

CHISHOLM'S PARADOX REVISITED: PUZZLES REGARDING CONTRARY-TO-DUTY OBLIGATIONS AND A DYNAMIC SOLUTION

YUNA WON

Indiana University, Bloomington

A contrary-to-duty obligation (CTD obligation) is a type of conditional obligation that tells us what to do when a primary duty is violated. Chisholm's Paradox is one of the most famous deontic puzzles about CTD obligations. It is widely believed that Chisholm's Paradox does not arise for ordering semantics, today's orthodox semantics for modals and conditionals. In this paper, I propose a new puzzle, *the CTD Trilemma*, to show that ordering semantics still has difficulties in adequately representing natural reasoning with CTD obligations. I argue that to solve the CTD Trilemma a formal account must attend to two different functions played by *ought*-statements in our normative reasoning and discourse: *ought*-statements as normative rules and normative judgments. To formally capture this distinction I develop a new dynamic account of *ought*-statements and normative reasoning inspired by Frank Veltman's update semantics for default reasoning. Finally, I show how my update semantics for normative reasoning provides a simple and elegant solution to the CTD Trilemma and explains seemingly inconsistent data about inferences using *ought*-statements in normative reasoning.

1. Introduction

A contrary-to-duty obligation (CTD obligation) is a type of conditional obligation that tells us what to do when a primary obligation is violated. Consider the following pair of sentences.

Contact: Yuna Won <yunawon@iu.edu>

- (p1) You ought to keep your promise.
- (p2) If you do not keep your promise, you ought to apologize.

(p2) expresses a CTD obligation, when paired with the non-conditional obligation expressed by (p1). This CTD obligation generates a new obligation to apologize for a person who has violated the primary obligation expressed by (p1). CTD obligations are not only used routinely by ordinary speakers, but also indispensable across moral, social and legal discourses. This familiar concept gives rise to one of the most difficult puzzles for Standard Deontic Logic, called *Chisholm's Paradox*. Here's Chisholm's (1963) original example.

The Chisholm Set

- (c1) It ought to be that Jones goes to help his neighbors.
- (c2) It ought to be that if he does go he tells them he is coming.
- (c3) If he does not go then he ought not to tell them he is coming.
- (c4) He does not go.

It is not hard to imagine Jones's normative predicament described by this set of statements. In Standard Deontic Logic, however, the Chisholm set is either inconsistent or informationally redundant. In particular, the Chisholm set is inconsistent because the following two inferences are valid in Standard Deontic Logic: (c3) and (c4) entail (c5) by Factual Detachment, and (c1) and (c2) entail (c6) by Deontic Detachment.

- (c5) Jones ought not to tell his neighbors he is coming.
- (c6) Jones ought to tell his neighbors he is coming.

Apparently, (c5) and (c6) express two incompatible obligations. This is why Chisholm's Paradox is often interpreted as a puzzle revealing the difficulty of having both detachment rules for conditional obligations, and motivates many attempts to devise a formal system that validates only one detachment rule, or validates both without the paradoxical results. Thus, it has been widely believed that ordering semantics, which does not validate both detachment rules, is free from Chisholm's Paradox, and thus Chisholm's Paradox has not been discussed much in the ordering semantic framework.

In this paper, I challenge this mainstream diagnosis and widely held belief about Chisholm's Paradox. Given the Chisholm set, the most natural interpretation of Jones's normative predicament is this: he violates his primary duty to go

help his neighbors, since he does not go, and thus he ought not to tell them he is coming. By proposing a new puzzle, *the CTD Trilemma*, I will show that ordering semantics cannot adequately capture Jones's normative predicament and that the real challenge raised by Chisholm's example is not the choice between Factual Detachment and Deontic Detachment.

I argue that accurately understanding two different ways of using normative *ought*-statements is not only the key to solving the CTD Trilemma, but also the first step to developing an adequate formal semantics for normative reasoning and discourse. Based on the two principal observations about our normative reasoning and discourse using *ought*-statements, I motivate and develop a version of *dynamic semantics*, update semantics for normative reasoning. Finally, I show that my dynamic system solves the CTD Trilemma and adequately explains our seemingly inconsistent data about inference patterns using conditional and unconditional *ought*-statements and why the different readings of Chisholm's example are possible.

In sum, this paper aims at making the following contributions:

- (1) To appreciate the real challenge raised by CTD obligations by constructing a new puzzle, the CTD Trilemma;
- (2) To critically evaluate the ordering semantic account of deontic *ought*-statements and conditional obligations by using the CTD Trilemma;
- (3) To recognize that (i) normative rules and normative judgments play distinct roles in our normative reasoning, and (ii) both are expressed by using *ought*-statements;
- (4) To develop a new dynamic semantics that adequately captures two different semantic functions of *ought*-statements in normative discourse and reasoning, and thereby to solve the CTD Trilemma.

In the next two sections, I briefly introduce the original puzzle, Chisholm's Paradox for Standard Deontic Logic, and show why ordering semantics is free from this classic puzzle.

2. Standard Deontic Logic (SDL) and Chisholm's Paradox

Chisholm's Paradox is one of the most serious challenges to Standard Deontic Logic (SDL). SDL is a normal modal logic with a deontic necessity operator 'OB' for obligations, and axiomatized as follows.

- | | | |
|----|--|-----------|
| A1 | All tautologies | |
| A2 | $\mathbf{OB}(\phi \supset \psi) \supset (\mathbf{OB}\phi \supset \mathbf{OB}\psi)$ | (OB-K) |
| A3 | $\mathbf{OB}\phi \supset \neg\mathbf{OB}\neg\phi$ | (OB-Dual) |
| R1 | If $\vdash \phi$ and $\vdash \phi \supset \psi$ then $\vdash \psi$. | (MP) |
| R2 | If $\vdash \phi$ then $\vdash \mathbf{OB}\phi$. | (OB-Nec) |

Since in SDL ordinary *if*-statements are expressed by material conditionals, there are two possible combinations of a material conditional and a deontic necessity operator to express conditional obligations: $\mathbf{OB}(\phi \supset \psi)$ and $\phi \supset \mathbf{OB}(\psi)$. Therefore, the most straightforward way of representing Chisholm's example is the following.

The Chisholm Set

- | | | |
|------|--|--------------------------------------|
| (c1) | It ought to be that Jones goes to help his neighbors. | $\mathbf{OB}(h)$ |
| (c2) | It ought to be that if he does go he tells them he is coming. | $\mathbf{OB}(h \supset t)$ |
| (c3) | If he does not go then he ought not to tell them he is coming. | $\neg h \supset \mathbf{OB}(\neg t)$ |
| (c4) | He does not go. | $\neg h$ |

We can easily imagine the situation where Jones has duties and obligations expressed by (c1)–(c3), but fails to meet his primary obligation (c1). In such a case, all four statements of the Chisholm set are true, and each provides independent information about his situation. However, this formal representation of the Chisholm set in SDL leads to a contradiction. Here is a simple proof. (c3) and (c4) entail (c5) by (MP), and (c1) and (c2) entail (c6) by (OB-K) and (MP).

- | | | |
|------|--|-----------------------|
| (c5) | Jones ought not to tell them he is coming. | $\mathbf{OB}(\neg t)$ |
| (c6) | Jones ought to tell them he is coming. | $\mathbf{OB}(t)$ |

Since (OB-Dual) holds, (c6) is equivalent to $\neg\mathbf{OB}(\neg t)$, which contradicts (c5). Therefore, (c1)–(c4) is *inconsistent*, in the sense that Jones has two mutually incompatible obligations.

Even if (c2) and (c3) are read as $h \supset \mathbf{OB}(t)$ and $\mathbf{OB}(\neg h \supset \neg t)$ respectively, SDL does not faithfully represent the Chisholm set. Since $\phi \supset \psi$ is equivalent to $\neg\phi \vee \psi$, and $\neg\phi$ entails $\phi \supset \psi$, (c4) $\neg h$ entails (c2') $h \supset \mathbf{OB}(t)$, and (c1) $\mathbf{OB}(h)$ entails (c3') $\mathbf{OB}(\neg h \supset \neg t)$. In these cases, the Chisholm set turns out to be *informationally redundant* in that one of the conditional obligations is entailed by one of the non-conditional statements in the Chisholm set. As we can see in the following table, each possible combination of a material conditional and a deontic modal operator in SDL is either inconsistent or informationally redundant.

A	B	C	D
(c1) OB (h) (c2) OB ($h \supset t$) (c3) $\neg h \supset \mathbf{OB}(\neg t)$ (c4) $\neg h$	(c1) OB (h) (c2') $h \supset \mathbf{OB}(t)$ (c3) $\neg h \supset \mathbf{OB}(\neg t)$ (c4) $\neg h$	(c1) OB (h) (c2) OB ($h \supset t$) (c3') OB ($\neg h \supset \neg t$) (c4) $\neg h$	(c1) OB (h) (c2') $h \supset \mathbf{OB}(t)$ (c3') OB ($\neg h \supset \neg t$) (c4) $\neg h$
Inconsistent: (c1) and (c2) entails OB (t); (c3) and (c4) entails OB ($\neg t$).	Informationally Redundant: (c4) entails (c2').	Informationally Redundant: (c1) entails (c3').	Informationally Redundant: (c1) entails (c3'); (c4) entails (c2').

Here's the paradoxical result with Chisholm's example in SDL.

Chisholm's Paradox (for Standard Deontic Logic)

Intuitively, the Chisholm set is consistent, and each statement in the set tells us independent information. However, the Chisholm set is either inconsistent or informationally redundant in Standard Deontic Logic.

The derivation of a new obligation from a conditional obligation and a factual statement, such as the inference of (c5) from (c3) and (c4), is called *Factual Detachment*. On the other hand, the derivation of a new obligation from a conditional obligation and a non-conditional obligation, such as the derivation of (c6) from (c1) and (c2), is called *Deontic Detachment*. Many proposed solutions of Chisholm's Paradox focus on resolving the tension between these two detachment rules to avoid inconsistency.¹ However, I will challenge this widely accepted interpretation of Chisholm's Paradox by showing that ordering semantics that does not validate Factual Detachment still fails to adequately represent Jones's normative predicament given the Chisholm set.

1. Åqvist (1984) calls this the dilemma on commitment and detachment. This interpretation of Chisholm's Paradox motivates the attempts to devise deontic logic systems that validate Factual Detachment without paradoxical results. Al-Hibri (1978) and Mott (1973) independently develop formal systems in which Factual Detachment is valid while Deontic Detachment is not. Arregui (2008) proposes a new account of deontic *should*-conditionals that validates Factual Detachment in the ordering semantic framework by carefully attending to the contrast between *should* and *should have* and thereby solves the *should* version of Chisholm's Paradox. Loewer and Belzer (1983) propose a system with both detachment rules without contradiction by introducing temporal elements. Straßer (2014) develops an adaptive logic validating both detachment rules by introducing different notions of obligations. Silk (2014) validates both detachment rules without tensions by understanding an *ought* in a conditional statements relative to a contextually salient ordering source. Willer (2014) proposes a dynamic solution to Chisholm's Paradox that validates both detachment rules in his non-monotonic system.

3. Ordering Semantics and Chisholm's Paradox

3.1. The Ordering Semantic Account of If and Ought

Ordering semantics is a broad family of semantic accounts that interpret conditionals and modals by using quantification over possible worlds and orderings over possible worlds. Ordering semantics is now an orthodox semantics in linguistics and philosophy of language, most famously developed by Angelika Kratzer and David Lewis.² The basic idea of ordering semantics is to understand modal claims as descriptions of the possibilities that are relevant in a given conversational context, and of how these possibilities are ranked relative to a given purpose.

In Kratzer's ordering semantics, the set of accessible possible worlds and a preference ordering on them are determined by two contextual parameters: a *modal base*, a function f mapping an index world i to a set of propositions, $f(i)$, and an *ordering source*, a function mapping i to a set of propositions, $g(i)$, that determines which possible worlds are regarded as ideal as far as that conversational context is concerned. Conditionals in ordering semantics are also modal claims. An *if*-clause plays the role of restricting the set of accessible possible worlds. A deontic modal claim like an *ought*-statement describes what is the case in the deontologically best possible worlds among the accessible possible worlds. Naturally, a conditional *ought*-statement describes what is the case in the deontically ideal worlds in the set of possible worlds that is restricted by the conditional's antecedent.³

Roughly, "it ought to be that ψ if ϕ ($\mathbf{O}(\psi/\phi)$)" is true if and only if ψ holds in the deontically ideal worlds among accessible ϕ -worlds in the given context.⁴ Here is the formal presentation of the truth condition of a conditional obligation ($\mathbf{O}(\psi/\phi)$) in Kratzer semantics.⁵

2. Different versions of ordering semantics have independently been developed by many linguists and philosophers to deal with different flavors of modals and conditionals: Kratzer (1977; 1981; 1991), Lewis (1973; 1974; 1981), Stalnaker (1968), von Wright (1951; 1970; 1965), Hansson (1969) and Føllesdal and Hilpinen (1971).

3. Here I assume the standard, simplified ordering semantic account of conditional *ought*-statements. On Kratzer semantics, the semantic function of the antecedent of a conditional is to restrict the domain of the covert modal in the consequent. However, this account is not free of controversy. Alternatively, some argue that a covert epistemic necessity modal takes scope over the deontic modal in order to deal with the Miner's Paradox and the trivial truth of "if ϕ , then $\mathbf{O}(\phi)$ " in Kratzer semantics (Kolodny and MacFarlane 2010; see also Bronfman and Dowell 2018; Carr 2014).

4. Following the convention of dyadic deontic logic, I will represent a conditional *ought*-sentence with a two-place deontic operator ' $\mathbf{O}(\psi/\phi)$ ' and ignore the stylistic difference between "it ought to be that if ϕ then ψ " and "if ϕ then it ought to be that ψ " in a natural language. An unconditional *ought*-statement, ' $\mathbf{O}(\phi)$ ' is an abbreviation for ' $\mathbf{O}(\phi/\top)$ ' — \top stands for a tautology.

5. For the sake of simplicity, I make the Limit Assumption throughout this paper, according to which there are always some closest possible worlds to the deontic ideal. Without the Limit Assumption, the truth condition of $\mathbf{O}(\psi/\phi)$ is given as follows: $\llbracket \mathbf{O}(\psi/\phi) \rrbracket_i^{f,g} = 1$ iff for $u, v, z \in \cap f^+(i)$,

For a language \mathcal{L} and $\mathcal{M} = \langle \mathcal{W}, f, g, \llbracket \cdot \rrbracket \rangle$,

$\llbracket \mathbf{O}(\psi/\phi) \rrbracket_i^{f,g} = 1$ iff $\{w : w \in \cap f^+(i) \wedge \neg \exists v \in \cap f^+(i) : v \prec_{g(i)} w\} \subseteq \llbracket \psi \rrbracket$,

where $\cap f^+(i) = \{w : w \in \cap f(i) \wedge w \in \llbracket \phi \rrbracket\}$.⁶

This definition reads: a conditional obligation “it ought to be that ψ if ϕ ($\mathbf{O}(\psi/\phi)$)” is true at an index world i if and only if in the subset of ϕ -worlds accessible from i the closest possible worlds to the deontic ideality determined by $g(i)$ are ψ -worlds. For the sake of convenience, let's use the following simplified gloss:

$\mathbf{O}(\psi/\phi)$ is true in i iff the deontically best worlds among ϕ -worlds accessible from i are ψ -worlds.

3.2. Ordering Semantics and Chisholm's Paradox

Given the ordering semantic account of *ought* and *if*, (c1)–(c3) in the Chisholm set have the following truth conditions as deontic modal statements.

- (c1) $\mathbf{O}(h)$ is true iff the deontically best worlds are h -worlds.
- (c2) $\mathbf{O}(t/h)$ is true iff the deontically best worlds among h -worlds are t -worlds.
- (c3) $\mathbf{O}(\neg t/\neg h)$ is true iff the deontically best worlds among $\neg h$ -worlds are $\neg t$ -worlds.

It is obvious that in ordering semantics each statement of the Chisholm set provides independent information: *no informational redundancy*. More importantly, there is *no inconsistency* in this understanding of the Chisholm set. (c1) and (c2) entail $\mathbf{O}(t)$, whereas (c3) and (c4) do not entail $\mathbf{O}(\neg t)$, because Deontic Detachment is valid, but Factual Detachment is *not valid* in this model.

Deontic Detachment for $\mathbf{O}(\psi/\phi)$: $\mathbf{O}(\psi/\phi) \wedge \mathbf{O}(\phi) \models \mathbf{O}(\psi)$

Factual Detachment for $\mathbf{O}(\psi/\phi)$: $\mathbf{O}(\psi/\phi) \wedge \phi \models \mathbf{O}(\psi)$

$\forall u \exists v [v \preceq_{g(i)} u \wedge \forall z : \text{if } z \preceq_{g(i)} v, \text{ then } z \in \llbracket \psi \rrbracket]$. Here $\cap f^+(i) = \{w : w \in \cap f(i) \wedge w \in \llbracket \phi \rrbracket\}$. And $w \preceq_{g(i)} w'$ means that w is at least as good as w' given $g(i)$; in other words, w meets at least as many propositions in $g(i)$ as w' does.

6. A model \mathcal{M} for a language \mathcal{L} is a quadruple, $\langle \mathcal{W}, f, g, \llbracket \cdot \rrbracket \rangle$ where \mathcal{W} is a non-empty set of possible worlds, f is a function from the index world, i , to a set of propositions, $\preceq_{g(i)}$ orders possible worlds given the set of propositions, $g(i)$, and a valuation function $\llbracket \cdot \rrbracket$ from a well-formed sentence to a set possible worlds; the truth of p at i is represented by $\llbracket p \rrbracket_i = 1$. $\cap f(i)$ is a set of possible worlds which satisfy the propositions in $f(i)$. $g(i)$ is a set of the propositions describing the deontic ideality. $w \prec w'$ means that w is strictly better than w' .

When $\mathbf{O}(\phi)$ and $\mathbf{O}(\psi/\phi)$ are true, $\mathbf{O}(\psi)$ cannot be false in this system. Suppose that the deontically best worlds are ϕ -worlds and among ϕ -worlds the deontically best worlds are ψ -worlds. Then, the deontically best worlds are ψ -worlds as well. Thus, Deontic Detachment, the inference from $\mathbf{O}(\phi)$ and $\mathbf{O}(\psi/\phi)$ to $\mathbf{O}(\psi)$, is valid in ordering semantics. The validity of Deontic Detachment has intuitive appeal. It is natural to infer a new obligation ($\mathbf{O}(\psi)$) from a conditional obligation ($\mathbf{O}(\psi/\phi)$) when its antecedent is obligatory ($\mathbf{O}(\phi)$).

On the other hand, Factual Detachment is not valid for deontic modals. When it comes to deontic modals, it is obvious that an index world, i , is not one of the deontically best possible worlds. For example, in the actual world that we live in, filicide is morally and legally forbidden, and people ought not to commit filicide; however, the actual world is not completely free of it. The actual world is not an ideal world from a moral and legal perspective. Thus it is better not to make the assumption, called *Weak Centering* for deontic modals, according to which the index world i is always one of the best worlds in a relevant sense.⁷ Without *Weak Centering*, the fact that ϕ and $\mathbf{O}(\psi/\phi)$ hold in i does not guarantee that $\mathbf{O}(\psi)$ holds in i —it is possible that $\mathbf{O}(\phi)$ is false in i . Therefore, Factual Detachment is invalid in ordering semantics. As a result, (c5) does not follow from (c3) and (c4), whereas (c6) does follow from (c1) and (c2). There is no inconsistency in the Chisholm set.

- | | |
|---|----------------------|
| (c5) Jones ought not to tell them he is coming. | $\mathbf{O}(\neg t)$ |
| (c6) Jones ought to tell them he is coming | $\mathbf{O}(t)$ |

In sum, ordering semantics is free from the paradoxical results of the Chisholm set that we saw in SDL thanks to its new understanding of the meaning of conditionals and the invalidity of Factual Detachment.

4. The CTD Trilemma : The New Puzzle

Chisholm's four sentences entail (c6), not (c5), in ordering semantics. That means, Jones ought to tell his neighbors he is coming when he does not go. This result is deeply unsatisfactory. The faithful formal representation of Jones's predicament in Chisholm's example seems to require more than just presenting Chisholm's four sentences without inconsistency and redundancy.

7. Formally, *Weak Centering* amounts to the following assumption about an ordering source $g(i)$: $i \in \{w : w \in \cap f^+(i) \wedge \neg \exists v \in \cap f^+(i) : v \prec_{g(i)} w\}$.

To identify the real challenge raised by Chisholm's example, I am going to examine three types of situations the Chisholm set can naturally call to mind and their characteristic features: the According-to-Duty (ATD) Situation, the Overriding Situation, and the CTD Situation. Noting that there are three different ways of interpreting Jones's predicament in Chisholm's example, we will be able to get clear on the constraints and desiderata of a satisfactory account of CTD obligations and normative reasoning.

4.1. Three Interpretations of Chisholm's Example

First, the According-to-Duty (ATD) Situation of the Chisholm's example is a situation in which Jones has not gone help his neighbors yet, but it is still practically possible for him to go help them. Suppose, for example, that Jones is out of town and knows that his weekend will be well-spent if he goes help his neighbors move tomorrow. Imagine that he can still go help them tomorrow if he catches the train back to town tonight. In this situation, what he ought to do is to go help his neighbors and let them know he is coming. In the ATD situation, (c1)–(c4) hold and the derived ought-statement (c6) holds, but not (c5).

Second, the Overriding Situation is a situation where it would be a good thing for him to go help his neighbors, but it is outside his ability; so it is not the case that he ought to go help them. For example, suppose that Jones is out of town and knows that his neighbors need his help tomorrow, but he has already missed all trains back to town (or, there's no train back to town due to the bad weather); so he won't be there to help them in time. Given Jones's predicament it cannot be said that Jones ought to go help his neighbors; he ought not to tell them he is coming. In the Overriding Situation, (c2)–(c5) hold in this situation, but (c1) does not. It is easy to see that the critical feature that distinguishes the ATD Situation and the Overriding Situation is Jones's practical ability to go help his neighbors.

Third, the CTD Situation is a situation in which Jones violates his duty and his CTD obligation takes effect. That means he ought not to tell them he is coming, but the duty he has failed to fulfill is still there as his duty. Imagine, for example, that Jones has a duty to help his neighbors as a community member or as part of his job (e.g., as a firefighter) and his neighbors need his help. However, Jones missed the last train that would get him back to town in time. He won't be able to help his neighbors. So he ought not to tell them he is coming. However, it is true at the same time that he ought to go help his neighbors since that's his duty. Although (c1) and (c2) express Jones's duties and obligations, we should not infer that Jones ought to tell his neighbors he is coming. Therefore, (c1)–(c5) hold in the CTD Situation, but not (c6).

The CTD Situation is importantly different from the Overriding Situation in that (c1) is read as expressing Jones's duty or obligation in the CTD Situation, which persists regardless of his practical ability of fulfilling the duty in the current circumstance. The Overriding Situation is the interpretation of the Chisholm set when reading (c1) as a judgment that tells Jones which act or course of action is the best thing to do in a relevant sense given his current situation, like recommendation or advice. And, that is why (c1) stops holding when (c4) is true.

Clarifying these three possible interpretations of Chisholm's example permits us to see that depending of how we interpret Jones's normative predicament Chisholm's example presents different challenges with different sets of statements to work on in a formal account.

The ATD Situation	The Overriding Situation	The CTD Situation
(c1) $O(h)$		(c1) $O(h)$
(c2) $O(t/h)$	(c2) $O(t/h)$	(c2) $O(t/h)$
(c3) $O(\neg t/\neg h)$	(c3) $O(\neg t/\neg h)$	(c3) $O(\neg t/\neg h)$
(c4) $\neg h$	(c4) $\neg h$	(c4) $\neg h$
(c6) $O(t)$	(c5) $O(\neg t)$	(c5) $O(\neg t)$

There have been many attempts to deal with Chisholm's Paradox by carefully attending to temporal or agential components of the example. These attempts, roughly, argue that since it is time-sensitive whether Jones can fulfill (c1), what Jones ought to do is different before and after a certain critical point of time that determines whether he can go help his neighbors or not: the ATD Situation before this point and the Overriding Situation after the point. Therefore, Jones's predicament in Chisholm's example is either the According-To-Duty (ATD) Situation or the Overriding Situation.

Given this dichotomic interpretation of Chisholm's example, neither the clash between (c5) and (c6) nor the tension between (c1) and (c5) arise, because they simply do not hold at the same time for Jones. If this interpretation of Chisholm's example is correct, formal frameworks for *ought*s and *if*s are lifted from the burden of reconciling Factual Detachment and Deontic Detachment, and the challenge raised by Chisholm's example is to devise a formal system that is sensitive to such temporal aspects or agent's abilities and can capture both the ATD and Overriding Situations.

Proponents of ordering semantics would welcome this interpretation of Chisholm's example, since the relation between the truth of (c1) and Jones's ability to go is easy to capture using the ordering semantic framework. Here's how that goes. It is a logical truth that either Jones can go help his neighbors, or Jones cannot go help his neighbors: $\diamond h \vee \neg \diamond h$. According to the standard ordering

semantic account of modals and conditionals, $\diamond h$ is true at w iff there is at least one h -possible world accessible from w . On the other hand, $\neg\diamond h$ is true at w iff there is no h -world among all the possible worlds accessible from w . Therefore, the two possible situations are captured by different modal bases: $\neg h \notin f_1$ and $\neg h \in f_2$, respectively. Suppose that the relevant ordering source is $g = \{h, h \supset t, \neg h \supset \neg t\}$. In the situation where Jones can go ($\diamond h$), the relevant contextual parameters are f_1 and g . Given f_1 and g , there are some h -worlds in the relevant domain, and h -worlds are the deontically best worlds. Therefore, Jones ought to go help his neighbors and tell them he is coming : (c1) \wedge (c6). On the other hand, in the situation where Jones cannot go to help his neighbors ($\neg\diamond h$), the relevant contextual parameters are f_2 and g . Given f_2 and g , there are no h -worlds in the relevant domain. Among $\neg h$ -worlds, $\neg h \wedge \neg t$ -worlds are better than $\neg h \wedge t$ -worlds. Therefore, Jones ought not to tell them he is coming ($\mathbf{O}(\neg t)$); but, importantly, it is not the case that he still has an obligation to go help his neighbors ($\neg\mathbf{O}(h)$).⁸

Although the dichotomic interpretation is possible for Chisholm's example, it is easy to construct examples of Chisholm's Paradox without temporal and agential components. Here's Prakken and Sergot's (1996) example:

- (d1) There ought to be no dog.
- (d2) If there is no dog, there ought to be no warning signs.
- (d3) If there is a dog, there ought to be a warning sign.
- (d4) There is a dog.

Unlike (c1) in the Overriding Situation, (d1) is neither defeated nor overridden in this timeless example. From (d1)–(d4) it follows that there ought to be a warning sign, but (d1) still holds as a holiday cottage regulation. Regardless of whether or not a dog can be removed, (d1) holds as a holiday cottage regulation, and (d1) is violated given the truth of (d4). It seems that the most natural reading of (d1)–(d4) is the CTD Situation. So we should conclude that the real puzzle invoked by Chisholm's example is not about temporal or agential components in normative discourse.

Moreover, both the ATD Situation and the Overriding Situation lack the characteristic twofold mechanism of a CTD obligation: the violation of a primary obligation and the generation of a new obligation upon the violation. In the ATD Situation, the CTD obligation does not take effect. In the Overriding Situation, there is no violation of Jones's primary obligation or duty since (c1) stops

8. This modal solution to Chisholm's Paradox in the ordering semantic framework is given and briefly discussed in Portner (2009: Ch. 3). Apparently, the modal solution is closely related to classic temporal solutions to Chisholm's Paradox. For such temporal approaches to solving Chisholm's Paradox, see Feldman (1986) and Loewer and Belzer (1983).

being relevant as being overridden by (c4), in other words, the situation where (c4) is an exceptional situation for (c1), where (c1) does not hold. However, the situation in which a CTD obligation takes effect is a sub-ideal situation, not an exceptional circumstance. In other words, when his CTD obligation takes effect, Jones's primary obligation is violated, not replaced by some other new obligation. Among the three interpretations of Chisholm's example, the CTD Situation is the only interpretation that captures the characteristic twofold mechanism of our reasoning with a CTD obligation. Let's take a close look at the CTD Situation.

4.2. The CTD Situation and The CTD Trilemma

Here are three characteristic features of the CTD Situation: (I) and (II) are the two features that distinguish the CTD Situation from the ATD Situation and the Overriding Situation, and (III) says that the CTD Situation is not an aberrant normative situation like a dilemma situation.

I. Violation of a Primary Obligation

In the CTD Situation, the relevant agent's primary obligation holds, but it is violated: $O(\phi) \wedge \neg\phi$.

II. Inference of a New *Ought*-Statement

In the CTD Situation, a new *ought*-statement that tells the agent what to do in her current circumstance is derived from a CTD obligation and factual statement: $O(\psi)$ from $\neg\phi$ and $O(\psi/\neg\phi)$.

III. No Tension among *Ought*-Statements

Conditional and unconditional *ought*-statements in the CTD Situation do not generate contradictory or difficult normative choices for the agent; the agent in the CTD Situation does not face a practically difficult choice.

These three characteristic features of the CTD Situation are also the three requirements a satisfactory formal account of *ought* and *if* must meet: (I) the Violation Requirement, (II) the New *Ought*-Statement Requirement and (III) the No Tension Requirement. In Chisholm's example, together with the first two requirements, (III) the No Tension Requirement put constraint on what we should not infer from the Chisholm set in the CTD Situation.

- (i) (c6) does not follow from (c1) and (c2), or even if (c6) follows from (c1) and (c2), it is not the case that Jones ought to *tell* and *not tell* he is coming;
- (ii) (c1) and (c5) hold, but it is not the case that Jones ought to *go help* and *not tell* them he is coming.

Let's call the challenge of meeting all three features of the CTD Situation *the CTD Trilemma*. Meeting all three of them amounts to representing Jones's CTD situation as the situation in which (c1)–(c5) hold but (c6)–(c8) do not hold. Let's call this set of statements (c1)–(c5) that represent Jones's CTD Situation, *the Extended Chisholm set*.

- | | | |
|--------|---|----------------------|
| (c1) | It ought to be that Jones goes to help his neighbors. | $O(h)$ |
| (c2) | It ought to be that if he does go he tells them he is coming. | $O(t/h)$ |
| (c3) | If he does not go then he ought not to tell them he is coming. | $O(\neg t/\neg h)$ |
| (c4) | He does not go. | $\neg h$ |
| (c5) | He ought not to tell them he is coming. | $O(\neg t)$ |
| # (c6) | Jones ought to tell them he is coming. | $O(t)$ |
| # (c7) | Jones ought to go help his neighbors and not to tell them he is coming. | $O(h \wedge \neg t)$ |
| # (c8) | Jones ought to and ought not to tell them he is coming. | $O(t \wedge \neg t)$ |

The CTD Trilemma specifies the three desiderata a formal account of normative language must meet. They strongly suggest the validity of Factual Detachment and the invalidity of Deontic Detachment and Agglomeration of *ought*-statements. Any satisfactory formal account must be able to explain these puzzling data points and formally incorporate them. Now the question is whether this new puzzle can be solved in the ordering semantic framework.

First, a version of ordering semantics that rejects Weak Centering can express the violation of Jones's primary obligation by using the conjunction $O(h) \wedge \neg h$. However, the rejection of Weak Centering results in the invalidity of Factual Detachment as we have seen in Section 3.2.⁹ Thus, this version of ordering semantics has difficulty in explaining why it is natural to infer the new *ought*-statement (c5) in the CTD Situation. Therefore, this approach meets (I) the Violation Requirement and (III) the No Tension Requirement, but fails to meet (II) the New *Ought*-Statement Requirement.

Second, a version of ordering semantics that assumes Weak Centering for deontic modals runs a different risk. This version of ordering semantics

9. The rejection of Weak Centering means that the actual world (or an index world) is not always one of the ideal possible worlds; thus, $O(h)$ could be true in the actual world, where h does not hold.

validates Factual Detachment since it assumes that the actual world is one of the ideal possible worlds when evaluating deontic modals. So it meets (II).¹⁰ However, the conjunction $\mathbf{O}(h) \wedge \neg h$ is not possible in this system, since the truth of $\mathbf{O}(h)$ entails the truth of h in i given Weak Centering. Therefore, this version of ordering semantics with Weak Centering meet (II) and (III) but fail to meet (I). In sum, whether or not Weak Centering is assumed, ordering semantics faces unwanted results: if Weak Centering is assumed, then it fails to meet (I) the Violation Requirement; and if not, it fails to meet (II) the New *Ought*-Statement Requirement.

Since the rejection of Weak Centering entails the invalidity of Factual Detachment, it would look like a problem of how to make sense of Factual Detachment in the ordering semantic framework without giving up on our strong intuition that the actual world is not a deontically ideal world. Let's check this route. Inspired by Stalnaker (1975), one might argue that Factual Detachment is technically invalid, but *pragmatically valid* in that the inference of $\mathbf{O}(\psi)$ from the assertions of $\mathbf{O}(\psi/\phi)$ and ϕ is *reasonable*.¹¹ Since Deontic Detachment is valid in ordering semantics, the pragmatic validity of Factual Detachment brings the familiar tension between (c5) and (c6) back. At least, this approach must explain why (c5) guides Jones's action, while (c6) does not.

Maybe we can concoct a version of ordering semantic account for *oughts* and *ifs* that invalidates Deontic Detachment but validates Factual Detachment, and get the positive data right: (c5) follows from (c1)–(c4), but not (c6). This will meet both (I) and (II). Unfortunately, even this contrived system cannot meet (III). Just (c1) and (c5) together generate a tension in an ordering semantic framework, which does not exist in our intuitive understanding of the CTD Situation. To say that both (c1) $\mathbf{O}(h)$ and (c5) $\mathbf{O}(\neg t)$ are true means that in the deontically ideal possible worlds Jones *goes help* his neighbors but *does not tell* them he is coming in the ordering semantic framework. Although (c1) and (c5) are not logically inconsistent with each other, $\mathbf{O}(h \wedge \neg t)$ does not hold in the CTD Situation given that (c1) and (c2) hold in the CTD Situation; according to (c1) and (c2), ideally Jones is supposed to go help his neighbors and tell them he is coming. Prakken and Sergot (1996; 1997) call this puzzling result in ordering semantics *Pragmatic Oddity*. Saint-Croix and Thomason (2016) also draw attention to the tension between (c1) and (c5) in Chisholm's example. They argue that with or without

10. When Jones does not go ($\neg h$) and his CTD obligation ($\mathbf{O}(\neg t/\neg h)$) holds, he ought not to tell them he is coming ($\mathbf{O}(\neg t)$), because the actual world is one of the deontically ideal possible worlds in which $\neg h$ holds when Weak Centering is assumed.

11. In the same vein, Kolodny and Macfarlane (2010) argue that the inference from $\mathbf{O}(\psi/\phi)$ and ϕ to $\mathbf{O}(\psi)$ is not valid, but *quasi-valid* in that it is a reliable inference when $\mathbf{O}(\psi/\phi)$ and ϕ are *known* to be true. Saint-Croix and Thomason (2016) argues that Factual Detachment is pragmatically valid to explain the inference of (c5) from (c3) and (c4).

conditional *ought*-statements in the Chisholm set (c1) and (c5) generate a puzzle; both (c1) and (c5) hold in Chisholm's example, but the conjunction of them, (c7), "is clearly false" (Saint-Croix & Thomason 2016: 193). They call this puzzle, *an unconditional version of Chisholm's Paradox*.¹²

In sum, the three characteristic features of the CTD Situation generate the following trilemma for ordering semantics:

The CTD Trilemma

The ordering semantic account of *oughts* and *ifs* can meet two of the three features of the CTD Situation, but not all three of them. Therefore, it fails to provide a faithful formal representation of the CTD Situation.

The discussion so far and the exploration of possible solutions to the CTD Trilemma in the ordering semantic framework naturally prompts us to suspect that (c1) and (c5) are different types of *ought*-statements and nudge us to explore alternative approaches.

5. The New Approach

To solve the CTD Trilemma I propose a new approach employing the following two observations about our normative reasoning and the dual nature of *ought*-statements in our normative discourse.

- (i) Normative rules and normative judgments have different functions in our normative reasoning. Accepted normative rules are prescriptive in that they define which acts or states of affairs an agent prefers or regards

12. Although I agree with their observation, I don't think that the reason that (c7) sounds marked has nothing to do with the conditional obligations in the Chisholm's example. (c7) is false or infelicitous to utter because both (c1) and (c2) hold for Jones in the CTD Situation. Since (c2) holds for Jones in the CTD Situation, the conjunction of (c1) and (c5) does not hold. Consider the following example:

- (j1) Jones ought to look after his child.
- (j3) If Jones does not look after his child, he ought to hire a nanny.
- (j4) But he will not.
- (j5) So he ought to hire a nanny.

The conjunction of (j1) and (j5) is not "clearly false" as Saint-Croix & Thomason argue. It sounds perfectly fine to say in Jones's normative predicament described by (j3)–(j5) and (j1): Jones ought to look after his child and hire a nanny. There is not even pragmatic oddity without the corresponding ATD obligation (j2) "If Jones looks after his child, then he ought not to hire a nanny." So it seems that an unconditional version of Chisholm's Paradox is not a general recipe for the puzzle Saint-Croix & Thomason put forward.

as more desirable than the others. On the other hand, a normative judgment describes which course of action or state of affairs is most desirable or preferred among alternatives in the given circumstance.

- (ii) In our normative discourse, *ought*-statements can be used to express both normative rules and normative judgements.

I will first explain each observation and how they help us explain the seemingly inconsistent linguistic data on the CTD Situation in the CTD Trilemma.

5.1. Normative Rules and Normative Judgements

The distinction between normative rules and normative judgments is straightforward. By normative rules, I mean regulative rules in our lives that offer guidance for judgments about what to do in various circumstances. Moral, civil, or legal duties and obligations, requirements of etiquette and social norms all are considered normative rules. On the other hand, a normative judgment is a practical judgment inferred from a set of normative rules and facts telling us which course of action is most desirable in the given circumstance.

Normative rules are prescriptive in that they generate commitments and expectations to perform, avoid, or permit certain types of acts. A set of normative rules that are accepted by an agent forms a preference pattern regarding which courses of action or states of affairs are most desirable. For example, for those who have learned and internalized a set of dinner manners, eating asparagus with their fingers is considered more desirable than eating asparagus with a knife and fork. For Kantian deontologists, being completely honest is always morally preferable to lying.

Normative rules are importantly different from empirical rules that express causal or empirical regularities. A case that does not conform to a certain empirical rule challenges or even disproves the empirical rule, but a particular case that does not conform to a normative rule is a violation of the rule and does not weaken the rule. For example, the observation of white ravens would challenge the empirical rule that all ravens are black; on the other hand, lying to a murderer outside your door to disallow him entrance into your home would be a violation of the moral principle not to lie but does not challenge or disprove it.¹³

The contrast between normative rules and empirical rules reveals another important attribute of normative rules. Normative rules remain relevant as regulative rules even when an agent violates them, unless an agent is exempted from

13. For more discussions on the nature of normative rules and their explanatory roles, see Sayre-McCord (1992).

them or stops being the subject of the rules. As a legally recognized subject of her country, for example, Jane has an obligation to serve on a jury when called. Suppose that Jane was called. "Jane ought to serve on a jury" as a normative rule holds, and does so even when she fails to show up for the duty. She can be exempt from her jury service if she meets certain conditions, but that exemption does not dissolve her duty as a citizen.¹⁴

In contrast to a normative rule, a normative judgment tells us what action or state of affairs is more desirable than its alternative. Normative judgments as practical judgments are sensitive to the agent's contingent circumstances and her practical abilities in a given situation and guide her actions, and can change as the circumstances change. An agent makes a normative judgment based on her factual judgments (what is known) about her current circumstance and the conditional and unconditional normative rules accepted by her. A devoted utilitarian would make a normative judgment "I ought to give more money to charity" from a combination of his moral principles and the facts regarding his income and necessary living expenses. When I fail to keep my promise to my partner, "I ought to apologize to my partner" is the normative judgment I infer from a given set of normative rules I accept and the fact that I failed to keep my promise.

Now we have an intuitive model of a simple normative reasoning using normative rules, normative judgments, and factual judgments. Conditional or unconditional normative rules structure the pattern of preferences by imposing obligations, granting permissions, and forbidding certain types of actions, and facts (or what is known) about the given circumstance determines what are practically feasible choices in the circumstance. A normative judgment is made by using both the pattern of preferences defined by normative rules and the set of choices limited by what is known about the given circumstance. For example, my moral duty to keep my promise to my partner generates the preference ranking over two states: keeping my promise over not keeping my promise. My conditional obligation to apologize when failing to keep my promise articulates the preference ranking further in the following order : 1) keeping my promise; 2) failing to keep my promise and apologizing for it, and 3) failing to keep my promise and not apologizing. When keeping my promise is still a possibility, I ought to keep my promise. When I fail to keep my promise, I have only two options: apologizing for it or not apologizing for it. And, my normative judgment has to be: "I ought to apologize." Even in this sub-ideal situation, my duty to keep my promise holds. In other words, my primary moral duty "I ought to keep my promise" holds.

14. Note that being exempt from a duty or obligation is different from failing to meet or violating one's duty or obligation. In this example, Jane would be free of her duty if she is qualified or her circumstance meets some pre-established conditions, such as being a student or a non-citizen.

Now is the time to talk about the second principle of this approach: even though they serve different functions in our normative reasoning, both normative rules and judgments can be expressed by *ought*-statements in our normative discourse.¹⁵ Imagine Jane is called for jury duty, and suppose it turns out that she was a victim in a similar crime and knows that it will be hard for her to maintain objectivity. In this case, she can reasonably make a normative judgment “I ought not to serve on a jury,” although this judgment is opposed to the relevant normative rule, which also could be stated in a form of *ought*-statement: “Jane ought to serve on a jury.”

Here I argue that *ought*-statements used in normative discourse that are commonly believed to be homogeneous are actually of two kinds: one expressing normative rules and the other expressing normative judgments. And their functions in our normative discourse and reasoning are distinct enough to be taken into our semantic account of *ought*-statements. The question of how to disambiguate two meanings of *ought*-statements in normative discourse in various contexts is beyond the scope of this paper. For the purposes of my project, I am going to presume this ambiguity is resolved at the pragmatic level and focus on analyzing their semantic contents. From now I am going to use two different operators, one for *ought*-statements expressing normative rules (O_R) and the other for *ought*-statements for normative judgments (O_J). In texts, I will use the subscripts to indicate different kinds of *oughts*: “ought_R” and “ought_J.”

5.2. A Sketch for a New Solution to the CTD Trilemma

Using the three components of normative reasoning and the dual nature of *ought*-statements we can give a natural explanation of Jones’s normative predicament in the CTD Situation. In the CTD Situation, (c1)–(c3) are the normative rules for Jones. (c1) expresses Jones’s duty as a community member; (c2) expresses Jones’s conditional obligation stemming from common etiquette; and (c3) expresses Jones’s CTD obligation related to a general obligation not to lie or deceive others. (c4) is the factual judgment about Jones’s contingent situation. (c5) is the normative judgment derived from his CTD obligation and factual judgment, (c3) and (c4).¹⁶ So the Extended Chisholm Set describing Jones’s normative


15. This principle seems to hold for other deontic modal auxiliary verbs, like *should* and *must*.

16. Alternatively, we can read (c1) as Jones’s duty and (c2) and (c3) as conditional normative judgments in the CTD Situation. The preference pattern for which such (c1)–(c3) hold is the same as the preference pattern defined by (c1)–(c3) as normative rules. Therefore, these readings will be represented by the same preference pattern of an information state. I will return to this point after articulating the technical distinctions between normative rules and normative judgments.


predicament in the CTD Situation actually consists of the three different types of statements as follows:

- (c1_R) It *ought*_R to be that Jones goes to help his neighbors. $O_R(h)$
- (c2_R) It *ought*_R to be that if he does go he tells them he is coming. $O_R(t/h)$
- (c3_R) If he does not go then he *ought*_R not to tell them he is coming. $O_R(\neg t/\neg h)$
- (c4) He does not go. $\neg h$
- (c5_J) He *ought*_J not to tell them he is coming. $O_J(\neg t)$

We can informally illustrate how the normative rules, the factual statement and the normative judgments in the CTD Situation interact in the CTD Situation with diagrams. The three normative rules, (c1)–(c3), define Jones's normative preference over possible worlds (or possible states of affair) as follows: (Solid-line circles represent possible worlds that are not ruled out, while dashed-line circles represent possible worlds ruled out by what is known; the arrow connecting two circles or boxes indicates which possibility is preferred; what an arrow is pointing to is the preferred possibility between two possibilities that the arrow connects.)



The factual statement (c4) defines what are practically possible options for Jones in his current predicament by ruling out all *h*-worlds from the set of possible worlds as in the following diagram.



Given (c1)–(c4), the normative judgment (c5) “Jones *ought*_J not to tell them he is coming” accurately describes the best Jones can do in the given circumstance. This new reading of the Extended Chisholm set suggests an intuitive solution to the CTD Trilemma. The simple conjunction of (c1) and (c4), $O_R(h) \wedge \neg h$, expresses the violation of Jones's primary obligation to help his neighbors. The normative judgment $O_J(\neg t)$ follows from the Chisholm set in that $\neg h \wedge \neg t$ -worlds are the most preferred worlds when (c4) holds, as in the diagram. Finally, the normative rule (c1) $O_R(h)$ and the normative judgment (c5) $O_J(\neg t)$ are not in conflict and

thus can be true in the same situation. As we have seen, what is expected by a normative rule does not always coincide with the best we can do in a particular situation, and a normative rule holds even when violated.

This new approach provides a simple intuitive explanation of why (c6)–(c8) all are false normative judgments.

#	(c6) _J	Jones <i>ought</i> _J to tell them he is coming.	$O_J(\neg t)$
#	(c7) _J	Jones <i>ought</i> _J to go help his neighbors and not to tell them he is coming.	$O_J(h \wedge \neg t)$
#	(c8) _J	Jones <i>ought</i> _J to and <i>ought</i> _J not to tell them he is coming.	$O_J(t \wedge \neg t)$

From the normative rules (c1)–(c3) and the factual judgment (c4), the normative judgment (c6) simply does not follow and is neither one of Jones’s normative rules nor an adequate normative judgment in the CTD Situation. Since (c1) is a normative rule and (c5) is a normative judgment, we cannot, using Agglomeration, combine them into one *ought*-statement, which explains the falsity of (c7) without invalidating Agglomeration.

In addition to this natural, intuitive account of the CTD Situation, this new reading can explain why different interpretations of Chisholm’s example are possible. In particular, when Chisholm’s example is read as the Overriding Situation, we read (c1) as a normative judgment overridden by a newly derived normative judgment (c5), because a normative judgment is sensitive to a factual judgment. As a situation changes, an adequate normative judgment may change.

In this section, above, I have proposed a new reading of the Extended Chisholm set and a new approach to the CTD Trilemma, using two key observations: (i) normative rules and normative judgments have distinct functions in our normative reasoning and (ii) *ought*-statements can express both normative rules and judgments in our normative discourse. With these two principles, we can provide a natural explanation of the CTD Situation and meet the three desiderata of the CTD Trilemma in a natural and unforced manner. In the next section, I will develop a new semantic account for normative reasoning that formally captures these principles, thus solving the CTD Trilemma.

6. Dynamic Semantics for Normative Reasoning

The new account I sketch here is a dynamic semantics formally inspired by Veltman’s (1996) update semantics for default reasoning. Here are two reasons that

I build a new semantic account for normative reasoning using Veltman's update semantics: (i) as a dynamic semantics it allows us to define the way normative rules contribute to normative discourse and keep track of them even when they are violated or cannot be fulfilled; and (ii) there is a clear structural parallelism between default reasoning and normative reasoning, which permits a formal adaptation of the update semantics for default reasoning to normative reasoning. I will first informally introduce Veltman's update semantics for default reasoning and show that some of the key features common in default reasoning and normative reasoning. Then, I will formally sketch a new update semantics for normative reasoning.

6.1. Veltman's Update Semantics

The core idea of dynamic semantics is that the meaning of a sentence is how it changes the context or what is believed or accepted by its participants. This context or conversational common ground is called an *information state*, which is formally represented by a set of possible worlds.¹⁷ In static systems, to know the meaning of "it's raining" is to know when "it's raining" is true. In the dynamic framework, to know the meaning of "it's raining" is to know what kind of changes it brings to a relevant agent's information state who accepts the information conveyed by the sentence. Since what is believed or accepted as true is a central part of a context, the dynamic semantic meaning of a sentence is often called a "context change potential" (CCP), which is a function from one information state to another.¹⁸

Without much formalism, we can easily understand how update semantics defines the meaning of simple statements (ϕ) and epistemic *might*-statements ($\diamond_e \phi$) and introduce some core concepts of dynamic semantics. Suppose Ben is wondering if it's snowing outside and naturally says "it might be snowing outside ($\diamond_e p$)." Abby tells him "it's not snowing ($\neg p$)." Her utterance changes what

17. An information state models what is believed or accepted by a relevant agent(s), relevant discourse participants, or a loosely define relevant social group. For the sake of argument, I will leave the notion of information state as an abstract and idealized concept, like a context in ordering semantics, and will not consider the questions on the nature of information states in the dynamic semantic framework.

18. The origin of this dynamic understanding of meaning traces back to Stalnaker's work on presupposition and assertion (1978). The core idea of dynamic semantics is endorsed and developed by many authors to provide a better explanation of various linguistic phenomena: Kamp's (1988, 2011) Discourse Representation Theory, Groenendijk and Stokhof's (1991) Dynamic Predicate Logic, Heim's (1983) File Change Semantics, Veltman's (1996) Update Semantics, and recent dynamic approaches such as Beaver (2001), von Stechow and Gillies (2007), Starr (2016), and Willer (2013; 2014).

Ben believes in that Ben stops believing what is incompatible with what he has learned. So after accepting Abby's utterance it is inadequate for him to say or accept "it might be snowing outside ($\diamond_e p$)." This process of incorporating new information into a current information state to evolve a new information state is called an *update*. When a certain information state (σ) is already compatible with ϕ , ϕ is accepted in σ (or, σ supports ϕ). In this example, "it might be snowing outside" was accepted in Ben's information state before Abby's utterance of $\neg p$, but not after the update with $\neg p$. This change shows that the process of updating new information is not monotonic. While accepting Abby's utterance "it's not snowing ($\neg p$)" amounts to ruling out the possibilities of snowing (p) from Ben's information state, accepting an epistemic *might*-statement does not require a change in the information state. The utterance of "it might be snowing ($\diamond_e p$)" would be accepted only when the possibility of snowing is not completely ruled out. The different behaviors of simple statements and epistemic *might*-statements reveal the important distinction between two kinds of updates in a dynamic account: "additive propositional *updates* and non-classical *tests*" (Veltman 1996: 5). An utterance of ϕ in a discourse requires updating the relevant information state with ϕ by ruling out the possibilities incompatible with it, while an utterance of *might*(ϕ) is an invitation to perform a test on the information state, checking whether the information state is compatible with the possibility of ϕ .¹⁹

Here is a formal presentation of this basic update semantics for simple statements and epistemic *might*-statements.

Definition 1 (Languages, L_0^A and L_1^A) Let \mathcal{A} be a set of atomic sentences, p, q, r, \dots . The language L_0^A has \mathcal{A} as its non-logical vocabulary and the logical operators, such as \neg , \vee , and \wedge , and parentheses,) and (, as its basic logical vocabulary.

A string of symbols χ is a sentence of L_0^A iff

- (i) for some atomic sentence $p \in \mathcal{A}$, $\chi = p$;
- (ii) for some sentence ϕ of L_0^A , $\chi = \neg\phi$;
- (iii) for some sentences ϕ and ψ of L_0^A , $\chi = \phi \wedge \psi$;
- (iv) for some sentences ϕ and ψ of L_0^A , $\chi = \phi \vee \psi$.

19. In this paper it is assumed that all utterances are sincere and felicitously made. Also, this paper sets aside the important issues regarding acceptability and assertability that could be relative to information states.

A language L_1^A has \mathcal{A} and its logical vocabulary and an additional unary operator \diamond_e for *might*.

A string of symbols χ is a sentence of L_1^A iff for some sentence ϕ of L_0^A $\chi = \phi$ or $\chi = \diamond_e \phi$

Definition 2 (The Space of Possible Worlds, W) Let W be the set of all possible worlds, w . Each possible world w is an atomic valuation, a function from every atomic sentence in \mathcal{A} to one truth value ($\mathbf{1}, \mathbf{0}$).

Definition 3 (Truth Sets/Propositions) For $p \in \mathcal{A}$, $\llbracket p \rrbracket^M = \{w \in W \mid w(p) = \mathbf{1}\}$. $\llbracket p \rrbracket$ is a set of possible worlds in which p is true. (A model M will be omitted for simplicity.)

Definition 4 (Information State) (i) σ is an *information state* iff $\sigma \subseteq W$; (ii) $\mathbf{0}$, the *minimal state*, is the information state given by W ; (iii) $\mathbf{1}$, the *absurd state*, is the information state given by the empty set.

For any sentence ϕ , $\mathbf{1}[\phi] = \mathbf{1}$. That means, in the absurd state every sentence is accepted, but no sentence is acceptable.

Definition 5 (Update System) An update system is a triple $\langle L, \Sigma, [*] \rangle$, where L is a language, Σ is a set of information states, σ , each of which is a subset of W , and $[*]$ is an update function. For any system M , an update function $[*]$ is a function from a sentence of L to functions from information states to information states in Σ .

Definition 6 (Update Function) For $p \in \mathcal{A}$ and ϕ and ψ of L_0^A , any sentence χ of L_1^A is such that $\chi = p; \neg\phi; \phi \wedge \psi; \phi \vee \psi$; or $= \diamond_e \phi$. For any χ of L_1^A and an information state σ , $\sigma[\chi]$ is determined as follows:

$$\sigma[p] = \sigma \cap \llbracket p \rrbracket (p \in \mathcal{A})$$

$$\sigma[\neg\phi] = \sigma - \sigma[\phi]$$

$$\sigma[\phi \wedge \psi] = \sigma[\phi][\psi]$$

$$\sigma[\phi \vee \psi] = \sigma[\phi] \cup \sigma[\psi]$$

$$\sigma[\diamond_e \phi] = \begin{cases} \sigma, & \text{if } \sigma[\phi] \neq \mathbf{1}; \\ \mathbf{1}, & \text{if } \sigma[\phi] = \mathbf{1}. \end{cases}$$

While updating the information state σ with p has the effect of excluding all $\neg p$ -worlds from σ , the function of epistemic modal statement behaves like a test. Updating σ with $\diamond_e p$ returns σ , if there is at least one p -world in σ ; otherwise, it is an absurd state, $\mathbf{1}$, given by the empty set.

Another central idea of a dynamic framework is the notion of acceptance. ϕ is accepted in an information state σ when updating σ with ϕ does not change σ : $\sigma[\phi] = \sigma$. In other words, such σ supports ϕ ; σ is called *the fixed point* of ϕ .

Definition 7 (Acceptance)

$\sigma \Vdash \phi$ iff $\sigma[\phi] = \sigma$

Consequence in a dynamic system can be defined in terms of acceptance as follows.

Definition 8 (Dynamic Consequence)

$\phi_1, \phi_2, \dots, \phi_n \models \psi$ iff $\forall \sigma \in \Sigma$: if $\sigma[\phi_1][\phi_2] \dots [\phi_n] = \sigma$, then $\sigma[\psi] = \sigma$.

This dynamic definition of consequence means that $\phi_1, \dots, \phi_n \models \psi$, just in case every state of information which is a fixed point of $[\phi_1] \dots [\phi_n]$ also supports ψ . The dynamic notion of consequence as the preservation of acceptance is non-monotonic in that it is sensitive to the order of premises.

6.2. Default Reasoning and Normative Reasoning

Veltman (1996) expands this simple model of update semantics to analyze default reasoning with *default rules* expressed by “normally (or usually) ϕ ” and *expectations* expressed by “presumably ϕ .” Default rules, such as “normally, ravens are black” and “normally, cold sores heal in a week by themselves,” are descriptive rules that serve to express empirical regularities. They help us get to reasonable expectations or guesses about what will happen. For example, “presumably, Julia’s new pet raven is black”; and “presumably, Jake’s cold sore will heal in a week.” Restricted default rules expressed in the conditional form allow us to make more precise expectations or guesses; for example, “if the patients do not get enough lysine from their diet, cold sore remedies normally take more than a week.”

Suppose that *normally* the campus bookstore is busy in the afternoon. Given this knowledge, it is reasonable for one to conclude that *presumably* the campus bookstore is busy this afternoon. Suppose that you have found that it’s pretty empty this afternoon. Still you can say, “normally, the campus bookstore is busy in the afternoon,” but once seeing it’s not busy, it is not apt anymore to say,

“presumably, the campus bookstore is busy.” This example instantiates the following characteristic inference patterns between “normally ϕ ,” “presumably ϕ ” and “ ϕ .”

- (d1) *Normally* $\phi \models$ *Presumably* ϕ
- (d2) *Normally* $\phi, \neg\phi \not\models$ *Presumably* ϕ
- (d3) *Normally* $\phi, \neg\phi \models$ *Normally* ϕ

Suppose that the following conditional default rule holds about the campus bookstore as well: “if the campus bookstore is not busy in the afternoon, *normally* there is a school event.” You have found that the bookstore is not busy this afternoon. Given these default rules and finding, you can infer: “presumably, there is a school event.” Again, the default rule, “Normally, the campus bookstore is busy in the afternoon” holds. The following utterance is true and felicitous to utter:

“The bookstore is not busy this afternoon. Presumably, there is a school event. However, normally, the campus bookstore is busy in the afternoon.”

Thus, the following inferences regarding non-conforming cases and conditional default rules are warranted in default reasoning.

- (d4) *Normally* $\phi, \neg\phi, \text{Normally } \psi \text{ if } \neg\phi \models$ *Presumably* ψ
- (d5) *Normally* $\phi, \neg\phi, \text{Normally } \psi \text{ if } \neg\phi \models$ *Normally* ϕ

The interplay among default rules, expectations and factual judgments in default reasoning is parallel to the one among normative rules, normative judgments and factual judgments in normative reasoning. Accepted default rules generate a pattern of expectations regarding what will more likely happen. Likewise, accepted normative rules generate a pattern of preferences regarding what actions or state of affairs is more desirable. A default rule holds even when a particular occasion or circumstance diverges from what the default rule predicts. Similarly, a normative rule holds even when it is violated on a particular occasion.

The following inferences using normative rules, normative judgments and factual judgments are naturally warranted in normative reasoning.

- (n1) $\mathbf{O}_R(\phi) \models \mathbf{O}_J(\phi)$
- (n2) $\mathbf{O}_R(\phi), \neg\phi \not\models \mathbf{O}_J(\phi)$

(n3) $O_R(\phi), \neg\phi \models O_R(\phi)$

(n4) $O_R(\phi), \neg\phi, O_R(\psi/\neg\phi) \models O_J(\psi)$

(n5) $O_R(\phi), \neg\phi, O_R(\psi/\neg\phi) \models O_R(\phi)$

These two sets of inferences, (d1)–(d5) and (n1)–(n5), manifest the structural parallelism between default reasoning and normative reasoning and show why it is natural to adapt Veltman’s update semantics for default reasoning to model normative reasoning using normative rules and normative judgments.

6.3. The Dynamic Meanings of Ought-Statements

To formally represent the distinctive functions of two kinds of *ought*-statements in dynamic semantics, we need a *structured* information state that can capture two things: what an agent knows and an agent’s normative preference pattern over possibilities. These two attributes of a structured information state can be modeled by a set of possible worlds and a preference ordering pattern over the possible worlds. Now we can articulate the meanings of $O_R(\psi/\phi)$ and $O_J(\psi/\phi)$ in terms of how they contribute to a structured information state σ . The meaning of a normative rule expressed by $O_R(\psi/\phi)$ is its potential to change the preference ordering of possible worlds in the set of all ϕ -worlds of σ in a way that all $\phi \& \psi$ -worlds are preferred over (or ranked higher than) $\phi \& \neg\psi$ -worlds. On the other hand, the meaning of a normative judgment expressed by $O_J(\psi/\phi)$ is its potential to test whether all $\phi \& \psi$ -worlds are preferred over $\phi \& \neg\psi$ -worlds in the current information state σ .

Before presenting formal definitions let me illustrate what their updates would look like with a diagram.




Figure 1

Suppose that a structured information set σ_0 is represented by a set of all possible worlds W and no preference ordering. Updating σ_0 with a normative rule $\mathbf{O}_R(\phi)$ and a conditional normative rule $\mathbf{O}_R(\psi/\neg\phi)$ in order has the effect of generating the preference pattern as in σ_2 , and updating $\neg\phi$ has the effect of ruling out all ϕ -worlds as in σ_3 . (Again, in diagrams solid-line circles represent sets of possible worlds; the circle an arrow points to is a preferred possibility; and gray dashed-line circles are the sets of possibilities that are ruled by propositional updates.)

We can easily understand the different types of updates by $\mathbf{O}_R(\psi/\phi)$, $\mathbf{O}_J(\psi/\phi)$ and ϕ in this diagram. Updating with a conditional normative rule $\mathbf{O}_R(\psi/\neg\phi)$ generates a new preference pattern in the set of $\neg\phi$ -worlds. These information statements support different normative judgments; in other words, different normative judgments hold as the information state changes. Both σ_1 and σ_2 support $\mathbf{O}_J(\phi)$, but not σ_3 . Both σ_2 and σ_3 support the conditional normative judgment $\mathbf{O}_J(\psi/\neg\phi)$; σ_3 even supports the unconditional normative judgment $\mathbf{O}_J(\psi)$ because ψ -worlds are the best worlds in σ_3 . That means an agent in an information state σ_1 or σ_2 makes a normative judgment that it ought to be that ϕ ($\mathbf{O}_J(\phi)$); however, an agent in σ_3 does not make the normative judgment because the possibility of ϕ is ruled out by what is known ($\neg\phi$). While adequate normative judgments change as the information state evolves, the normative rule $\mathbf{O}_R(\phi)$ remains accepted in the information state; that is, all three information states σ_1 , σ_2 , and σ_3 support $\mathbf{O}_R(\phi)$.

In the next subsection, I will formally present my update semantics for normative reasoning that captures the dynamic meanings of the three types of statements and their interplay summarized in (n1)–(n5). The formal strategy of my update semantics for normative reasoning largely follows Veltman's approach with some stylistic modification and simplification because in this paper I do not deal with rules with exceptions.

6.4. Update Semantics for Normative Reasoning

The major work we need to do to expand the simple update semantics into update semantics for normative reasoning is to construct a structured information state that can keep track of the changes by both normative rules and factual statements. With a structured information state, we will be able to define the two new operators $\mathbf{O}_R(\psi/\phi)$ and $\mathbf{O}_J(\psi/\phi)$ in a way that captures their meanings illustrated with the diagrams and the inference patterns (n1)–(n5).

Definition 9 (Normative Language L_2^A) Let \mathcal{A} be a set of atomic sentences, p, q, \dots . The language L_2^A has \mathcal{A} , as its non-logical vocabulary and the logical

operators, such as \neg , \vee , and \wedge , and parentheses, $)$ and $($, as its basic logical vocabulary.

A string of symbols ϕ is a sentence of L_0^A iff

- (i) for some atomic sentence $p \in \mathcal{A}$, $\phi = p$;
- (ii) for some sentence ψ of L_0^A , $\phi = \neg\psi$;
- (iii) for some sentences ψ and χ of L_0^A , $\phi = \psi \wedge \chi$;
- (iv) for some sentences ψ and χ of L_0^A , $\phi = \psi \vee \chi$.

A language L_2^A has two additional dyadic operators: $\mathbf{O}_R(-/-)$ and $\mathbf{O}_J(-/-)$.

A string of symbols χ is a sentence of L_2^A iff for sentences ϕ and ψ of L_0^A ,

- (i) $\chi = \phi$; or (ii) $\chi = \mathbf{O}_R(\psi/\phi)$ or $\chi = \mathbf{O}_J(\psi/\phi)$.

In order to introduce a structured information state, we need to define two formal devices: a frame N on W and a preference pattern $N(d)$ on d .

Definition 10 (Frame) Let W be the power set of the set A of atomic sentences. A frame N on W is a function assigning a preference pattern $N(d)$ on d for every $d \subseteq W$.

Definition 11 (Preference Pattern) $N(d)$ is a preference pattern on $d \subseteq W$ iff $N(d)$ is a reflexive and transitive relation on d .

A non-trivial normative preference pattern $N(d)$ on d systematically encodes the accepted normative rules relevant to d . The accepted normative rules relevant to d gives a set of propositions that determines the preference ranking over possible worlds in d like an ordering source g does in ordering semantics. The set of relevant normative rules to d includes all normative rules regarding the supersets of d . Consider a normative preference pattern $N(\llbracket \phi \rrbracket)$ on $\llbracket \phi \rrbracket$ which is determined by the set of accepted normative rules relevant to situations where ϕ holds. The set of normative rules that are relevant to the situation where ϕ holds are: $\mathbf{O}_R(\psi_1/\chi)$, $\mathbf{O}_R(\psi_2/\chi)$, \dots , $\mathbf{O}_R(\psi_n/\chi)$ for any χ such that $\llbracket \chi \rrbracket \supseteq \llbracket \phi \rrbracket$. For example, both Sue's moral duty not to lie and her conditional obligation to apologize if she breaks a promise are the relevant normative rules to the situation where she breaks her promise.

Definition 12 (Normative Preference Pattern on $\llbracket \phi \rrbracket$) For χ such that $\llbracket \chi \rrbracket \supseteq \llbracket \phi \rrbracket$, $\mathbf{O}_R(\psi_1/\chi)$, $\mathbf{O}_R(\psi_2/\chi)$, \dots , $\mathbf{O}_R(\psi_n/\chi)$ are the normative rules applied to ϕ -situations. Let P_ϕ be the set of all propositions that are expected to be fulfilled given the normative rules relevant to ϕ -situations: $P_\phi = \{\psi_1, \psi_2, \dots, \psi_n\}$.

$N(\llbracket \phi \rrbracket) = \{\langle w, v \rangle \mid \text{for } w, v \in \llbracket \phi \rrbracket, \text{ every proposition in } P_\phi \text{ that holds in } v \text{ also holds in } w\}$

A preference pattern $N(d)$ is defined by a set of pairs of possible words; when $\langle w, v \rangle \in N(d)$, let's write $w \preceq_{N(d)} v$. It reads: w fulfills as many relevant normative rules as v does.

Definition 13 (Structured Information State σ) σ is a structured information state defined by a pair of a frame N and a set of possible worlds, $s \subseteq W$: $\sigma = \langle N, s \rangle$.

- (i) An information state $\sigma = \langle N, s \rangle$ is coherent, where s is a non-empty subset set of W and N is a coherent frame on W ;
- (ii) The minimal state $\mathbf{0} = \langle N_0, W \rangle$, where $N_0(d) = \{\langle w, v \rangle \mid w, v \in d\}$ for every $d \subseteq W$;
- (iii) The absurd state $\mathbf{1} = \langle N_1, \emptyset \rangle$ where $N_1(d) = \{\langle w, w \rangle \mid w \in d\}$ for every $d \subseteq W$.²⁰

Interestingly, the notion of a structured information state in my update semantics is comparable to the way that ordering semantics uses two contextual parameters. In this system, an information state σ is defined by a pair of a preference frame N and a set of accessible possible worlds s : $\sigma = \langle N, s \rangle$. s is comparable to $\bigcap f(i)$ and the role of a preference frame N is comparable to $g(i)$ in ordering semantics.

Definition 14 (Refinement of Preference Pattern) Let $N(d)$ and $N'(d)$ be preference patterns on $d \subseteq W$.

- (i) $N'(d)$ is a refinement of $N(d)$ iff $N'(d) \subseteq N(d)$;
- (ii) $N(d) \circ \llbracket p \rrbracket$ is a refinement of $N(d)$ with the proposition $\llbracket p \rrbracket$;
 $N(d) \circ \llbracket p \rrbracket = \{\langle w, v \rangle \in N(d) \mid \text{if } v \in \llbracket p \rrbracket, \text{ then } w \in \llbracket p \rrbracket\}$.

The notion of refinement allows us to capture the change in a preference pattern when accepting a normative rule expressed by $\mathbf{O}_R(\psi/\phi)$. Consider, for example, Sue's moral duty to keep her promise ($\mathbf{O}_R(p/\top)$) and her conditional moral duty to apologize if she does not keep her promise ($\mathbf{O}_R(a/\neg p)$). Accepting $\mathbf{O}_R(p/\top)$ has the effect of ranking all p -worlds over $\neg p$ -worlds. Accepting $\mathbf{O}_R(a/\neg p)$ has the effect of ranking all a & $\neg p$ -worlds over $\neg a$ & $\neg p$ -worlds. These

20. Here N_1 is an incoherent frame. See the definition of Coherence below.

effects on a preference pattern by accepting normative rules are expressed in terms of the refinements of a preference pattern.

Accepting $\mathbf{O}_R(p/\top) : N(W) \circ \llbracket p \rrbracket = \{\langle w, v \rangle \in N(W) \mid \text{if } v \in \llbracket p \rrbracket, w \in \llbracket p \rrbracket\}$.

Accepting $\mathbf{O}_R(a/\neg p) : N(\llbracket \neg p \rrbracket) \circ \llbracket a \rrbracket = \{\langle w, v \rangle \in N(\llbracket \neg p \rrbracket) \mid \text{if } v \in \llbracket a \rrbracket, w \in \llbracket a \rrbracket\}$.

We can also define the notion of refinement of a frame and its coherence as follows.

Definition 15 (Refinement of Frame) Let N, N' and $N_{\psi/\phi}$ be frames on W .

- (i) The frame N' is a refinement of N iff $N'(d) \subseteq N(d)$ for every $d \subseteq W$.
- (ii) The frame $N_{\psi/\phi}$ is the refinement of N such that for any $d \subseteq W$,
 - (a) if $d \neq \llbracket \phi \rrbracket$, $N_{\psi/\phi}$ assigns $N(d)$ to d ; and
 - (b) if $d = \llbracket \phi \rrbracket$, $N_{\psi/\phi}$ assigns $N(\llbracket \phi \rrbracket) \circ \llbracket \psi \rrbracket$ to d .

Definition 16 (Ideal and Optimal Worlds)

- (i) w is an *ideal* world in $\sigma = \langle N, s \rangle$ iff $w \in W$ and $w \preceq_{N(w)} v$ for every $v \in W$.
- (ii) **Ideal** $(\langle N, s \rangle)$ is the set of all ideal worlds in $\langle N, s \rangle$.
- (iii) w is an *optimal* world in $\sigma = \langle N, s \rangle$ iff $w \in s$ and there is no $v \in s$ such that $v \prec_{N(s)} w$.
- (iv) **Optimal** $(\langle N, s \rangle)$ is the set of all optimal worlds in $\langle N, s \rangle$.

Definition 17 (Coherence) Let N be a frame on W , and $d \subseteq W$.

- (i) w is a *best* world in $N(d)$ iff $w \in d$ and $w \preceq_{N(d)} v$ for every $v \in d$.
- (ii) **Best** $(N(d))$ is the set of all best worlds in $N(d)$.
- (iii) N is *coherent* iff for every non-empty $d \subseteq W$, **Best** $(N(d)) \neq \emptyset$.

Using these new devices now we can define the dynamic meanings of sentences in L_2^A .

Definition 18 (Update Function for Structured Information States) Any sentence χ of L_2^A is such that $\chi = \phi; = \mathbf{O}_R(\psi/\phi)$; or $= \mathbf{O}_J(\psi/\phi)$ for ϕ and ψ in L_0^A . For every sentence χ of L_2^A and an information state $\sigma = \langle N, s \rangle$, $\sigma[\chi]$ is determined as follows:

- (i) If $\chi = \phi$, then $\sigma[\chi] = \langle N, s[\chi] \rangle$, provided that $s[\chi] \neq \emptyset$.
- (ii) If $\chi = \mathbf{O}_R(\psi/\phi)$, then $\sigma[\chi] = \langle N_{\psi/\phi}, s \rangle$, where $N_{\psi/\phi}$ is coherent and $\llbracket \phi \rrbracket \cap \llbracket \psi \rrbracket \neq \emptyset; \mathbf{1}$, otherwise.²¹
- (iii) If $\chi = \mathbf{O}_J(\psi/\phi)$, then $\sigma[\chi] = \begin{cases} \sigma, & \text{if } \mathbf{Optimal}(\langle N, s[\phi] \rangle) \subseteq \llbracket \psi \rrbracket; \\ \mathbf{1}, & \text{otherwise} \end{cases}$

Note that updating $\sigma = \langle N, s \rangle$ with $\mathbf{O}_R(\psi/\phi)$ amounts to refining the frame N . Updating σ with $\mathbf{O}_J(\psi/\phi)$ amounts to testing whether the optimal worlds (most preferred worlds) in $s[\phi]$ given $N(s)$ are ψ -worlds. The update semantics is designed to capture an information state σ changes when an agent in an information state σ accepts ϕ . If $\sigma[\phi] \neq \mathbf{1}$, then ϕ is acceptable in σ ; if $\sigma[\phi] = \mathbf{1}$, then ϕ is not acceptable in σ . Likewise, a normative rule $\mathbf{O}_R(\psi/\chi)$ is acceptable only when the resulting refinement of a frame N is coherent.²²

If an agent is in the information state $\sigma[\mathbf{O}_J(q/p)] = \sigma$, it is adequate for the agent to accept or make the normative judgment that it ought to be the case that q if p ($\mathbf{O}_J(q/p)$). If an agent is in the information state $\sigma[\mathbf{O}_R(q/p)] = \sigma$, it is adequate for the agent to accept or utter a normative rule can be expressed by $\mathbf{O}_R(q/p)$.

Acceptance, Consequence and Consistency for a structured information state σ are defined similarly as before.

Definition 19 (Acceptance/Support)

$\sigma \Vdash \phi$ iff $\sigma[\phi] = \sigma$

Definition 20 (Dynamic Consequence)

$\phi_1, \phi_2, \dots, \phi_n \models \psi$ iff $\forall \sigma \in \Sigma$: if $\sigma[\phi_1][\phi_2] \dots [\phi_n] = \sigma$, then $\sigma[\psi] = \sigma$.

Definition 21 (Consistency)

A sequence of sentences, ϕ_1, \dots, ϕ_n is consistent iff $\exists \sigma : \sigma[\phi_1] \dots [\phi_n] \neq \mathbf{1}$.

This updated semantics for normative reasoning licenses the the characteristic inferences in normative reasoning (n1)–(n5). This can be shown as follows using acceptance: (n1*)–(n5*).

21. Given the definition of refinement, when $\llbracket \phi \rrbracket \cap \llbracket \psi \rrbracket \neq \emptyset$, $N_{\psi/\phi}$ is vacuously coherent given that N is coherent. Also, $\mathbf{O}_R(\psi/\phi)$ does not make much sense when $\llbracket \phi \rrbracket \cap \llbracket \psi \rrbracket \neq \emptyset$.

22. This does not mean that the agent in σ must reject new information ϕ when $\sigma[\phi] = \mathbf{1}$ and reject a new normative rule $\mathbf{O}_R(\psi/\phi)$ when the refinement of the frame with $\mathbf{O}_R(\psi/\phi)$ is not coherent. For the purpose of this paper and the sake of simplicity, however, accommodation in update semantics will not be discussed in this paper.

$$(n1^*) \mathbf{0}[\mathbf{O}_R(\phi)] \Vdash \mathbf{O}_J(\phi)$$

$$(n2^*) \sigma[\mathbf{O}_R(\phi)][\neg\phi] \not\vdash \mathbf{O}_J(\phi)$$

$$(n3^*) \sigma[\mathbf{O}_R(\phi)][\neg\phi] \Vdash \mathbf{O}_R(\phi)$$

$$(n4^*) \mathbf{0}[\mathbf{O}_R(\phi)][\neg\phi][\mathbf{O}_R(\psi/\neg\phi)] \Vdash \mathbf{O}_J(\psi)$$

$$(n5^*) \sigma[\mathbf{O}_R(\phi)][\neg\phi][\mathbf{O}_R(\psi/\neg\phi)] \Vdash \mathbf{O}_R(\phi)$$

These inferences show that acceptance is non-monotonic with normative and factual judgments, while normative rules are not sensitive to propositional updates.

The proposed update semantics validates some versions of Factual Detachment and Deontic Detachment for a coherent information state σ that allow us to infer new normative judgments, while new normative rules are not derived from either detachment rules.

Factual Detachment

$$\mathbf{O}_J(\psi/\phi), \phi \models \mathbf{O}_J(\psi)$$

$$\mathbf{O}_R(\psi/\phi), \phi \not\models \mathbf{O}_J(\psi)$$

Deontic Detachment

$$\mathbf{O}_J(\psi/\phi), \mathbf{O}_J(\phi) \models \mathbf{O}_J(\psi)$$

$$\mathbf{O}_R(\psi/\phi), \mathbf{O}_R(\phi) \not\models \mathbf{O}_J(\psi)$$

$$\mathbf{O}_R(\psi/\phi), \mathbf{O}_R(\phi) \not\models \mathbf{O}_R(\psi)$$

$$\mathbf{O}_J(\psi/\phi), \mathbf{O}_J(\phi) \not\models \mathbf{O}_R(\psi)$$

Agglomeration does not generally hold, but a conjunction of normative judgments is inferable in limited occasions.

Agglomeration

$$\mathbf{O}_J(\phi), \mathbf{O}_J(\psi) \models \mathbf{O}_J(\phi \wedge \psi)$$

$$\mathbf{O}_R(\phi), \mathbf{O}_R(\psi) \not\models \mathbf{O}_R(\phi \wedge \psi)$$

$$\mathbf{0}[\mathbf{O}_R(\phi)][\mathbf{O}_R(\psi)] \Vdash \mathbf{O}_J(\phi \wedge \psi)$$

The two kinds of *oughts* in this update semantics for normative reasoning play a critical role in explaining conflicting intuitions about normative inferences.

There is no general inference pattern in normative reasoning that allows us to derive new normative rules, while Factual Detachment, Deontic Detachment, and Limited Agglomeration allow us to infer new normative judgments. This result makes sense because an agent can infer what she practically ought to do from a given set of accepted normative rules and what is known, but an agent does not acquire new rules from what is already known to or accepted by her.

So far I have sketched a simple version of the update semantics for normative reasoning.²³ There are some details left to be filled in, but the basic picture of this dynamic approach must be clear enough to illustrate the intuitive and natural distinction between normative rules and normative judgments and solve the CTD Trilemma.


6.5. A Dynamic Solution to the CTD Trilemma

Let's see how the update semantics for normative reasoning handles Chisholm's example and solves the CTD Trilemma. We have seen that Jones's predicament in the CTD Situation is most naturally illustrated by the following five statements.

- (c1_R) $O_R(h/\top)$
- (c2_R) $O_R(t/h)$
- (c3_R) $O_R(\neg t/\neg h)$
- (c4) $\neg h$
- (c5_J) $O_J(\neg t/\top)$

In the update semantics for normative reasoning, accepting the normative rules (c1)–(c3) and the factual statement (c4) in the Chisholm set brings about the changes illustrated in the following diagram.

23. This is a simple version because it is a semantics for a limited language that does not deal with compound *ought*-statements. Also, it lacks some formal apparatus to deal with normative rules for exceptions. In this paper, this technical complication is suppressed since CTD obligations are not rules for exceptions. However, this system is easily expandable to deal with normative rules for exceptions. The relevant technical details are fully developed and discussed regarding default rules for exceptions in Veltman (1996).



In this diagram, σ_0 is the minimal state $\mathbf{0} = \langle N_0, W \rangle$. σ_4 is the resulting information state of updating (c1)–(c4) in the CTD situation:

$$\sigma_0 [\mathbf{O}_R(h/\top)] [\mathbf{O}_R(t/h)] [\mathbf{O}_R(\neg t/\neg h)] [t] = \sigma_4.$$

The resulting information state σ_4 supports not only (c1)–(c4), but also (c5) $\mathbf{O}_J(\neg t)$. $\sigma_4 [\mathbf{O}_J(\neg t)] = \sigma_4$ since $\mathbf{Optimal}(\sigma_4) \subseteq \llbracket \neg t \rrbracket$.²⁴ However, (c6) $\mathbf{O}_J(t)$, (c7) $\mathbf{O}_J(h \wedge \neg t)$ and (c8) $\mathbf{O}_J(t \wedge \neg t)$ are not supported by σ_4 . For $\sigma_4 = \langle N, s \rangle$, $\mathbf{Optimal}(\langle N, s \rangle) \not\subseteq \llbracket t \rrbracket$, $\not\subseteq \llbracket h \wedge \neg t \rrbracket$ and $\not\subseteq \llbracket t \wedge \neg t \rrbracket$.

$$\sigma_4 \Vdash \mathbf{O}_R(h), \mathbf{O}_R(t/h), \mathbf{O}_R(\neg t/\neg h), \text{ and } \neg h$$

$$\sigma_4 \Vdash \mathbf{O}_J(\neg t/\top)$$

$$\sigma_4 \not\Vdash \mathbf{O}_J(t), \mathbf{O}_J(h \wedge \neg t) \text{ and } \mathbf{O}_J(t \wedge \neg t)$$

The structured information state σ_4 , which is the result of updating the minimal state with (c1)–(c4), adequately represents Jones’s predicament in the CTD Situation and meets all three desiderata of the CTD Trilemma. Jones’s violation of his duty to go to help is expressed by the conjunction $\mathbf{O}_R(h) \wedge \neg h$; and this conjunction is supported by σ_4 : $\sigma_4 [\mathbf{O}_R(h)] [\neg h] = \sigma_4$. Second, σ_4 supports $\mathbf{O}_J(\neg t)$ since $\sigma [\mathbf{O}_R(\psi/\phi)] [\phi] \Vdash \mathbf{O}_J(\psi)$.²⁵ Finally, it is clear why (c6) $\mathbf{O}_J(t)$ is not supported

24. As we have stipulated before, $\mathbf{O}_R(\phi)$ and $\mathbf{O}_J(\phi)$ are abbreviations for dyadic notations $\mathbf{O}_R(\phi/\top)$ and $\mathbf{O}_J(\phi/\top)$.

25. In this system, both versions of Factual Detachments using normative rules and judgments are licensed for $\sigma = \langle N, s \rangle$ such that $s \cap \llbracket \phi \rrbracket \neq \emptyset$: $\mathbf{O}_R(\psi/\phi)$, $\phi \models \mathbf{O}_J(\psi)$ and $\mathbf{O}_J(\psi/\phi)$, $\phi \models \mathbf{O}_J(\psi)$.

by σ_4 and why conjoining $\mathbf{O}_R(t)$ and $\mathbf{O}_J(-t)$ is not legitimate in the update semantics for normative reasoning.

This dynamic account nicely explains why different normative judgments hold before and after accepting $\neg h$. Using this non-monotonic feature of the update semantics, we can give the temporal and agential explanations to Chisholm's Paradox in the dynamic semantic framework. Before being updated with $\neg t$, σ_3 supports $\mathbf{O}_J(h)$, $\mathbf{O}_J(t)$ and $\mathbf{O}_J(h \wedge t)$. But $\sigma_4 = \sigma_3[-t]$ supports none of these three normative judgments. Rather, it supports $\mathbf{O}_J(-t)$ and $\mathbf{O}_J(\neg h \wedge \neg t)$. Using this non-monotonic feature of the update semantics, we can also explain why (c1) as a normative judgment $\mathbf{O}_J(h/\top)$ is overridden in the Overriding Situation. Suppose an information state σ_6 that supports $\mathbf{O}_J(h/\top)$: $\sigma_6[\mathbf{O}_J(h/\top)] = \sigma_6$. This information state σ_6 stops supporting $\mathbf{O}_J(h/\top)$ once it is updated with $\neg h$.

$$\sigma_6 \Vdash \mathbf{O}_J(h/\top)$$

$$\sigma_6[-h] \not\Vdash \mathbf{O}_J(h/\top)$$

In the Overriding Situation (c1) is read as $\mathbf{O}_J(h)$ and thus overridden by (c4).

Unlike normative judgments, the accepted normative rules in an information state are not affected by additional propositional updates. They stay the same as we can see with σ_3 and σ_4 unlike factual and normative judgments.

$$\sigma_3 \Vdash \mathbf{O}_R(h/\top)$$

$$\sigma_3[-h] \Vdash \mathbf{O}_R(h/\top)$$

$$\sigma_3[-h] \not\Vdash \mathbf{O}_J(h/\top)$$

Finally, the information state σ_4 also supports some alternative readings of the CTD Situation. So far I have read all three *ought*-statements (c1)–(c3) in the CTD Situation as normative rules. But it is possible to read (c2) and (c3) as normative judgments: $\mathbf{O}_J(t/h)$ and $\mathbf{O}_J(\neg t/\neg h)$, and σ_4 supports these two normative judgments as well. The critical feature of the CTD Situation is that (c1) is a normative rule, and (c5) is a normative judgment following from (c3) and (c4). Maybe a slight hesitation in reading (c2) as a normative judgment would come from this formal result: the information state that supports $\neg h$ supports $\mathbf{O}_J(t/h)$ vacuously in that for $\sigma = \langle N, s \rangle$, $\mathbf{Optimal}(s[-h][h]) \subseteq \llbracket \neg t \rrbracket$ holds vacuously. Nevertheless, both readings of conditional *ought*-statements (c2) and (c3) are permitted in the CTD Situation, and our observations about the CTD Situation remain the same in either case.

6.6. Comparisons

The key move of my update semantics for normative reasoning is in understanding the meanings of *ought*-statements in normative discourse through the distinct functions of normative rules and judgments in normative reasoning. The prescriptive nature of normative rules and the descriptive nature of normative judgments are captured in two types of *ought*-statements in this dynamic semantics. And this two-*ought* account offers the most straightforward solution to the tension between (c1) and (c5) in Chisholm's example. The solution to the CTD Trilemma is not *ad hoc*. It not only captures the positive data (c1)–(c5), but also explains the negative data (c6)–(c8) systematically. With the dynamic meanings of the two types of *ought*-statements $O_R(\psi/\phi)$ and $O_J(\psi/\phi)$, we can gain both Factual Detachment and Deontic Detachment without generating a contradiction with Chisholm's example.

Before closing this paper, let me compare my new approach with two recently proposed solutions to Chisholm's Paradox that share my observations about Chisholm's example: Saint-Croix and Thomason's (2016; S&T hereafter) contextualist solution and Willer's (2014) dynamic solution. Specifically, they both aim to explain that (c1)–(c5) hold, but not (c6); and both (c1) and (c5) hold, but not (c7).

S&T (2016) present an unconditional version of Chisholm's Paradox, highlighting the oddity in the conjunction of (c1) and (c5). In Chisholm's example, both (c1) and (c5) are true, but the conjunction of (c1) and (c5) "is clearly false" (S&T 2016: 193). They solve this puzzling phenomenon by appealing to the context-sensitivity of conditionals and *ought*-statements in the ordering semantic framework. They claim that for the *ought*-statements (c1)–(c3) and (c5) to be non-vacuously true and felicitous they must be evaluated against different contexts. The context we assign to (c1) and the one we assign to (c5) are different and thus "cannot be unified into a single one" (S&T 2016: 206). The agglomeration of (c1) and (c5) is therefore not warranted. This radical contextualist approach concludes that "there is no single context that satisfies them all [(c1)–(c5)] appropriately" (S&T 2016: 205). It is a puzzling outcome given that (c1)–(c5) altogether faithfully describe Jones's predicament in the CTD Situation accurately at the same time. S&T justify the context-shifts in reading (c1)–(c5) by appealing to two pragmatic phenomena known as Lewis's Accommodation (1979) and Stalnaker's Assertion (1978). Another kink in this approach is the invalidity of Factual Detachment. The contextualist treatment of each sentence of the Chisholm set has nothing to do with explaining why Factual Detachment seems to be valid in Chisholm's example. To deal with the invalidity of Factual Detachment in ordering semantics, S&T appeal to the notion of pragmatic validity, concluding that Factual Detachment is invalid in ordering semantics, but pragmatically

valid; so (c5) is pragmatically derived from (c3) and (c4). Since this result brings back the familiar tension between (c5) and (c6), contextualists again must explain why (c5) and (c6) are not in conflict and explain why (c5) guides Jones's action but (c6) does not.²⁶

Now let's turn to Willer's dynamic semantics for *ought*-statements. Willer (2014) proposes a version of update semantics for *ought*-statements to solve Chisholm's Paradox without giving up on either Factual Detachment or Deontic Detachment. Willer constructs his semantic account of deontic necessity modal ($\Box_d\phi$) using update semantics as well, but his account is quite different from mine in that it does not see the two kinds of *ought*-statements in Chisholm's example. In his dynamic semantics there is only one type of deontic *ought*-statement ($\Box_d\phi$), which is comparable to $\mathbf{O}_J(\phi/\top)$ in my account.

$$\sigma[\Box_d\phi] = \{w \in \sigma : \sigma_d \models \phi\}$$

This reads: $\Box_d\phi$ is accepted in σ given d just in case ϕ is accepted by the subset of σ , which is the set of deontically ideal possible worlds in σ given d . Willer's non-monotonic update semantics validates both Factual Detachment and Deontic Detachment and thus explains why (c6) follows from (c1) and (c2), but not from (c1)–(c4); instead, (c5) follows from (c1)–(c4).

$$\Box_d h, h \Rightarrow \Box_d t \models \Box_d t$$

$$\Box_d h, h \Rightarrow \Box_d t, \neg h \Rightarrow \Box_d \neg t, \neg h \not\models \Box_d t$$

$$\Box_d h, h \Rightarrow \Box_d t, \neg h \Rightarrow \Box_d \neg t, \neg h \models \Box_d \neg t$$

However, this simple update semantics for deontic *ought*-statements has a problem with the persistent nature of (c1). The information state updated with (c1)–(c4) supports (c5), but not (c1) anymore due to its non-monotonicity.

$$\Box_d h, h \Rightarrow \Box_d t, \neg h \Rightarrow \Box_d \neg t, \neg h \not\models \Box_d h$$

To address this undesired result, Willer introduces a formal device *downdating* and proposes a revised updating rule for $\Box_d\phi$.

$$\sigma[\Box_d\phi] = \left\{w \in \sigma : \left(\sigma_d \downarrow \phi \downarrow \neg\phi\right)_d \models \phi\right\}$$

26. These issues are not discussed by S&T, but it is not hard to conjecture contextualist explanations for them.

Basically, downdating (\downarrow) forces an information state to make a necessary expansion to include both ϕ - and $\neg\phi$ -worlds in order to test whether ϕ is accepted by the set of deontically ideal possible worlds in the expanded σ given d . To justify this revision using downdating, Willer claims that sentences such as the following conditional should not be trivially true in a formal account: “If Jones goes to help his neighbors, then he ought to help his neighbors.” Given this update rule, it is possible for $\sigma[\neg\phi]$ to support $\Box_d\phi$. Therefore, there is a consistent information state that supports all (c1)–(c5): $\sigma[\Box_d h][h \Rightarrow \Box_d t][\neg h \Rightarrow \Box_d \neg t][\neg h] \models \Box_d \neg t$. With this revision using downdating, this account gets the positive data right and validates both Factual Detachment and Deontic Detachment.

A contextualist approach seems to be the best one can do to formally represent (c1)–(c5) without inconsistency in the ordering semantic framework, but at the price of theoretical elegance. Presenting (c1)–(c5) as true statements in the CTD situation and validating Factual Detachment do not seem to be two separate issues. But the contextualist treatment of individual *ought*-statements has nothing to do with the intuitive validity of Factual Detachment, and they have to be handled separately in the ordering semantic framework.

On the other hand, in Willer’s non-monotonic account of deontic modals the validity of both Factual Detachment and Deontic Detachment is naturally explained in terms of the update function of deontic modals and the dynamic consequence. And the validity of both detachment rules does not generate a contradiction with Chisholm’s example in this system. Due to its non-monotonicity, Willer’s dynamic account needs a major revision with the update function of deontic modals in order to capture the truth of (c1) in the CTD Situation. It seems that the consequences of using downdating may need to be more carefully examined. *Ought*-statements one might utter as a last resort such as “I ought to stay put because there are no options” seem to be simply false given the revised update function.²⁷

The update semantics for normative reasoning presented in this paper does not require some further device to explain why (c1) holds when (c4) holds because the persistent nature of *ought*-statements as normative rules is well encoded in the definition of $\mathbf{O}_R(\psi/\phi)$ in my account. Moreover, the introduction of two *ought*-statements in my update semantics are well-motivated by the natural understanding of how normative rules and judgments in our normative

27. Another technical concern I have about downdating is that it seems that this revision invites a problem with conditional *ought*-statements expressing exceptions. Consider, for example, this set of statements: (f1) Jones ought to go help his neighbors; (f2) If Jones has the flu, he ought not to go help his neighbors, and (f3) Jones has the flu. Given his revised update rule, there is neither an information state that supports all (f1)–(f3) nor an information state updated with these three statements that supports “Jones ought not to help his neighbors.”

reasoning interplay. Finally, here are a couple of interesting theoretical implications of the update semantics for normative reasoning. My update semantics for normative reasoning can be seen as a dynamic implementation of the ordering semantic account of *oughts* and *ifs*. As an information state σ in my expansion is formally represented by a set of possible worlds and preference ordering over them in the dynamic semantic framework, it is comparable to a context in ordering semantics defined by a pair of modal base and ordering source: $\langle f, g \rangle$. With that in mind, we can see that my account explains the mechanism through which a preference ordering changes as a discourse evolves. Any version of ordering semantics or dynamic semantics using the notion of ordering source like Willer's dynamic account are silent about how to form an ordering source and how it grows. The evolution of the ordering source is part of my update semantics, and the way $\mathbf{O}_R(\psi/\phi)$ refines a frame N is more fine-grained than simply adding $\phi \supset \psi$ into $g(i)$.

7. Conclusion

There have been many interpretations of Chisholm's example and discussions on what exactly are the problems in the example. In this paper, I have devised a new puzzle I call the CTD Trilemma focusing on the function of CTD obligations in our normative reasoning. This new puzzle not only reveals the limitations of the ordering semantic account of normative *ought*-statements but also guides us to a novel semantic account for normative reasoning. The CTD Situation where a CTD obligation takes effect is the situation where both the violated duty and a new obligation hold. So it turns out that a faithful representation of our reasoning in the CTD Situation requires not only an accurate account of conditional *ought*-statements, but also an acknowledgment of the dual nature of *ought*-statements in our normative reasoning. The update semantics for normative reasoning formally captures the two functions of *ought*-statements in a simple, unforced manner and handles the seemingly inconsistent observations about Chisholm's example. As a result, it provides a seamless formal account of the meaning of *ought*-statements and valid inferences in normative reasoning.

Some might think that the departure from the mainstream static semantics to a dynamic semantics is a big leap. But the proposed update semantics for normative reasoning preserves the core ideas in the ordering semantic account of *oughts* and *ifs*; that means it can be seen as a dynamic implementation of ordering semantics. More importantly, the theoretical benefits of my new formal approach and its intuitive framework should not be overlooked. Update semantics for normative reasoning has the potential to solve other deontic puzzles,

like Forrester's (1984) Paradox of Gentle Murder, Sartre's Dilemma (Lemmon 1962), and Jørgensen's (1937) Dilemma. In addition, the significance of dynamic approach deserves to be evaluated in the larger context of linguistic data, moral theories, and metaethics. It is widely recognized that the normative use of language demonstrates both prescriptive and descriptive features. Given this dual nature of normative language, a dynamic approach is an ideal semantic framework for normative language in general.

Acknowledgements

I thank Will Starr for reading numerous drafts of this paper, discussing many earlier ideas and various related topics with me and providing astute feedback on them. I also thank Janice Dowell, Matti Eklund, Julia Markovits, and Ergo's two anonymous referees for their extremely helpful feedback and suggestions that helped this paper take its current shape. Early ideas and ancestral versions of this paper were shared in many different forms. I would like to extend my gratitude to my friends, colleagues, commentators and audiences of workshops and conferences for their constructive feedback and discussion with me: Ana Arregui, Cory Brown, Mark Brown, Eric Epstein, Harold Hodes, Teresa Kouri, Thomas McKay, Ted Sider, Bokai Yao and the audiences at Cornell Semantics Group, ESSLLI 2015 Student Session, Cornell Philosophy Department Workshop, Logic Colloquium 2016, Conference on Contemporary Philosophy in East Asia 2016, "Logic and Language" Graduate Conference 2016 at University of Calgary and "Colloquium: Paradox" at Pacific APA 2017. Special thanks to Ted for all of his support.

References

- Al-hibri, Azizah Y. (1978). *Deontic Logic: A Comprehensive Appraisal and a New Proposal*. University Press of America.
- Åqvist, Lennart (1984). Deontic Logic. In Dov M. Gabbay and Franz Guenther (Eds.), *Handbook of Philosophical Logic. Volume II: Extensions of Classical Logic* (605–714). Springer Netherlands.
- Arregui, Ana (2008). Chisholm's Paradox in "Should"-Conditionals. In Tova Friedman and Satoshi Ito (Eds.), *Proceedings of SALT 18* (55–72). Linguistic Society of America.
- Beaver, David I. (2001). *Presupposition and Assertion in Dynamic Semantics* (Vol. 16). CSLI.
- Bronfman, Aaron and J. L. Dowell (2018). The Language of "Ought," and Reasons. In Daniel Star (Ed.), *Oxford Handbook of Reasons and Normativity* (85–112). Oxford University Press.
- Carr, Jennifer (2014). The If P, Ought P Problem. *Pacific Philosophical Quarterly*, 95(4), 555–583.

- Chisholm, Roderick M. (1963). Contrary-to-Duty Imperatives and Deontic Logic. *Analysis*, 24(2), 33–36.
- Feldman, Fred (1986). *Doing the Best We Can: An Essay in Informal Deontic Logic*. D. Reidel.
- Føllesdal, Dagfinn and Risto Hilpinen (1971). *Deontic Logic: An Introduction*. Springer.
- Forrester, James William (1984). Gentle Murder, or the Adverbial Samaritan. *The Journal of Philosophy*, 81(4), 193–197.
- Groenendijk, Jeroen and Martin Stokhof (1991). Dynamic Predicate Logic. *Linguistics and Philosophy*, 14(1), 39–100.
- Hansson, Bengt (1969). An Analysis of Some Deontic Logics. *Noûs*, 3(4), 373–398.
- Heim, Irene (1983). File Change Semantics and the Familiarity Theory of Definiteness. In Javier Gutiérrez-Rexach (Ed.), *Semantics: Critical Concepts in Linguistics* (108–135). Psychology Press.
- Jørgensen, Jörgen (1937). Imperatives and Logic. *Erkenntnis*, 7(4), 288–296.
- Kamp, Hans (1988). Discourse Representation Theory: What It Is and Where It Ought to Go. *Natural Language at the Computer*, 320(1), 84–111.
- Kamp, Hans, Josef Van Genabith, and Uwe Reyle (2011). Discourse Representation Theory. In Dov M. Gabbay (Ed.), *Handbook of Philosophical Logic* (125–394). Springer.
- Kolodny, Niko and John Macfarlane (2010). Ifs and Oughts. *The Journal of Philosophy*, 107(3), 115–143.
- Kratzer, Angelika (1977). What 'Must' and 'Can' Must and Can Mean. *Linguistics and Philosophy*, 1(3), 337–355.
- Kratzer, Angelika (1981). The Notional Category of Modality. In Hans-Jürgen Eikmeyer and Hannes Rieser (Eds.), *Words, Worlds, and Contexts: New Approaches in Word Semantics* (38–74). Walter de Gruyter.
- Kratzer, Angelika (1991). Modality. In A. von Stechow and D. Wunderlich (Eds.), *Semantics: An International Handbook of Contemporary Research* (Vol. 7, 639–650). Berlin.
- Lemmon, E. J. (1962). Moral Dilemma. *The Philosophical Review*, 71, 139–158.
- Lewis, David (1973). *Counterfactuals*. Harvard University Press.
- Lewis, David (1974). Semantic Analyses for Dyadic Deontic Logic. In Sören Stenlund, Ann-Mari Henschen-Dahlquist, Lars Lindahl, Lennart Nordenfelt, and Jan Odelstad (Eds.), *Logical Theory and Semantic Analysis* (1–14). Springer.
- Lewis, David (1979). Scorekeeping in a Language Game. In Rainer Bäuerle, Urs Egli, and Arnim von Stechow (Eds.), *Semantics from Different Points of View* (172–187). Springer.
- Lewis, David (1981). Ordering Semantics and Premise Semantics for Counterfactuals. *Journal of Philosophical Logic*, 10(2), 217–234.
- Loewer, Barry and Marvin Belzer (1983). Dyadic Deontic Detachment. *Synthese*, 54(2), 295–318.
- Mott, Peter L. (1973). On Chisholm's Paradox. *Journal of Philosophical Logic*, 2(2), 197–211.
- Portner, Paul (2009). *Modality* (Vol. 1). Oxford University Press.
- Prakken, Henry and Marek Sergot (1996). Contrary-to-Duty Obligations. *Studia Logica*, 57(1), 91–115.
- Prakken, Henry and Marek Sergot (1997). Dyadic Deontic Logic and Contrary-to-Duty Obligations. In Donald Nute (Ed.), *Defeasible Deontic Logic* (223–262). Springer.
- Saint-Croix, Catharine and Richmond H. Thomason (2016). Chisholm's Paradox and Conditional Oughts. *Journal of Logic and Computation*, 29(3), 369–386.
- Sayre-McCord, Geoffrey (1992). Normative Explanations. *Philosophical Perspectives*, 6, 55–71.

- Silk, Alex (2014). Why 'Ought' Detaches: Or, Why You Ought to Get with My Friends (If You Want to Be My Lover). *Philosophers' Imprint*, 14(7), 1–16.
- Stalnaker, Robert C. (1968). A Theory of Conditionals. In William L. Harper, Robert Stalnaker, and Glenn Pearce (Eds.), *Ifs* (41–55). Springer.
- Stalnaker, Robert C. (1975). Indicative Conditionals. *Philosophia*, 5(3), 269–286.
- Stalnaker, Robert C. (1978). Assertion. In Peter Cole (Ed.), *Pragmatics* (315–332). Brill.
- Starr, William B. (2016). Dynamic Expressivism about Deontic Modality. In Nate Charlow and Matthew Chrisman (Eds.), *Deontic Modality* (355–394). Oxford University Press.
- Straßer, Christian (2014). A Deontic Logic Framework Allowing for Factual Detachment. In *Adaptive Logics for Defeasible Reasoning: Applications in Argumentation, Normative Reasoning and Default Reasoning* (297–333). Springer.
- Veltman, Frank (1996). Defaults in Update Semantics. *Journal of Philosophical Logic*, 25, 221–261.
- Von Fintel, Kai and Anthony Gillies (2007). An Opinionated Guide to Epistemic Modality. In Tamar Szabo Gendler and John Hawthorne (Eds.), *Oxford Studies in Epistemology* (Vol. 2, 32–62). Oxford University Press.
- Von Wright, Georg Henrik (1951). Deontic Logic. *Mind*, 60(237), 1–15.
- Von Wright, Georg Henrik (1965). A Correction to a New System of Deontic Logic. *Danish Yearbook of Philosophy*, 2, 103–107.
- Von Wright, Georg Henrik (1970). A New System of Deontic Logic. In Risto Hilpinen (Ed.), *Deontic Logic: Introductory and Systematic Readings* (105–120). Springer.
- Willer, Malte (2013). Dynamics of Epistemic Modality. *Philosophical Review*, 122(1), 45–92.
- Willer, Malte (2014). Dynamic Thoughts on Ifs and Oughts. *Philosophers' Imprint*, 14(28), 1–30.