamond\psi$. So by the transitivity of entailment $\varphi \vee \psi$ implies $\Diamond\varphi \wedge \Diamond\psi$. Now suppose $\varphi \vee \psi \models \Diamond\varphi$. Then $\Diamond(\varphi \vee \psi) \models \Diamond\Diamond\varphi$ by upwards monotonicity, which by the 4 axiom implies $\Diamond\varphi$. □

Definition 3.16. $s[\varphi \vee_t \psi] = s[\Diamond\varphi][\Diamond\psi]$

Fact 3.6. If $*$ validates Free Choice, then $*$ is at least as strong as $\vee_t$.

Proof. By the T axiom for $\Diamond$, $\varphi * \psi \models \Diamond(\varphi * \psi)$. Suppose FC 1 is valid for $*$; then $\Diamond(\varphi * \psi) \models \Diamond\varphi$ and $\Diamond(\varphi * \psi) \models \Diamond\psi$. This is sufficient to show that $\Diamond(\varphi * \psi) \models \varphi \vee_t \psi$, and hence that $\varphi * \psi \models \varphi \vee_t \psi$ □

Definition 3.18. $*$ is factorizable iff there is some distributive operator $\bowtie$ such that $[\varphi * \psi] = s[\varphi \bowtie \psi][\varphi \vee_t \psi]$.

Definition 3.19. $s[\varphi \vee_u \psi] = s[\varphi] \cup s[\psi]$

Definition 3.20. $s[\varphi \vee_{ut} \psi] = s[\varphi \vee_u \psi][\varphi \vee_t \psi]$.

Fact 3.7. if $*$ is factorizable and satisfies Free Choice, and if $\varphi \not\models \psi$ and $\psi \not\models \varphi$, then $\varphi \vee_{ut} \psi \models \varphi * \psi$.

Proof. Suppose that $s[\varphi \vee_u \psi][\varphi \vee_t \psi] = s$. We must show that $s \models \varphi * \psi$. We can suppose that s is not $\emptyset$, since the proof would be trivial in that case. Since $s[\varphi \vee_u \psi][\varphi \vee_t \psi] = s$, we know that (i) $s[\varphi] \neq \emptyset$ and $s[\psi] \neq \emptyset$, and (ii) s is made up exclusively of worlds w where either $\{w\} \models \varphi$ or $\{w\} \models \psi$. From (i), we know that the possibility tests required by $*$ are satisfied. To complete our proof, we can use (ii) to show that $s[\varphi \bowtie \psi] = s$.

Since $\bowtie$ is distributive, we know that $s[\varphi \bowtie \psi] = \{w \in s \mid \{w\} \models \varphi \bowtie \psi\}$. So to prove that $s[\varphi \bowtie \psi] = s$, we can take an arbitrary $w \in s$ and show that $\{w\} \models \varphi \bowtie \psi$. From (ii) we know that either $\{w\} \models \varphi$ or $\{w\} \models \psi$. We can show that in either case $\{w\} \models \varphi \bowtie \psi$.

So suppose the former. Either $\{w\} \models \psi$ or not. Suppose the former. Then since $\{w\} \models \varphi$ and $\{w\} \models \psi$, we have that $\{w\} \models \Diamond\varphi$ and $\{w\} \models \Diamond\psi$. So since $*$ satisfies Free Choice, we know that $\{w\} \models \Diamond(\varphi * \psi)$. Since $\{w\}$ is maximal, this means that $\{w\} \models \varphi * \psi$ and hence $\{w\} \models \varphi \bowtie \psi$.

But now suppose instead that $\{w\} \models \varphi$ and $\{w\} \not\models \psi$. Now let v be such that $\{v\} \models \psi$ and $\{v\} \not\models \varphi$. Here, we rely on the assumption that φ and ψ are not ordered by strength. First, this means that $\{w, v\} \models \Diamond\varphi$ and $\{w, v\} \models \Diamond\psi$. So by Free Choice $\{w, v\} \models \Diamond(\varphi * \psi)$. But this means that $\{w, v\}[\varphi \bowtie \psi] \neq \emptyset$. This in turn implies that $\{w, v\}[\varphi \bowtie \psi] = \{w, v\}$. For otherwise one of w or v would be removed, in which case one of $\Diamond\varphi$ or $\Diamond\psi$ would not be supported by $\{w, v\}[\varphi \bowtie \psi]$, and so $\{w, v\} \not\models \Diamond(\varphi * \psi)$. So since $\{w, v\} \models \varphi \bowtie \psi$ and $\bowtie$ is distributive, $\{w\} \models \varphi \bowtie \psi$. □

Definition 3.23. $*$ is quasi-conjunctive relative to R iff for any A and B, if A and B are expressible via R, then $A * B = A \cap B$.

Fact 3.8. If Wide Free Choice is valid for $\Diamond_i$ and disjunction is propositional, then disjunction is quasi-conjunctive relative to R_i.

Proof. Suppose Wide Free Choice is valid for $\Diamond_i$ and disjunction is propositional, expressing the operation $*$. Now suppose that A and B are expressible via R_i. That is, suppose $[\![\Diamond_i\varphi]\!] = A$ and suppose $[\![\Diamond_i\psi]\!] = B$, for some φ and ψ. By Wide Free Choice, $\Diamond_i\varphi \vee \Diamond_i\psi$ implies each of $\Diamond_i\varphi$ and

$\Diamond_i\psi$, and vice versa. But this means that $[\![\Diamond i\varphi \vee \Diamond i\psi]\!] = [\![\Diamond_i\varphi]\!] \cap [\![\Diamond_i\psi]\!]$. Since $[\![\Diamond i\varphi \vee \Diamond i\psi]\!] = A * B$, this means that $A * B = A \cap B$. □

Fact 3.9. Wide Free Choice is valid for $\Diamond_i$ and $\vee_t$ iff for any w, v, u: if wR_iv then uR_iv.

Proof. Suppose Wide Free Choice is valid for $\Diamond_i$ and $\vee_t$. Then whenever $s \models \Diamond\Diamond_i\varphi$ and $s \models \Diamond\Diamond_i\psi$, we have that $s \models \Diamond_i\varphi$ and $s \models \Diamond_i\psi$. This means that if there is some $w \in s$ where wR_iv and $v \in A$, we are guaranteed that for every $w \in s$ there is an accessible v world where A. This means that if wR_iv, then uR_iv for any u. Otherwise, s could contain just w and u, and then $\{v\}$ would be a counterexampling choice of A. Conversely, suppose that for any w, v, u: if wR_iv then uR_iv. Now suppose $s \models \Diamond\Diamond_i\varphi$ and $s \models \Diamond\Diamond_i\psi$. Then $\exists w \in s : \{v \mid wR_iv\}[\varphi] \neq \emptyset$. Now let u be an arbitrary world in s. By the constraint above, we know that $\{v' \mid uR_iv'\}[\varphi] \neq \emptyset$, since $\{v \mid wR_iv\} = \{v' \mid uR_iv'\}$. So $s \models \Diamond_i\varphi$. □

Fact 3.10. Free Choice, Dual Prohibition, Contraposition, and Transitivity imply Explosion.

Proof. By Free Choice, $\Diamond(\varphi \vee \psi) \models \Diamond\psi$. So by Contraposition $\neg\Diamond\psi \models \neg\Diamond(\varphi \vee \psi)$. By Dual Prohibition, $\neg\Diamond(\varphi \vee \psi) \models \neg\Diamond\varphi$. So by Transitivity, $\neg\Diamond\psi \models \neg\Diamond\varphi$ and hence by Contraposition again $\Diamond\varphi \models \Diamond\psi$. □

Fact 3.11. If Free Choice and Dual Prohibition are valid, then Explosion is valid.

Proof. Suppose $s \models \Diamond\varphi$. Suppose that s is non-empty (otherwise the proof is trivial). We can now show that $s \models \Diamond\psi$. First, $s[\varphi] \neq \emptyset$. Since Dual Prohibition is valid, this means that $s[\varphi \vee \psi] \neq \emptyset$. But since Free Choice is valid, this last implies that $s[\psi] \neq \emptyset$. So $s \models \Diamond\psi$. □

5.3 A theory of conditional assertion

Definition 4.3. $\rightarrow$ is a conditional assertion operator iff:

1. $s[\varphi \rightarrow \psi][\varphi] \models \psi$
2. $s[\varphi \rightarrow \psi][\neg\varphi] = s[\neg\varphi]$

Fact 4.1. Suppose $\models$, $[\cdot]$, and $\neg$ are well-behaved, φ and $\neg\psi$ are persistent, and $\neg\varphi$ is idempotent. Then if $\rightarrow$ is a conditional assertion operator, then:

1. $\varphi; \neg\psi \models \neg(\varphi \to \psi)$

2. $\neg\varphi \models \varphi \to \psi$.

Proof. Suppose for reductio that $\varphi; \neg\psi \not\models \neg(\varphi \to \psi)$, and suppose that $\to$ is a conditional assertion operator. Then there is some s where $s[\varphi][\neg\psi] \not\models \neg(\varphi \to \psi)$. So by the well-behavedness of negation, we know that $s[\varphi][\neg\psi][\varphi \to \psi]$ is not absurd. But $s[\varphi] \models \varphi$, and so by the persistence of φ, we know that $s[\varphi][\neg\psi][\varphi \to \psi] \models \varphi$. So $s[\varphi][\neg\psi][\varphi \to \psi][\varphi] = s[\varphi][\neg\psi][\varphi \to \psi]$. So by condition (i) of $\to$ being a conditional assertion operator we can infer that $s[\varphi][\neg\psi][\varphi \to \psi] \models \psi$. But $s[\varphi][\neg\psi] \models \neg\psi$, and so by the persistence of $\neg\psi$ $s[\varphi][\neg\psi][\varphi \to \psi] \models \neg\psi$. So $s[\varphi][\neg\psi][\varphi \to \psi]$ is absurd after all. Contradiction.

For the last condition, take some context s. Since $\neg\varphi$ is idempotent, we know that $s[\neg\varphi] \models \neg\varphi$. But by what we have proved above, we know that whenever $s \models \neg\varphi$, we have that $s \models \varphi \to \psi$. So $s[\neg\varphi] \models \varphi \to \psi$. □

Fact 4.2. If $\to$ is propositional and $\neg$ is well behaved, then $\to$ is a conditional assertion operator iff:

1. $[\![\varphi]\!] \cap [\![\neg\psi]\!] \cap [\![\varphi \to \psi]\!] = \emptyset$

2. $[\![\neg\varphi]\!] \subseteq [\![\varphi \to \psi]\!]$

Proof. If $\to$ is propositional, then the two requirements of conditional assertion operators are as follows. First, for any set of worlds s, $s \cap [\![\varphi \to \psi]\!] \cap [\![\varphi]\!]$ is a subset of $[\![\psi]\!]$. Second, for any set of worlds s, $s \cap [\![\varphi \to \psi]\!] \cap [\![\neg\varphi]\!]$ is equal to $s \cap [\![\neg\varphi]\!]$. Further, if negation is well behaved then $[\![\neg\varphi]\!] = W - [\![\varphi]\!]$. Given this assumption, the two requirements here are equivalent to those above. □

Fact 4.3. If $\to$ is a test operator, then $\to$ is not a conditional assertion operator.

Proof. Let c be a context $\{w, v\}$, where $w(\varphi) = 0$, $v(\varphi) = 1$, and $v(\psi) = 0$. If $\to$ is a test operator, then $c[\varphi \to \psi]$ is either s or $\emptyset$. Suppose the former. Then $c[\varphi \to \psi][\varphi] \not\models \psi$, so $\to$ is not a conditional assertion operator. Suppose the latter. Then $c[\varphi \to \psi][\neg\varphi] = \emptyset \neq \{w\} = c[\neg\varphi]$. So $\to$ is not a conditional assertion operator. □

Definition 4.15.

1. A generalized selection function f is a function from a context and two *ccp*s to a new context, where $f(s, [\varphi], [\psi]) \subseteq s[\varphi][\psi]$.

2. $\to$ is selective iff $s[\varphi \to \psi] = s[\neg\varphi] \cup f(s, [\varphi], [\psi])$.

Fact 4.4. $\to$ is a conditional assertion operator iff $\to$ is selective.

Proof. Suppose $\to$ is selective. Then $s[\varphi \to \psi][\varphi] \models \psi$, since the only φ worlds in $s[\varphi \to \psi]$ are ψ worlds. In addition, $s[\varphi \to \psi][\neg\varphi] = s[\neg\varphi]$ because $s[\varphi \to \psi]$ includes all of $s[\neg\varphi]$.

Now suppose $\to$ is a conditional assertion operator. It suffices to show that for any s, $s[\varphi \to \psi] \subseteq s[\neg\varphi] \cup s[\varphi][\psi]$. In that case, we can construct a selector f so that $f(s, [\varphi], [\psi]) = s[\varphi \to \psi] - s[\neg\varphi]$. So suppose $w \in s[\varphi \to \psi]$ but $w \notin s[\neg\varphi] \cup s[\varphi][\psi]$. Then $w \in s[\varphi][\neg\psi]$. But this means that $s[\varphi \to \psi][\varphi] \not\models \psi$. □

Definition 4.16.

1. f is testlike iff $f(s, [\varphi], [\psi]) = s[\varphi][\psi]$ whenever $s[\varphi] \models \psi$.

2. $\to$ is testlike iff $s[\varphi \to \psi] = \begin{cases} s & \text{if } s[\varphi] \models \psi \\ s[\neg\varphi] \cup f(s, [\varphi], [\psi]) & \text{otherwise.} \end{cases}$

Fact 4.5. $\to$ is testlike just in case $\to$ is selective and f is testlike.

Proof. Suppose $\to$ is selective and f is testlike. Now suppose $s[\varphi] \models \psi$. Then $f(s, [\varphi], [\psi]) = s[\varphi][\psi]$, and so $s[\varphi \to \psi] = s$. □

Fact 4.6. If $\to$ is testlike and $[\to] \neq [\supset]$, then $\to$ is not propositional.

Proof. Suppose $\to$ is testlike and not the material conditional. Then $f(s, [\varphi], [\psi])$ must eliminate a $\varphi \wedge \psi$ world w in some case. But since $\to$ is testlike, $\{w\} \models \varphi \to \psi$. So $w \in f(\{w\}, [\varphi], [\psi])$. So $s[\varphi \to \psi] \neq \bigcup \{\{w\}[\varphi \to \psi] \mid w \in s\}$. □

Definition 4.17.

1. $s[\varphi \vee \psi] = s[\varphi] \cup s[\psi]$

2. $s[\varphi \barwedge \psi] = f(s, [\varphi], [\psi])$

Fact 4.7. If $\rightarrow$ is selective, then $[\varphi \rightarrow \psi] = [\neg\varphi \vee (\varphi \barwedge \psi)]$.

Proof. $s[\varphi \rightarrow \psi] = s[\neg\varphi] \cup f(s, [\varphi], [\psi]) = s[\neg\varphi] \cup s[\varphi \barwedge \psi] = s[\neg\varphi \vee (\varphi \barwedge \psi)]$. □

Definition 4.18. $\rightarrow$ is Ramseyan just in case $s \models \varphi \rightarrow \psi$ iff $s[\varphi] \models \psi$.

Fact 4.8. If $\rightarrow$ is selective, then $\rightarrow$ is testlike iff $\rightarrow$ is Ramseyan.

Proof. Suppose $\rightarrow$ is selective, and that $\rightarrow$ is testlike. Suppose $s[\varphi] \models \psi$. Then $s \models \varphi \rightarrow \psi$ by the first case of being a testlike operator. Now suppose $s[\varphi] \not\models \psi$. In that case $s[\varphi \rightarrow \psi] = s[\neg\varphi] \cup f(s[\varphi][\psi])$. Since $s[\varphi] \not\models \psi$, we know that $s[\varphi][\neg\psi] \neq \emptyset$, and so $s[\neg\varphi] \cup s[\varphi][\psi] \subset s$. Since $f(s[\varphi][\psi]) \subset s[\varphi][\psi]$, this implies that $s[\varphi \rightarrow \psi] = s[\neg\varphi] \cup f(s[\varphi][\psi]) \subset s$, and so $s \not\models \varphi \rightarrow \psi$.

Suppose that $\rightarrow$ is selective and Ramseyan. Now suppose that $s[\varphi] \models \psi$. We must show that $f(s, [\varphi], [\psi]) = s[\varphi][\psi]$. Since $\rightarrow$ is Ramseyan and $s[\varphi] \models \psi$, we know that $s \models \varphi \rightarrow \psi$. This means that $s = s[\neg\varphi] \cup f(s, [\varphi], [\psi])$. But if $\rightarrow$ is not testlike, then $f(s, [\varphi], [\psi])$ is not $s[\varphi][\psi]$, which means that some $\varphi \wedge \psi$ world is eliminated from s by $f(s, [\varphi], [\psi])$. This last step assumes that $s[\varphi]$ is nonempty; but if this condition fails then $f(s, [\varphi], [\psi])$ is trivially identical to $s[\varphi][\psi]$. But if some $\varphi \wedge \psi$ world is eliminated from s by $f(s, [\varphi], [\psi])$, it cannot be contained in $s[\neg\varphi]$. So it is not in $s[\neg\varphi] \cup f(s, [\varphi], [\psi])$, and so $s \not\models \varphi \rightarrow \psi$. □

Fact 4.9. If $\rightarrow$ is testlike, then $\rightarrow$ is idempotent.

Proof. If $\rightarrow$ is selective, then $s[\varphi \rightarrow \psi]$ contains only worlds where either $\neg\varphi$, or $\varphi \wedge \psi$. This means that $s[\varphi \rightarrow \psi][\varphi] \models \psi$. Now suppose $\rightarrow$ is testlike. This then implies that $s[\varphi \rightarrow \psi] \models \varphi \rightarrow \psi$, since $s[\varphi \rightarrow \psi]$ satisfies the condition for support required by testlike conditionals. Since $s[\varphi \rightarrow \psi] \models \varphi \rightarrow \psi$ for arbitrary s, φ, and ψ, we know that $\rightarrow$ is idempotent. □

Definition 4.28.

1. $c[\alpha] = \langle \bigcup \sigma_{c[\alpha]}, \{s \in \sigma_c \mid \forall w \in s : w(\alpha) = 1\} \rangle$

2. $c[\neg\varphi] = \begin{cases} c[\neg^{s}_{A}\varphi] & \text{if } [\varphi] \text{ is an assertion} \\ c[\neg^{w}_{CA}\varphi] & \text{otherwise} \end{cases}$

3. $c[\varphi \wedge \psi] = c[\varphi][\psi]$

4. $c[\varphi \rightarrow \psi] = \langle s_c, \sigma_{c[\neg\varphi]} \cup \sigma_{c[\psi]} \rangle$

5. $c[\Box\varphi] = \langle s_c, \sigma_{c[\varphi]} \rangle$

6. $c[\Diamond\varphi] = \langle s_c, \sigma_c - \sigma_{c[\neg\varphi]} \rangle$

Definition 4.29.

1. $s[\alpha]_2 = \{w \in s \mid w(\alpha) = 1\}$

2. $s[\neg\varphi]_2 = s - s[\varphi]_2$

3. $s[\varphi \wedge \psi]_2 = s[\varphi]_2[\psi]_2$

4. $s[\varphi \rightarrow \psi]_2 = \{w \in s \mid s[\varphi]_2 \models \psi\}$

5. $s[\Box\varphi]_2 = \{w \in s \mid s \models \varphi\}$

6. $s[\Diamond\varphi]_2 = \{w \in s \mid s[\varphi]_2 \neq \emptyset\}$

Definition 4.30. $c[\varphi]_2^{\uparrow} = \langle s_c, \{s \in \sigma_c \mid s[\varphi]_2 = s\} \rangle$

Fact 4.10.

1. If φ does not contain $\rightarrow$, $\Box$, or $\Diamond$, then for every c, $c[\varphi]_3 = c[\varphi]_2^{\uparrow}!$.

2. If φ contains $\rightarrow$, $\Box$, or $\Diamond$, then $[\varphi]_3 = [\varphi]_2^{\uparrow}$.

Proof. For simplicity, let $[\cdot]$ denote $[\cdot]_3$. In addition, we'll focus in what follows on σ_c exclusively, proving that $\sigma_{c[\varphi]} = \sigma_{c[\varphi]_2^{\uparrow}}$. We can do this because $s_{c[\varphi]} = \bigcup \sigma_{c[\varphi]}$ whenever φ is non-modal. We proceed by induction on $\mathscr{L}$. First, suppose $s \in \sigma_{c[\alpha]}$. Then s is made up exclusively of α worlds. So $s \in \sigma_{c[\alpha]_2^{\uparrow}}$. Similarly, if $s \in \sigma_{c[\alpha]_2^{\uparrow}}$, then s is made up exclusively of α worlds, and so s is in $\sigma_{c[\alpha]}$.

Now suppose φ satisfies the equation above, and consider $\neg\varphi$. Here, it's worth considering two cases separately: where $[\varphi]$ is an assertion (and $[\neg\varphi]$ is $[\neg_A^s \varphi]$), and where $[\varphi]$ is not (and $[\neg\varphi]$ is $[\neg_{CA}^w \varphi]$). Consider the former case first. Then $s \in \sigma_{c[\neg\varphi]}$ iff $\langle s, \{s' \in \sigma_c \mid s' \subseteq s\} \rangle[\varphi] = \bot$ iff (by the inductive hypothesis) $s'[\varphi]_2 \neq s'$ for every $s' \subseteq s$. This in turn holds just in case $s' - s'[\varphi]_2 = s'$

for every $s' \subseteq s$, because when φ is an assertion, $s[\varphi]_2$ being empty implies that $s'[\varphi]_2$ is empty for each $s' \subseteq s$. This last condition is then equivalent to $s'[\neg\varphi]_2 = s'$ for every $s' \subseteq s$, and hence to $s \in \sigma_{c[\neg\varphi]_2^{\uparrow}}$. Now suppose instead that $[\varphi]$ is a conditional assertion. Then $s \in \sigma_{c[\neg\varphi]}$ iff $s \in \sigma_c$ and $s \notin \sigma_{c[\varphi]}$. This implies by our inductive hypothesis that $s[\varphi]_2 \neq s$, which, since φ is a conditional assertion and therefore is a test in $[\cdot]_2$, is equivalent to the condition $s[\neg\varphi]_2 = s$. So $s \in \sigma_c$ and $s \notin \sigma_{c[\varphi]}$ is equivalent to $s \in \sigma_{c[\neg\varphi]_2^{\uparrow}}$.

Now suppose φ and ψ satisfy the equation above, and consider $\varphi \wedge \psi$. $s \in \sigma_{c[\varphi\wedge\psi]}$ just in case $s \in \sigma_{c[\varphi][\psi]}$, which holds just in case $s \in \sigma_{c[\varphi\wedge\psi]_2^{\uparrow}}$ by the inductive hypothesis.

Now suppose φ and ψ satisfy the equation above, and consider $\varphi \rightarrow \psi$. $s \in \sigma_{c[\varphi\rightarrow\psi]}$ just in case $s \in \sigma_{c[\neg\varphi]}$ or $s \in \sigma_{c[\psi]}$. By the inductive hypothesis, this in turn holds just in case $s[\neg\varphi]_2 = s$ or $s[\varphi]_2 = s$, which in turn holds just in case $s[\varphi \rightarrow \psi]_2 = s$, iff $s \in \sigma_{c[\varphi\rightarrow\psi]_2^{\uparrow}}$.

Now suppose φ satisfies the equation above, and consider $\Box\varphi$ and $\Diamond\varphi$. $s \in \sigma_{c[\Box\varphi]}$ just in case $s \in \sigma_{c[\varphi]}$, just in case $s \in \sigma_{c[\varphi]_2^{\uparrow}}$. $s \in \sigma_{c[\Diamond\varphi]}$ just in case $s \in \sigma_c$ and $s \notin \sigma_{c[\neg\varphi]}$. By the inductive hypothesis, this last holds just in case $s[\neg\varphi]_2 \neq s$, just in case $s[\Diamond\varphi]_2 = s$, which all holds just in case $s \in \sigma_{c[\Diamond\varphi]_2^{\uparrow}}$. $\Box$

Bibliography

Maria Aloni. Free choice, modals, and imperatives. *Natural Language Semantics*, 15(1):65--94, 2007.

Luis Alonso-Ovalle. *Disjunction in Alternative Semantics*. PhD thesis, University of Massachusetts Amherst, 2006.

Pranav Anand and Valentine Hacquard. Epistemics and attitudes. *Semantics and Pragmatics*, 6 (8):1--59, 2013.

Andrew Bacon. Stalnaker's thesis in context. *The Review of Symbolic Logic*, 8(1):131--163, March 2015.

Chris Barker. Free choice permission as resource-sensitive reasoning. *Semantics and Pragmatics*, 3(10):1--38, 2010.

Stephen J. Barker. Towards a pragmatic theory of 'if'. *Philosophical Studies*, 79(2):185--211, 1995.

David Beaver. *Presupposition and Assertion in Dynamic Semantics*. CSLI Publications, Stanford, CA, 2001.

Nuel Belnap. Conditional assertion and restricted quantification. *Noûs*, 4(1):1--12, February 1970.

Nuel Belnap. Restricted quantification and conditional assertion. In Hugues Leblanc, editor, *Truth, Syntax and Modality*, pages 48--75. North-Holland Publishing Company, Amsterdam, 1973.

Jonathan Bennett. Counterfactuals and possible worlds. *Canadian Journal of Philosophy*, 4:381--402, 1974.

Daniel Büring. Identity, modality, and the candidate behind the wall. In Devon Strolovitch and Aaron Lawson, editors, *SALT*, volume VIII, pages 36--54, Ithica, NY, 1998. Cornell University.

Nate Charlow. Triviality for restrictor conditionals. *Noûs*, 2015.

Simon Charlow. The scope of alternatives: Indefiniteness and islands. March 2017.

Roderick Chisholm. The contrary-to-fact conditional. *Mind*, 55:289--307, 1946.

Keith DeRose. Can it be that it would have been even though it might not have been? *Philosophical Perspectives*, 13:387--413, 1999.

Michael Dummett. *Frege: Philosophy of Language*. Duckworth, London, 1973.

Dorothy Edgington. On conditionals. *Mind*, 104(414):235--329, 1995.

Kit Fine. Critical notice of *Counterfactuals*. *Mind*, 84(335):451--458, 1975.

Kit Fine. The impossibility of vagueness. November 2016.

Danny Fox. Free choice and the theory of scalar implicatures. MIT 2009.

Melissa Fusco. Deontic modals and the semantics of choice. *Philosophers' Imprint*, 15(28):1--27, 2015.

Dov M. Gabbay. A general theory of the conditional in terms of a ternary operator. *Theoria*, 38: 97--104, 1972.

Bart Geurts. Entertaining alternatives: Disjunctions as modals. *Natural Language Semantics*, 13: 383--410, 2005.

Allan Gibbard. Two recent theories of conditionals. In William L. Harper, Robert Stalnaker, and Glenn Pearce, editors, *IFS: Conditionals, Belief, Decision, Chance and Time*, volume 15 of *The University of Western Ontario Series in Philosophy of Science*, pages 211--247. Springer Netherlands, 1981.

Anthony S. Gillies. Epistemic conditionals and conditional epistemics. *Nous*, 38(4):585--616, 2004.

Anthony S. Gillies. On truth-conditions for *if* (but not quite only *if*). *Philosophical Review*, 118 (3):325--349, 2009.

Nelson Goodman. The problem of counterfactual conditionals. *The Journal of Philosophy*, 44: 1137--128, 1947.

Jeroen Groenendijk and Martin Stokhof. Dynamic montague grammar. In Laszlo Kalman and Laszlo Polos, editors, *Papers from the Second Symposium on Logic and Language*, pages 3--48, Budapest, 1990. Akademiai Kiado.

Jeroen Groenendijk and Martin Stokhof. Dynamic predicate logic. *Linguistics and Philosophy*, 14 (1):39--100, February 1991a.

Jeroen Groenendijk and Martin Stokhof. Two theories of dynamic semantics. In Jan van Eijck, editor, *Lecture Notes in Computer Science*, volume 478, pages 55--64. Springer, 1991b.

Jeroen Groenendijk, Martin Stokhof, and Frank Veltman. Coreference and modality. In Shalom Lappin, editor, *Handbook of Contemporary Semantic Theory*, pages 179--213. Blackwell, 1996.

Alan Hájek. Staying regular? 2013.

C.L. Hamblin. Questions in montague english. In *Foundations of Language*, volume 10, pages 41--53. 1973.

David Harel, Dexter Kozen, and Jerzy Tiuryn. *Dynamic Logic*. MIT Press, 2000.

Chaoan He. Conjunction, connection and counterfactuals. *Erkenntnis*, pages 1--15, 2016.

Irene Heim. *The Semantics of Definite and Indefinite Noun Phrases*. PhD thesis, University of Massachusetts, Amherst, MA, 1982.

Irene Heim. On the projection problem for presuppositions. *WCCFL*, 2:114--125, 1983.

Janneke Huitink. Partial semantics and iterated *if*-clauses. In A. Riester and T. Solstad, editors, *Proceedings of Sinn and Bedeutung*, volume 13, pages 203--216, 2009a.

Janneke Huitink. Domain restriction by conditional connectives. July 2009b.

Richard Jeffrey. On indeterminate conditionals. *Philosophical Studies*, 14:37--43, 1963.

Hans Kamp. Free choice permission. *Proceedings of the Aristotelian Society*, pages 57--74, 1974.

Lauri Karttunen. *Possible* and *Must*. In J. Kimball, editor, *Syntax and Semantics*, volume 1, pages 1--20. New York Academic Press, 1972.

Lauri Karttunen. Presuppositions and linguistic context. *Theoretical Linguistics*, 1:181--194, 1974.

Magdalena Kaufmann and Stefan Kaufmann. *The Handbook of Contemporary Semantic Theory*, chapter Conditionals and Modality. Number 8. Wiley and Sons, 2015.

Angelika Kratzer. Conditionals. *Chicago Linguistics Society*, 22:1--15, 1986.

Angelika Kratzer. Modality. In Arnim von Stechow and Dieter Wunderlich, editors, *Semantics: an international handbook of contemporary research*, pages 639--650. de Gruyter, 1991.

Angelika Kratzer. *Modals and Conditionals*. Oxford University Press, Oxford, 2012.

Angelika Kratzer. Modality and the semantics of embedding. December 2013.

Angelika Kratzer and Junko Shimoyama. Indeterminate pronouns: The view from japanese. In Y. Otsu, editor, *The Proceedings of the Third Tokyo Conference on Psycholinguistics*, Tokyo, 2002. Hituzi Syobo.

Saul Kripke. Semantical considerations on modal logic. *Acta Philosophica Fennica*, 16:83--94, 1963.

Daniel Lassiter. *Must*, knowledge, and (in)directness. *Natural Language Semantics*, 24(2):117--163, June 2016.

David Lewis. *Counterfactuals*. Blackwell, 1973.

David Lewis. Probabilities of conditionals and conditional probabilities. *The Philosophical Review*, 85:297--315, 1976.

William Lycan. Conditional-assertion theories of conditionals. In Judith Jarvis Thompson and Alex Byrne, editors, *Content and Modality: Themes from the Philosophy of Robert Stalnaker*, pages 148--164. Oxford University Press, 2006.

J. L. Mackie. Counterfactuals and causal laws. In R.J. Butler, editor, *Analytical Philosophy*. Basil Blackwell, 1962.

J. L. Mackie. *Truth, Probability and Paradox*. Clarendon Press, Oxford, 1973.

Ruth Manor. A semantic analysis of conditional assertion. *Journal of Philosophical Logic*, 3(1/2):37--52, March 1974.

Michael McDermott. On the truth conditions of certain 'if' sentences. *The Philosophical Review*, 105(1):1--37, 1996.

Michael McDermott. True antecedents. *Acta Analytica*, 22:333--335, 2007.

Vann McGee. A counterexample to modus ponens. *The Journal of Philosophy*, 82(9):462--471, 1985.

Aidan McGlynn. The problem of true-true counterfactuals. *Analysis*, 72:276--85, 2012.

Zachary Miller. The reformulation argument: reining in gricean pragmatics. *Philosophical Studies*, 173(2):525--546, 2016.

Peter Milne. Bruno de finetti and the logic of conditional events. *British Journal for the Philosophy of Science*, 48:195--232, 1997.

Sarah Moss. On the semantics and pragmatics of epistemic modals. *Semantics and Pragmatics*, 2015.

Sarah Murray and William B. Starr. The structure of communicative acts. *Linguistics and Philosophy*, Forthcoming.

Dilip Ninan. Relational semantics and domain semantics for epistemic modals. *Journal of Philosophical Logic*, 2016.

Alan Penczek. Counterfactuals with true components. *Erkenntnis*, 46:79--85, 1997.

Paul Portner. Imperatives and modals. *Natural Language Semantics*, 15(4):351--383, 2007.

W. V. Quine. *Methods of Logic*. Holt, 1959.

Frank Plumpton Ramsey. *Philosophical Papers*. Cambridge University Press, 1990.

Nicholas Rescher. *Hypothetical Reasoning*. North-Holland Publishing Company, Amsterdam, 1964.

Craige Roberts. Information structure in discourse: towards an integrated formal theory of pragmatics. *Semantics and Pragmatics*, 5(6):1--69, 2012.

Daniel Rothschild. Expressing credences. *Proceedings of the Aristotelian Society*, 112(1):99--114, 2012.

Daniel Rothschild and Seth Yalcin. On the dynamics of conversation. *Nous*, September 2015a.

Daniel Rothschild and Seth Yalcin. On the dynamics of conversation. *Noûs*, 2015b.

Jeff Russell and John Hawthorne. General dynamic triviality theorems. *Philosophical Review*, 125 (3):307--339, 2016.

Mandy Simons. Dividing things up: The semantics of *or* and the modal/*or* interaction. *Natural Language Semantics*, 13:271--316, 2005.

Robert Stalnaker. Presuppositions. *Journal of Philosophical Logic*, 2:447--457, 1973.

Robert Stalnaker. Indicative conditionals. *Philosophia*, 5:269--86, 1975.

Robert Stalnaker. Assertion. In *Syntax and Semantics*, volume 9, pages 315--332. New York Academic Press, 1978.

Robert C. Stalnaker. Conditional propositions and conditional assertions. In Andy Egan and Brian Weatherson, editors, *Epistemic Modality*. Oxford University Press, 2011.

William B. Starr. Expressing permission. In *Proceedings of SALT 26*, Ithaca, NY, 2016. CLC Publications.

William B. Starr. Indicative conditionals, strictly. *Philosophers' Imprint*, Forthcoming.

Tamina Stephenson. Judge dependence, epistemic modals, and predicates of personal taste. *Linguistics and Philosophy*, 30(4):487--525, 2007.

Eric Swanson. How not to theorize about the language of subjective uncertainty. In Andy Egan and B. Weatherson, editors, *Epistemic Modality*. Oxford University Press, 2011.

Eric Swanson. The application of constraint semantics to the language of subjective uncertainty. *Journal of Philosophical Logic*, pages 1--28, 2012.

Johan van Benthem. *Essays in Logical Semantics*. Springer, 1986.

Johan van Benthem. Semantic parallels in natural language and computation. *Studies in Logic and the Foundations of Mathematics*, 129:331--375, 1989.

Johan van Benthem. *Exploring Logical Dynamics*. CSLI Publications, 1996.

Frank Veltman. *Logics for Conditionals*. PhD thesis, University of Amsterdam, 1985.

Frank Veltman. Defaults in update semantics. *Journal of Philosophical Logic*, 25(3):221--261, 1996.

Kai von Fintel and Anthony S. Gillies. *Must* …stay …strong! *Natural Language Semantics*, 18(4): 351--383, 2010.

G. H. von Wright. *Logical Studies*. The Humanities Press, New York, 1957.

Georg Henrik von Wright. Deontic logic and the theory of conditions. *Critica: Revista Hispanoamericana de Filosofia*, 2(6):3--31, 1968.

Malte Willer. Dynamics of epistemic modality. *Philosophical Review*, 122(1):44--92, 2013.

Malte Willer. Simplifying counterfactuals. In *Proceedings of the 20th Amsterdam Colloquium*, pages 428--437, 2015.

Timothy Williamson. *Knowledge and its Limits*. Oxford University Press, Oxford, 2000.

Seth Yalcin. Epistemic modals. *Mind*, 116(464):983--1026, 2007.

Seth Yalcin. Bayesian expressivism. *Proceedings of the Aristotelian Society*, 112(2):123--160, 2012.

Seth Yalcin. Epistemic modality *De Re*. *Ergo*, 2(19), 2015.

Thomas Ede Zimmermann. Free choice disjunction and epistemic possibility. *Natural Language Semantics*, 8:255--290, 2000.