

# Chisholm's Paradox and Conditional Oughts

Catharine Saint Croix and Richmond H. Thomason

Philosophy Department, University of Michigan, Ann Arbor, MI, USA

**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 Chisholm's Paradox

[1] formulates the problem of contrary-to-duty obligations with the following example.<sup>1</sup>

- (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. This is only considered to be paradoxical because such examples are difficult to formalize.

We begin with an unconditional version of the Chisholm quartet. Having shed some complication, we argue, it becomes clear that 'oughts' are context-sensitive. Bearing that in mind, we shift to conditional obligation in Section 5. There, we offer considerations in favor of factual, rather than deontic, detachment. In Section 5.3, we show that a narrow-scope deontic conditional is unsuitable for the formalization of reparational obligations unless the conditional is contextualized. This contextualization avoids the problematic inference. We then discuss how generic constructions may be used to recover deontic detachment for wide-scope conditional 'ought's where necessary. Finally, we return to Chisholm's paradox in Section 6, where we demonstrate our proposed solution.

---

<sup>1</sup> This is Chisholm's exact wording, except that we have substituted 'Jones' for 'a certain man' in (1a).

## 2 An Unconditional Version

Steps (1b) and (1c) are conditional oughts, conditionals whose main clause involves an 'ought'. Like Chisholm himself, many authors ([2], [3], [4], [5], 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 [6], 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.

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  $\langle \text{Jack}, 3, 9 \rangle$  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  $\langle \text{Queen}, \text{King}, 4 \rangle$ , 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 ([7,8]),  $\bigcirc(\phi)$  is true just in case there is an outcome  $u$  satisfying  $\phi$  such that any outcome at least as good as  $u$  also satisfies  $\phi$ .

Under the  $\langle \text{Queen, King, 4} \rangle$  layout, then, let  $\phi = \text{'choose left or middle'}$ . In this case, both outcomes satisfying  $\phi$ , left and middle, reward \$50 and the only other possible outcome of Jones' choice rewards \$0. So,  $\bigcirc(\phi)$  is true.  $\bigcirc(\text{choose left})$ , however turns out to be false. This is because there is another outcome (middle), that doesn't satisfy 'choose left' but has a reward as good as that of 'choose left'.

We can formalize this by augmenting a boolean propositional logic with  $\bigcirc$  as a modal operator. 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  $W$ . 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  $\langle \text{Queen, King, 4} \rangle$ , we don't need to check irrelevant possibilities in which Jones chooses the left 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  $\langle \text{Queen, 5, 2} \rangle$ . Perhaps surprisingly, it is still natural to say that Jones ought to choose the leftmost card, even though he doesn't *know* this.

If the dealer were to say, "Jones, there's no card you ought to choose" in this case, the natural interpretation would be that there is no unique face card on the table, rather than that there may be a unique face card although Jones doesn't know this. Practical oughts act, in fact, as if they want to ignore the agent's epistemic situation, even though agents must find the best choice in light of their knowledge.

We can, however, find cases where practical oughts are relativized to the agents' knowledge. If, for instance, Jones detects a bias for queens on the left in the face-down version of the game he might say to himself, as he chooses, "I ought to choose the leftmost card." But this epistemically conditioned usage is evanescent. If the cards are turned over after Jones plays and the layout is  $\langle 5, 2, \text{Queen} \rangle$ , Jones could well say "Damn! I ought to have chosen the card on the right!" He can't, however, say "Damn! I ought to have chosen the card on

the left, but it wasn't a queen!" The "miners' paradox" scenario presented in [9] illustrates the point: although these knowledge-relative uses can be exhibited, doing so takes a certain amount of work.

## 4.2 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, 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  $\langle \text{Queen}, 2 \rangle$ . Since choosing left allows Jones to reach the optimal outcome, \$500,  $\circ(L_1)$  is true in our model: (a)  $\models \circ(L_1)$ . If he does choose left, (b)  $\models \circ(L_2)$ : 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  $L_2$  or \$350 for  $R_2$ . With these options, then, (c)  $\models \circ(R_2)$ . Thus, we have:

- (3a)  $\circ(L_1)$  is true.
- (3b) But  $\neg L_1$  is true.
- (3c) So,  $\circ(R_2)$  is true

At this point, it has become clear that the satisfaction relation is context dependent. To represent examples like this, 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  $X$ , representing the set of alternatives presumed to be open to choice.

**Definition 2.** Contextualized satisfaction in a model for  $\circ(\phi)$ .

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  $w \in W$ ,  $\mathfrak{M}, X, w \models \circ(\phi)$  **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)  $\mathfrak{M}, X, w \models \circ(L_1)$
- (4b) But  $\mathfrak{M}, X, w \models \neg L_1$
- (4c) So,  $\mathfrak{M}, X \cap \llbracket \neg L_1 \rrbracket, w \models \circ(R_2)$

Here  $\llbracket \neg L_1 \rrbracket$  is the set of worlds satisfying  $\neg L_1$ .

Thus, the assertion in (4b) serves to add its content to the context set used to interpret the next step, (4c). This is similar to the effect of assertion on presupposition noted in Stalnaker's [10] 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 [6] and has become a standard part of linguistic theories of modals; see [11,12]. 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  $\circ(\psi/\phi)$ , 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 generally to be preferred. 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 ' $\bullet\bullet\rightarrow$ ', making no assumptions about its logic for the moment. Assuming compositionality, there are two options for formalizing 'If  $\phi$  then ought  $\psi$ ': (i) wide scope  $\circ, \circ(\phi\bullet\bullet\rightarrow\psi)$  and (ii) narrow scope  $\circ, \phi\bullet\bullet\rightarrow\circ(\psi)$ .

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)  $\circ(\text{Help})$
- (6b)  $\circ(\text{Help}\bullet\bullet\rightarrow\text{Tell})$
- (6c)  $\neg\text{Help}\bullet\bullet\rightarrow\circ\neg\text{Tell}$
- (6d)  $\neg\text{Help}$

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 [13].)

### 5.1 Detachment

In Example (6) we want to conclude that Jones ought to tell his neighbors he isn't coming to help,  $\circ\neg\text{Tell}$ . Now, with few exceptions, logics of the conditional

allow us to infer  $\psi$  from  $\phi \bullet\bullet \rightarrow \psi$  and  $\phi$ . And with fewer exceptions, deontic logics with a conditional  $\bullet\bullet \rightarrow$  validate the inference from  $\bigcirc(\phi \bullet\bullet \rightarrow \psi)$  and  $\bigcirc\phi$  to  $\bigcirc\psi$ . There are, then, two ways to conclude an 'ought' statement.

**Definition 3.** *Factual deontic detachment (FDD)*

Infer  $\bigcirc\psi$  from  $\phi \bullet\bullet \rightarrow \bigcirc\psi$  and  $\phi$ .

**Definition 4.** *Deontic deontic detachment (DDD)*

Infer  $\bigcirc\psi$  from  $\bigcirc(\phi \bullet\bullet \rightarrow \psi)$  and  $\bigcirc\phi$ .

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

## 5.2 Factual 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.

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

## 5.3 Factual 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 ([10]).<sup>2</sup> This invokes a "selection function"  $s$  from propositions or sets of worlds to sets of worlds. This function satisfies the following conditions.

- (8a) For some  $u \in W$ ,  $s(Y, w) \subseteq \{u\}$
- (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  $Y = Y'$ .

<sup>2</sup> We mention at this point [4], which also proposes  $\phi \bullet\bullet \rightarrow \bigcirc\psi$  as a formalization of conditional oughts but without mentioning contextual effects.

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 with 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  $s(\llbracket \phi \rrbracket, w)$ .

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)  $\neg\text{Help} > \bigcirc\neg\text{Tell}$

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 is 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 the best alternatives, he does. Of course, in those alternatives,  $\neg\text{Tell}$  is false. So,  $u$  isn’t an  $\bigcirc\neg\text{Tell}$  world. Thus, (9) turns out to be false and the rather pointless (10) is true.

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

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.<sup>3</sup>

**Definition 6.** 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  $\mathfrak{M}, X \cap \llbracket \phi \rrbracket, u \models \psi$ .

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,

<sup>3</sup> [14] argues that this condition provides an improvement on Stalnaker’s semantics for the conditional.

because in the closest world  $u$  where John doesn't help his neighbors,  $\bigcirc\neg\text{Tell}$  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  $\mathfrak{M}, X, v \not\models \bigcirc\neg\text{Tell}$  for the best alternatives,  $v$  in  $X$ . Under the second we use  $X \cap \llbracket \neg\text{Help} \rrbracket$  instead of  $X$ , and (9) is true.

On the other hand, *modus ponens* is no longer valid on this interpretation of  $>$  (see [14] 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 [10]. Although the inference

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

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  $\phi$  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 providing a wide scope logical form  $\bigcirc(\phi > \psi)$ . 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. There is very little linguistic evidence for *any* cases in which modals take wide scope over conditionals. If such cases occur at all, they rarely occur naturally.

As [13] point out, deontic statements should figure in bodies of rules and maxims, and DDD does seem to play a useful role in reasoning in these domains. If employees ought to be paid for every day they work, and employees ought to work on weekdays that aren't holidays, then it might be useful to conclude that employees ought to be paid on weekdays that aren't holidays. We turn now to the formalization of DDD, and to how we might be able to capture inferences such as this, and begin by turning to another difficulty with FDD.

#### 5.5 Worries about Strengthening the Antecedent

Our treatment of conditional oughts as narrow scope conditionals allows for left-nonmonotonicity effects—failure of Strengthening the Antecedent. For instance, suppose that Jones is looking at  $\langle \text{Queen}, 4, \text{King} \rangle$ . Then, 'If the right card were a



6, I ought to choose the left card' is clearly true. And equally clearly, 'If the right card were a 6 and the middle card were a jack and the left card were an eight, I ought to choose the left card' is false. This is intended, but it creates problems for treating these conditional oughts as standing rules that can be applied in new circumstances.

When we lack Strengthening the Antecedent, it is harder to apply FDD than one might think. Suppose we are in a world  $u$  satisfying  $p$ . Jones knows this, and has reason to believe  $p > \circ q$  and  $(p \wedge r) > \circ q$ . For instance, Jones may know he has received a bill for cleaning a carpet, and think that he ought to pay the bill if he's received it, but also that he ought *not* to pay it if he's received it and the carpet has not, in fact, been cleaned. Intuitively, the following deontic conditionals are true in this scenario:

(12a)  $p > \circ q$  ['I ought to pay the bill if I received it.']

(12b)  $(p \wedge \neg r) > \neg \circ q$  ['I ought not to pay the bill if I received it and the carpet wasn't cleaned.']

Using FDD, Jones can reach the following conclusions.

(13a)  $p, p > \circ q \therefore \circ q$

(13b)  $p, \circ q, (p \wedge r) > \neg \circ q \therefore \neg r$

(13c)  $p, r, (p \wedge r) > \neg \circ q \therefore \neg(p > \circ q)$

In our example, Jones can't use FDD to conclude  $\circ(q)$  without also committing to  $\neg r$ . If, on the other hand, he finds that the rug hasn't been cleaned, he must, in view of (13c), give up the conditional obligation  $p > \circ q$ . In this common situation, it seems that we can't use typical, presumably action-guiding conditional oughts without checking many background facts (and, in particular, all the possible defeaters). It seems that we can't at the same time take conditional obligations to be *standing* oughts—true in the worlds in which we deliberate—and have FDD.

## 5.6 Standing Oughts as Generics

If at all possible, we should avoid having to resolve a conflict between FDD and standing maxims by choosing one at the expense of the other. In practical reasoning, we need to combine conditional oughts with beliefs about our current circumstances to obtain unconditional, immediate oughts—so we need FDD. But we also need standing oughts, and these too must be applicable to circumstances. This we take to be the real challenge of accommodating DDD.

We propose appealing to *generic constructions*, and in particular, generic conditionals, as a solution to this problem. Generics are quite generally available in the world's languages. 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'. The critical feature to observe is that generic claims tolerate exceptions: there's no contradiction in saying 'Even though I jog home from work, I think I'll take the bus today'.

Linguists postulate a GEN operator to provide a logical form for generics, so that  $\text{GEN}(\phi)$  is the generic of  $\phi$ . (See [15].)

While there isn't an overt tense marker signaling a generic in the case of 'ought' or 'should', generic uses of modal statements would not be surprising. Suppose you are discussing the monthly bills with a significant other, and say,

(14a) "You ought to pay this bill."

(14b) "Sure, I ought to, but I can't pay it till I have the money."

(14c) ?"No, I ought not to pay it; I don't have the money to pay it."

(14b) is the natural response in this case, rather than (14c). Without a generic interpretation of (14a), this would create a problem for the plausible idea that 'ought' implies 'can'. If (14a) is interpreted as a *generic* ought, however, (14b)'s characteristic 'but' would signal that the clause following it marks an exception to a general rule. Conditional oughts can have the same flavor, as the following variation on (14a–b) illustrates.

(15a) "If you got a bill from the carpet cleaner, you ought to pay it."

(15b) "Yes, if I got the bill I ought to pay it, but I don't have the money."

Generic constructions like (15a), involving a conditional and a deontic modal, do not appear to have been studied much by linguists. However, it's natural to treat them just like other generics, assigning them the same form  $\text{GEN}(\phi)$ , but where  $\text{GEN}$  is now operating on a conditional ought.

We committed ourselves in Section 5.3 to  $\phi > \circ \psi$  as a formalization of statements of conditional obligation. This means that (15a) will have a fully compositional formalization along the following lines.

(16)  $\text{GEN}(\text{BillReceived} > \circ \text{PayBill})$

Unfortunately, this suggestion offers less help with satisfaction conditions than one would like. The semantic interpretation of  $\text{GEN}$  is chronically problematic. Although linguists have proposed satisfaction conditions,<sup>4</sup> there is no general agreement about what they should be. Even the postulation of  $\text{GEN}$  is to some extent controversial.<sup>5</sup> Fortunately, however, we do not need to commit ourselves to any specific account of  $\text{GEN}$  to see that formalizations like (16) can serve the purposes that have made wide scope deontic conditionals and DDD seem attractive to some deontic logicians.

First, like all generics, statements like (16) are standing generalizations, ready to be used in any deontic deliberation.<sup>6</sup> Second, generics support defeasible instantiation. For instance, with no reason to the contrary, from (15a) one can infer that I ought to pay the bill if I received it.

Thus, we can replicate DDD with a two-stage inference: first, from (15a) we defeasibly infer  $\text{BillReceived} > \circ \text{PayBill}$ . This and  $\text{BillReceived}$  pragmatically imply

<sup>4</sup> See, for instance, [16].

<sup>5</sup> Also, it's natural to think that a singular generic like (14a) is in fact derived from a more general formulation using the generic plural, such as 'People ought to pay their bills'. The semantic interpretation of generic plurals is even more controversial.

<sup>6</sup> In [17] 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.

$\circ$ PayBill. Thus, a combination of defeasible inference and pragmatic implication delivers the desired consequence.

According to this theory, (12a–b) would now appear as

- (12a')  $\text{GEN}(\text{BillReceived} > \circ\text{PayBill})$  ['I ought (normally) to pay the bill if I received it.']  
 (12b')  $\text{GEN}((\neg\text{CarpetCleaned} \wedge \text{BillReceived}) > \circ\neg\text{PayBill})$  ['I ought not (normally) to pay the bill if I received it but the carpet wasn't cleaned.']

These formalizations do not cause the difficulties we saw in (13a–d), because  $\text{BillReceived} > \circ\text{PayBill}$  is only a defeasible consequence of  $\text{GEN}(\text{BillReceived} > \circ\text{PayBill})$ . This consequence is defeated by  $\text{BillReceived} \wedge \neg\text{CarpetCleaned}$ , in the presence of (12b').

## 6 Revisiting Chisholm

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

- (17a)  $\circ\text{Help}$   
 (17b)  $\circ(\text{Help} > \text{Tell})$   
 (17c)  $\neg\text{Help} > \circ\neg\text{Tell}$   
 (17d)  $\neg\text{Help}$

Here we have used a wide-scope  $\circ$  to formalize (17b). Although we have said that  $\circ$  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

- (18a)  $\circ\text{Help}, \circ(\text{Help} > \text{Tell})$  imply  $\circ\text{Tell}$ , and  
 (18b)  $\neg\text{Help}, \neg\text{Help} > \circ\neg\text{Tell}$  imply  $\circ\neg\text{Tell}$

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

While (17a–d) do imply this deontic contradiction, we don't need to accept the conclusion. The context sensitivity of  $\circ$  provides a natural account of how we can, in a sense, accept (17a–d) without thereby accepting  $\circ \perp$ . In fact, the solution to paradox in its original form mirrors the contextual solution we gave in Section 4.2 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 4.2. 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 (King, 6). We'll call the four outcomes  $w_1$  (\$500),  $w_2$  (\$0),  $w_3$  (\$250), and  $w_4$

(\$350). 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:

(19)  $\text{OR}_2$  ['Jones ought to choose right for the second pick']

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

The observer could also take a more remote 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:

(20)  $\text{OL}_2$  ['Jones ought to choose left for the second pick']

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

(21) Jones ought to be choosing left right now.<sup>7</sup>

If there are semantic differences between (20) and (21), 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 the reason that (21) is more felicitous than (20) in the impractical sense is pragmatic, and is similar in many respects to the difference between indicative and subjunctive mood.

Now, there is an apparent contradiction between (19) on the one hand, and (20/21) 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 (19) 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 right on the second turn instead of left in  $w_4$ . To evaluate this practical ought, we set the context,  $X$ , to  $\llbracket \text{R}_1 \rrbracket$ , or  $\{w_3, w_4\}$ . Then,  $\mathfrak{M}, X, w_3 \models \text{OR}_2$ , since  $\{w_4\}$  is preferable to  $\{w_3\}$ .

But the impractical 'ought' in (20/21) 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,  $\text{OL}_2$  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  $\text{L}_2$  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

<sup>7</sup> For some reason, usages with 'be' seem to go better with impractical contexts. This seems related to the difference that philosophers like Castañeda (See [18]) have noted between 'ought to be' and 'ought to do'.

‘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, (19) is practical and (20/21) are impractical. The contradiction between the two is therefore only apparent, since the appropriate contexts for them are different. In fact, although  $\circ R_2$  is true relative to a context  $\{w_3, w_4\}$  of alternatives, and  $\circ L_2$  is true relative to a context  $\{w_1, w_2\}$  of alternatives, the two formulas and the deontic contradiction  $\circ \perp$  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 (19) and (20/21)—apparently contradictory, but perfectly compatible if you take the context shift into account—is, directly analogous to the main issue in Chisholm’s Paradox. Consider this *Heads Up* paradox:

- (22a) It ought to be that Jones chooses left initially.
- (22b) It ought to be that if he chooses left initially he chooses left next.
- (22c) If Jones does not choose left initially, he ought not to choose left next.
- (22d) Jones does not choose left initially.

We propose to formalize these as follows.

- (23a)  $\circ L_1$
- (23b)  $L_1 > \circ L_2$ .
- (23c)  $\neg L_1 > \circ \neg L_2$
- (23d)  $\neg L_1$ .<sup>8</sup>

Again we imagine our bystander uttering (22a–d). This time, she speaks as Jones is about to make his first choice, and in the world  $w_4$ . We will evaluate (23a–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 two 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, (23a) 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 4.2, a context like total ignorance, which doesn’t alter the alternatives, is not a felicitous context for a Chisholm-style premise set.<sup>9</sup>

---

<sup>8</sup> Chisholm may well have intended a wide-scope formalization of (22b), but, as we have argued, such formulations are implausible, and are better treated as generics. We take (22b) to be a plausible formulation, however, given the foregoing discussion.

<sup>9</sup> There are also reasonable interpretations of  $\circ$  that make (23a) false in this context, but that issue is beyond the scope of the present paper.

**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, (23a) is true, since  $w_1$  is better than  $w_4$  and  $w_1$  is an  $L_1$  world. Furthermore, (23b) is true, but vacuously so since  $w_1$  is the only alternative under which we can evaluate  $\circ L_2$  under the antecedent and this context set. This suggests that  $\{w_1, w_4\}$  is not a felicitous context for that formula. This leaves (23c), 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, (23a) is true;  $\mathfrak{M}, \{w_1, w_2\}, w_4 \models \circ L_1$ . It's vacuously true, however, since the context assumes he will choose left. (23b) is also true, since  $w_1$  is better than  $w_2$ . (23c) is also true, but it's vacuous because the antecedent of the conditional is inconsistent with the context set. Finally, (23d) 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  $\circ \perp$ . To that extent, we've shown that the Chisholm premise set is not paradoxical on our account. In general,  $\circ p, p > \circ q, \neg p > \circ \neg q$ , and  $\neg p$  do not imply  $\circ \perp$ . We lost this implication at the point where we added contexts to the interpretation of  $\circ$  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, 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 (23a–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 [19], 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  $\circ \phi$ , context sets that entail  $\phi$  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 (22a). At any rate, the totally ignorant context falsifies (22a), and contexts that determine the first choice make (22a) vacuous. But this context makes (23b) and (23c) 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 (22b) and (22c). 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  $\circ$ , entail (23a)).

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, (22d), because any context is compatible with its truth. But things are different if we imagine (22d) to have been asserted. Dynamic theories of assertion, such as the one presented in [20], take assertion to add context to the context. When the initial context for an assertion of  $\phi$  is  $X$ , the subsequent context is  $X \cap \llbracket \phi \rrbracket$ .

If we take (20d) to have been asserted, then, we get quite a different picture of the premisses (22a–d). The order of the premisses 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 (22d) has been asserted salient, let's revise the order of premisses as follows.

(24a) Jones does not choose left initially.

(24b) It ought to be that Jones chooses left at first.

(24c) It ought to be that if he chooses left initially he chooses left next.

(24d) If Jones does not choose left initially, he ought not to choose left next.

The assertion of (24a), restricts the context to  $\{w_3, w_4\}$ . This forces a *counterfactual* interpretation of (24b) and (24c), in which this restriction is temporarily suspended and replaced with the totally ignorant context. Chisholm's wording, with 'ought to be' in both (24b) and (24c), 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 premisses, 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 the naturalness of this paradox. Beyond solving the paradox, these ideas provide a theory of deontic operators and conditionals that is linguistically motivated and intuitive.

## References

1. Chisholm, R.M.: Contrary-to-duty imperatives and deontic logic. *Analysis* 24, 33–36 (1963)
2. Bentham, J., Grossi, D., Liu, F.: Priority structures in deontic logic. *Theoria* 80 (2013)
3. Cantwell, J.: Changing the modal context. *Theoria* 74, 331–351 (2008)
4. Mott, P.L.: On Chisholm's paradox. *Journal of Philosophical Logic* 2, 197–211 (1973)
5. Sellars, W.: Reflections on contrary-to-duty imperatives. *Noûs* 1, 303–344 (1967)
6. 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)
7. Lewis, D.K.: *Counterfactuals*. Harvard University Press, Cambridge (1973)
8. Lewis, D.K.: Counterfactuals and comparative possibility. *Journal of Philosophical Logic* 2, 418–446 (1973)
9. Kolodny, N., MacFarlane, J.: Ifs and oughts. *The Journal of Philosophy* 107, 115–143 (2010)
10. Stalnaker, R.C.: Indicative conditionals. *Philosophia* 5, 269–286 (1975)
11. Portner, P.: *Modality*. Oxford University Press, Oxford (2009)
12. Kratzer, A.: *Modals and Conditionals*. Oxford University Press, Oxford (2012)
13. Jones, A.J., Carmo, J.: Deontic logic and contrary-to-duties. In: Gabbay, D.M., Guenther, F. (eds.) *Handbook of Philosophical Logic*, 2nd edn., vol. VIII, pp. 265–344. Kluwer Academic Publishers, Amsterdam (2002)
14. Thomason, R.H.: The semantics of conditional modality (2012), <http://web.eecs.umich.edu/~rthomaso/documents/modal-logic/lkminf.pdf>
15. Carlson, G.N., Pelletier, F.J. (eds.): *The Generic Book*. Chicago University Press, Chicago (1995)
16. Hacquard, V.: *Aspects of Modality*. Ph.d. dissertation, Linguistics Department, Massachusetts Institute of Technology, Cambridge, Massachusetts (2006)
17. 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 (2013)
18. Castañeda, H.N.: Paradoxes of moral reparation: Deontic foci versus circumstances. *Philosophical Studies* 57, 1–21 (1989)
19. Lewis, D.K.: Scorekeeping in a language game. *Journal of Philosophical Logic* 8, 339–359 (1979)
20. Stalnaker, R.C.: Assertion. In: Cole, P. (ed.) *Syntax and Semantics 9: Pragmatics*. Academic Press, New York (1981)