Trivalent Semantics for Conditional Obligations

Paul Égré

IRL Crossing, CNRS, Adelaide
ENS-PSL, Paris

Lorenzo Rossi

University of Turin
Palazzo Nuovo, Via Sant'Ottavio 20
10124 Torino, Italy

Jan Sprenger

University of Turin
Palazzo Nuovo, Via Sant'Ottavio 20
10124 Torino, Italy

Abstract

This paper provides a new framework for formalizing conditional obligations in natural language: it pairs a unary deontic operator with trivalent semantics for the indicative conditional and the assumption that the antecedents of conditionals restrict the scope of modals in the consequent. Combining these three ideas, we obtain a fully compositional theory of "if" and "ought" that validates plausible principles for deontic reasoning. Moreover, it addresses classical challenges such as the "if A then ought A" problem, the paradox of the miners, and the modeling of contrary-to-duty obligations (viz. Chisholm's quartet). All in all, our proposal provides a unified account of deontic modals and conditionals thatsquares well with general theories of natural language reasoning.

Keywords: conditional obligations, deontic modals, deontic reasoning, trivalent conditionals, Chisholms' quartet, paradox of the miners

1 Introduction

In natural language we often express *conditional obligations* by means of conditionals, as in

(1) If you help your neighbors, you should tell them.

Despite their surface form, conditional obligations are frequently formalized via a dyadic operator $O(B|A)$, which is then interpreted in a Kripke-style semantics, so that $O(B|A)$ is true (at the world of evaluation) if B is obligatory if we restrict attention to A-worlds [12,20,28,33]. Despite its widespread use, this approach has significant methodological and linguistic limitations [5,35].

First, a theory of compositional meaning should be *modular* and *parsimonious*: it should separate the contributions of "if" and "ought" and explain their interaction compositionally, instead of introducing an *ad hoc* dyadic operator [39].

Second, as argued by Kratzer [25,26,27], conditionals frequently restrict the scope of modals that occur in the consequent. A reasonable working hypothesis is therefore that there is a general mechanism whereby conditionals restrict modals (including deontic ones), and that we can derive an adequate account of conditional obligations from such a general mechanism. This would also help us to analyze the meaning of complex conditional sentences such as

(2) If John helps his neighbors, he will just show up, but actually, he should tell them.

In this paper, we offer a fully compositional account of the interaction of "if" and "ought" based on a trivalent semantics of indicative conditionals [14,17,16], a Kratzerian restriction mechanism and the evaluation of deontic modals relative to a context or informational state. For the rest, our semantics of "ought" remains standard.

The paper is structured as follows. In Section 2, we explain the trivalent semantics of conditionals and its interactions with deontic modals. Section 3 develops an account of logical consequence and observes some key results. Section 4 and 5 apply our account to the paradox of the miners and to Chisholm's puzzle about contrary-to-duty obligations. Section 6 draws the balance and explains the advantages of our approach over competing accounts.

All in all, combining trivalent, truth-functional semantics for conditionals with a Kratzerian restriction mechanism for modals in the consequent, we obtain a systematic, powerful and unified theory of conditional obligations and deontic reasoning that connects well with established semantic theories of conditionals and modals in natural language.

2 The Trivalent Theory of Conditionals

We work with a propositional language $\mathcal{L}$, whose vocabulary includes finitely many propositional variables (p_0, p_1, ...), the Boolean connectives $\neg$, $\wedge$, and $\vee$, and auxiliary symbols. $\mathcal{L}^{\rightarrow}$ is the extension of $\mathcal{L}$ with a conditional connective $\rightarrow$ representing the natural language indicative conditional "if ... then".

It has been frequently observed that the terms "true" and "false" have no clear ordinary sense when applied to indicative conditionals with false antecedents [1]. For example, it is not clear what determines the truth value of

(3) If Mary goes to the party, John will go, too.

when Mary does not go the party. The trivalent account gives a principled response to this question: the sentence is neither true nor false, but *void*—a third semantic value symbolized by $1/2$. Instead, if Mary goes to the party, the sentence takes the semantic value of the consequent. Symbolically: $A \to B$ is $1/2$ if A is 0, and takes the value of B otherwise.

In other words, we interpret conditionals as *conditional assertions*, akin to conditional bets. Suppose that Alice and Bob bet on the conditional (3): will John go to the party if Mary does? In order to be able to settle the bet, the precondition that Mary goes to the party has to be satisfied. In that case, the bet is on: Alice wins it John goes, and Bob wins otherwise. However, if Mary does not attend the party, the bet is off, and no-one can claim victory.

This basic intuition can be developed into a truth-functional theory of conditionals [10,13,30]. We use the trivalent truth tables proposed by Cooper in [10] for the conditional and the Boolean connectives $\neg$ and $\wedge$ as given in Table 1, with disjunction defined in the standard way as $A \vee B := \neg(\neg A \wedge \neg B)$.[1]

$f_{\to}$	1	$1/2$	0
1	1	$1/2$	0
$1/2$	1	$1/2$	0
0	$1/2$	$1/2$	$1/2$

	$f_{\neg}$
1	0
$1/2$	$1/2$
0	1

$f_{\wedge}$	1	$1/2$	0
1	1	1	0
$1/2$	1	$1/2$	0
0	0	0	0

Table 1
Cooper's truth table for the indicative conditional, paired with Strong Kleene negation and a modification of Strong Kleene conjunction ("quasi-conjunction").

The "quasi-conjunction" we adopt modifies the more familiar strong Kleene conjunction, by letting the conjunction of a classical and a non-classical value always take the classical value. So, for instance, $f_{\wedge}(1, 1/2) = 1$. This feature is required if we want to make sentences such as

(4) If the sun shines tomorrow, John goes to the beach; and if it rains, he goes to the museum.

true (in some possible worlds). We defer to previous publications for extensive philosophical defense of this semantics, and for its integration with theories of probability and epistemic modals [15,16,17].

The above tables generate the following valuation functions for formulae of $\mathcal{L}^{\to}$:

[1] Void antecedents are aligned with true rather than false antecedents for reasons discussed in [16]. However, since the conditional is the only source of the third truth value (we assume atom-classical valuations), this only affects expressions of the form $(A \to B) \to C$, which are outside the scope of this paper.

Definition 2.1 [Cooper valuations] A *Cooper valuation* is a valuation function $[\![]\!] : \mathcal{L}^{\rightarrow} \times W \longmapsto \{0, 1/2, 1\}$ that respects the truth tables from Table 1 and assigns classical truth values from the set $\{0, 1\}$ to all atomic formulas of $\mathcal{L}$ at a world $w \in W$.

Cooper valuations make convincing predictions for the evaluation of non-modal conditionals and their compounds. Moreover, we have developed a general template for extending this account to modal operators and their interaction with conditionals [17]. This extension agrees with Lewis's and Kratzer's observations that "if" often serves to restrict modal operators [26,27,29]. For example, in sentences such as

(5) If Richard is not in the office, then he will be at home.

(6) If Sarah leaves the party early, then she might be tired.

(7) If Teddy is not too busy, then he should help his neighbors.

the "will", the "might" and the "should" seem to quantify over the set of worlds specified by the antecedent. This is arguably the most natural way of expressing conditional modality in natural language. In [17], we show for epistemic modals how this restriction operation can be implemented while treating "if" as a trivalent binary connective. That is, we analyze the previous sentences as being of the form $A \rightarrow \mathsf{M}(B)$, with M standing for the modal operator. In this paper, we apply this template to deontic modals.

Let us introduce a unary operator O standing for "it is obligatory that" or simply "ought". The dual operator $\mathsf{P}(A)$ for "it is permissible that" is defined as $\neg\mathsf{O}(\neg A)$. To provide a semantics for a language containing O and P, we need three semantic parameters:

(1) a set of possible worlds W;
(2) a context or information state s indicating the *epistemically* possible worlds;
(3) a *deontic selection function* $d : \mathcal{P}(W) \longmapsto \mathcal{P}(W)$ that maps every context s to a *deontic context* $d(s)$ relevant for the evaluation of deontic modals at s, with the constraint $d(\emptyset) = \emptyset$.

While (1) is the basis of every modal semantics, (2) is standard in informational state semantics for epistemic modals and conditionals [19,44], and (3) expresses the idea that deontic modals are evaluated uniformly in a given context s. The idea is that your practical obligations depend on what you know about the world, i.e., the context you are in, but not on the exact world you happen to be in (unbeknownst to me). While not being completely uncontroversial, this approach has been defended by a number of prominent authors under the name of **Perspectivism** [11,22,23,37]. It is also often assumed, if implicitly, by the baseline theory of deontic modals in linguistics, in particular by authors who have followed Kratzer [25,26].

We evaluate deontic modals for obligation and permission in a standard way as quantification over a set of possible worlds. For $A \in \mathcal{L}$ and context s, an agent ought to A if A is true in the entire deontic context $d(s)$ (i.e., in all deontically acceptable worlds), and the agent may A at s if and only if A is true in at least one deontically acceptable world. Since the deontic selection function d is fixed for each model, we write $[\![\mathsf{O}(A)]\!]^{s,w}$ rather than $[\![\mathsf{O}(A)]\!]^{s,w,d}$, and similarly for other sentences:

$$[\![\mathsf{O}(A)]\!]^{s,w} = \begin{cases} 1 & \text{if } \forall w' \in d(s) : [\![A]\!]^{s,w'} = 1 \\ 0 & \text{otherwise.} \end{cases} \tag{1}$$

$$[\![\mathsf{P}(A)]\!]^{s,w} = \begin{cases} 1 & \text{if } \exists w' \in d(s) : [\![A]\!]^{s,w'} = 1 \\ 0 & \text{otherwise.} \end{cases} \tag{2}$$

When we let $A \in \mathcal{L}^{\rightarrow}$, these conditions expand to three values in the following way, mimicking a standard understanding of modal operators as expressing universal and existential quantification generalizing Cooper's trivalent semantics for conjunction and disjunction.

$$[\![\mathsf{O}(A)]\!]^{s,w} := \begin{cases} 1 & \text{if for all } w' \in d(s) : [\![A]\!]^{s,w'} \geq 1/2 \\ & \text{and for at least one } w'' \in d(s) : [\![A]\!]^{s,w''} = 1 \\ 1/2 & \text{if for all } w' \in d(s) : [\![A]\!]^{s,w'} = 1/2 \\ 0, & \text{if for a } w' \in d(s) : [\![A]\!]^{s,w'} = 0. \end{cases} \tag{3}$$

$$[\![\mathsf{P}(A)]\!]^{s,w} := \begin{cases} 1 & \text{if for a } w' \in d(s) : [\![A]\!]^{s,w'} = 1; \\ 1/2 & \text{if for all } w' \in d(s) : [\![A]\!]^{s,w'} = 1/2; \\ 0 & \text{if for all } w' \in d(s) : [\![A]\!]^{s,w'} \leq 1/2 \text{ and} \\ & \text{for at least one } w'' \in d(s) : [\![A]\!]^{s,w''} = 0. \end{cases} \tag{4}$$

Henceforth, $\mathcal{L}^{\rightarrow}_{\mathsf{O}}$ denotes the language generated by formulae of $\mathcal{L}^{\rightarrow}$ and their Boolean and conditional compounds with expressions of the form $\mathsf{O}(A)$ and $\mathsf{P}(A)$, for non-modal $A \in \mathcal{L}^{\rightarrow}$. For the sake of simplicity we exclude iterated deontic modals. [2]

Next we describe how deontic modals and conditionals interact. Following the template from [17], we hard-wire the Kratzerian restriction operation—antecedents of conditionals restrict the scope of modal operators—into the truth conditions of modal-conditional sentences. For this, we first need a definition of restricting and updating context:

Definition 2.2 [Restricted and Updated Contexts] Let s be a context and $A \in \mathcal{L}^{\rightarrow}_{\mathsf{O}}$. Then the *restriction* of s to A is defined as

$$s/A := \{w \in s : [\![A]\!]^{s,w} \geq 1/2\} \tag{5}$$

[2] An extension to all formulae of the modal-conditional language can be done following the general recipe in [17].

and the *update* of s on A is defined as

$$s[A] := \begin{cases} s/A & \text{if } \forall w \in s/A : [\![A]\!]^{s/A,w} \geq 1/2 \\ \emptyset & \text{if } \exists w \in s/A : [\![A]\!]^{s/A,w} = 0 \end{cases} \tag{6}$$

The reason why update and restriction have to be defined separately is that some modal sentences A may change their truth value after restricting the context to s/A. This is because deontic modals which were true relative to $d(s)$ may turn out false in the shifted deontic context $d(s/A)$. A minimal example is $A = p \wedge \mathsf{O}(\neg p)$ in a context $s = \{w_1, w_2\}$ with $[\![p]\!]^{s,w_1} = 1$, $[\![p]\!]^{s,w_2} = 0$ and $d(s) = \{w_2\}$. In this case $s/A = \{w_1\}$, but the restriction may shift the deontic context, so if $d(s/A) = s$, then $[\![\mathsf{O}(\neg p)]\!]^{s/A,w_1} = 0$, and consequently, $[\![A]\!]^{s/A,w_1} = 0$. We handle such degenerate cases by letting the updated context $s[A]$ collapse.

We can now define the general valuation clauses for conditional sentences with deontic operators ($A, B \in \mathcal{L}_{\mathsf{O}}^{\rightarrow}$). We evaluate $A \rightarrow B$ as void if A is false and as having the truth value of B if A is non-false. This is essentially the truth table for the trivalent conditional from Table 1. However, in the latter case, *we evaluate B at the deontic context $d(s[A])$ that corresponds to the original context s updated with the information that A is not false* (=essentially, restricted to the A-worlds). The distinction between these cases is made by the condition $w \in s[A]$, saying that w is an A-world and that the update on A does not collapse. Formally:

$$[\![A \rightarrow B]\!]^{s,w} = \begin{cases} [\![B]\!]^{s[A],w} & \text{if } w \in s[A] \\ 1/2 & \text{if } w \notin s[A] \end{cases} \tag{7}$$

With respect to proposals that model conditional obligations with dyadic deontic operators, an advantage is that we can straightforwardly evaluate sentences such as

(8) If John helps his neighbors, he just shows up, but actually, he should tell them.

The logical form of (8) is $H \rightarrow (S \wedge \mathsf{O}(T))$. Such sentences are hard to evaluate when conditional obligations are formalized with a dyadic operator, without reanalyzing the sentence into the more complex form $(H \rightarrow S) \wedge \mathsf{O}(T|H)$. On our semantics, this sentence is evaluated as void if John does not help his neighbors, and it takes the semantic value of $S \wedge \mathsf{O}(T)$ relative to the updated context $s[H]$ if John helps his neighbors. Similarly, by incorporating the evaluation clauses for epistemic modals [17], we could evaluate sentences mixing epistemic and deontic modals such as

(9) If John helps his neighbors, then he must have a lot of free time, but he should tell them that he comes.

Even though our semantics is not dynamic, the evaluation clause for conditionals mimics update operations that are typical of dynamic approaches to conditionals and modality. At the same time, it emulates the restricting function of Kratzer-style semantics without having to deny that "if" is a proper sentential connective (and it avoids postulating a covert modal in bare conditionals, as Kratzerian approaches do).

A classical problem of Kratzer-style semantics for conditionals and deontic modals is that they seem to validate the is-ought inference "if A, then ought A" [6]. This is because restricting the scope of the deontic modal in the consequent to A-worlds makes $\mathsf{O}(A)$ come out true. But clearly we do not want "if A, then ought A" to be a tautology of our semantics. (Consider the sentence "if I drink three beers a day, then I ought to drink three beers a day.") This problem affects Kratzer's baseline theory, and in principle any theory where conditionals restrict a semantic parameter that is relevant for evaluating a modal in the consequent [4,24,42,43].

A way out is to take the iffy ought in a sentence like $A \to \mathsf{O}(B)$ to be *doubly modalized* [27, pp. 106–107]: the antecedent A restricts the scope of a covert epistemic modal, and this epistemic modal regulates the scope of the deontic modal in the consequent. However, this solution does not explicate the mechanism for restricting the deontic modal in the consequent, and it may yield wrong predictions for a variety of sentences that express conditional obligations [6]. Something more needs to be said on how the scope of deontic modals depends on covert epistemic modals.

The above observations have two major consequences for our account. First, we should *not* require that in general, $d(s) \subseteq s$: in this case, since $s[A] \models A$ and by assumption $d(s[A]) \subseteq s[A]$, it would follow that $d(s[A]) \models A$, hence $\mathsf{O}(A)$ would be true at any world in $s[A]$ and the conditional $A \to \mathsf{O}(A)$ would never be false at any world. The requirement $d(s) \subseteq s$ is therefore too strong to square well with our trivalent restrictor semantics, for it would impose an "is-ought" collapse.

Second, we should avoid that $d(s) \cap s = \emptyset$: this would imply that sentences of the form $A \to O(\neg A)$ can be true at a world $w \in s$. But "if A then ought $\neg A$" expresses a conditional obligation that is impossible to satisfy given the supposition in the antecedent. Such sentences should not be true. Obligations seem to depend at least partially on the epistemic context. If there is no epistemically possible world where Jones *can* help his neighbors, then he is apparently *not* obliged to help them. We should not allow obligations to detach completely from knowledge and evidence, as in objectivist theories of "ought" [41]. Deontic selection functions with these features will be called *regular*.

Definition 2.3 [Regularity] A deontic selection function $d : \mathcal{P}(W) \longmapsto \mathcal{P}(W)$ is *regular* if for any $s \neq \emptyset$, $d(s) \cap s \neq \emptyset$.

Regularity implies in particular the seriality axiom schema D of deontic logic: $\neg(\mathsf{O}(A) \wedge \mathsf{O}(\neg A))$, barring obligation to contradictory actions. So in any non-empty context, obligation does not trivialize. To see that $A \to \mathsf{O}(A)$ can indeed come out false when assuming regularity, let $W = \{w_1, w_2\}$, where w_1 is a p-world, and w_2 a $\neg p$-world, with $s = \{w_1\}$ and $d(s) = \{w_1, w_2\}$. Then d is regular and $[\![p \rightarrow \mathsf{O}(p)]\!]^{s,w_1} = 0$. We can thus avoid the "if A, ought A" problem in a simple and elegant way, without increasing the complexity of our approach. At the same time, regularity guarantees that sentences of the form $A \rightarrow O(\neg A)$ can never be true, barring any contrary-to-possibility obligations.

This completes the exposition of our semantic framework. The models that satisfy the valuation clauses introduced so far, and that are based on regular deontic selection functions, are called *deontic Cooper-Kratzer models*.

Definition 2.4 [Deontic Cooper-Kratzer Models] An $\mathcal{L}_{\mathsf{O}}^{\rightarrow}$*-model* $\mathcal{M}$ is a quadruple $\langle W, s, d, [\![]\!] \rangle$ where W is a set of possible worlds, $s \subseteq W$ a context, $d : \mathcal{P}(W) \longmapsto \mathcal{P}(W)$ a regular deontic selection function and $[\![]\!]$ is a valuation function that assigns classical truth values to atomic formulae of $\mathcal{L}$ at any world, and is extended to $\mathcal{L}_{\mathsf{O}}^{\rightarrow}$ by means of the compositional rules in Table 1 and equations (3), (4) and (7).

3 Logical Consequence

How should logical consequence be defined when the underlying semantics is non-classical, as in our case? *A priori*, there are no obvious answers since a trivalent semantics allows for (strict) truth preservation, non-falsity preservation, or a mixture of both, and these logics may have different strengths and weaknesses [9,14].

We adopt a view of logical consequence that is tailor-made for natural language reasoning: whenever we accept all premises of an argument (i.e., we are certain of them or willing to assert them), we should also accept the conclusion. This notion of logical consequence was introduced under the label of "reasonable inference" by Robert Stalnaker:

> an inference [...] is *reasonable* just in case, in every context in which the premises could appropriately be asserted or supposed, it is impossible for anyone to accept the premises without committing himself to the conclusion ([38], p. 271)

Stalnalker's notion of reasonable inference treats consequence in terms of acceptance preservation, rather than truth preservation. This pragmatic view of logical consequence is now rather popular among semanticists and philosophers of language working on conditionals, especially in accounts that evaluate conditionals relative to information states and tie the role of logic to the preservation of structural features of information [3]. Santorio even claims that defining logical consequence along these lines is "the obvious notion of consequence for assessing consistency and validity for asserted claims in natural language" ([36], p. 81). In our trivalent framework, this idea corresponds to preservation of non-falsity *uniformly in a context*. For every sentence A and state s, let's say that s satisfies A in a given model (in symbols: $s \models A$) if and only if, for every $w \in s$, $[\![A]\!]^{s,w} \geq 1/2$. Then we define logical consequence by

Definition 3.1 [Logical Consequence] For $\Gamma \subseteq \mathcal{L}_{\mathsf{O}}^{\rightarrow}$ and $B \in \mathcal{L}_{\mathsf{O}}^{\rightarrow}$: $\Gamma \models B$ if and only if in all deontic Cooper-Kratzer models $\mathcal{M} = \langle W, s, d, [\![]\!] \rangle$:

$$s \models A \text{ for every } A \in \Gamma \qquad \Rightarrow \qquad s \models B. \tag{8}$$

This semantic notion of consequence can also be given a purely probabilistic presentation in terms of preserving probability 1, highlighting its informational motivation and the link to acceptance preservation (for details, see [17,16]). Furthermore, this consequence relation has the usual structural properties, i.e. Reflexivity ($A \models A$), Monotonicity (if $\Gamma \models B$, then $\Gamma, A \models B$) and Transitivity (if $\Gamma \models A$ for all $A \in \Sigma$ and $\Sigma, \Delta \models B$, then $\Gamma, \Delta \models B$). What about its valid inferences?

Our first observation is that any non-modal, non-conditional A does not imply $\mathsf{O}(A)$, but will imply $\mathsf{P}(A)$, i.e., it is not obligatory that $\neg A$. This is the inferential analogue of our observations about the truth values of $A \to \mathsf{O}(A)$ and $A \to \mathsf{O}(\neg A)$ in the previous section. Note that the proposition may fail when A is not Boolean, but a conditional.[3]

Proposition 3.2 (No Is-Ought Inference) *For Boolean $A \in \mathcal{L}$:*

$$A \not\models \mathsf{O}(A) \;\; \textit{and} \;\; A \models \mathsf{P}(A) \tag{9}$$

Second, by and large, the deontic operators behave like in standard deontic logic (SDL) [7,31].[4] Tautologies are obligatory ($\mathsf{O}(\top)$) and no contradiction is obligatory ($\neg\mathsf{O}(\bot)$). Crucially, obligations aggregate under conjunction. This implies that our logic does not permit conflicting obligations.

Fact 3.3 (Obligation Aggregation) *For non-modal $A, B \in \mathcal{L}^{\to}$:*

$$\mathsf{O}(A), \mathsf{O}(B) \models \mathsf{O}(A \wedge B) \tag{10}$$

Third, deontic reasoning is closed under logical consequence, but not under the conditional connective:

Proposition 3.4 (Closure under Consequence) *For $A, B \in \mathcal{L}^{\to}_{\mathsf{O}}$:*

$$\textit{If } A \models B \quad \textit{then} \quad \mathsf{O}(A) \models \mathsf{O}(B) \tag{11}$$

$$\mathsf{O}(A), A \to B \not\models \mathsf{O}(B) \tag{12}$$

Principle (11) also holds in SDL and is often called Monotonicity. It expresses the idea that an agent should "take moral responsibility for the logical consequences of what he/she has committed to do" [32]. By contrast, the inference stated in (12) fails, and reasonably so. To use a well-known example

[3] Suppose $A = p \to q$, $s = \{w\}$, $d(s) = \{w, w'\}$ with $[\![p]\!]^{s,w} = 0$, $[\![p]\!]^{s,w'} = 1$ and $[\![q]\!]^{s,w'} = 0$. Then $[\![p \to q]\!]^{s,w} = 1/2$ and $[\![p \to q]\!]^{s,w'} = 0$, and therefore $[\![\mathsf{P}(A)]\!]^{s,w} = 0$. While $s \models A$, it is not the case that $s \models \mathsf{P}(A)$.

[4] We do not have sufficient space to discuss the differences between our framework and SDL in detail. They appear most notably when considering the full language, including the conditional connective and iterated modals.

[40], if you ought to save as many lives as possible, but doing so implies, *in the actual epistemic context*, killing an innocent bystander, it does not follow that you ought to kill an innocent bystander. Indeed, SDL too fails the version of (12) with the material conditional, that is $\mathsf{O}(A), A \supset B \not\models \mathsf{O}(B)$.

Another important observation is that our logic validates Modus Ponens:

Fact 3.5 (Modus Ponens) *For all* $A, B \in \mathcal{L}_{\mathsf{O}}^{\rightarrow}$*:*

$$A \rightarrow B, A \models B \qquad \text{(Modus Ponens)}$$

As an immediate consequence, we observe that $\models$ also respects Factual and Deontic Detachment, two principles that have an important role in the discussion of Chisholm's puzzle about contrary-to-duty obligations [8]:

Corollary 3.6 (Factual and Deontic Detachment) *For non-modal* $A, B \in \mathcal{L}^{\rightarrow}$*:*

$$A \rightarrow \mathsf{O}(B), A \models \mathsf{O}(B) \qquad \text{(Factual Detachment)}$$

$$\mathsf{O}(A \rightarrow B), \mathsf{O}(A) \models \mathsf{O}(B) \qquad \text{(Deontic Detachment)}$$

This is a notable difference to Kratzer's baseline semantics, which invalidates Factual Detachment.

However, not all classical propositional inferences are valid: for example, Modus Tollens fails for conditional and modal consequents (for discussion, see [17,16]).

Fact 3.7 (Limited Modus Tollens) *For* $A, B \in \mathcal{L}$*:* $A \rightarrow B, \neg B \models \neg A$*, but this entailment fails conditional or modal* A *or* B*.*

This fact deserves a brief justification. In the modal-free fragment, a counterexample is the inference $p \rightarrow (q \rightarrow r), \neg(q \rightarrow r) \not\models \neg p$ (consider $[\![p]\!]^{s,w} = 1$ and $[\![q]\!]^{s,w} = 0$). Leaving aside iterated conditionals and turning to modals, it has been observed by Yalcin [45] that Modus Tollens is a problematic inference pattern for the interaction of conditionals and modals. Adapting one of his examples to the deontic setting, consider the following sentences:

(a) If Jones does not assist his neighbors, he ought to have a good reason.

(b) Jones does not need a good reason for not assisting his neighbors.

(c) Jones assists his neighbors.

It looks unwarranted to infer the factual conclusion (c) from accepting both the conditional (a) and the deontic modal (b).

Our conditional-deontic logic thus validates a surprisingly high number of plausible inferences and meta-inferences. We will now apply our theory to two classical test cases for theories of conditional obligation: the paradox of the miners and Chisholm's quartet.

4 Application 1: The Paradox of the Miners

Following [24], the paradox of the miners, originally due to Parfit, can be presented as follows:

	BA	BB	NB
A	$w_{A,BA}$	$w_{A,BB}$	$w_{A,NB}$
B	$w_{B,BA}$	$w_{B,BB}$	$w_{B,NB}$

Table 2
A minimal model of the paradox of the miners, which contains six possible worlds (=W) corresponding to combinations of the location of the miners (A, B) and our actions (BA, BB, NB). We assume $s = W$ and the abbreviation $NB := \neg BA \wedge \neg BB$ denotes that neither shaft is blocked. For the deontic selection function d we assume that $d(s) = \{w_{A,NB}, w_{B,NB}\}$, $d(s[A]) = \{w_{A,BA}\}$ and $d(s[B]) = \{w_{B,BB}\}$

> Ten miners are trapped either in shaft A or in shaft B, but we do not know which. Flood waters threaten to flood the shafts. We have enough sandbags to block one shaft, but not both. If we block one shaft, all the water will go into the other shaft, killing any miners inside it. If we block neither shaft, both shafts will fill halfway with water, and just one miner, the lowest in the shaft, will be killed.

Kolodny and MacFarlane assume that the following factual and deontic statements hold in this situation:

(1) We ought to block neither shaft. ($\mathsf{O}(\neg BA) \wedge \mathsf{O}(\neg BB)$)[5]

(2) If the miners are in shaft A, we ought to block shaft A. ($A \to \mathsf{O}(BA)$)

(3) If the miners are in shaft B, we ought to block shaft B. ($B \to \mathsf{O}(BB)$)

(4) The miners are in shaft A or shaft B. ($A \vee B$)

They argue that as long as we have a context-insensitive *ought*-operator, (2)–(4) entail (5), using the rules of Disjunction Introduction, Disjunction Elimination and Modus Ponens for the indicative conditional.

(5) We ought to block either shaft A or shaft B. ($\mathsf{O}(BA) \vee \mathsf{O}(BB)$)

But (5) contradicts (1), assuming our deontic logic is standard and does not allow for conflicting obligations. According to Kolodny and MacFarlane, rejecting Modus Ponens for the indicative conditional is the best way out of the paradox. This conclusion has not garnered much support in the literature (see, for example, [3,4,5]). And for us, it is not viable, since our conditional operator satisfies Modus Ponens (see Fact 3.5).

Our semantics avoids the paradox because the modal in $A \to \mathsf{O}(BA)$ is evaluated “Kratzer-style”, i.e., with respect to the deontic context $d(s[A])$ generated by the restricted context $s[A]$. The minimal model in Table 2 demonstrates the consistency of (1), (2), (3) and (4). The six possible worlds correspond to combinations of the location of the miners (A, B) and our actions (BA, BB, NB), and the subscripts indicate which sentences are true and which are false at a world. The abbreviation $NB := \neg BA \wedge \neg BB$ denotes that neither shaft is blocked.

[5] We can conjoin these obligations if desired, since our logic satisfies the principle of aggregating obligations.

Assuming that $s = W$, the four premises can be explicated as follows: (1) says that $d(s) \models \neg BA \wedge \neg BB$. (2) and (3) say that $d(s[A]) \models BA$ and $d(s[B]) \models BB$. Finally, (4) claims that $s \models A \vee B$. The above model satisfies (4) by construction and it is easy to see that it also satisfies (1)—(3) if we choose the deontic selection d such that $d(s) = \{w_{A,NB}, w_{B,NB}\}$, $d(s[A]) = \{w_{A,BA}\}$ and $d(s[B]) = \{w_{B,BB}\}$:

Proposition 4.1 (Consistency of the premises of the miners' paradox) *The deontic Cooper-Kratzer model $\langle W, s, d, [\![]\!] \rangle$ described in Table 2 and the above paragraph satisfies (1)–(4) at any world in $s = W$. In particular:*

$$\mathsf{O}(\neg BA), \mathsf{O}(\neg BB), A \to \mathsf{O}(BA), B \to \mathsf{O}(BB), A \vee B \not\models \bot. \tag{13}$$

One reason why we obtain consistency despite validating Factual Detachment (i.e., Modus Ponens for deontic consequents) is that we *invalidate* Reasoning by Cases (i.e., Disjunction Elimination).[6] This is typical of any "global" consequence relation that quantifies over contexts instead of worlds: we cannot infer from $s \models A \vee B$ that either $s \models A$ or $s \models B$, and so, from $A \models C$ and $B \models C$ it does not follow that $A \vee B \models C$. Applied to our example: we can infer $\mathsf{O}(BA) \vee \mathsf{O}(BB)$ both from A and from B, but we cannot infer that conclusion from $A \vee B$. This means that we cannot infer (5) from (2), (3) and (4) in the miners' paradox, similar to [18].

By contrast, this split is possible for "local" consequence relations that track truth at a single world of evaluation: $[\![A \vee B]\!]^w = 1$ implies either $[\![A]\!]^w = 1$ or $[\![B]\!]^w = 1$, and so Reasoning by Cases works.

Kolodny and MacFarlane reject this solution, because they argue that (1), (2) and (3) alone suffice to derive a contradiction, using Modus Tollens. But in our approach, Modus Tollens fails for modal consequents, as shown in Fact 3.7.[7] This failure can be motivated independently [17,45]. If we maintain the conditional obligation to block shaft A if the miners are there, and somebody tells us that we ought *not* to block shaft A, then this can be for two reasons: either this person believes that the miners are in shaft B, or she has no clue about their location and tells us not to block shaft A for this reason. But the latter case is perfectly compatible with the miners being in shaft A. So Modus Tollens *should* actually fail, as argued at the end of the preceding section.[8]

Finally, Kolodny and MacFarlane object that a context-sensitive interpretation of "ought" faces the following problem:

> (1), (2), and (3) will naturally occur in a single episode of deliberation. Why should it be that, in our paradoxical argument, (1) is used relative to the agent's current evidence, while (2) and (3) are used relative to a more informed body of evidence?

---

[6] For a related strategy, see also [18].

[7] Having unrestricted Modus Tollens would allow us to infer $\neg A$ from (1) and (2), and $\neg B$ from (1) and (3), using D, thereby contradicting (4).

[8] Moreover, we agree with Bledin ([4], p. 76) that "we must exercise caution when using proof by cases in languages that contain informational modals and the indicative."

> The contextualist owes an explanation of why in such cases there should always be a shift in the contextually relevant evidence.

Our answer is that (1), (2) and (3) are actually evaluated with respect to the same evidence. The truth conditions of the conditional imply that the antecedent restricts the scope of modals in the consequent, but the deontic conditional $A \to \mathsf{O}(BA)$ is still evaluated at a specific world w in the original context s. The context shifts for the purpose of evaluating the consequent, but not for evaluating the entire conditional.

5 Application 2: Contrary-to-duty obligations

Chisholm's quartet highlights a problem of adequately capturing contrary-to-duty obligations in a standard semantics for deontic modals. It consists of the following four sentences, which we provide with the standard formalizations [8,34]:

(a) $\mathsf{O}(G)$: Jones ought to go to assist his neighbors.

(b) $\mathsf{O}(G \to T)$: It ought to be that if he goes he tells them he is coming.

(c) $\neg G \to \mathsf{O}(\neg T)$: If he does not go, then he ought not to tell them he is coming.

(d) $\neg G$: He does not go.

When interpreting $A \to B$ as the material conditional $\neg A \vee B$, and sticking to unary deontic operators obeying the SDL principles, either independence or consistency of the four premises fails, regardless of whether the conditional obligations in (b) and (c) are taking wide or narrow scope [31]. Since our semantics for deontic operators and conditionals preserves SDL features like Factual/Deontic Detachment and agglomeration of obligations, one may wonder how it can handle the puzzle. Kratzer's baseline semantics for conditionals and modals avoids collapse because it invalidates Factual Detachment (i.e., the inference from (a) and (b) to $\mathsf{O}(T)$), but it has been argued that Chisholm-style paradoxes can be produced without assigning a central role to Factual Detachment [2,5].

In the standard analysis, independence and consistency are understood in terms of truth values *at a world*. That is, either the four sentences are not simultaneously satisfiable, or satisfying three of them enforces the truth of the remaining sentence.

We proceed in two steps. First, we show that (a)–(d) can simultaneously be true at a single world, and moreover, satisfying an arbitrary subset of them at w does not fix the truth values of the remaining sentences at w. This is arguably the original challenge posed by the paradox. Second, we consider (a)–(d) from the point of view of our notion of logical consequence that tracks acceptance in a context rather than truth at a world. We obtain that $\neg G$ and $\mathsf{O}(G)$ are inconsistent with each other, but we argue that this is a sensible result.

For the first step, consider the regular deontic Cooper-Kratzer model with a context that contains exactly four worlds: $s = \{w_{GT}, w_{G\neg T}, w_{\neg GT}, w_{\neg G\neg T}\}$, corresponding to the four combinations of the truth values of G and T.

Proposition 5.1 (Joint satisfiability and independence) *Assume that $s = W = \{w_{GT}, w_{G\neg T}, w_{\neg GT}, w_{\neg G\neg T}\}$, corresponding to valuations of the atomic sentences G and T. Then, there is a deontic selection function $d : \mathcal{P}(W) \longmapsto \mathcal{P}(W)$ such that in the deontic Cooper-Kratzer model $\langle W, s, d, [\![]\!]\rangle$, (a)–(d) are simultaneously true at any $w \in \{w_{\neg GT}, w_{\neg G\neg T}\}$:*

$$[\![\mathsf{O}(G)]\!]^{s,w} = [\![\mathsf{O}(G \to T)]\!]^{s,w} = [\![\neg G \to \mathsf{O}(\neg T)]\!]^{s,w} = [\![\neg G]\!]^{s,w} = 1 \qquad (14)$$

Moreover, the constraints expressed on actual world, context and deontic selection function by any subset of (a)–(d) are not sufficient to decide the semantic value of the remaining sentences.

The proof of satisfiability is insightful and so we provide it here. (a) imposes that $d(s) \models G$. (b) imposes that $w_{G\neg T} \notin d(s)$ (because $[\![G \to T]\!]^{s,w_{G\neg T}} = 0$). Since $d(s)$ must be non-empty, we infer that $d(s) = \{w_{GT}\}$, and this is sufficient to ensure the strict truth of (a) and (b). (c) imposes that $d(s[\neg G]) \models \neg T$ and (d) imposes that the actual world of evaluation $w_@$ is a $\neg G$-world: $w_@ = w_{\neg GT}$ or $w_@ = w_{\neg G\neg T}$. As a result, all sentences of Chisholm's quartet will be true at $(s, w_@)$. The reason why consistency can be maintained is that (a) and (b) make claims about what ought to be the case relative to the context s, (c) makes a claim about what ought to be the case in a *shifted context* $s[\neg G]$, and (d) makes a claim about the actual world in s.[9, 10]

By contrast, the second part of the proposition shows that the truth of three sentences among (a)—(d) does not fix the truth value of the fourth sentence. For example, we may assume that (a), (c) and (d) are true, which means that $d(s) \models G$, $d(s[\neg G]) \models \neg T$ and $w_@ \in \{w_{\neg GT}, w_{\neg G\neg T}\}$. This is compatible with $[\![\mathsf{O}(G \to T)]\!]^{s,w} = 1$ (as by the above model), but also with $[\![\mathsf{O}(G \to T)]\!]^{s,w} = 0$. It suffices to assume that $d(s) = \{w_{G\neg T}\}$. This small modification is sufficient to falsify (b) while maintaining the truth of the other three premises. Similarly, the modification $d(s) = \{w_{\neg G\neg T}, w_{GT}\}$ falsifies (a) but maintains the truth of (b)–(d). Finally, $d(s[\neg G]) = \{w_{\neg G\neg T}, w_{GT}\}$ falsifies (c) but maintains the truth of (a), (b) and (d) while $w_@ = w_{GT}$ falsifies (d) but maintains the truth of the deontic claims (a)–(c).

Summing up: there are deontic Cooper-Kratzer models where all four premises are satisfied, at a world, but also models where only three premises

[9] This diagnosis does not change if we interpret the deontic modal in (b) as having narrow scope, e.g., because it might look unnatural to let a deontic modal scope over a conditional [35,5]. In this case, (b) imposes the constraint $d(s[G]) \models T$, which is consistent with the other three premises.

[10] In general, wide and narrow scope deontic conditionals $\mathsf{O}(A \to B)$ and $A \to \mathsf{O}(B)$ do not always take the same semantic value—even if A and B are factual sentences. To see this, suppose that $[\![A]\!]^{s,w} = 1$. Then $[\![\mathsf{O}(A \to B)]\!]^{s,w} = [\![\mathsf{O}(B)]\!]^{s,w}$, but $[\![A \to \mathsf{O}(B)]\!]^{s,w} = [\![\mathsf{O}(B)]\!]^{s[A],w}$. However, what is obligatory in s need not coincide with what is obligatory in $s[A]$.

are satisfied and the fourth is false. This is a reasonable sense in which the premises are at the same time consistent and independent of each other.

However, we can also look at Chisholm's quartet from the point of view of valid inference (analyzed as *uniform* preservation of non-falsity at a context). Then premise (a) and (d) are sufficient to produce a contradiction (equation 15): $\neg G$ rules out $\mathsf{O}(G)$ because "ought" implies "can" (compare Proposition 3.2). On the other hand, the deontic and conditional-deontic claims in (a), (b) and (c) are compatible with any factual decision that Jones makes—i.e., they neither imply that Jones goes, nor that he does not go. We reproduce this partial logical indepedence in equation (16) and (17) below.

Proposition 5.2 (Partial Logical Consistency and Independence) *For factual and non-contradictory $G, T \in \mathcal{L}$:*

$$\neg G, \mathsf{O}(G) \models \bot \tag{15}$$
$$\mathsf{O}(G), \mathsf{O}(G \to T), \neg G \to \mathsf{O}(\neg T) \not\models G \tag{16}$$
$$\mathsf{O}(G), \mathsf{O}(G \to T), \neg G \to \mathsf{O}(\neg T) \not\models \neg G \tag{17}$$

Some readers may find (15) hard to swallow. To make it more palatable, note first that unlike in classical logic, $\mathsf{O}(G), \neg G \models \bot$ does *not* imply $\mathsf{O}(G) \models G$. "Ought" implies "can", but does not imply "is". That we cannot simultaneously accept $\neg G$ and $\mathsf{O}(G)$ in an epistemic context does not imply that we can reason from deontic to factual claims.

Second, equation (15) does *not* rule out that $\neg G$ and $\mathsf{O}(G)$ are simulatenously true at a world w in a context s. I may not help my neighbors, but have an obligation to do so. The consistency intuition in Chisholm's quartet arguably aims at joint satisfiability *at a world*, rather than at valid inference.

What equation (15) indeed rules out is that we jointly accept, in a context, that something is the case and that we have an *actual* obligation to do the contrary. Having an actual obligation presupposes the epistemic possibility of satisfying it, in line with Perspectivism. Once it is settled that Jones does not go, there is no meaningful (evidence-based) sense in which he should go. Going has simply ceased to be an option. Our obligations may be contrary to duty, but not contrary to possibility.

Finally, Proposition 5.2 implies that the contrary-to-duty obligation expressed by (c) does not clash with the unconditional obligations expressed by (a) and (b), or vice versa. Thus, we obtain a meaningful and non-trivial account of contrary-to-duty obligations.[11] In particular, we maintain central principles of standard deontic logic, such as agglomeration of obligations or Factual and Deontic Detachment. This is a notable difference to the major accounts in the literature where one of them is typically given up [2,25,26,35]. All in all, we can model obligations as context-dependent, explain why some conditional obligations depend on what is the case and assign a non-trivial role to contrary-to-duty oughts.

[11] We repeat that our logic is *not* conflict-tolerant: while we can reason from (a) and (b) to $\mathsf{O}(T)$ and from (c) and (d) to $\mathsf{O}(\neg T)$, already (a) and (d) are sufficient to produce a contradiction.

6 Conclusion

Our account of conditional obligations draws its inspiration from various sources: (i) Cooper's trivalent semantics for indicative conditionals; (ii) Lewis's and Kratzer's idea that "if"-sentences restrict the scope of modal operators in the consequent; (iii) the evaluation of deontic modals relative to a context rather than the precise world of evaluation; (iv) explicating logical consequence in terms of preserving acceptance at a context rather than truth at a world. Putting these four ideas together, which can be motivated independently, we obtain accurate and plausible predictions for two classical paradoxes of conditional obligation (the paradox of the miners and Chisholm's quartet). Furthermore, the "if A then ought A" problem is resolved. Moreover we retain the standard principles for deontic reasoning such as Factual Detachment / Modus Ponens, and the aggregation of obligations.

In other words, we offer an integrated logical theory of "iffy oughts" (i.e., conditional obligations, and the interaction of conditionals and modals) that squares well with our linguistic practices and with established ideas about the semantics of modals and conditionals in natural language. It is fully modular and compositional: the truth conditions of conditionals with deontic modals are derived from the truth conditions of conditionals and non-embedded modals.

This means that our account has numerous attractive features compared to rivaling approaches that model the natural language interaction of deontic and conditional structures. It does not require dyadic conditional obligations, but explains conditional obligation compositionally. The proposed mechanism for the interaction of conditionals and modals applies to various kinds of modals [17], while alternatives such as preference-based deontic logic do not offer such an integration [12,20,21]. Finally, with respect to Kratzer's baseline theory, our account retains "if" as a sentential connective (faithful to the linguistic form of conditionals), it preserves Factual Detachment (faithful to entrenched ideas about valid inference), and it proceeds without covert epistemic modals to give a smoother resolution of the classical deontic paradoxes. Future work might aim at an axiomatization of our account, and this would be helpful for detailed comparisons with alternative proposals, too.

Appendix

Proof [Proof of Proposition 3.2] Consider a context s where $s \models A$. Since d is regular, $d(s) \cap s \neq \emptyset$ and so, there is a world w in $d(s)$ where $[\![A]\!]^{s,w} = 1$. By the truth conditions for P, it follows that $s \models \mathsf{P}(A)$. Assume now that A is false anywhere outside s and that $d(s) \not\subseteq s$. In this case, $d(s) \not\models A$ and hence the is-ought inference $A \models \mathsf{O}(A)$ fails. □

Proof [Proof of Fact 3.3] Consider a context s with $s \models \mathsf{O}(A)$ and $s \models \mathsf{O}(B)$. Then $d(s) \models A$ and $d(s) \models B$. By the truth tables for conjunction, it follows that $d(s) \models A \wedge B$. This is equivalent to $s \models \mathsf{O}(A \wedge B)$. □

Proof [Proof of Proposition 3.4] To see the validity of (11), consider a context s and suppose that $s \models \mathsf{O}(A)$. Then $d(s) \models A$, and since $A \models B$, also $d(s) \models B$. But this means that $s \models \mathsf{O}(B)$ and hence $\mathsf{O}(A) \models \mathsf{O}(B)$. For the failure of (12), consider $W = \{w_1, w_2\}$, $s = \{w_1\}$, and $d(s) = W$ with $A := \top$ and $B := p$, with $[\![p]\!]^{s,w_1} = 1$, $[\![p]\!]^{s,w_2} = 0$. Then $s \models \mathsf{O}(\mathsf{A})$ (because $d(s) \models \top$) and $s \models A \to B$, but clearly, $s \not\models \mathsf{O}(B)$. □

Proof [Proof of Fact 3.5] Suppose that we have a deontic Cooper-Kratzer model $\langle W, s, d, [\![]\!] \rangle$ with $s \models A \to B$ and $s \models A$. Then also $s[A] = s$ and $1/2 \leq [\![A \to B]\!]^{s,w} = [\![B]\!]^{s[A],w} = [\![B]\!]^{s,w}$ for all $w \in s$. Therefore $s \models B$. □

Proof [Proof of Corollary 3.6] Factual Detachment is an immediate consequence of the validity of Modus Ponens. For Deontic Detachment, suppose that $s \models \mathsf{O}(A)$ and $s \models \mathsf{O}(A \to B)$. This means that $d(s) \models A$ and $d(s) \models A \to B$. Therefore, by Modus Ponens also $d(s) \models B$, and $s \models \mathsf{O}(B)$, as requested. □

Proof [Proof of Fact 3.7] Take $W = s = d(s) = \{w_1, w_2\}$ with $[\![p]\!]^{s,w_1} = 1$, $[\![p]\!]^{s,w_2} = 0$ and $d(s[p]) = s[p] = \{w_1\}$. In this case, $s \models p \to \mathsf{O}(p)$ and $s \models \neg\mathsf{O}(p)$, but $s \not\models \neg p$. □

Proof [Proof of Proposition 5.2] For $\neg G, \mathsf{O}(G) \models \bot$, note that $d(s) \cap s \neq \emptyset$ and therefore, if $s \models \neg G$, then $\mathsf{O}(G)$ will scope over some $\neg G$-worlds. Therefore $s \not\models \mathsf{O}(G)$.

For the two remaining claims, take the model from Proposition 5.1 and note that $s \models \mathsf{O}(G)$, $s \models \mathsf{O}(G \to T)$, and $s \models \neg G \to \mathsf{O}(\neg T)$. The first two claims are purely deontic claims and therefore uniformly true at s, while the third claim is void at G-worlds and true at $\neg G$-worlds (by construction, $d(s[\neg G]) \models \neg T$). But clearly, s contains both G- and $\neg G$-worlds and therefore neither $s \models G$ nor $s \models \neg G$. □

References

[1] Adams, E. W., *The Logic of Conditionals*, Inquiry **8** (1965), pp. 166–197.

[2] Arregui, A., *Detaching if-clauses from should*, Natural Language Semantics **18** (2010), pp. 241–293.

[3] Bledin, J., *Logic informed*, Mind **123** (2014), pp. 277–316.

[4] Bledin, J., *Modus ponens defended*, Journal of Philosophy **112** (2015), pp. 57–83.

[5] Cariani, F., *Deontic logic and natural language*, in: D. Gabbay, R. van der Meyden, J. Horty, X. Parent and L. van der Torre, editors, *The Handbook of Deontic Logic, Vol. 2*, College Publications, 2021 pp. 499–548.

[6] Carr, J., *The "If P, Ought P" problem*, Pacific Philosophical Quarterly **95** (2014), pp. 555–583.

[7] Chellas, B. F., "Modal Logic: An Introduction," Cambridge University Press, Cambridge, 1980.

[8] Chisholm, R. M., *Contrary-to-duty imperatives and deontic logic*, Analysis **24** (1963), pp. 33–36.

[9] Cobreros, P., P. Égré, D. Ripley and R. van Rooij, *Tolerant, classical, strict*, Journal of Philosophical Logic **41** (2012), pp. 347–85.

[10] Cooper, W. S., *The propositional logic of ordinary discourse*, Inquiry **11** (1968), pp. 295–320.

[11] Dancy, J., "Practical Reality," Oxford University Press, Oxford, 2000.

[12] Danielsson, S., "Preference and Obligation. Studies in the Logic of Ethics," Filosofiska Faereningen, 1968.

[13] de Finetti, B., *La logique de la probabilité*, in: *Actes du congrès international de philosophie scientifique*, Hermann, Paris, 1936 pp. 565–573.

[14] Égré, P., L. Rossi and J. Sprenger, *De Finettian Logics of Indicative Conditionals. Part I: Trivalent Semantics and Validity*, Journal of Philosophical Logic **50** (2021), pp. 187–213.

[15] Égré, P., L. Rossi and J. Sprenger, *Probability for Trivalent Conditionals* (2024), manuscript available at `https://philpapers.org/rec/GRTCTM`.

[16] Égré, P., L. Rossi and J. Sprenger, *Certain and Uncertain Inference with Trivalent Conditionals*, Australasian Journal of Philosophy (2025).

[17] Égré, P., L. Rossi and J. Sprenger, *Trivalent conditionals, Kratzer style* (2025), manuscript available at `https://philarchive.org/rec/EGRTCK`.

[18] Gabbay, D., L. Robaldo, X. Sun, L. van der Torre and Z. Baniasadi, *Toward a linguistic interpretation of deontic paradoxes*, in: F. Cariani, D. Grossi, J. Meheus and X. Parent, editors, *Deontic Logic and Normative Systems* (2014), pp. 108–123.

[19] Gillies, A. S., *On Truth-Conditions for If (but Not Quite Only If)*, Philosophical Review **118** (2009), pp. 325–349.

[20] Hansson, B., *An analysis of some deontic logics*, Noûs **3** (1969), pp. 373–398.

[21] Hansson, S. O., *Preference-Based Deontic Logic (PDL)*, Journal of Philosophical Logic **19** (1990), pp. 75–93.

[22] Kiesewetter, B., *Ought and the perspective of the agent*, The Journal of Ethics and Social Philosophy **5** (2011), pp. 1–24.

[23] Kiesewetter, B., "The Normativity of Rationality," Oxford University Press, Oxford, 2017.

[24] Kolodny, N. and J. MacFarlane, *Ifs and oughts*, Journal of Philosophy **107** (2010), pp. 115–143.

[25] Kratzer, A., *Partition and revision: The semantics of counterfactuals*, Journal of Philosophical Logic **10** (1981), pp. 201–216.

[26] Kratzer, A., *Conditionals*, Chicago Linguistics Society **22** (1986), pp. 1–15.

[27] Kratzer, A., "Modals and Conditionals. New and Revised Perspectives," Oxford University Press, 2012.

[28] Lewis, D., *Semantic analyses for dyadic deontic logic*, in: S. Kanger and S. Stenlund, editors, *Logical theory and semantic analysis: essays dedicated to Stig Kanger on his 50thbirthday*, Reidel, Boston, 1974 pp. 1–14.

[29] Lewis, D., *Adverbs of quantification*, in: E. Keenan, editor, *Formal Semantics of Natural Language*, Cambridge University Press, 1975 pp. 3–15.

[30] McDermott, M., *On the truth conditions of certain if-sentences*, Philosophical Review **105** (1996), pp. 1–37.

[31] McNamara, P. and F. van de Putte, *Deontic logic*, in: E. N. Zalta and U. Nodelman, editors, *The Stanford Encyclopedia of Philosophy*, Metaphysics Research Lab, Stanford University, 2022, Fall 2022 edition .

[32] Nute, D. and X. Yu, *Introduction*, in: D. Nute, editor, *Defeasible deontic logic*, Kluwer Academic Publishers, Dordrecht pp. 1–16.

[33] Prakken, H. and M. Sergot, *Dyadic deontic logic and contrary-to-duty obligations*, in: D. Nute, editor, *Defeasible deontic logic*, Springer, 1997 pp. 223–262.

[34] Rönnedal, D., *Contrary-to-duty paradox*, Internet Encyclopedia of Philosophy (2025).

[35] Saint Croix, C. and R. Thomason, *Chisholm's paradox and conditional oughts*, Lecture Notes in Computer Science **8554** (2014), pp. 192–207.

[36] Santorio, P., *Path semantics for indicative conditionals*, Mind **131** (2022), pp. 59–98.

[37] Scanlon, T. M., "Moral Dimensions: Permissibility, Meaning, Blame," Belknap Press, Cambridge/MA, 2008.

[38] Stalnaker, R., *Indicative Conditionals*, Philosophia **5** (1975), pp. 269–286.

[39] Thomason, R., *Deontic logic as founded on tense logic*, in: R. Hilpinen, editor, *New Studies in Deontic Logic: Norms, Actions, and the Foundations of Ethics*, Wiley-Blackwell, 1981 pp. 165–176.

[40] Thomson, J. J., *Killing, letting die, and the trolley problem*, The Monist (1976), pp. 204–217.

[41] Thomson, J. J., *Imposing risks*, in: W. Parent, editor, *Rights, Restitution, and Risk*, Harvard University Press, Cambridge/MA, 1986 pp. 173–191.

[42] Willer, M., *Dynamic thoughts on ifs and oughts*, Philosophers' Imprint **14** (2014), pp. 1–30.

[43] Willer, M., *Dynamic foundations for deontic logic*, in: N. Charlow and M. Chrisman, editors, *Deontic Modality*, Oxford University Press, Oxford, 2016 pp. 324–354.
[44] Yalcin, S., *Epistemic modals*, Mind **116** (2007), pp. 983–1026.
[45] Yalcin, S., *A counterexample to modus tollens*, Journal of Philosophical Logic **41** (2012), pp. 1001–1024.