Chisholm's Paradox and Conditional Oughts

Catharine Saint-Croix and Richmond H. Thomason

[1] Philosophy Department, University of Michigan
[2] Philosophy Department, University of Michigan

Abstract. Since it was presented in 1963, Chisholm's paradox has attracted constant attention in the deontic logic literature, but without the emergence of any definitive solution. We claim this is due to its having no single solution. The paradox actually presents many challenges to the formalization of deontic statements, including (1) context-sensitivity of unconditional oughts, (2) formalizing conditional oughts, and (3) distinguishing generic from nongeneric oughts. Using the practical interpretation of 'ought' as a guideline, we propose a linguistically motivated logical solution to each of these problems, and explain the relation of the solution to the problem of contrary-to-duty obligations.

1 Introduction: Chisholm's paradox

[5] formulates the problem of contrary-to-duty obligations with the following four-sentence example.[3]

(1a) It ought to be that Jones go to the assistance of his neighbors.
(1b) It ought to be that if he does go he tells them he is coming.
(1c) If he does not go, then he ought not to tell them he is coming.
(1d) He does not go.

Although the language in (1a) and (1b) is somewhat stilted and unnatural, there is nothing uncommon about the situation it describes. Frequently, a secondary obligation results from the violation of a primary obligation. Such situations are only called paradoxical because they are difficult to formalize.

We begin with an unconditional version of the Chisholm quartet. We argue that the central problem consists in the fact that 'ought' is context-sensitive, rather than conditionals. Nevertheless, and independently, the interaction of 'ought' with conditionals remains to be accounted for. With this in mind, we shift in Section 5 to conditional obligation. Contrasting two types of detachment—factual and deontic—we argue that the former is much more important in practical reasoning. This accords well with the idea that conditional oughts are narrow-scope deontic conditionals. In Section 5.3, we remove a problem for this analysis by contextualizing the conditional. We then discuss how the contextualization of conditionals and 'ought' interacts in the interpretation of conditional oughts.

The difficulty remains of accounting for deontic detachment. This inference is useful for some important purposes, and usually is explained by allowing 'ought'

[3] This is Chisholm's exact wording, except that we have substituted 'Jones' for 'a certain man' in (1a).

to have wide scope over the conditional. We argue that there is independent evidence that deontic statements can be generic, and that generic conditional oughts can support deontic detachment. There is, therefore, no need to postulate scope ambiguities in conditional oughts.

Finally, in Section 6 we illustrate our account of context, 'ought', and conditionals by returning to Chisholm's paradox.

2 An unconditional version

Steps (1b) and (1c) of Chisholm's example are conditional oughts—conditionals whose main clause involves an 'ought'. Like Chisholm himself, many authors ([1], [2], [17], [20], for instance) have felt that this example reveals the inadequacy of naive formalizations of conditional obligation, and that an adequate logic of conditional obligation will solve the problem.

This can't be entirely right, because of examples like (2), which are similar to (1) but do not involve the conditional.

(2a) Jones ought to assist his neighbors.
(2b) But he will not go.
(2c) So he ought to tell his neighbors he is not going to assist them.
(2d) Therefore, he ought to assist his neighbors and to tell them he is not going to assist them.

(2d) follows from (2a) and (2c) in standard deontic logics. But, although (2a–c) appear to be mutually consistent, and provide a plausible description of the Chisholm scenario, (2d) is clearly false.

We don't deny that the conditional version of Chisholm's paradox illustrates logical difficulties having to do with conditional obligation. But the unconditional version reveals a more fundamental problem that needs to be cleared up before turning to the conditional case.

3 A methodology

Work in deontic logic tends to concentrate on examples with moral overtones, like promise-keeping. But 'ought' has many uses. If we assume, with [11], that these uses differ only in the sort of possibilities that are in play, these differences will not affect the underlying logic. Practical or prudential uses of 'ought' provide intuitions that in general are crisper than moral uses, and moreover are readily restricted to simple domains or scenarios.

With these considerations in mind, we propose to use a game that we'll call *Heads Up* as a laboratory for testing deontic intuitions. A number of playing cards are set down side by side. A player, Jones, gets to choose a card. The player's payoffs are dependent on whether he chooses a face card. Simple versions of this game will involve just one choice, while more complicated versions will involve successive choices.

4 Unconditional oughts

We begin with a simple three-card version of *Heads Up*. In this game, there will always be at least one face card and at least one non-face card on the table. If Jones chooses a face card, he gets \$50. Otherwise, he gets nothing.

Suppose Jones is presented with $\langle \mathsf{Jack}, 3, 9\rangle$ as a layout. Clearly, he ought to choose the leftmost card, ought not to choose the middle card, and ought not to choose the rightmost card. Similarly, if presented with $\langle \mathsf{Queen}, \mathsf{King}, 4\rangle$, Jones ought not to choose the rightmost card. And it's not the case that he ought to choose the leftmost card, since choosing the middle card would result in the same payoff. So, he ought to choose *either* left or middle.

Following Lewis' semantics for deontic operators ([14, 15]), $\bigcirc\phi$ is true just in case there is an outcome u satisfying ϕ such that any outcome at least as good as u also satisfies ϕ. Under the $\langle \mathsf{Queen}, \mathsf{King}, 4\rangle$ layout, consider $\mathsf{ChooseLeftOrMiddle}$. Both outcomes satisfying this formula reward \$50, while the only other outcome rewards \$0. So $\bigcirc\mathsf{ChooseLeftOrMiddle}$ is true. $\bigcirc\mathsf{ChooseLeft}$, however, turns out to be false. This is because there is another outcome (middle), that doesn't satisfy $\mathsf{ChooseLeft}$ but has a reward that is just as good.

Deontic logicians routinely formalize such things by augmenting a boolean propositional logic with $\bigcirc$ as a modal operator, so that for any formula ϕ, $\bigcirc\phi$ is also a formula. Let a Lewis frame be a structure $\mathcal{F} = \langle W, \preceq, f\rangle$, where W is a non-empty set (of outcomes), $\preceq$ is a preorder over W, and f is a function taking members of W to subsets of W (picking out a set of deontic alternatives from w). A model over $\mathcal{F}$ is a structure $\langle W, \preceq, f, V\rangle$, where V is a function taking propositional atoms to subsets of W.

Definition 1. Satisfaction in a model for $\bigcirc\phi$.

Given a model $\mathfrak{M}$ and a world $w \in W$, $\mathfrak{M}, w \models \bigcirc\phi$ **iff** there is some $u \in f(w)$ such that $\mathfrak{M}, v \models \phi$ for all $v \in f(w)$ such that $v \preceq u$.

By allowing f to pick out deontic alternatives from w, we can restrict the outcomes evaluated to those that are relevant given the situation. When the layout is $\langle \mathsf{Queen}, \mathsf{King}, 4\rangle$, we don't need to check irrelevant possibilities in which Jones chooses the left card and it isn't a Queen.

4.1 Knowledge of circumstances and uncertainty

So far, we have said nothing about what Jones knows: *Heads Up* doesn't specify whether the cards are dealt face down. If they are dealt face up, Jones will know all of the relevant information, but otherwise we can't ignore his epistemic state. Suppose now that the cards are dealt face down and that the layout is $\langle \mathsf{Queen}, 5, 2\rangle$. Perhaps surprisingly, it is still natural to say that Jones ought to choose the leftmost card, even though he doesn't *know* this.

Suppose the dealer were to say, "Jones, there's no card you ought to choose." The only natural interpretation of this would be that there is no unique face card on the table. The dealer *can't mean* that although there is a unique face card, Jones doesn't know this, or doesn't know where it is. Practical oughts act, in fact, as if they want to ignore the agent's epistemic situation, even though, of course, agents are trapped by the limits of their knowledge in making decisions.

4.2 Oughts and optimality

The conflicting thought that oughts should track something like expected utility is tempting. To get a sense of this, consider a case in which, after many rounds of a face down version of *Heads Up*, Jones has begun to sense a bias in the dealer—it seems significantly more likely that there will be a face card on the left than in any other position. In such situations, it is natural for Jones to think something like, "I ought to choose the leftmost card."

Now, suppose Jones plays the next round. He turns over the leftmost card and finds that in fact, it's a two. The comment we would then expect is "Damn! I ought to have chosen a different card!" rather than, "Damn! I ought to have chosen this card, but it was a loser!"

Even when expected utility guides our choices, it seems that it's still natural to use 'ought' as if it's tracking the optimal action.

If this is so, how can Jones think to himself, "I ought to choose the leftmost card" as he makes the decision? Actually, this is expected and natural, and does not in itself require a different, expected utility tracking sense of 'ought'. When Jones says this, he believes he ought to choose the card on the left. When the cards are revealed, he realizes that this belief was false: he did what he incorrectly *thought* he ought to do. Thus, we typically don't describe lapses of this sort as cases in which, although we did what we ought to have done, this turned out to be suboptimal. We describe them as cases in which we have mistaken—perhaps with good reason—what we ought in fact to do.

With a certain amount of effort, however, by raising the stakes and heightening the impact of ignorance, we can manage to detach practical 'ought's from optimality and relativize them to the agent's knowledge. The "miners' paradox" scenario presented in [10] illustrates the point.

4.3 Solving the unconditional paradox

Confine attention now to objective oughts that ignore the agent's epistemic state. In the *Heads Up* domain this means, among other things, that if $w' \in f(w)$ then w and w' do not differ in their factual circumstances (that is, in the layout of cards), although they can differ in the choices that the agent makes. Now, consider another version of *Heads Up*, with just two cards in the layout and where the player will make two successive choices. The payoffs are: \$500 if a face card is chosen in both rounds, \$0 if a face card is chosen in the first round but not the second, \$250 if a face card is chosen only in the second round, and, finally, \$350 if no face card is chosen in either round. The relevant choices are now

L_1: Choose the left card on the first turn
L_2: Choose the left card on the second turn
R_1: Choose the right card on the first turn
R_2: Choose the right card on the second turn

Suppose the layout is $\langle \mathsf{Queen}, 2 \rangle$. Since choosing left allows Jones to reach the optimal outcome, \$500, $\bigcirc\mathsf{L}_1$ is true in our model: $\models \bigcirc\mathsf{L}_1$. If he does choose left, $\models \bigcirc\mathsf{L}_2$: Jones should choose left in the second round, since the outcomes available after L_1 are \$500 for L_2 and \$0 for R_2.

But now suppose that Jones performs poorly in the first round and chooses the card on the right. Then the outcomes available to Jones are different. In light of his poor first choice, Jones' payoffs are \$250 for L_2 and \$350 for R_2. With these options, then, $\models \bigcirc\mathsf{R}_2$. Thus, we have:

(3a) $\bigcirc\mathsf{L}_1$ is true.
(3b) But $\neg\mathsf{L}_1$ is true.
(3c) So, $\bigcirc\mathsf{R}_2$ is true

But of course, $\bigcirc\mathsf{L}_1$ and $\bigcirc\mathsf{R}_2$ aren't *simultaneously* true: that would lead to the very worst \$0 payoff. The assertion of (3b) has changed the context.

At this point, it has become clear that the satisfaction relation is context-dependent. To represent examples like (3a–c), we must make this explicit. We modify the function f of a Lewis frame to take a context set X of worlds into account: $f(w, X) \subseteq X$. We then define *contextualized satisfaction* by relativizing satisfaction to this set X. For practical oughts, X represents the set of alternatives presumed to be open to choice.

Definition 2. Contextualized satisfaction in a model for $\bigcirc\phi$.
Given a Lewis frame $\mathcal{F} = \langle W, \preceq, f \rangle$, a model $\mathfrak{M}$ on $\mathcal{F}$, a subset X of W, and a world $w \in W$, $\mathfrak{M}, X, w \models \bigcirc\phi$ **iff** there is some $u \in f(w, X)$ such that $\mathfrak{M}, X, v \models \phi$ for all $v \in f(w, X)$ such that $v \preceq u$.

Using contextualized satisfaction, (3a–c) can now be stated as follows.

(4a) $\mathfrak{M}, X, w \models \bigcirc\mathsf{L}_1$
(4b) But $\mathfrak{M}, X, w \models \neg\mathsf{L}_1$
(4c) So, $\mathfrak{M}, X \cap [\![\neg\mathsf{L}_1]\!], w \models \bigcirc\mathsf{R}_2$

Here $[\![\neg\mathsf{L}_1]\!]$ is the set of worlds satisfying $\neg\mathsf{L}_1$.

Thus, the assertion in (4b) adds its content to the context set used to interpret the next step, (4c). This dynamic effect is similar to the effect of assertion on presupposition noted in Stalnaker's [22] and similar writings.

The scenario in (4) closely resembles (2), the unconditional version of Chisholm's paradox. This is easier to see from the following English version of (4).

(5a) Jones ought to choose left on round one.
(5b) But he will not.
(5c) So he ought to choose right on round two.

The sensitivity of 'ought' and other modals to a context set was noted by Kratzer in [11] and has become a standard part of linguistic theories of modals;

see [18, 12]. Thus, this contextual solution to the unconditional Chisholm paradox is well motivated in terms of linguistic theories of modals and assertion.

5 Conditional oughts

We now turn to the conditional version of Chisholm's paradox. Here, the crucial issue is how to formalize statements of conditional obligation.

At the outset, there are two approaches to this issue: (i) take conditional oughts to be primitive 2-place modalities $\bigcirc(\psi/\phi)$, and (ii) take them to be compositional combinations of 'ought' and 'if'. Although many linguists and logicians take the first approach, a compositional theory is always to be preferred, when it's compatible with the evidence.

With this in mind, we will not consider Approach (i) here, but will explore Approach (ii), with the thought that a linguistically adequate compositional theory will render Approach (i) unnecessary. Our hypothesis, then, is that conditional oughts involve the same conditional that figures generally in other conditional constructions, with or without modals.

We begin with a neutral conditional '$\dashrightarrow$', making no assumptions about its logic for the moment. Assuming compositionality, there are two options for formalizing 'If ϕ then ought ψ': (i) wide-scope $\bigcirc$, $\bigcirc(\phi \dashrightarrow \psi)$ and (ii) narrow-scope $\bigcirc$, $\phi \dashrightarrow \bigcirc\psi$.

Chisholm's somewhat tortured phrasing in (1) suggests that both are involved, with (1b) taking wide-scope and (1c) narrow. With this in mind, we might formalize (1) as follows:

(6a) $\bigcirc\mathsf{Help}$
(6b) $\bigcirc(\mathsf{Help} \dashrightarrow \mathsf{Tell})$
(6c) $\neg\mathsf{Help} \dashrightarrow \bigcirc\neg\mathsf{Tell}$
(6d) $\neg\mathsf{Help}$

This formalization uses both the wide-scope formalization of conditional 'ought' (6b) and the narrow-scope (6c). Ultimately, we will question the formalization of (6b), but this contrast provides a useful way to frame the important issues, which have to do with detachment. (See [9].)

5.1 Detachment

In Example (6) we want to conclude that Jones ought to tell his neighbors he isn't coming to help, $\bigcirc\neg\mathsf{Tell}$. Now, with few exceptions, logics of the conditional allow us to infer ψ from $\phi \dashrightarrow \psi$ and ϕ. And with fewer exceptions, deontic logics validate the inference from $\bigcirc(\phi \dashrightarrow \psi)$ and $\bigcirc\phi$ to $\bigcirc\psi$. There are, then, two ways to conclude an 'ought' statement:[4]

Definition 3. *Factual deontic detachment* (FDD)
Infer $\bigcirc\psi$ from $\phi \dashrightarrow \bigcirc\psi$ and ϕ.

[4] Following [7], these are often referred to as simply *factual detachment* and *deontic detachment*. We prefer this somewhat awkward scheme because it emphasizes that both are types of deontic detachment.

Definition 4. *Deontic deontic detachment* (DDD)
Infer $\bigcirc\psi$ from $\bigcirc(\phi \dashrightarrow \psi)$ and $\bigcirc\phi$.

Factual deontic detachment provides the inference of $\bigcirc\neg\mathsf{Tell}$ from (6c) and (6d), and this inference is welcome in the example. The difficulty is that deontic deontic detachment allows us to infer $\bigcirc\mathsf{Tell}$ from (6b) and (6a), and this conclusion is not so welcome.

5.2 Factual deontic detachment in *Heads Up*

Let's refine our intuitions by returning to a simple version of *Heads Up*. In this version, we still have just two cards in the layout, exactly one of which will be a face card, and there is only a single round.

Intuitively, the following deontic conditionals are true in this scenario:

(7a) If the left card is a queen, Jones should choose it.
(7b) If the right card is a queen, Jones should choose it.
(7c) If the left card is a two, Jones should not choose it.
(7d) If the right card is a two, Jones should not choose it.

Here, the intuitions in favor of FDD are very powerful. As soon as we learn, for instance, that the left card is a queen, we think that Jones should choose it. This supports formalizations of (7a–d) as narrow-scope deontic conditionals $\phi \dashrightarrow \bigcirc\psi$, as indeed the language suggests.

5.3 Factual deontic detachment as *modus ponens*

Superficially, FDD may appear to be very simple if it merely amounts to using *modus ponens* with a narrow-scope deontic conditional. But if we look more carefully at this matter in model-theoretic terms, the matter is more complex.

For definiteness, we will work from now on with Stalnaker's semantics for the conditional ([22]).[5]

And at the same time we adopt Stalnaker's notation, $>$, for the conditional.

[5] We mention at this point [17] and [25], which also propose $\phi \dashrightarrow \bigcirc\psi$ as a formalization of conditional oughts, but in quite different ways. On the one hand, [17] does not allow at all for informational effects on $\bigcirc$, and incorporates a world-shifting approach to conditionals, borrowed from [14]. But [25], on the other hand, concentrates on information state updates, and ignores world-shifting entirely.

Our own approach seeks to motivate independent semantic accounts of the conditional and of 'ought' and to formalize conditional obligation by simply combining these two accounts. Following the work of Stalnaker and Lewis, and beginning with non-informational conditionals like 'If this match is struck it will light', we believe that conditionals must involve a world-shifting element. Like Stalnaker, we favor a conditional that validates Conditional Excluded Middle. But we add an informational ingredient, which also is shifted by conditional antecedents and can affect the interpretation of information-sensitive consequents. Unlike the currently popular dynamic approaches of [6] and [25], we do not endorse a logic with a dynamic consequence relation. Instead, information dynamics is embedded in the satisfaction conditions for conditionals and deontic $\bigcirc$.

Stalnaker invokes a "selection function" s from propositions or sets of worlds to sets of worlds. This function satisfies the following conditions.

(8a) For all $w \in W$ and $Y \subseteq W$ there is a $u \in W$ such that $s(Y, w) \subseteq \{u\}$.
(8b) $s(Y, w) \subseteq Y$.
(8c) $s(Y, w) = \{w\}$ if $w \in Y$.
(8d) If $s(Y, w) = \emptyset$ then $Y = \emptyset$.
(8e) If $s(Y, w) \subseteq Y'$ and $s(Y', w) \subseteq Y$ then $Y = Y'$.

To interpret the conditional $>$, we add the function s to our frames and add the following satisfaction clause. (Because in Section 4 we decided to relativize satisfaction to contexts, this definition contains a parameter 'X'.)

Definition 5. Naive satisfaction for $>$.
Given a model $\mathfrak{M}$, a context set X, and a world $w \in W$,
$\mathfrak{M}, X, w \models \phi > \psi$ **iff** if $s(\phi, w) = \{u\}$ then $\mathfrak{M}, X, u \models \psi$.

Here, $s(\phi, w)$ is $s([\![\phi]\!], w)$.

It is easy to verify that this definition validates *modus ponens*, so that it supports FDD with narrow-scope formalizations of conditional obligation. This depends crucially on Stalnaker's "centering" condition (8c).

But Definition 5 creates problems in formalizing reparational obligation. Suppose, to return to (1), that we formalize (1c), 'If he does not go, then he ought not to tell them he is coming,' as follows.

(9) $\neg\mathsf{Help} > \bigcirc\neg\mathsf{Tell}$

According to the naive satisfaction condition in Definition 5, (9) is true in a world w where, say, Jones has promised to help his neighbors, iff in the closest world u where Jones does not help his neighbors he ought to tell them he will not come to help them. But in this world u, the factual circumstances, including Jones' promise, remain the same—only Jones' choice has changed. (Otherwise, u would not be the closest world.) So in this world u, just as in w, Jones ought to help his neighbors because, in view of his promise, he helps them in the best alternatives to u. Of course, in those alternatives, $\neg\mathsf{Tell}$ is false. So, u isn't an $\bigcirc\neg\mathsf{Tell}$-world. Thus, (9) turns out to be false, and instead the rather pointless (10) is true.

(10) $\neg\mathsf{Help} > \bigcirc\mathsf{Help}$

This conditional, amounting to 'If Jones does not help his neighbors, then he (still) ought to help them' may make some sense as an admonition, but it is impractical and certainly doesn't correspond to our intuitions concerning secondary obligations.

To solve this problem, we replace the naive satisfaction clause for $>$ with a more sophisticated version that contextualizes the conditional as well as $\bigcirc$. First, we make the selection function s sensitive to context, so that s now inputs a set of worlds (the antecedent proposition), another set of worlds X (the context), and a world, and, as before, returns a unit set of worlds or the empty set. Our new satisfaction condition is this:[6]

Definition 6. Contextualized satisfaction for $>$.
Given a model $\mathfrak{M}$, a context set X, and a world $w \in W$,
$\mathfrak{M}, X, w \models \phi > \psi$ **iff** if $s(\phi, X, w) = \{u\}$ then $\mathfrak{M}, X \cap [\![\phi]\!], u \models \psi$.

This cumulative satisfaction clause adds the antecedent proposition to the context in which the consequent is evaluated. Only worlds satisfying the antecedent are to be taken into account in evaluating a consequent $\bigcirc\psi$. And this solves the problem of formalizing secondary obligations. (9), for instance, is true, because in the closest world u where John doesn't help his neighbors, $\bigcirc\neg\mathsf{Tell}$ is true in the best options in which the background facts are assumed, as well as the proposition that Jones will not help his neighbors. In other words, under the first definition, (9) failed because $\mathfrak{M}, X, v \not\models \bigcirc\neg\mathsf{Tell}$ for the best alternatives, v in X. Under the second we use $X \cap [\![\neg\mathsf{Help}]\!]$ instead of X, and (9) is true.

On the other hand, *modus ponens* is no longer valid on this interpretation of $>$ (see [24] for details), so that FDD is threatened.

It may seem at this point that it is difficult or even impossible to retain (i) a compositional account of the conditional oughts involved in FDD, (ii) a compositional account of the conditional oughts involved in stating secondary obligations, and (iii) a logical endorsement of FDD.

However, we do not believe that things are as bad as this. FDD is *pragmatically valid*, in a sense that was first introduced in [22]. Although the inference

(11a) ϕ
(11b) $\phi > \bigcirc\psi$
(11c) $\therefore\ \bigcirc\psi$

is invalid when all three terms are evaluated with respect to the same context X, (11c) will be true if the first step adds its proposition to the context in which the subsequent steps are validated. In other words, $\bigcirc\psi$ follows from *the assertion* of ϕ and $\phi > \bigcirc\psi$. We feel that pragmatic validity provides an adequate account of the very strong intuitions that favor FDD, and also serves the purposes of FDD in practical reasoning, allowing detachment when the minor premise has been learned and added to the background context.

5.4 Deontic deontic detachment

Treating conditional oughts compositionality has the apparent advantage of providing a natural formalization for DDD, by making a wide-scope logical form $\bigcirc(\phi > \psi)$ available. If the underlying conditional validates *modus ponens* and $\bigcirc$ is a modal operator, then DDD with wide-scope is just the modal principle K.

But in fact, this idea is not well supported on linguistic grounds. We know of no convincing evidence that in normal conditional sentences with 'ought' or 'should' in the main clause, the modal takes wide scope over the conditional. The

[6] [24] argues that—independently of deontic considerations—this condition provides an improvement on Stalnaker's semantics for the conditional.

problem is most evident with permission. For instance, contrast the following two formalizations of ‘If you are sick, you may take the day off’.

(12a) $\mathsf{Sick} > \mathsf{PDayOff}$
(12b.) $\mathsf{P}(\mathsf{Sick} > \mathsf{DayOff})$

(12a) captures the intended meaning well, making permission conditional on illness. It is hard, though, to see what sensible meaning, if any (12b) could have, or how to render this formula at all in English.

Rather than multiplying scope ambiguities, we prefer to adopt the working assumption that ‘ought’ and ‘should’ rarely take wide scope over the conditional. The view, then, is that there is nothing special about the use of ‘if’ in statements of conditional obligation—these differ from other conditionals only in the deontic content of the main clause.

5.5 Oughts for practical occasions

In practical deliberation, an agent is faced with a set of choices: the purpose of deliberation is to restrict these alternatives. If the deliberation process is successful, it will deliver a set of alternatives that are indifferent and optimal, so that one of them can be selected arbitrarily.[7]

Many kinds of preferences can figure in deliberation: idealizing, we suppose that the agent has prioritized these dimensions of preference and has identified among these certain core preferences that can serve as absolute criteria for acceptability: any alternative that is not optimal with respect to these core preferences is unacceptable. Conditional and unconditional oughts based on these core preferences can structure the deliberative process by filtering the initial set of alternatives: an agent can use them to eliminate at the outset of deliberation any possibility that doesn’t belong to this set.

In decisions with full information—cases where the agent has complete knowledge of the relevant factors—only unconditional oughts are needed. That is illustrated by our first *Heads up* example, at the begining of Section 4. Jones can see the layout $\langle \mathsf{Queen}, \mathsf{King}, 4 \rangle$. He may have whimsical preferences—say, for red rather than black cards—but these are dominated by the payoffs of the game and can be set aside at the outset. Because the payoffs satisfy $\bigcirc\mathsf{ChooseLeftOrMiddle}$, the alternative $\mathsf{ChooseRight}$ is eliminated at the outset of deliberation.

Conditional oughts come into play when there is lack of knowledge or uncertainty. Imagine a face-down version of *Heads up*, in which it’s understood that there is one and only one face card in each layout. The cards are turned over successively, and Jones is allowed to choose or reject each card in turn.

If he wishes to plan in advance, his plans must be contingent, because of uncertainty. Before any cards are turned, he can use conditionals and say:

(13a) “If the first card is a face card, I should choose it.”
(13b) “If the first card is not a face card, I should reject it.”

[7] This account, of course, idealizes away from considerations having to do with limited rationality.

Armed with these conditionals, he begins to play the game. When the first card is turned over, FDD creates an unconditional obligation that can guide the first choice.

Even when an agent has complete information, conditional oughts can be useful in formulating *policies*. In these formulations, deontic conditionals typically are in subjunctive mood.[8] Return to our original version of *Heads Up* and suppose, for instance, that Jones is looking at $\langle \mathsf{Queen}, 4, \mathsf{King} \rangle$, but is interested not only in the decision at hand but in what he should do under slightly different circumstances. He might say

(14) "If the right card were a six, I ought to choose the left card."

Entertaining possibilities even further from the actual layout, he could also say

(15) "If the right card were a six and the middle card were a jack and the left card were an eight, I ought to choose the left card."

Left-nonmonotonicity effects such as this—failures of Strengthening the Antecedent—are features of the classical logics of conditionals, and follow from our claim that conditional oughts are conditionals. These effects are intended and, we believe, unavoidable. But they do mean that a certain amount of caution is needed in selecting the conditional oughts that apply to a given practical situation.

For instance, knowing that he has received a bill for cleaning a carpet, Jones may entertain two conditional oughts: he ought to pay the bill if he has received it, but he ought *not* to pay it if he has received it and the carpet has not, in fact, been cleaned.[9]

We formalize the conflicting conditional oughts as narrow-scope deontic conditionals:

(16a) $\mathsf{BillRcvd} > \bigcirc\mathsf{PayBill}$
(16b) $(\mathsf{BillRcvd} \wedge \neg\mathsf{CarpetCleaned}) > \bigcirc\neg\mathsf{PayBill}$

The following validities constrain the possibilities in this case.[10]

(17a) $\mathsf{BillRcvd}, \mathsf{BillRcvd} > \bigcirc\mathsf{PayBill} \;\therefore\; \bigcirc\mathsf{PayBill}$
(17b) $\mathsf{BillRcvd}, \bigcirc\mathsf{PayBill}, (\mathsf{BillRcvd} \wedge \neg\mathsf{CarpetCleaned}) > \bigcirc\neg\mathsf{PayBill}$
$\therefore\; \mathsf{CarpetCleaned}$
(17c) $\mathsf{BillRcvd}, \neg\mathsf{CarpetCleaned}, (\mathsf{BillRcvd} \wedge \neg\mathsf{CarpetCleaned}) > \bigcirc\neg\mathsf{PayBill}$
$\therefore\; \neg(\mathsf{BillRcvd} > \bigcirc\mathsf{PayBill})$

[8] The differences between subjunctive and indicative conditional oughts—which, following Stalnaker, we assume to be pragmatic—are not important for our present purposes.

[9] We'll discuss how conditional oughts such as these might come to be entertained below, in Section 5.6.

[10] (17a) is an instance of FDD and, as we explained, is pragmatically valid. (17b–c) are valid.

Either the carpet was cleaned, both conditional oughts are true, and Jones ought to pay the bill, or the carpet wasn't cleaned, (16a) is false, and Jones ought not to pay the bill.

In view of these considerations, care needs to be exercised in using action-guiding conditional oughts in practical situations. In our example, Jones can't use FDD to conclude $\bigcirc\mathsf{PayBill}$ without also committing to $\mathsf{CarpetCleaned}$. Conditional oughts can have *defeaters*. And when they do have defeaters, an agent who uses them is presuming that these defeaters are false. A cautious agent will take care to make sure that these defeaters in fact are false.

5.6 Standing oughts as generics

The view that statements of conditional obligation are conditionals is challenged by authors who claim that deontic deontic detachment (DDD) can play an important role in reasoning about obligation. As [9] points out, we need to formulate maxims and bodies of rules, and DDD seems to be needed in applying these materials.

If employees ought to be paid normal wages for every weekday they work that isn't a holiday, and employees ought to work on weekdays that aren't holidays, then it might be useful to conclude that employees ought to be paid normal wages on weekdays that aren't holidays. It is appealing to formalize this inference using wide-scope $\bigcirc$:

(18) $\bigcirc(\mathsf{Work} \to \mathsf{Paid}),\ \bigcirc\mathsf{Work}\ \therefore\ \bigcirc\mathsf{Paid}.$

Moreover, our account in Section 5.5 of practical deliberation assumed that appropriate circumstantial oughts would somehow become available in deliberative situations. Surely these circumstantial oughts must originate in a body of *standing oughts*, and must be obtained from this body by an inferential process of some sort. In rejecting wide-scope conditional oughts, we seem to have deprived ourselves of the inferential mechanisms we need to make practical use of general deontic knowledge.

There is a plausible and simple answer to both of these worries: our proposal is that standing oughts are *generic*.

It has often been observed that general, standing statements of obligation are defeasible. And projects in the deontological tradition such as [19] find it necessary to appeal to defeasible rules in attempting to produce systematic bodies of standing norms. Generic constructions are the natural expression of such standing oughts.

Generics are quite generally available in the world's languages. There are in fact many ways to express genericity; see [3, pp. 2–14]. In English, non-progressive, non-perfective present tense sentences are quite likely to be generics: 'She jogs home from work', 'It rains in Seattle', 'He likes red wine better than white wine'. It is such constructions that are our particular focus of attention here; linguists sometimes classify these under the heading "generic tense." Sometimes they are called "habituals," but they often express other sorts of regularities. Generic uses of present tense often contrast with "occasional" uses. For instance, ‘I prefer standing to sitting’ could be used to express either a standing preference, or a one-time preference, actuated on a particular occasion.

Such generic claims have three important logical characteristics. (1) They tolerate exceptions: there’s no contradiction in saying ‘Even though I jog home from work, I think I’ll take the bus today’. (2) They are context-sensitive, where the context often is a type of situation. (3) They support default inference: instances can be inferred from them in appropriate contexts, when the conclusions are not undermined by other considerations.[11]

Linguists have proposed several tests for recognizing generic constructions. According to one of the best and most generally applicable tests, a generic sentence can be paraphrased without significant change of meaning by adding ‘usually’.

According to this test, for instance, (19a) is a generic. But by exactly the same reasoning, so is the very similar (19b), which, of course, involves ‘ought’.

(19a) At a dinner party, the hostess will seat her guests.
(19b) At a dinner party, the hostess ought to seat her guests.

These examples also illustrate the context-sensitivity of generics. In a paragraph where the topic is dinner party goings-on, the qualifying phrase can be dropped. Despite the lack of explicit qualification, the last sentence is implicitly restricted to dinner parties.

(20) At the beginning of a dinner party, drinks and hors d’oeuvres will be served. After an interval, the group will move to the dining room. The hostess will seat her guests.

There is solid evidence, then, that (1) sentences involving ‘ought’ can be generic, and (2) that generics can be context-dependent. But this applies to conditional oughts as well as to unconditional ones. There is no great difference, for instance, between (19b) and (21).

(21) At a dinner party, if the occasion is at all formal, the hostess ought seat her guests.

The semantics of generics is poorly understood. But provisionally, and for purposes of illustration, we can assume that the generic constructions in which we are interested can be explained using preferential models.[12] Briefly, a generic sentence ϕ is true if it is true in all worlds that are maximally preferred according to some ordering, representing normality. This readily explains two of our logical features: (1) exceptions can occur because not all worlds maximize normality and (2) default inference is based on the assumption that the actual world is normal. We can explain context dependence by saying that the normal worlds are restricted to a contextually given background set of worlds, representing the background assumptions in play.

[11] This formulation is vague: making it precise is a difficult task, but work in nonmonotonic logic has shed light on this matter. See, in particular, [8].

[12] See [21].

We assume that a generic conditional ought has the doubly narrow-scope form

(22) $\phi > \mathrm{Gen}\bigcirc\psi$.

The antecedent ϕ, as before, shifts to a close world in which the antecedent ϕ is true, and restricts the background context worlds to ones in which ϕ is true. The world shift will typically have little or no semantic effect in these constructions, because we can expect obligations and normalities to be unchanged in a sufficiently close world. But the restriction to worlds where the antecedent is true is crucial: it affects the subsequent interpretation of both Gen and $\bigcirc$.

The interpretation of a generic conditional such as 'If you receive a bill, you ought to pay for it', then, amounts to this: in all normal worlds w in which you receive a bill, you pay the bill in all worlds w' that are normatively best in the set of deontic alternatives to w. We illustrate this idea, and provide further support for treating standing oughts as generics, with an example.

Suppose you are discussing the monthly bills with a significant other, who says:

(23a) "You ought to pay this bill."

You might reply:

(23b) "Sure, I ought to, but I can't pay it till I have the money."

But you would never say:

(23c) ?"No, I don't have to pay it; I don't have enough money."

The unnaturalness of (23c) supports the interpretation of (23a) as generic. If (23a) were an occasional, rather than a generic ought, (23c) would be the correct response.

Generic claims, however, characteristically tolerate exceptions. So if (23a) is interpreted, as we suggest, as generic, the use of 'but' in (23b) would mark the subsequent clause as an exception to a general rule.

Conditional oughts can have the same generic flavor, as the following variation on (23a–b) illustrates.

(24a) "If you got a bill from the carpet cleaner, you ought to pay it."
(24b) "Yes, if I got the bill I ought to pay it, but I don't have the money."

First, like all generics, statements like (24a) are standing generalizations, available to be used as appropriate in any deontic deliberation.[13] Second, generics support defeasible instantiation. For instance, with no reason to the contrary, from (24a) one can infer that I ought to pay the bill if I received it.

[13] In [13] Adam Lerner and Sarah-Jane Leslie explicitly identify ethical maxims with deontic generic constructions. This paper goes into some detail about how such constructions enter into reasoning about what one ought to do.

Thus, we can replicate DDD with a two-stage inference: first, from (24a) we defeasibly infer $\mathsf{BillReceived} > \bigcirc\mathsf{PayBill}$. This and $\mathsf{BillReceived}$—in the absence of defeaters—pragmatically imply $\bigcirc\mathsf{PayBill}$. Importantly, normal worlds in which you receive the bill are worlds in which you *ought* to receive the bill. So, while the inference in question does not proceed by way of $\bigcirc\mathsf{BillReceived}$ (as it would were DDD employed directly), it is nonetheless true. Thus, a combination of defeasible inference and pragmatic implication delivers the desired consequence.

6 Revisiting Chisholm

With the formal apparatus developed in the preceding sections, we can return to (1a–d), formalizing the example as follows:

(25a) $\bigcirc\mathsf{Help}$
(25b) $\bigcirc(\mathsf{Help} > \mathsf{Tell})$
(25c) $\neg\mathsf{Help} > \bigcirc\neg\mathsf{Tell}$
(25d) $\neg\mathsf{Help}$

Here we have used a wide-scope $\bigcirc$ to formalize (25b). Although we have said that $\bigcirc$ rarely takes wide scope over a conditional, we can take Chisholm's awkward 'it ought to be that if' to force such a reading, as Chisholm probably intended.
The apparent problem, then, is that

(26a) $\bigcirc\mathsf{Help}, \bigcirc(\mathsf{Help} > \mathsf{Tell})$ imply $\bigcirc\mathsf{Tell}$, and
(26b) $\neg\mathsf{Help}, \neg\mathsf{Help} > \bigcirc\neg\mathsf{Tell}$ imply $\bigcirc\neg\mathsf{Tell}$

which together *pragmatically* imply a deontic contradiction, $\bigcirc\perp$.

While (25a–d) do imply this deontic contradiction, we don't need to accept the conclusion. The context-sensitivity of $\bigcirc$ provides a natural account of how we can, in a sense, accept (25a–d) without thereby accepting $\bigcirc\perp$. In fact, the solution to paradox in its original form mirrors the contextual solution we gave in Section 4.3 to the simpler, unconditional form of the paradox.

Below, in Section 6.2, we provide a more detailed analysis of of the contextual solution, this time with a narrow-scope formalization of the second premise.

6.1 *Heads Up*, one more time

Recall the two-move version of *Heads Up* introduced in Section 4.3.

Here, we can find what might be called *impractical oughts* arising from a speaker's choice of context. Suppose the cards are dealt face up and the layout Jones faces is $\langle\mathsf{King}, 6\rangle$. We'll call the four outcomes w_1 (\$500), w_2 (\$0), w_3 (\$250), and w_4 (\$350).

We can picture the possibilities in the two-choice scenario using a branching temporal graph. The possible worlds correspond to histories or paths through the tree, and are shown at the top of the picture, along with the payoffs.

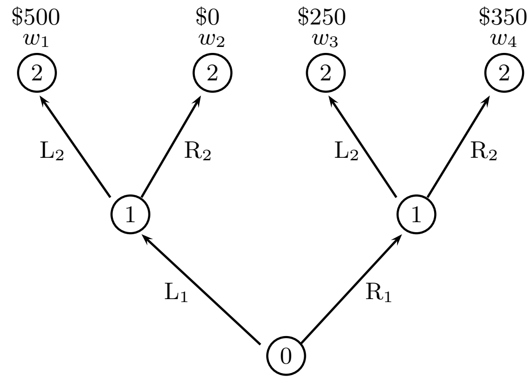

We imagine an observer watching over Jones' shoulder as he makes his choices in the world w_3. A king is on the left and a six is on the right, and Jones has chosen right. This first choice has left him in a suboptimal position; he can no longer get \$500, but by choosing right again he can get \$350. Choosing left will get him \$250. The observer might well say:

(27) $\bigcirc\mathsf{R}_2$ ['Jones ought to choose right in the second round]

The observer has made a good point: Jones will get less if he chooses left.

The observer could also take a more contrary-to-fact and less practical perspective. If at the outset Jones had gone for the best outcome, he would have chosen left. At the second round, then, the correct choice would also be left. With this perspective (and a bit of wishful thinking) the observer could also say:

(28) $\bigcirc\mathsf{L}_2$ ['Jones ought to choose left in the second round']

This is a bit unnatural; this impractical perspective is better expressed as:

(29) Jones ought to be choosing left right now.[14]

If there are semantic differences between (28) and (29), they're subtle. At any rate, Jones can't be choosing left unless he chooses left, and if he chooses left he is choosing left. It may be that (29) is more felicitous in the impractical employment of (28) for pragmatic reasons, and that this use of 'be' is similar in many respects to subjunctive mood.

Now, there is an apparent contradiction between (27) on the one hand, and (28–29) on the other, both said at the same time in w_3. But the contradiction is only apparent. There is a contextual element in play: the background of possibilities considered to be open alternatives.

The practical 'ought' in (27) should be interpreted with respect to the practical set of possibilities that are open at w_3. This will be a set consisting of two worlds: w_3 and a world w_4 which is like w_3 except for the fact that Jones chooses

[14] For some reason, usages with 'be' seem to go better with impractical contexts. This seems related to the difference that philosophers like Castañeda have noted between 'ought to be' and 'ought to do'. See [4].

right on the second turn instead of left in w_4. To evaluate this practical ought, we set the context, X, to $[\![\mathsf{R}_1]\!]$, or $\{w_3, w_4\}$. Then, $\mathfrak{M}, X, w_3 \models \bigcirc\mathsf{R}_2$, since $\{w_4\}$ is preferable to $\{w_3\}$.

But the impractical 'ought' in (28–29) requires a different set of possibilities. Here, the observer is meddling with the open alternatives by supposing that Jones had made the correct first choice. We therefore want the set of possibilities to be $\{w_1, w_2\}$. For this set of possibilities, $\bigcirc\mathsf{L}_2$ is true, even after Jones' first choice at w_3. This is because $f(w_3, \{w_1, w_2\}) \subseteq \{w_1, w_2\}$ and w_1, which is a L_2 world, is better than w_2.

The main point may have been lost in these formal details. It's this: a context, in the form of a set of background possibilities, contributes to the interpretation of an 'ought'. For practical oughts (and this is the default), these are the possibilities that vary according to exogenous chance factors and the agent's choice of an action. But 'ought' can also be used impractically, with respect to a counterfactual set of possibilities; such usages are often associated with the verb 'be'. The truth of an 'ought' statement will depend, among other things, on the context that is used to interpret it.

In our examples from the two-choice version of the game, (27) is practical and (28–29) are impractical. The contradiction between the two is therefore only apparent, since the appropriate contexts for them are different. In fact, although $\bigcirc\mathsf{R}_2$ is true relative to a context $\{w_3, w_4\}$ of alternatives, and $\bigcirc\mathsf{L}_2$ is true relative to a context $\{w_1, w_2\}$ of alternatives, the two formulas and the deontic contradiction $\bigcirc\bot$ that they entail are never true relative to the same context set of alternatives. So, there is no "paradox" here.

6.2 The Chisholm quartet

What we concluded about (27) and (28–29)—apparently contradictory, but perfectly compatible if you take the context shift into account—is, directly relevant to Chisholm's Paradox. Consider this *Heads Up* paradox:

(30a) It ought to be that Jones chooses left initially.
(30b) It ought to be that if he chooses left initially he chooses left next.
(30c) If Jones does not choose left initially, he ought not to choose left next.
(30d) Jones does not choose left initially.

We propose to formalize these as follows.

(31a) $\bigcirc\mathsf{L}_1$
(31b) $\mathsf{L}_1 > \bigcirc\mathsf{L}_2$.
(31c) $\neg\mathsf{L}_1 > \bigcirc\neg\mathsf{L}_2$
(31d) $\neg\mathsf{L}_1$.[15]

[15] Chisholm may well have intended a wide-scope formalization of (30b), but, as we have argued, such formulations are implausible, and are better treated as generics. In view of the foregoing discussion, we take (30b) to be a better formulation.

Again we imagine our bystander uttering (30a–d). This time, she speaks as Jones is about to make his first choice, and in the world w_4. We will evaluate (31a–d) from this standpoint. As we saw above, we also must take context into account, in the form of a set of alternatives. Many sets could be in play at this point. Let's consider three possibilities.

Context 1. Total ignorance. Suppose our observer uses a context set in which all possibilities are open, $X = \{w_1, w_2, w_3, w_4\}$. Here, (31a) is true, since L_1 is true in w_1 and w_1 is the best alternative to w_4. In fact, this context-world combination renders all four premises true. However, as we discussed in section 4.3, a context like total ignorance, which doesn't alter the alternatives, is not a felicitous context for a Chisholm-style premise set.[16]

Context 2. Optimal second choice. Suppose the observer believes that Jones may lapse in making his first choice, but will regain his senses in the second round. (Perhaps he appears to be temporarily distracted.) Then, the second choice will be optimal, meaning that her context is $\{w_1, w_4\}$.

Relative to this context, (31a) is true, since w_1 is better than w_4 and w_1 is an L_1 world. Furthermore, (31b) is true, but vacuously so since w_1 is the only alternative under which we can evaluate $\bigcirc\mathsf{L}_2$ under the antecedent and this context set. This suggests that $\{w_1, w_4\}$ is not a felicitous context for that formula.

This leaves (31c), which also is vacuously true.

Context 3. Bold first choice. Assuming that the first choice is bold, keeping the best possibility available, gives us the set $\{w_1, w_2\}$. Here, (31a) is true; $\mathfrak{M}, \{w_1, w_2\}, w_4 \models \bigcirc\mathsf{L}_1$. It's vacuously true, however, since the context assumes he will choose left. (31b) is also true, since w_1 is better than w_2. (31c) is also true, but it's vacuous because the antecedent of the conditional is inconsistent with the context set. Finally, (31d) is true as well, since the world is w_4.

6.3 Plausibility of the Chisholm premise set

Each of these contexts satisfies all of the premises of Chisholm's Paradox without sacrificing consistency. Part of the problem presented by the paradox was that these premises are supposed to imply $\bigcirc\perp$. To that extent, we've shown that the Chisholm premise set is not paradoxical on our account. In general, $\bigcirc p$, $p > \bigcirc q$, $\neg p > \bigcirc\neg q$, and $\neg p$ do not imply $\bigcirc\perp$. We lost this implication at the point where we added contexts to the interpretation of $\bigcirc$ and the conditional. While this is an interesting feature of the contextualized satisfaction conditions, it isn't really a solution to the paradox. Contexts two and three both employ vacuous oughts,

[16] There are also reasonable interpretations of $\bigcirc$ that make (31a) false in this context, but that issue is beyond the scope of the present paper.

which signify an inappropriate context. This is a problem because the four terms of Chisholm's Paradox not only seem true, but are meant to seem natural.

We obtain a more satisfactory solution if we say that each of (31a–d) is satisfied and appropriate in some context, but there there is no single context that satisfies them all appropriately. There are at least two linguistically well-motivated ways we might account for this context change.

Accommodation. We know from pragmatics that an utterance attracts an appropriate context. When a sentence is uttered (within limits) a context for interpreting it is selected that makes it a sensible thing to say. Following [16], this phenomenon is called "accommodation."

Aside from general rules, such as "Try to make the utterance true," some special rules seem to apply to the interpretation of 'ought'.

(i) All things equal, prefer an indicative or practical use, in which what is beyond the agent's control is supposed, but what depends on actions under the agent's control is allowed to vary.
(ii) Vacuous cases are to be avoided, and in particular, in interpreting $\bigcirc\phi$, context sets that entail ϕ or $\neg\phi$ are to be avoided.

These rules make $\{w_1, w_3\}$ (Context 2, the optimal second-choice context) the most plausible context for (30a). At any rate, the totally ignorant context falsifies (30a), and contexts that determine the first choice make (30a) vacuous. But Context 2 makes (31b) and (31c) vacuous.

On the other hand, $W = \{w_1, w_2, w_3, w_4\}$ (Context 1, the total ignorance context) is the most plausible for (30b) and (30c). This context makes both conditionals true, and it entails neither the antecedents of the conditionals nor their negations (though it does, on this account of $\bigcirc$, entail (31a)).

Together, these different preferences for contexts may help to explain the plausibility of Chisholm-style paradoxes. The premises seem true and felicitous because, when they're accommodated to their respective appropriate contexts, each *is* true and felicitous. But these different contexts cannot be unified into a single one that makes all the premises true and felicitous.

Assertion. It wasn't necessary to say anything about the last premise, (30d), because any context is compatible with its truth. But things are different if we imagine (30d) to have been asserted. Dynamic theories of assertion, such as the one presented in [23], take assertion to add context to the context. When the initial context for an assertion of ϕ is X, the subsequent context is $X \cap [\![\phi]\!]$.

If we take (20d) to have been asserted, then, we get quite a different picture of the premises (30a–d). The order of the premises in Chisholm's formulation, of course, doesn't invite this interpretation, since (1d) is his last premise. But the interpretation is still available, we think, even with Chisholm's order, and his wording of (1a) and (1b) actually encourages this interpretation.

To make the case where (30d) has been asserted salient, let's revise the order of premises as follows.

(32a) Jones does not choose left initially.
(32b) It ought to be that Jones chooses left at first.
(32c) It ought to be that if he chooses left initially he chooses left next.
(32d) If Jones does not choose left initially, he ought not to choose left next.

The assertion of (32a) restricts the context to $\{w_3, w_4\}$. This forces a *counterfactual* interpretation of (32b) and (32c), in which this restriction is temporarily suspended and replaced with the totally ignorant context. Chisholm's wording, with 'ought to be' in both (32b) and (32c), encourages this subjunctive interpretation. With a return to the restricted context at the last premise, FDD can be applied, enabling the conclusion 'Jones ought not to choose left next'.

This provides another natural, non-paradoxical interpretation of the Chisholm premises, in which (1a) and (1b) are taken to be subjunctive.

7 Conclusions

Chisholm's paradox is not merely the byproduct of a naive theory of conditional obligation. We have shown that, by integrating linguistic ideas, such as contextual effects and their interaction with assertion and accommodation, we are able to solve the paradox's *unconditional*, equally troublesome variation. Further, by combining this operator with factual deontic detachment, we rendered the Chisholm set consistent. Suggesting that some deontic constructions are generic, we also offered an account of standing obligations within this framework. Finally, citing the vacuous satisfaction of certain premises in contexts allowing all four Chisholm premises to be true, we provided an account of how the paradox can be seductive: the premises are all true, but not in the same context.

Beyond solving the paradox, these ideas provide a theory of deontic operators and conditionals that is linguistically motivated and intuitive, and that can, we think, be incorporated in a plausible picture of their role in practical reasoning.

References

1. Benthem, J., Grossi, D., Liu, F.: Priority structures in deontic logic. Theoria 80(1) (2013)
2. Cantwell, J.: Changing the modal context. Theoria 74(4), 331–351 (2008)
3. Carlson, G.N., Pelletier, F.J. (eds.): The Generic Book. Chicago University Press, Chicago, IL (1995)
4. Castañeda, H.N.: Paradoxes of moral reparation: Deontic foci versus circumstances. Philosophical Studies 57(1), 1–21 (1989)
5. Chisholm, R.M.: Contrary-to-duty imperatives and deontic logic. Analysis 24(2), 33–36 (1963)
6. Gillies, A.S.: On truth conditions for *If* (but not quite only *If*). The Philosophical Review 118(3), 325–345 (2009)
7. Greenspan, P.: Conditional ought and hypothetical imperatives. Journal of Philosophy (72), 259–276 (1975)

8. Horty, J.F.: Reasons as Defaults. Oxford University Press, Oxford (2012)
9. Jones, A.J., Carmo, J.: Deontic logic and contrary-to-duties. In: Gabbay, D.M., Guenthner, F. (eds.) Handbook of Philosophical Logic, Volume VIII, pp. 265–344. Kluwer Academic Publishers, Amsterdam, 2 edn. (2002)
10. Kolodny, N., MacFarlane, J.: Ifs and oughts. The Journal of Philosophy 107(3), 115–143 (2010)
11. Kratzer, A.: The notional category of modality. In: Eikmeyer, H.J., Rieser, H. (eds.) Words, Worlds and Contexts, pp. 38–74. Walter de Gruyter, Berlin (1981), http://semanticsarchive.net/Archive/Tc2NjA1M/notional-category-modality.pdf
12. Kratzer, A.: Modals and Conditionals. Oxford University Press, Oxford (2012)
13. Lerner, A., Leslie, S.J.: Generics, generalism, and reflective equilibrium: Implications for moral theorizing from the study of language. In: Hawthorne, J., Turner, J. (eds.) Philosophical Perspectives 27: Philosophy of Language, pp. 366–403. Wiley Periodicals, Malden, Massachusetts (2013)
14. Lewis, D.K.: Counterfactuals. Harvard University Press, Cambridge, Massachusetts (1973)
15. Lewis, D.K.: Counterfactuals and comparative possibility. Journal of Philosophical Logic 2(4), 418–446 (1973)
16. Lewis, D.K.: Scorekeeping in a language game. Journal of Philosophical Logic 8(3), 339–359 (1979)
17. Mott, P.L.: On Chisholm's paradox. Journal of Philosophical Logic 2(2), 197–211 (1973)
18. Portner, P.: Modality. Oxford University Press, Oxford (2009)
19. Ross, W.D.: The Right and the Good. Oxford University Press, Oxford (1930)
20. Sellars, W.: Reflections on contrary-to-duty imperatives. Noûs 1(4), 303–344 (1967)
21. Shoham, Y.: Reasoning about Change: Time and Causation From the Standpoint of Artificial Intelligence. The MIT Press, Cambridge, Massachusetts (1988)
22. Stalnaker, R.C.: Indicative conditionals. Philosophia 5, 269–286 (1975)
23. Stalnaker, R.C.: Assertion. In: Cole, P. (ed.) Syntax and Semantics 9: Pragmatics. Academic Press, New York (1981)
24. Thomason, R.H.: The semantics of conditional modality (2012), available at http://web.eecs.umich.edu/ rthomaso/documents/modal-logic/lkminf.pdf
25. Willer, M.: Dynamic thoughts on ifs and oughts. Philosophers' Imprint 14(28) (2014)