Coarse Deontic Logic (short version)

Frederik Van De Putte [1]

Ghent University
Blandijnberg 2, B-9000 Ghent, Belgium
frederik.vandeputte@ugent.be

Abstract

In recent work, Cariani has proposed a semantics for *ought* that combines two features: (i) it invalidates Inheritance in a principled manner; (ii) it allows for coarseness, which means that $ought(\varphi)$ can be true even if there are specific ways of making φ true that are (intuitively speaking) impermissible. We present a group of multi-modal logics based on Cariani's proposal. We study their formal properties and compare them to existing approaches in the deontic logic literature — most notably Anglberger et al.'s logic of obligation as weakest permission, and deontic stit logic.

Keywords: Deontic logic; contrastivism; modal inheritance; Ross paradox; deontic STIT logic; coarseness

1 Introduction

Contrastivism about "ought" says that claims using this modality can only be understood relative to a (usually implicit) contrast class. [2] So according to this view, "you ought to take the bus" is shorthand for "given the set of alternatives $\mathcal{A}$ under consideration, you ought to take the bus". Here $\mathcal{A}$ may consist of various ways of getting somewhere (say, the university).

In recent work, Cariani has proposed a formal semantics which starts from a contrastivist reading of ought [5]. This proposal is interesting for at least two reasons. First, it gives a principled account of why Inheritance [3] fails in cases like the Ross paradox, which makes it more insightful than most existing semantics for non-normal modalities. [4] Second, it allows for what Cariani calls *coarse* ought-claims, which means that $ought(\varphi)$ can be true even if there are specific ways of making φ true that are (intuitively speaking) impermissi-

[1] We are indebted to Mathieu Beirlaen and three anonymous referees for incisive comments on previous versions.

[2] See [20, footnote 1] for some key references to contrastivism in deontic logic.

[3] By *Inheritance* we mean here: from $ought(\varphi)$ and $\varphi \vdash \psi$, to infer $ought(\psi)$. This property is also often called *monotony*.

[4] As Cariani [5, p. 537] remarks, such semantics are "often purely algebraic", in the sense that they just translate rules for *ought* into conditions on neighbourhood functions.

ble.[5] This unusual combination – coarseness *without* Inheritance – is possible precisely because of the way the alternatives are modeled: rather than single worlds, they are (mutually exclusive) *sets* of worlds.

Before one can argue for or against Cariani's proposal, one has to study the logics obtained from it. We do this here. In Section 2, we present Cariani's proposal, both informally and in terms of a possible-worlds semantics. We discuss the most salient properties of the resulting logic. Next, we consider variants of this semantics that are defined over the same modal language (Section 3). Section 4 provides a map of the various logics obtained and presents their axiomatization.[6] Finally, we show how they relate to existing work in the deontic logic field, and where one can draw on this link in order to solve existing problems and puzzles (Section 5).

Preliminaries We use $p, q, \ldots$ for arbitrary propositional variables. The boolean connectives are denoted by $\neg, \vee, \wedge, \supset, \equiv$ (only the first two are primitive) and occasionally we will use the falsum and verum constants ($\bot$, resp. $\top$). $\varphi, \psi, \ldots$ are metavariables for formulas and $\Gamma, \Delta, \ldots$ for sets of formulas. *ought* refers to operators proposed as formal counterparts of the natural language "ought". Given an expression of the type *ought*(φ), φ is the *prejacent* of this formula.

2 Cariani's Semantics

In this section, we introduce and illustrate Cariani's semantics for *ought*. We first present the semantics informally in our own terms, after which we indicate the relation with Cariani's original presentation (Section 2.1). Next, we define a formal semantics which implements Cariani's ideas (Section 2.2) and discuss the most salient properties of the resulting logic (Section 2.3).

2.1 Cariani's proposal, informally

Our Version Cariani's *ought* is defined in terms of various more basic concepts. To spell these out, we need three parameters:

(a) a set of (mutually exclusive) *alternatives* or *options* $\mathcal{A}$
(b) a set $\mathcal{B} \subseteq \mathcal{A}$ of "optimal" or "best" options
(c) a set $\mathcal{I} \subseteq \mathcal{A}$ of "impermissible" options

For instance, in a context where we are deliberating about how Lisa ought to get to the university, her options may be represented by the following set:

$$\mathcal{A}_{\mathsf{ex}} = \{\text{walk}, \text{bike}, \text{bus}, \text{car}\}$$

indicating that she may walk to the university, drive her bike, take the bus, or drive by car. Some of these options may be optimal – e.g. biking or taking the

[5] We explain and illustrate Cariani's notion of coarseness in Section 2.1.

[6] In the full version of this paper [23] (available on request), we show that our axiomatizations are sound and (strongly) complete, and we establish the finite model property for each of the logics.

bus. Driving may well be impermissible (since she may not yet have obtained her driver's licence) and walking may be suboptimal (since given the distance, she risks getting late) but nevertheless permissible. So we have:

$$\mathcal{B}_{\mathsf{ex}} = \{\text{bike}, \text{bus}\}$$

$$\mathcal{I}_{\mathsf{ex}} = \{\text{car}\}$$

Each of the options in $\mathcal{A}_{\mathsf{ex}}$ can be carried out in many different ways; e.g. Lisa may drive her bike in a blue dress or in a green dress; she may drive her bike in a hazardous way or very cautiously. In Cariani's terms, this means the alternatives are *coarse-grained*. In other words, they correspond to generic action-types or general properties (sets of worlds in a Kripke-model), in contrast to action-tokens or maximally specific descriptions of a state of affairs (worlds in a Kripke-model). [7]

This explains at once how it is possible, in Cariani's framework, that there are (intuitively) impermissible instances of an optimal (or permissible) alternative. Even if Lisa ought to drive her bike or take the bus, this does not imply that every way of doing so is normatively ok. Indeed, relative to a more fine-grained set of alternatives, it may turn out that some ways of driving her bike are impermissible. Mind that the framework does not explicitly represent the impermissibility of such specific actions – hence, they are only impermissible "intuitively speaking". The point is exactly that, by choosing one specific level of granularity in a certain context, we decide to leave those more specific (impermissible) actions out of the picture. Once we make them explicit, the level of granularity changes, and with it the truth of any given *ought*-claim. [8]

Since options are coarse-grained, they do not fix every property of the world. Still, some propositions are fixed by taking one option rather than the other. If Lisa takes her bike, she is definitely not taking the bus or driving her car. In general, we say that an option $X \in \mathcal{A}$ *guarantees* a proposition φ iff following that option ensures that φ is the case.

We are now ready to spell out an informal version of Cariani's proposal. That is, where φ is a proposition, $ought(\varphi)$ is true (relative to $\mathcal{A}, \mathcal{B}, \mathcal{I}$) iff each of the following hold:

(i) φ is *visible*, i.e. for all $X \in \mathcal{A}$: X guarantees φ or X guarantees $\neg\varphi$
(ii) φ is *optimal*, i.e. for all $X \in \mathcal{B}$: X guarantees φ
(iii) φ is *strongly permitted*, i.e. for all $X \in \mathcal{A}$ that guarantee φ, $X \notin \mathcal{I}$.

For instance, in our example, it is true that Lisa ought to ride her bike or take the bus. It is false that she ought to ride her bike, take the bus or take the car, since taking the car is impermissible. It is equally false that she ought

[7] See [5, pp. 544-545] for a more detailed discussion of the link between action types/tokens and Cariani's semantics.

[8] This of course raises the question how oughts concerning such more fine-grained $\mathcal{A}'$ relate to the coarse-grained $\mathcal{A}$ – we return to this point in Section 5.

to ride her bike or take the bus in a green dress, since that proposition is not visible. [9]

This shows us at once that Inheritance is invalid on Cariani's semantics. It is in fact blocked in two different ways – see (i) and (iii) above. As a result, also the Ross paradox is blocked: "you ought to mail the letter" may be true while "you ought to mail the letter or burn it" is false. This will either be the case because burning the letter is invisible, or if we do take it to be a visible option, because it is impermissible.

Ranking and threshold In Cariani's original proposal, instead of $\mathcal{B}$ and $\mathcal{I}$, a "ranking" of $\mathcal{A}$ is used together with a "threshold" t on that ranking. The idea is that the "best" options are those that are maximal (according to the ranking), and the impermissible ones are those that are below the threshold. Although Cariani is not very explicit about the formal properties of his ranking and threshold, it seems that his ranking is a modular pre-order, in the sense that it distinguishes different layers of "ever better" options. [10] In other words, it can be defined as a function $r : \mathcal{A} \to \mathbb{R}$, where intuitively, X is better than X' (for $X, X' \in \mathcal{A}$) iff $r(X) > r(X')$. The threshold is then simply a $t \in \mathbb{R}$, such that whenever $r(X) < t$, X is impermissible.

It is easy enough to check that, once such an r and t are given, we can obtain $\mathcal{B}$ and $\mathcal{I}$ from them as follows: (i) $\mathcal{B} = \{X \in \mathcal{A} \mid r(X) = \mathsf{max}_<(\{r(Y) \mid Y \in \mathcal{A}\})\}$, and (ii) $\mathcal{I}$ is the set of all $X \in \mathcal{A}$ such that $r(X) < t$. Hence our simplified version of Cariani's semantics is at least as general as his original version.

Given fairly weak assumptions, we can also show the converse. That is, consider an arbitrary $\langle \mathcal{A}, \mathcal{B}, \mathcal{I} \rangle$ and suppose that each of the following hold:

(D) $\mathcal{B} \neq \emptyset$
(C∩ $\mathcal{B} \cap \mathcal{I} = \emptyset$

In other words, there are best options, and every best option is permissible. Define the function $r : \mathcal{A} \to \{1, 2, 3\}$ as follows:

(1) if $X \in \mathcal{B}$, then $r(X) = 3$
(2) if $X \in \mathcal{I}$, then $r(X) = 1$
(3) if $X \in \mathcal{A} \setminus (\mathcal{B} \cup \mathcal{I})$ then $r(X) = 2$

Let $t = 2$. It can easily be checked that (i) and (ii) hold. So if we assume (D) and (C∩), the two formats are equivalent (deontically speaking).

In the current section, we will leave restrictions (D) and (C∩) aside. In Section 3.1 we consider variants of our base logic in which these restrictions are added to the semantics.

---

[9] As the reader may note, "Lisa ought to take her bike, take the bus, or walk to the university" is also true in our example, which might strike one as odd. We return to this point in Section 3.2.

[10] At least it is in all the examples he gives. Also, this seems to be presupposed by the way he uses the notion of a threshold, viz. as a single member X of $\mathcal{A}$ such that any option below X is impermissible.

2.2 The formal semantics of $\mathbf{CDL^c}$

Our language $\mathcal{L}$ is obtained by closing the set of propositional variables $\mathcal{S} = \{p, q, \ldots\}$ under the Boolean connectives and the modal operators U (necessary/holds in every possible world), A (is guaranteed by the chosen alternative), B (is best/is guaranteed by all optimal alternatives), and P (is strongly permitted).

Two comments are in place here. First, Cariani does not explicitly mention the operators U and A. However, both are fairly natural modalities in this context. U is just a global modality – see [11] for a systematic study. A expresses the concept of being guaranteed by a given option, which Cariani uses in the semantic clause of his *ought*-operator. Moreover, adding both modalities to the language allows us to obtain a sound and complete axiomatization of the logic – see Section 4. [11]

Second, rather than taking it as primitive as Cariani does, we treat "is visible", V, as a defined operator:

$$\mathsf{V}\varphi =_{\mathsf{df}} \mathsf{U}(\mathsf{A}\varphi \vee \mathsf{A}\neg\varphi)$$

Likewise, O (Cariani's *ought*) is a defined operator:

$$\mathsf{O}\varphi =_{\mathsf{df}} \mathsf{V}\varphi \wedge \mathsf{B}\varphi \wedge \mathsf{P}\varphi$$

The following two definitions make the informal semantics from Section 2.1 exact: [12]

Definition 2.1 A $\mathbf{CDL^c}$-*frame* is a tuple $F = \langle W, \mathcal{A}, \mathcal{B}, \mathcal{I}\rangle$, where W is a non-empty set, $\mathcal{A} \in \wp(\wp(W))$ is a partition of W, $\mathcal{B} \subseteq \mathcal{A}$ is the set of *best options* in $\mathcal{A}$, and $\mathcal{I} \subseteq \mathcal{A}$ is the set of *impermissible options* in $\mathcal{A}$.

A $\mathbf{CDL^c}$-*model* M is a $\mathbf{CDL^c}$-frame $\langle W, \mathcal{A}, \mathcal{B}, \mathcal{I}\rangle$ augmented with a valuation function $v : \mathcal{S} \to \wp(W)$.

Since $\mathcal{A}$ is a partition of W, all worlds are by definition a member of some alternative in the contrast class. In other words, we exclude the possibility that some members of W are simply irrelevant for the deontic claims that are at stake. We leave the investigation of such a possibility for another occasion.

In line with the preceding, the members of $\mathcal{A}$ are interpreted as action types or options a given agent faces, whereas the members of W represent action tokens (specific ways of carrying out a given action or option). Formulas are evaluated relative to a given $w \in W$, in accordance with Definition 2.2. This means that in general, whether or not a formula is true may depend on the alternative that is chosen and on the specific way it is carried out or materializes. However, for purely normative claims, this is not the case (see our discussion of the property of Uniformity in Sections 2.3 and 3.3).

Definition 2.2 Let $M = \langle W, \mathcal{A}, \mathcal{B}, \mathcal{I}, v\rangle$ be a $\mathbf{CDL^c}$-model and $w \in W$. Where $w \in W$, let X^w denote the $X \in \mathcal{A}$ such that $w \in X$.

[11] It remains an open question whether one can obtain such an axiomatization without these modalities, and with V primitive.

[12] **CDL** is shorthand for "Coarse Deontic Logic". The superscript c refers to Cariani.

(SC1) $M, w \models \varphi$ iff $w \in v(\varphi)$ for all $\varphi \in \mathcal{S}$
(SC2) $M, w \models \neg\varphi$ iff $M, w \not\models \varphi$
(SC3) $M, w \models \varphi \vee \psi$ iff $M, w \models \varphi$ or $M, w \models \psi$
(SC4) $M, w \models \mathsf{U}\varphi$ iff $M, w' \models \varphi$ for all $w' \in W$
(SC5) $M, w \models \mathsf{A}\varphi$ iff $M, w' \models \varphi$ for all $w' \in X^w$
(SC6) $M, w \models \mathsf{B}\varphi$ iff for all $X \in \mathcal{B}$, for all $v \in X$, $M, v \models \varphi$
(SC7) $M, w \models \mathsf{P}\varphi$ iff for all $X \in \mathcal{A}$ s.t. (for all $v \in X$, $M, v \models \varphi$), $X \notin \mathcal{I}$

Note that $\mathsf{V}\varphi$ means (by our definition) that at every world w in the current model, either φ is guaranteed or $\neg\varphi$ is guaranteed. Since $\mathcal{A}$ is a partition of W, this is equivalent to saying that every option either guarantees φ or guarantees $\neg\varphi$, which corresponds to Cariani's original semantics for "is visible".

As usual, $\Gamma \Vdash_{\mathbf{CDL^c}} \varphi$ iff for all $\mathbf{CDL^c}$-models M and every world w in the domain of M, if $M, w \models \psi$ for all $\psi \in \Gamma$, then $M, w \models \varphi$.

2.3 Properties of $\mathbf{CDL^c}$

It can be easily verified that each of U, A and B are normal modal operators in $\mathbf{CDL^c}$. In fact, both U and A are **S5**-modalities. Second, P is a non-normal but classical modality (in the sense of Chellas [6]), which means it satisfies at least replacement of equivalents. As a result, also the defined operators V and O are classical.

Now for some more distinctive properties. Each of the following hold for $\Vdash = \Vdash_{\mathbf{CDL^c}}$:

$$\mathsf{O}(\varphi \wedge \psi) \nVdash \mathsf{O}\varphi \tag{1}$$

$$\mathsf{O}\varphi, \mathsf{O}\psi \Vdash \mathsf{O}(\varphi \wedge \psi) \tag{2}$$

$$\mathsf{O}\varphi, \mathsf{O}\psi \Vdash \mathsf{O}(\varphi \vee \psi) \tag{3}$$

$$\mathsf{O}\varphi, \mathsf{P}\psi \nVdash \mathsf{O}(\varphi \vee \psi) \tag{4}$$

$$\mathsf{O}\varphi, \mathsf{P}(\varphi \vee \psi) \nVdash \mathsf{O}(\varphi \vee \psi) \tag{5}$$

$$\mathsf{O}\varphi, \mathsf{P}\psi, \mathsf{V}\psi \Vdash \mathsf{O}(\varphi \vee \psi) \tag{6}$$

$$\mathsf{V}\varphi, \mathsf{V}\psi \Vdash \mathsf{V}(\varphi \vee \psi) \tag{7}$$

$$\mathsf{V}\varphi, \mathsf{V}\psi \Vdash \mathsf{V}(\varphi \wedge \psi) \tag{8}$$

$$\Vdash \mathsf{P}(\varphi \vee \psi) \supset (\mathsf{P}\varphi \wedge \mathsf{P}\psi) \tag{9}$$

$$\not\Vdash (\mathsf{P}\varphi \wedge \mathsf{P}\psi) \supset \mathsf{P}(\varphi \vee \psi) \tag{10}$$

$$\Vdash (\mathsf{P}\neg\mathsf{A}\neg\varphi \wedge \mathsf{P}\neg\mathsf{A}\neg\psi) \supset \mathsf{P}(\neg\mathsf{A}\neg\varphi \vee \neg\mathsf{A}\neg\psi) \tag{11}$$

Let us comment on these properties one by one. That O does not satisfy *Inheritance* – see (1) – was already explained above. Quite surprisingly, *Aggregation* (2) holds for O. In a context where the possibility of deontic conflicts is omitted, this is often considered a nice feature. It follows from the fact that the three operators B, V, and P are each aggregative – witness (8) and (9). For similar reasons, *Weakening* (3) also holds in $\mathbf{CDL^c}$ – this follows by the normality of B, (7), and (10).

Both Aggregation and Weakening deserve our attention here. As shown in [4], these properties fail on what is perhaps the most well-known contrastive semantics for *ought*, viz. the actualist semantics from [17], which has been worked out and axiomatized by Goble [8,9].

(4) and (5) tell us that, contrary to what one might expect, neither $\mathsf{P}\psi$ nor $\mathsf{P}(\psi \vee \varphi)$ suffice in order to derive $\mathsf{O}(\varphi \vee \psi)$ from $\mathsf{O}\varphi$. [13] The reason is that neither of those propositions warrant that $\varphi \vee \psi$ is visible, which is required for $\mathsf{O}(\varphi \vee \psi)$ to hold. Only if we add $\mathsf{V}\psi$ do we obtain a restricted form of Inheritance that is $\mathbf{CDL^c}$-valid – see (6).

Together with replacement of equivalents, (9) entails that P is "downward closed": whatever is stronger than something that is permitted, is itself also permitted. To see why this is so, note that $\mathsf{P}(\varphi \vee \psi)$ expresses that guaranteeing $\varphi \vee \psi$ implies that one is choosing a permissible option. Hence *a fortiori* guaranteeing φ (resp. ψ) is sufficient for permissibility. By the definition of O, this also means that $\mathsf{O}(\varphi \vee \psi) \Vdash_{\mathbf{CDL^c}} \mathsf{P}\varphi, \mathsf{P}\psi$: that $\varphi \vee \psi$ ought to be implies that $\varphi \vee \psi$ is strongly permitted, which in turn implies that both φ and ψ are strongly permitted. We return to this property in Section 3.2.

In view of (10), P is not an operator of "free choice permission" in the strict sense of [25]. To see why (10) holds, recall our example. Here, "Lisa takes the car in a green dress" ($car \wedge green$) is permissible in a vacuous way, since there is simply no option which guarantees that proposition. Likewise, "Lisa takes the car, but not in a green dress" ($car \wedge \neg green$) is permissible. However, *car* (which is equivalent to the disjunction of both propositions) is not permissible. [14]

(11) shows that for the more specific case where φ and ψ are of the form $\neg\mathsf{A}\neg\tau$, we do get the converse of (9). If it is permissible that (a) one leaves open

---

[13] Snedegar [20, pp. 217-218] refers to Goble [10, Note 49] who rejects such a rule. However, in Goble's case, the P-operator is one of weak permission, i.e. $\mathsf{P} =_{\mathsf{df}} \neg\mathsf{O}\neg$. Besides that, Goble's main concern is to accommodate deontic conflicts, a target which Cariani explicitly rules out – as Snedegar acknowledges.

[14] In view of this example, P seems to express only part of the meaning of "is permitted". A more plausible operator of (strong) permission can be defined as $\mathsf{P}^v = \mathsf{P}\varphi \wedge \mathsf{V}\varphi$. Note that $(\mathsf{P}^v\varphi \wedge \mathsf{P}^v\psi) \Vdash_{\mathbf{CDL^c}} \mathsf{P}^v(\varphi \vee \psi)$, but $\mathsf{P}^v(\varphi \vee \psi) \not\Vdash_{\mathbf{CDL^c}} \mathsf{P}^v\varphi \wedge \mathsf{P}^v\psi$. We leave the investigation of such definable operators for future work.

the possibility that φ, and it is also permissible that (b) one leaves open the possibility that ψ, then it is permissible that (c) one leaves open the possibility that φ or one leave open the possibility that ψ. Indeed, whenever (c) holds, either (a) or (b) hold and hence one is definitely taking one of the permissible options.

Other interesting validities concern the interaction between the alethic modalities U, A and the deontic modalities B, P, and O. These are of two types:

$$\text{where } \nabla \in \{\mathsf{B}, \mathsf{P}, \mathsf{O}\} : \Vdash \nabla\varphi \equiv \nabla\mathsf{A}\varphi \tag{12}$$

$$\text{where } \nabla \in \{\mathsf{B}, \mathsf{P}, \mathsf{O}\} : \Vdash \nabla\varphi \equiv \mathsf{U}\nabla\varphi \tag{13}$$

Contrast-sensitivity, (12), expresses that the deontic modalities really apply to alternatives $X \in \mathcal{A}$, rather than worlds $w \in W$. For instance, $\mathsf{B}\varphi$ can only be true if φ is true in all worlds that belong to an optimal alternative; but that is the same as saying that all optimal alternatives guarantee φ. This property is therefore essential for Cariani's constrastive approach.

Uniformity, (13), expresses that deontic claims are either always settled true or settled false (to use terminology from [2]). It follows from the fact that $\mathcal{B}$ and $\mathcal{I}$ are independent of the world w one happens to be at in a model. We return to this property in Section 3.3.

3 Some Variants

We now consider variants of the $\mathbf{CDL^c}$-semantics and motivate each of them independently. This will be useful in Section 5, where we compare Cariani's construction to existing work in deontic logic.

3.1 Conditions (D) and (C∩)

We first return to the conditions mentioned at the end of Section 2.1. (D) corresponds to the requirement in Standard Deontic Logic that the accessibility relation is *serial*, and hence, that there is at least one "ideal" or "optimal" world. It can be moreover easily checked that (D) is expressed by the familiar axiom schema $\mathsf{B}\varphi \supset \neg\mathsf{B}\neg\varphi$ within $\mathbf{CDL^c}$. This axiom schema (along with the failure of the T-schema, $\mathsf{B}\varphi \supset \varphi$) is traditionally seen as the distinctive feature of deontic logics.

Although it is a much debated property in the context of deontic logic in general, (D) does seem to have some intuitive power in the present context. After all, the idea is that we start from a fixed set of alternatives, one particular ranking r, and one threshold t.[15] Finiteness of $\mathcal{A}$ already entails (D). But even if we allow for a possibly infinite number of options, it seems sensible to say that we only consider finitely many of those as viable options, such that a ranking on them will always yield a non-empty set of best alternatives.

[15] As Cariani notes, one may generalize the entire setting to cases with multiple rankings and threshold functions; that seems to be his preferred way of allowing for deontic conflicts.

(C∩) is more difficult to interpret in the present context. It states that every best option is permissible. Interestingly, this condition is not definable in the language of $\mathbf{CDL^c}$. In fact, imposing it onto the semantics has no impact on the resulting logic. [16] This means in turn that, once we assume (D), and as far as the consequence relation is concerned, there really is no difference between Cariani's original semantics and our reformulation of it.

3.2 Putting the threshold at optimality

Bronfman & Dowell note that Cariani's use of a set of alternatives (as a set of sets of worlds) and a ranking on them does not conflict *per se* with the standard approach in modal logic [3, p. 6]. The distinctive feature of Cariani's semantics, according to them, is the use of the permissibility threshold in order to block Inheritance. It is this feature that they attack.

To understand their argument, we should briefly rehearse the pragmatic defense of Inheritance for *ought*. This defense says, roughly speaking, that although affirming $ought(\varphi \vee \psi)$ is rather pointless in cases where we also know $ought(\varphi)$, the former expression is nevertheless true whenever the latter is. It is much like affirming "John is either Dutch or Italian" when we actually know that John is Dutch: not maximally helpful, but also not plainly false or mistaken. What *is* false is the Gricean implicature that follows when we *only* state $ought(\varphi \vee \psi)$, viz. that $\varphi \vee \psi$ is the most specific necessary condition for optimality.

Cariani rejects this defense of Inheritance, since it cannot account for the way *ought* behaves in embeddings [5, pp. 549]. Such behavior, he argues, can only be explained by the following principle:

(Implicated) $ought(\varphi \vee \psi)$ communicates that one has two ways of doing as one ought, viz. by making φ true or by making ψ true.

In contrast, Cariani's account covers (Implicated) well: as we saw in Section 2.3, $\mathsf{O}(\varphi \vee \psi) \Vdash_{\mathbf{CDL^c}} \mathsf{P}\varphi, \mathsf{P}\psi$.

However, Bronfman & Dowell rightly remark that (Implicated) gives counterintuitive results when applied to Cariani's own semantics. That is, by taking an option that is suboptimal but permissible, the agent is also doing as (s)he ought – at least if (Implicated) holds. Let us illustrate this with our running example. The options *bus* and *bike* are the only two optimal ones. However, since *walk* is permissible, $ought(bus \vee bike \vee walk)$ comes out true. But, given (Implicated), this means that by walking to the university, Lisa is doing as she ought.

Bronfman & Dowell suggest that, if one really wants to satisfy Cariani's requirement, one should put the threshold at optimality. [17] There are two

[16] See [23] where these claims are proven.

[17] There remains a problem though. Suppose that "Lisa ought to go to the supermarket" is true. Since the semantics satisfies replacement of equivalents, it follows that "Lisa ought to either go to the supermarket and pay for whatever she buys or go to the supermarket and steal something." Given (Implicated), it follows that by going to the supermarket and stealing something, Lisa is doing as she ought. So whatever refinement one proposes of

ways to implement this suggestion. The first is to change the semantic clause for P, such that $M, w \models \mathsf{P}\varphi$ iff, whenever $X \in \mathcal{A}$ is such that $M, w' \models \varphi$ for all $w' \in X$, then $X \in \mathcal{B}$. This means that $\mathcal{I}$ becomes superfluous in the semantics of the logic.

Secondly, one may leave the semantic clause for P unaltered, but treat $\mathcal{I}$ simply as the set of all suboptimal alternatives. This means that we impose the following frame condition on $\mathbf{CDL^c}$-models:

(C+) $\mathcal{I} = \mathcal{A} \setminus \mathcal{B}$

The advantage of this second approach – which we will follow in the remainder – is that it allows for a smooth comparison with Cariani's original proposal. Note that (C+) is equivalent to the conjunction of condition (C$\cap$) (see Section 2.1) and the following:

(C$\cup$) $\mathcal{I} \cup \mathcal{B} = \mathcal{A}$

Henceforth, let M be a $\mathbf{CDL^{bd}}$-model iff it is a $\mathbf{CDL^c}$-model that satisfies (C+); we denote the associated consequence relation by $\Vdash_{\mathbf{CDL^{bd}}}$.

Obviously, $\mathbf{CDL^{bd}}$ is an extension of $\mathbf{CDL^c}$. But exactly what additional validities (in our language $\mathcal{L}$) do we get from imposing this condition? Each of (1)-(10) from Section 2.3 hold also for $\Vdash\ =\ \Vdash_{\mathbf{CDL^{bd}}}$, and hence not much seems to change to the deontic part of the language.

However, once we consider the interaction with U, we do get an important additional feature: if two ought-claims are both true, their prejacents have the same extension in the model. Following [7], we call this *Uniqueness*:

$$\mathsf{O}\varphi, \mathsf{O}\psi \Vdash_{\mathbf{CDL^{bd}}} \mathsf{U}(\varphi \equiv \psi) \tag{14}$$

This property fails for $\mathbf{CDL^c}$ – witness our example: both *ought*(*bike*$\vee$*bus*) and *ought*(*bike*$\vee$*bus*$\vee$*walk*) are true, but one is obviously more specific than the other. Note that by **S5**-properties of U, Uniqueness entails both Aggregation and Weakening for O.

Even if condition (C+) is well-motivated, *Uniqueness* may be hard to swallow from the viewpoint of natural language. One morning John may have to ensure that he gets to the office in time (p), but also that his children get to school in time (q). So *ought*(p) and *ought*(q) both seem true in this scenario. But John would be rather lucky if making q true would at once ensure that p also holds (or vice versa).

Still, this kind of critique misses the point behind Cariani's semantics. That is: once we fix a set of alternatives and a way to compare them, then (usually) there is no doubt that one or several of those alternatives are optimal. So in the above scenario, we are really looking at different sets of alternatives, or in more technical terms, different partitions of one and the same set of possible worlds.

Cariani's (or Kratzer's) semantics, pragmatic factors will anyway have to be called for at some point. (This example is a variant of Hansson's "vegetarian's free lunch" [12, p. 218].)

Another question is whether, if we do allow for several such partitions of the given W, there should be some interaction between the related *ought*-claims. We will not be able to tackle this important issue in the present paper and postpone it for future work.

3.3 Rejecting Uniformity

As we just saw, there are reasons for strengthening $\mathbf{CDL^c}$ in various ways. There are however also reasons for weakening $\mathbf{CDL^c}$, in the sense that it is no longer assumed that optimality and permissibility are uniform throughout a model. That is, rather than taking $\mathcal{B}$ and $\mathcal{I}$ as sets of alternatives, one may think of them as functions, taking as their argument worlds $w \in W$ (or alternatives $X \in \mathcal{A}$), and mapping these to sets of alternatives. This means in turn that the validities mentioned in (13) — see page 8 — are denied.

To motivate such a weakening, we can point to various arguments that have been put forth in the literature. First, from the viewpoint of game theory, it has been argued that which action of a given agent α is best, may depend on the actions other agents perform; hence, it will also depend on the specific world one happens to be at. See e.g. [1, Section 4.2] where this point is discussed and linked to some properties of the deontic operators.

Second, in [26], Wansing attacks specific constructions of deontic logic based on a branching-time framework, in which the truth of "obligation reports" (say, claims about what ought to be, what is best, what one ought to do, etc) depend only on the moment m of evaluation. This means that such claims are either true at all moment-history pairs m/h, or false at all m/h. In the present, more abstract framework, moments correspond to the entire set W, whereas moment/history-pairs correspond to single worlds $w \in W$.

Wansing's arguments for this claim are of two kinds: on the one hand, he says that certain obligations are simply of such a type that they depend on future contingents. For instance, if "you ought to give the prize to the winner of this race" is true, then depending on who actually wins (say a or b), it may be true that "you ought to give the prize to a" – but this will of course not be settled true. The other argument is more intricate, as it concerns the so-called Restricted Complement thesis from [2]. As Wansing shows, this thesis together with Uniformity trivializes nested ought-claims of the type "John ought to see to it that it is forbidden for Mary to eat the cake."

Third and last, Uniformity is typically rejected by actualist theories of *ought*. In contrast to possibilists, actualists argue that what ought to be depends on what is actually the case (now or in the future), rather than on what can be (or may become) the case.[18] Of course, the temporal dimension is not explicit in the simple $\mathbf{CDL^c}$-models we considered so far. Nevertheless, the fact that we abstract from the temporal dimension in our models seems a sufficient reason to remain neutral about those properties that would become

[18] See e.g. [15, Section 7.4.3] where the two views are briefly discussed and linked to two different notions of *ought* in stit logic. A more unified theory that encompasses both these notions is presented in [16].

problematic, once we add time back in.

4 Coarse Deontic Logics

Let us take stock. We first generalize Definition 2.1 from Section 2.2:

Definition 4.1 A **CDL**-*frame* is a tuple $F = \langle W, \mathcal{A}, \mathcal{B}, \mathcal{I} \rangle$, where W is a non-empty set, $\mathcal{A} \in \wp(\wp(W))$ is a partition of W, $\mathcal{B} : W \to \wp(\mathcal{A})$ maps every $w \in W$ to the set of w-*best options* in $\mathcal{A}$, and $\mathcal{I} : W \to \wp(\mathcal{A})$ maps every $w \in W$ to the set of w-*impermissible options* in $\mathcal{A}$.

The definition of a model and the semantic clauses remain the same, with the exception of the following:

(SC6') $M, w \models \mathsf{B}\varphi$ iff for all $X \in \mathcal{B}(w)$, for all $v \in X$, $M, v \models \varphi$
(SC7') $M, w \models \mathsf{P}\varphi$ iff for all $X \in \mathcal{A}$ s.t. (for all $v \in X$, $M, v \models \varphi$), $X \in \mathcal{A} \backslash \mathcal{I}(w)$

Table 1 gives a sound and (strongly) complete axiomatization of **CDL**. The first six axioms and rules in this table are standard. The axioms (G_{A}), (G_{B}) and (G_{P}) follow from the fact that U is a global modality. (C_{B}), (C_{P}), (P1) and (P2) were already discussed in Section 2.3. Finally, (EQ_{P}) is a strengthened version of replacement of equivalents for P.

(CL)	any complete axiomatization of classical propositional logic
(MP)	from $\varphi, \varphi \supset \psi$ to infer ψ
(NEC_{U})	from $\vdash \varphi$, to infer $\vdash \mathsf{U}\varphi$
(K_{B})	$\mathsf{B}(\varphi \supset \psi) \supset (\mathsf{B}\varphi \supset \mathsf{B}\psi)$
($S5_{\mathsf{A}}$)	**S5** for U
($S5_{\mathsf{A}}$)	**S5** for A
(G_{A})	$\mathsf{U}\varphi \supset \mathsf{A}\varphi$
(G_{B})	$\mathsf{U}\varphi \supset \mathsf{B}\varphi$
(G_{P})	$\mathsf{U}\varphi \supset \mathsf{P}\neg\varphi$
(C_{B})	$\mathsf{B}\varphi \equiv \mathsf{BA}\varphi$
(C_{P})	$\mathsf{P}\varphi \equiv \mathsf{PA}\varphi$
(P1)	$\mathsf{P}(\varphi \vee \psi) \supset (\mathsf{P}\varphi \wedge \mathsf{P}\psi)$
(P2)	$(\mathsf{P}\neg\mathsf{A}\neg\varphi \wedge \mathsf{P}\neg\mathsf{A}\neg\psi) \supset \mathsf{P}(\neg\mathsf{A}\neg\varphi \vee \neg\mathsf{A}\neg\psi)$
(EQ_{P})	$\mathsf{U}(\varphi \equiv \psi) \supset (\mathsf{P}\varphi \equiv \mathsf{P}\psi)$

Table 1
Axiomatization of **CDL**.

Table 2 provides an overview of the conditions on **CDL**-frames we have considered so far, and the axioms (if any) that correspond to these frame conditions. Where (C1), ..., (Cn) are frame conditions from Table 2, say M is an $\mathbf{CDL_{C1,\ldots,Cn}}$-model iff M is an **CDL**-model and M obeys these conditions. We use $\Vdash_{\mathbf{CDL_{C1,\ldots,Cn}}}$ to refer to the associated semantic consequence relation. Note that $\mathbf{CDL^c} = \mathbf{CDL_{U_B,U_P}}$ and $\mathbf{CDL^{bd}} = \mathbf{CDL_{U_B,U_P,C+}}$.

Not all of these conditions are independent. As noted, (C+) is equivalent to the conjunction of (C∪) and (C∩). In view of axiom (G_{A}), (U_{B}) implies (A_{B}), and (U_{P}) implies (A_{P}). Also, (C+) implies that (U_{B}) and (U_{P}) are equivalent,

(U_{B})	for all $w, w' \in W$, $\mathcal{B}(w) = \mathcal{B}(w')$	$\mathsf{B}\varphi \equiv \mathsf{UB}\varphi$
(U_{P})	for all $w, w' \in W$, $\mathcal{I}(w) = \mathcal{I}(w')$	$\mathsf{P}\neg\mathsf{A}\varphi \equiv \mathsf{UP}\neg\mathsf{A}\varphi$
(A_{B})	where $w, w' \in X$, $\mathcal{B}(w) = \mathcal{B}(w')$	$\mathsf{B}\varphi \equiv \mathsf{AB}\varphi$
(A_{P})	where $w, w' \in X$, $\mathcal{I}(w) = \mathcal{I}(w')$	$\mathsf{P}\neg\mathsf{A}\varphi \equiv \mathsf{AP}\neg\mathsf{A}\varphi$
(D)	for all $w \in W$, $\mathcal{B}(w) \neq \emptyset$	$\mathsf{B}\varphi \supset \neg\mathsf{B}\neg\varphi$
(C$\cup$)	for all $w \in W$, $\mathcal{B}(w) \cup \mathcal{I}(w) = W$	$(\mathsf{B}\varphi \wedge \mathsf{P}\neg\mathsf{A}\varphi) \supset \mathsf{U}\varphi$
(C$\cap$)	for all $w \in W$, $\mathcal{B}(w) \cap \mathcal{I}(w) = \emptyset$	-
(C+)	for all $w \in W$, $\mathcal{B}(w) = \mathcal{A} \setminus \mathcal{I}(w)$	-

Table 2
Frame conditions and axioms for **CDL**.

and that (A_{B}) and (A_{I}) are equivalent. We leave it to the reader to verify that this exhausts the dependencies between each of the frame conditions from Table 2.

Whereas most of the frame conditions are modally expressible, (C+) and (C$\cap$) are not. This means that there is no formula φ which is globally valid on all and only those **CDL**-frames that satisfy (C+). One can nevertheless give a sound and complete, Hilbert-style axiomatization of the logics in question, by adding the associated axioms from Table 2 to **CDL**. One can moreover show that all these logics satisfy the finite model property, and hence are decidable. We refer to [23] where each of these claims are spelled out in exact terms and proven.

5 Related Work

CDL and its extensions bear close resemblances to existing work in deontic logic. In fact, leaving some specific modeling choices aside, one could say that they are just a combination of two well-known constructions in the field. We only explain both of these here in a nutshell; a more detailed and exact comparison of the respective logics is left for future work.

5.1 Deontic necessity and sufficiency

The idea of combining a notion of necessity and sufficiency for modeling *ought* was proposed fairly recently in [1,19] under the name "obligation as weakest permission". The idea is that what one ought to do is that which is implied by every strongly permitted proposition, where a proposition is strongly permitted iff it is sufficient for optimality. The resulting *ought*-operator satisfies the same basic properties as our O in $\mathbf{CDL}_{\mathbf{C}+}$ does – Uniqueness, and hence also Aggregation and Weakening. Likewise, it does not satisfy Inheritance, Uniformity, and the rule of Necessitation.

In [24], richer logics are studied in which both deontic necessity and sufficiency are expressible, which can be traced back to an extended abstract by van Benthem [22]. As shown in [7, Section 3], the deontic action logic from [21] is a fragment of van Benthem's system, and hence belongs to the same family of logics.

The main difference between the aforementioned logics and the **CDL**-family is that the former speak about the optimality (permissibility) of single worlds or action-tokens, whereas the latter speak about sets of worlds or action-types. As a result, we can also express the additional condition that φ is visible whenever $ought(\varphi)$ is true. Also, because of this feature, the logic of obligation as weakest permission and its relatives do not allow for coarseness.

5.2 Deontic stit logics

Deontic logics for optimal actions, which are conceived as *sets* of worlds, are at least as old as Horty's [14], which he further developed in the 2001 book [15].[19] Roughly speaking, $ought_\alpha(\varphi)$ is true at a world w in a model of Horty's most basic semantics if and only if α sees to it that φ whenever it takes one of its best options at w. Horty further distinguishes between two ways to determine what the best options are; one is called *dominance act utilitarianism* and satisfies Uniformity; the other is called *orthodox act utilitarianism* and invalidates Uniformity.[20] Both satisfy the (D)-axiom (see Table 2). Horty's $ought_\alpha$-operators are hence much like the B-operator of $\mathbf{CDL}_{\mathbf{D},\mathbf{U}_{\mathsf{B}}}$ (resp. $\mathbf{CDL}_{\mathbf{D}}$), with the obvious difference that they refer explicitly to an agent or group of agents. Horty's systems lack an operator for strong permission (our P).

Apart from the usual benefits – the transfer of insights and results from one system to the other –, there is one particular sense in which this link can be highly useful. In [20], Snedegar considers the problem of *coarsening inferences*, i.e. inferences that involve sets of alternatives that differ in their degree of coarseness. Snedegar's question then is: how do ought-claims relative to $\mathcal{A}$ relate to ought-claims relative to a finer partition $\mathcal{A}'$?

In view of the preceding, this question is analogous to asking how the obligations of a group of agents relate to the obligations of subgroups of that group, within the framework of deontic stit-logic.[21] Indeed, the alternatives that are available to the group correspond exactly to a partition that refines the partition corresponding to the alternatives available to a subgroup.

6 Summary and Outlook

The main contribution of this paper consists in the formal study of different variants of Cariani's semantics for *ought*. Spelling out these variants in turn allowed us to point at links with existing work in deontic logic, most particularly the logic of obligation as weakest permission and deontic stit logic.

Many issues remain unsettled, such as a more exact comparison of these systems. As explained, the link with deontic stit logic suggests possible solu-

[19] In Horty's stit-based semantics, the points of evaluation are moment-history pairs rather than worlds, and the sets of worlds are rather sets of histories. There is however a one-to-one correspondence between such models and more regular Kripke-models – see e.g. [13, Section 2.1].

[20] In [16], Horty proposes a way to unify both accounts and hence overcome semantic ambiguity w.r.t. "the right action(s)".

[21] Horty discusses this relation in 6.2 of his book, showing that dominant act utilitarianism differs from orthodox act utilitarianism in this respect. See e.g. [18] for formal results on this matter.

tions to the problem of coarsening inferences; in future work we want to study this relation in more detail. Also, it is an open question whether deontic stit logic can be enriched with an operator for strong permission, and in particular, how such an operator will behave for group obligations.

References

[1] Albert J.J. Anglberger, Norbert Gratzl, and Olivier Roy. Obligation, free choice, and the logic of weakest permissions. *Review of Symbolic Logic*, 8(4):807–827, 2015.

[2] N. Belnap, Perloff M., Xu M., and Bartha P. *Facing the Future: Agents and Choice in Our Indeterminist World.* Oxford University Press, 2001.

[3] Aaron Bronfman and Janice Dowell. The language of reasons and oughts. In: Star, D. (ed.) Oxford Handbook of Reasons. Oxford University Press, Oxford (forthcoming).

[4] Fabrizio Cariani. Logical consequence for actualists (and other contrastivists). Unpublished manuscript, April 15, 2015.

[5] Fabrizio Cariani. "Ought" and resolution semantics. *Noûs*, 47(3):534–558, 2013.

[6] Brian Chellas. *Modal Logic: an Introduction.* Cambridge : Cambridge university press, 1980.

[7] Huimin Dong and Olivier Roy. Three deontic logics for rational agency in games. In *Studies in Logic*, volume 8, pages 7–31. College Publications, 2015.

[8] Lou Goble. A logic of "good", "should", and "would": Part I. *Journal of Philosophical Logic*, 19(2):169–199, 1990.

[9] Lou Goble. A logic of "good", "should", and "would": Part II. *Journal of Philosophical Logic*, 19(3):253–76, 1990.

[10] Lou Goble. Normative conflicts and the logic of "ought". *Noûs*, 43(3):450–489, 2009.

[11] Valentin Goranko and Solomon Passy. Using the universal modality: Gains and questions. *Journal of Logic and Computation*, 2(1):5 – 30, 1992.

[12] Sven Ove Hansson. *The Varieties of Permission*, chapter 3, pages 195–240. College Publications, 2013.

[13] Andreas Herzig and François Schwarzentruber. Properties of logics of individual and group agency. In *Advances in Modal Logic*, 2008.

[14] John F. Horty. *Deontic Logic, Agency and Normative Systems: ΔEON '96: Third International Workshop on Deontic Logic in Computer Science, Sesimbra, Portugal, 11 – 13 January 1996*, chapter Combining Agency and Obligation (Preliminary Version), pages 98–122. Springer London, London, 1996.

[15] John F. Horty. *Agency and Deontic Logic.* Oxford University Press, 2001.

[16] John F. Horty. Perspectival act utilitarianism. In Patrick Girard, Olivier Roy, and Mathieu Marion, editors, *Dynamic Formal Epistemology*, volume 351 of *Synthese Library*, pages 197–221. Springer Netherlands, 2011.

[17] Frank Jackson. On the semantics and logic of obligation. *Mind*, 94:177–195, 1985.

[18] Barteld Kooi and Allard Tamminga. Moral conflicts between groups of agents. *Journal of Philosophical Logic*, 37(1):1–21, February 2008.

[19] Olivier Roy, Albert J.J. Anglberger, and Norbert Gratzl. The logic of obligation as weakest permission. In Thomas Agotnes, Jan Broersen, and Dag Elgesem, editors, *Deontic Logic in Computer Science*, volume 7393 of *Lecture Notes in Computer Science*, pages 139–150. Springer Berlin Heidelberg, 2012.

[20] Justin Snedegar. Deontic reasoning across contexts. In Fabrizio Cariani, Davide Grossi, Joke Meheus, and Xavier Parent, editors, *Deontic Logic and Normative Systems*, volume 8554 of *Lecture Notes in Computer Science*, pages 208–223. Springer International Publishing, 2014.

[21] Robert Trypuz and Piotr Kulicki. A systematics of deontic action logics based on boolean algebra. *Logic and Logical Philosophy*, 18(3-4):253–270, 2010.

[22] Johan van Benthem. Minimal deontic logics (abstract). *Bulletin of the Section of Logic*, 8(1):36–41, 1979.

[23] Frederik Van De Putte. Coarse Deontic Logic (Extended Version). In preparation.
[24] Frederik Van De Putte. That will do: Logics of deontic necessity and sufficiency. Under review, 2016.
[25] Georg Henrik von Wright. *An Essay in Deontic Logic and the General Theory of Action.* North-Holland Publishing Company, 1968.
[26] Heinrich Wansing. Nested deontic modalities: Another view of parking on highways. *Erkenntnis*, 49(2):185–199, 1998.