

Priority Structures in Deontic Logic

by

JOHAN VAN BENTHEM

University of Amsterdam, and Stanford University

DAVIDE GROSSI

University of Liverpool

FENRONG LIU

Tsinghua University

Abstract: This article proposes a systematic application of recent developments in the logic of preference to a number of topics in deontic logic. The key junction is the well-known Hansson conditional for dyadic obligations. These conditionals are generalized by pairing them with reasoning about syntactic priority structures. The resulting two-level approach to obligations is tested first against standard scenarios of contrary-to-duty obligations, leading also to a generalization for the Kanger-Anderson reduction of deontic logic. Next, the priority framework is applied to model two intuitively different sorts of deontic dynamics of obligations, based on information changes and on genuine normative events. In this two-level setting, we also offer novel takes on vexed issues such as the Chisholm paradox and modelling strong permission. Finally, the priority framework is shown to provide a unifying setting for the study of operations on norms as such, in particular, adding or deleting individual norms, and even merging whole norm systems in different manners.

Keywords: deontic logic, preference logic, modal logic, dynamic logic, contrary-to-duty, norm change

1. Introduction

THERE IS A LONG-STANDING meta-ethical intuition that deontic notions of obligation, permission and prohibition involve a normative “ideality ordering”. This ordering ranks possible situations or courses of action according to how well they meet moral demands. Thus, Moore (1903) writes (as quoted in van Fraassen, 1973, p. 6):

[. . .] to assert that a certain line of conduct is [. . .] absolutely right or obligatory, is obviously to assert that more good or less evil will exist in the world, if it is adopted, than if anything else be done instead.

Accordingly, deontic logic has long considered models involving betterness ordering of worlds or states, going back at least to Hansson (1969). There, statements of dyadic obligation like “it ought to be the case that φ under condition ψ ” (in symbols, $\mathbf{O}(\varphi \mid \psi)$) were interpreted in terms of a binary relation $s \preceq t$ between states s, t – representing that t is *at least as ideal* as s – according to the following semantics:

$$\mathcal{M}, s \models \mathbf{O}(\varphi \mid \psi) \Leftrightarrow \max_{\preceq}(\llbracket \psi \rrbracket_{\mathcal{M}}) \subseteq \llbracket \varphi \rrbracket_{\mathcal{M}} \quad (1)$$

Here, \mathcal{M} is a modal model on a frame $\mathcal{F} = \langle S, \preceq \rangle$, $\llbracket \cdot \rrbracket_{\mathcal{M}}$ is the truth-set function of \mathcal{M} , and \max_{\preceq} is a ‘selection function’ picking out the \preceq -maximal elements in any given subset of S . Typically, the relation \preceq is constrained so that \max_{\preceq} always outputs a non-empty set.¹

This framework is quite flexible: depending on the properties of the relation \preceq , different logics can be obtained. Hansson (1969) starts with a \preceq which is only reflexive, moving eventually to total preorders that are reflexive, transitive and connected. Each step imposes stronger restrictions on how we make comparisons between worlds. Another type of variation makes the ordering \preceq dependent on the world of evaluation.² This semantics has sparked a literature that continued well into the 1990s (cf. van der Torre, 1997), though it has also been criticized for less than ideal fit with intuitive deontic reasoning. By now, maximality-based ordering semantics has emerged in many branches of philosophical logic, including conditional logic, doxastic logic and defeasible inference (Makinson, 1993).

In the present article, we extend maximality-based ordering models and their applications to deontic logic by tapping into recent developments of this semantics in the area of preference logic (in particular, Liu, 2008, 2011a; Girard, 2008). An outline of the article and of its results follows.

Outline of the article

Section 2 establishes a relation between the abovementioned ideality – or betterness – orderings on states and syntactic priority structures consisting of orderings on (propositional) formulae. Syntactic priority structures are seen as the explicit reasons determining the ideality ordering on states. Some technical results concerning this first section are reported in the Appendix.

Section 3 looks at the simple modal logics induced by priority structures and relates them to Hansson-style deontic logics. This yields a two-level perspective on reasoning about ideality orderings whose usefulness we illustrate by three applications to deontic logic.

Section 4 applies priority structures to the benchmark problem of contrary-to-duty obligations (cf. Prakken and Sergot, 1997; Carmo and Jones, 2001, for overviews) and the Kanger-Anderson reduction of the deontic modalities.

Sections 5 and 6 explain how our setting supports a variety of dynamic events that change obligations. Using techniques from Liu (2008, 2011a) – inspired by

1 In this article only the single-agent case will be considered. Extensions to multi-agent deontic scenarios, which would not pose any technical difficulty, require indexing the ideality relations by different agents.

2 Lewis (1974) is an overview of various moves in the early literature on dyadic obligation.

normative actions such as commands – we show how dynamics on ideality relations on states and dynamics on priority structures work in harmony, delivering a novel analysis of several deontic phenomena from the so-called factual detachment to norm change. This brings deontic logic closer to current dynamic logics of belief and preference change (van Benthem, 2011).

Finally, section 7 identifies further lines of research that now open up. Throughout the article we make use of, and elaborate on, well-known examples taken from the literature in deontic logic.

2. Priorities and Betterness

2.1 Reasons for orderings

The betterness relation between situations that grounds our obligations often stems from an explicit code for what is right or wrong. Let us start with this often cited rousing quote from St Paul’s First Letter to the Corinthians:

It is good for a man not to touch a woman. But if they cannot contain, let them marry: for it is better to marry than to burn. (cf. van Fraassen, 1973, p. 6)

This passage identifies three properties of states: most ideal is that men do not touch women, less ideal is the more permissive property of men either not touching women or marrying them, and most permissive, and hence least ideal, is ‘none of the above’, leading to the trivial property \top . This is not a scenario of deontic reasoning, but of an equally important process of norm giving. We can formalize the latter as a sequence of relevant propositions ordered by a binary priority relation:

$$(\neg t \vee m \vee \neg m) \prec (\neg t \vee m) \prec \neg t \quad (2)$$

where atoms m and t have their expected readings (‘marry’ and ‘touch’), and \prec represents a binary priority relation: $\neg t$ is strictly better than $(\neg t \vee m)$, which is strictly better than $(\neg t \vee m \vee \neg m)$ (i.e., \top).³

We call structures of this kind ‘priority sequences’ (Liu, 2008), and their more general variants – ‘priority graphs’ – will be the focus of this article.

2.2 Priority graphs and derived betterness ordering

Definition 1 (P-graphs). Let $\mathcal{L}(\mathbf{P})$ be a propositional language built on the set of atoms \mathbf{P} . A P -graph is a tuple $\mathcal{G} = \langle \Phi, \prec \rangle$ such that:

3 One can also construe the order of properties differently with not-touching on top, and marrying as second best. We return to options in extracting priorities below.




Figure 1. Hasse diagrams of P-graphs. Our convention is that higher formulae are strictly better than lower ones

- $\Phi \subset \mathcal{L}(\mathbf{P})$ with $|\Phi| < \omega$;
- \prec is a strict order on Φ such that, for all propositions $\phi, \psi \in \Phi$, if $\phi \prec \psi$ then ψ logically implies ϕ .

The maximal (resp., minimal) elements of a P-graph $\mathcal{G} = \langle \Phi, \prec \rangle$ are denoted by $\max(\mathcal{G})$ (resp., $\min(\mathcal{G})$).⁴

Intuitively, $\phi \prec \psi$ means that (the logically stronger) formula ψ is strictly better than (the logically weaker) formula ϕ . In particular, a P-graph is a finite graph of formulae from a propositional language, where each formula logically implies its immediate successor in the order.

P-graphs where \prec is a strict linear order, i.e., \prec -chains, are called *P-sequences*. They are referred to as sequences $\langle \phi_1, \dots, \phi_n \rangle$ of propositional formulae and are denoted by letter \mathcal{S} . The length of a sequence is its number of elements. Each P-graph $\langle \Phi, \prec \rangle$ defines in a natural way a set of P-sequences as the longest sequences that can be built with pairs of formulae from the relation \prec . For instance, the graph on the left of Figure 1 defines the P-sequences: $\langle p_1, p_1 \vee p_2, p_1 \vee p_2 \vee p_3 \rangle$ and $\langle p_1, p_1 \vee p_3, p_1 \vee p_2 \vee p_3 \rangle$. Given a P-graph $\mathcal{G} = \langle \Phi, \prec \rangle$, we denote by $S_{\mathcal{G}}$ the set of longest P-sequences defined by \mathcal{G} .

Here is how a given P-graph determines a betterness ordering on states:

Definition 2 (State betterness from P-graphs). Let $\mathcal{G} = \langle \Phi, \prec \rangle$ be a P-graph, S a non-empty set of states and $\mathcal{I}: \mathbf{P} \rightarrow 2^S$ a valuation function. The betterness relation $\preceq_{\mathcal{G}} \subseteq S^2$ is defined as follows:

$$s \preceq_{\mathcal{G}} s' := \forall \phi \in \Phi : s \in \llbracket \phi \rrbracket_{\mathcal{I}} \Rightarrow s' \in \llbracket \phi \rrbracket_{\mathcal{I}}. \quad (3)$$

The function outputting this preorder is called sub (from ‘subsumption’). We will sometimes refer to relations derived via sub explicitly as $\preceq_{\mathcal{G}}^{\text{sub}}$.

4 Notice that since Φ is finite and \prec is a strict order, $\max(\mathcal{G}) \neq \emptyset$ and $\min(\mathcal{G}) \neq \emptyset$ for any P-graph \mathcal{G} .

Definition 2 orders states in S according to which elements of the P-graph they satisfy. If a state satisfies a property in the graph, it also satisfies all \prec -worse properties in the graph.

Definitions 1 and 2 are illustrated at several places in the ensuing sections but the interested reader might already want to consult Examples 1 and 2 in section 4. Here we first mention a few important properties of P-graphs:

Fact 1 (Basic properties of \preceq_G) *Let $\mathcal{G} = \langle \Phi, \prec \rangle$ be a P-graph. For any valuation $\mathcal{I} : \mathbf{P} \rightarrow 2^S$ it holds that:*

1. *The relation \preceq_G is a preorder⁵ whose strict part \prec_G is upward well-founded;⁶*
2. *If $\varphi_i \prec \varphi_j$, then for all worlds $s \in \llbracket \varphi_i \rrbracket_{\mathcal{I}}$, $s' \in \llbracket \varphi_j \rrbracket_{\mathcal{I}}$: $s \preceq_G s'$;*
3. *If $\varphi_i \prec \varphi_j$, then for all worlds $s \in \llbracket \varphi_i \wedge \neg \varphi_j \rrbracket_{\mathcal{I}}$, $s' \in \llbracket \varphi_j \rrbracket_{\mathcal{I}}$: $s \prec_G s'$.*

Remark 1 (From inclusion to priority as primitive). We chose to make P-graphs consist of propositions ordered by (proper) inclusion (Definition 1). These induce a total preorder on states (Definition 2). An alternative definition (see Liu, 2011b) allows unrelated properties in the graph. An ordering on states is then induced by the following lexicographic map lex :

$$s \preceq_G^{\text{lex}} s' \Leftrightarrow \forall \varphi \in \Phi : [s \in \llbracket \varphi \rrbracket_{\mathcal{I}} \Rightarrow s' \in \llbracket \varphi \rrbracket_{\mathcal{I}} \text{ or } \exists \varphi' : [\varphi \prec \varphi' \text{ and } s \notin \llbracket \varphi' \rrbracket_{\mathcal{I}} \text{ and } s' \in \llbracket \varphi' \rrbracket_{\mathcal{I}}]]. \quad (4)$$

State s' is at least as good as s if and only if either all properties in the graph that are satisfied by s are also satisfied by s' or, if s satisfies a property φ that s' does not satisfy, then there is a better property φ' which is satisfied by s' but not by s . Relations obtained from P-graphs through lex can be proven equivalent to relations obtained through sub , as shown in the Appendix.

Remark 2 (The history of P-graphs). The representation of ideality relations through orderings of formulae – like P-graphs – has a long history in the deontic logic literature, where they occur in several papers as central tools to prove the completeness of Hansson-style deontic systems. Their first application can be found in Spohn (1975) where they are used to construct the canonical total preorder in the completeness proof for Hansson's logic known as DSDL3. This technique has later been applied also in Hansen (2001) and in the more general setting of preorders in Hansen (2005), where they are also shown to be related to the so-called multiplex semantics of Goble (2000). Our article will show that the usefulness of

⁵ A preorder is a reflexive and transitive binary relation.

⁶ Intuitively, this means that if one is to build chains of elements from strictly worse to strictly better, these chains are always bound to stop at some element, which is therefore a maximal element of \preceq .

syntactic preference structures in deontic logic goes way beyond merely technical applications. We will argue for them as a prolific and versatile means for the representation and analysis of normative systems, in particular once interfaced with a number of novel techniques recently developed in the preference logic literature.

3. Reasoning about Betterness

Semantic models, even enriched with syntactic priority structure, do not yet tell us how agents reason about their obligations. We now present a simple logic capturing reasoning with betterness orderings induced by priority graphs.

3.1 A modal logic of preorders

The definitions that follow here are standard in current modal logics of preference, and we refer to Boutilier (1994) and Girard (2008) for details and further motivation:

Language and semantics. The basic modal language of preference $\mathcal{L}(\mathbf{U}, \preceq)$ is built from a countable set \mathbf{P} of atoms according to the following inductive syntax:

$$\mathcal{L}(\mathbf{U}, \preceq) : \varphi ::= p \mid \top \mid \neg\varphi \mid \varphi \wedge \varphi \mid [\preceq]\varphi \mid [\mathbf{U}]\varphi$$

Here $p \in \mathbf{P}$. Existential modalities and Boolean operators are defined as usual. The intended meaning shows in the semantic structures of section 2:

Definition 3 (Models). A model for $\mathcal{L}(\mathbf{U}, \preceq)$ on the set of atoms \mathbf{P} is a tuple $\mathcal{M} = \langle S, \preceq, \mathcal{I} \rangle$ where:

- i) S is a non-empty set of states;
- ii) \preceq is a preorder over S ($s \preceq t$ stands for “ t is at least as good as s ”);
- iii) $\mathcal{I} : \mathbf{P} \rightarrow 2^S$.

A pointed model is one with a distinguished world $s : \langle \mathcal{M}, s \rangle$. We also define the strict suborder \prec (“strictly better than”) of \preceq as usual: $s \prec t$ iff $s \preceq t$ and $t \not\preceq s$.

Here are the key semantic clauses for the modal operators:

Definition 4 (Satisfaction). Let $\mathcal{M} = \langle S, \preceq, \mathcal{I} \rangle$ be a model. Truth for a formula $\varphi \in \mathcal{L}(\mathbf{U}, \preceq)$ in a pointed model $\langle \mathcal{M}, s \rangle$ is defined inductively:

$$\mathcal{M}, s \models p \Leftrightarrow s \in \mathcal{I}(p)$$

$$\mathcal{M}, s \models [\preceq]\varphi \Leftrightarrow \forall s' \in S \text{ s.t. } s \preceq s' : \mathcal{M}, s' \models \varphi$$

$$\mathcal{M}, s \models [\mathbf{U}]\varphi \Leftrightarrow \forall s' \in S : \mathcal{M}, s' \models \varphi$$

A useful notation is this: $\llbracket \varphi \rrbracket_{\mathcal{M}}$ denotes the truth-set of φ in \mathcal{M} .

[U]-formulae define global properties of a model, [\preceq]-formulae local properties true in all states at least as good as the current state.

Axiomatics. A complete axiomatic proof calculus for our system consists of the standard modal logic **S4** for betterness, **S5** axioms for the universal modality, and one inclusion axiom $[U]\varphi \rightarrow [\preceq]\varphi$. This logic is known to be sound and strongly complete for preorders. Its uses go back to Boutilier (1994).

Models from P-graphs. Having formulated the basic logic of preferential reasoning, let us now investigate our models a little more closely for their fine-structure. Given a P-graph \mathcal{G} and a propositional valuation function \mathcal{I} , Definition 2 yields models of the above type $\mathcal{M} = \langle S, \preceq_{\mathcal{G}}, \mathcal{I} \rangle$ where $\preceq_{\mathcal{G}}$ is the total preorder derived from \mathcal{G} .⁷

Expressive power: defining semantic ‘best’. Consider the formula

$$[U](\psi \rightarrow \langle \preceq \rangle (\psi \wedge [\preceq](\psi \rightarrow \varphi))) \quad (5)$$

It expresses the property that for each state that satisfies ψ there is a better state satisfying ψ and such that all of its better alternatives either do not satisfy ψ or satisfy both ψ and φ .

As already observed for instance in Boutilier (1994), Formula 5 expresses, on the models of interest for us here, precisely the notion of maximality, which, we have seen, lies at the heart of Hansson’s semantic for dyadic obligations (Formula 1):

Fact 2 *Let $\mathcal{M} = \langle S, \preceq, \mathcal{I} \rangle$ be a model where the preorder \preceq has an upward well-founded strict part \prec . It holds that:*

$$\mathcal{M}, s \models [U](\psi \rightarrow \langle \preceq \rangle (\psi \wedge [\preceq](\psi \rightarrow \varphi))) \Leftrightarrow \max_{\preceq}(\llbracket \psi \rrbracket_{\mathcal{M}}) \subseteq \llbracket \varphi \rrbracket_{\mathcal{M}}$$

This equivalence may be rephrased as follows. On upward well-founded models, all maximal states in ψ satisfy φ if and only if, for every ψ -state, there is a better state satisfying ψ and such that for each state upward all ψ -states are φ -states.

So, on the right kind of models, modal logics can also deal with global conditional operators and therefore express Hansson’s dyadic operators $\mathbf{O}(\varphi | \psi)$ via Formula 5. Moreover, it is easy to add expressive power to this semantics and

⁷ Similar structures occur in other areas. Each P-graph plus valuation induces a Lewis system of spheres in conditional logic, or a doxastic Grove model. In this respect, priority structures have analogues in default logic, and belief entrenchment (Nayak *et al.*, 1996).

axiomatics. For instance, one can add modalities for strict betterness (cf. van Benthem *et al.*, 2009) and extend both the models and the axiomatization to deal with these extensions. The current article will limit itself to the simple language introduced in this section.

Remark 3 (Permission). In this article we will not concern ourselves with the deontic notion of permission or, as it is sometimes called in the deontic logic literature, *weak permission* (although we will provide an analysis of the more specific notion of *strong permission* in section 6.4). Still, it is worth observing that the logical set up of the previous sections naturally supports a definition of permission of φ under condition ψ as $\neg \mathbf{O}(\neg \varphi \mid \psi)$. That is, φ is permitted under condition ψ if and only if among the most ideal ψ -states there exists some state satisfying φ .

3.2 Incorporating priority structures: fitting in syntactic ‘best’

Our modal logic talks about semantic ordering. What about the syntactic priority structure of section 2?

Expressive power: defining syntactic ‘best’. Hansson’s semantics for dyadic obligations (Formula 1) relies on an ideality ordering on states and on the notion of maximality. If the ideality ordering is taken to be defined by a P-graph \mathcal{G} , as in Definition 2, Hansson’s semantics gets a matching intuition in terms of the priority structure of P-graphs.

Let $\mathcal{G} = \langle \Phi, \prec \rangle$ be a P-graph and $\mathcal{M} = \langle S, \preceq_{\mathcal{G}}, \mathcal{I} \rangle$ a model. To say that in \mathcal{M} , given \mathcal{G} , “it ought to be the case that φ under condition ψ ” amounts to saying that all ψ -states that satisfy a \prec -maximal formula of Φ also satisfy φ . Given a graph \mathcal{G} , this can be expressed in our modal language as follows:

$$[\mathbf{U}] \left(\left(\bigvee_{\langle \varphi_1, \dots, \varphi_n \rangle \in S_{\mathcal{G}}} \bigwedge_{1 \leq i \leq n} (\langle \mathbf{U} \rangle (\varphi_i \wedge \psi) \rightarrow (\varphi_i \wedge \psi)) \right) \rightarrow \varphi \right) \quad (6)$$

The formula has the structure $[\mathbf{U}](\chi(\psi) \rightarrow \varphi)$ of a subsumption, where $\chi(\psi)$ denotes the antecedent as a function of condition ψ . It therefore expresses a set-theoretic inclusion property: the truth-set of the complex antecedent $\chi(\psi)$ is included in the truth-set of the consequent φ . So what does the antecedent say? It denotes the set of states s such that, for some P-sequence $\langle \varphi_1, \dots, \varphi_n \rangle$ among the P-sequences defined by \mathcal{G} (i.e., the P-sequences in $S_{\mathcal{G}}$), for all φ_i in the sequence, if φ_i is compatible with condition ψ then $\varphi_i \wedge \psi$ is the case at s . In other words, the antecedent denotes the set of states that satisfy ψ and, at the same time, the best among all the formulae in \mathcal{G} which are compatible with ψ . So, altogether, Formula 6 states that the best formulae of \mathcal{G} compatible with condition ψ are included in φ . Applications of Formula 6 to concrete scenarios are discussed in Examples 3 and 4 below.

It must be stressed that the property expressed by Formula 6, unlike Formula 5, does not make use of operators interpreted on the semantic ordering $\preceq_{\mathcal{G}}$. It just uses the ‘syntactic’ information encoded by the P-graph \mathcal{G} .

Remark 4 (Non-monotonicity of Formula (6)). It is worth noticing that Formula (6) encodes a non-monotonic behavior, even though its syntactic form $([U](\chi(\psi) \rightarrow \varphi))$, at first inspection, seems to suggest the contrary. The syntactic form $\chi(\psi)$ makes the strengthening of ψ in the antecedent in $[U](\chi(\psi) \rightarrow \varphi)$ unsound. The reason is simply that $\chi(\psi \wedge \psi')$ does not logically imply $\chi(\psi)$.

Correspondence of syntactic and semantic ‘best’. A simple inspection of Formulae 5 and 6 reveals deep differences: while Formula 5 expresses a notion of maximality by using operators (like $[\preceq]$) that are interpreted on the underlying ‘semantic’ ideality ordering ($\preceq_{\mathcal{G}}$), Formula 6 resorts only to universal quantification and the ‘syntactic’ information given by the P-graph. However, they both express precisely the very same notion, namely the conditional maximality underpinning Hansson’s semantics for dyadic obligations:

Theorem 1 (Correspondence). *Let $\mathcal{G} = \langle \Phi, \prec \rangle$ be a P-graph, $\mathcal{M}_{\mathcal{G}}$ a model derived by Definition 2 from \mathcal{G} , \mathcal{I} a valuation function and s a state. The following three statements are equivalent:*

- i) $\max_{\preceq}(\llbracket \psi \rrbracket_{\mathcal{M}}) \subseteq \llbracket \varphi \rrbracket_{\mathcal{M}}$;
- ii) s satisfies Formula 5;
- iii) s satisfies Formula 6.

Proof. First of all, observe that $\mathcal{M}_{\mathcal{G}}$ is built on a preorder with an upward well-founded strict component by Fact 1. That i) and ii) are equivalent follows therefore from Fact 2. The equivalence of i) and iii) follows directly once we show that the set of $\preceq_{\mathcal{G}}$ -maximal ψ -elements coincides with the truth-set of the antecedent of Formula 6:

$$\max_{\preceq_{\mathcal{G}}}(\llbracket \psi \rrbracket_{\mathcal{M}_{\mathcal{G}}}) = \left\llbracket \bigvee_{(\varphi_1, \dots, \varphi_n) \in S_{\mathcal{G}}} \bigwedge_{1 \leq i \leq n} (\langle U \rangle(\varphi_i \wedge \psi) \rightarrow (\varphi_i \wedge \psi)) \right\rrbracket_{\mathcal{M}_{\mathcal{G}}}$$

Proof from left to right. Suppose that s is a $\preceq_{\mathcal{G}}$ -maximal state in $\llbracket \psi \rrbracket_{\mathcal{M}_{\mathcal{G}}}$. Then by Definition 2 there is a formula $\xi \in \Phi$ such that $s \in \llbracket \psi \wedge \xi \rrbracket_{\mathcal{M}_{\mathcal{G}}}$ and there is no $\xi' \in \Phi$ such that $\llbracket \psi \wedge \xi' \rrbracket_{\mathcal{M}_{\mathcal{G}}} \neq \emptyset$ and $\xi \prec \xi'$. Now, ξ must belong to a P-sequence $\langle \varphi_1, \dots, \varphi_n \rangle \in S_{\mathcal{G}}$ in $S_{\mathcal{G}}$. Since \mathcal{G} is ordered by logical entailment (Definition 1), and ξ is \prec -maximal among the φ s which have a non-empty intersection with $\llbracket \psi \rrbracket_{\mathcal{M}_{\mathcal{G}}}$, we have that, for all φ_i in the sequence, s satisfies $\langle U \rangle(\varphi_i \wedge \psi) \rightarrow (\varphi_i \wedge \psi)$.

Making the quantification explicit we thus have that s satisfies $\bigvee_{\langle \varphi_1, \dots, \varphi_n \rangle \in S_G} \bigwedge_{1 \leq i \leq n} (\langle U \rangle (\varphi_i \wedge \psi) \rightarrow (\varphi_i \wedge \psi))$. The direction from right to left is similar and is left to the reader. \square

The correspondence given by Theorem 1 adds an interesting layer to the logic of betterness we have introduced above. It shows that the standard Hansson's semantics is captured by statements of the form $[U](\chi(\psi) \rightarrow \varphi)$ where no reference to an ideality relation is made. These statements capture a natural way of reasoning about obligations: first, use the priority structure to find best properties compatible with the given conditions, then apply a simple subsumption check. Examples 3 and 4 below will illustrate the theorem further.

The rest of the article takes this two-level view of betterness to deontic logic, where priority-based reasoning arguably plays a central role in the way we commonly conceptualize the notion of obligation.

4. First Application: Contrary-to-Duty Reasoning

The section applies P-graphs – in their special form of P-sequences – and the modal language introduced above to classic topics in deontic logic: contrary-to-duty obligations (CTDs) and the so-called Kanger-Anderson reduction of deontic logic.

4.1 Introducing priorities: contrary-to-duty obligations

Along the lines of the St Paul's example that opened section 2, we show how P-sequences naturally arise in so-called contrary-to-duty scenarios.

Example 1 (Gentle murder). Here is our first example (Forrester, 1984, p. 194):

Let us suppose a legal system which forbids all kinds of murder, but which considers murdering violently to be a worse crime than murdering gently. [...] The system then captures its views about murder by means of a number of rules, including these two:

1. It is obligatory under the law that Smith not murder Jones.
2. It is obligatory that, if Smith murders Jones, Smith murders Jones gently.

The scenario mentions two relevant sets of states: those in which Smith does not murder Jones, represented by the formula $\neg m$; and those in which either Smith does not murder Jones or he does murder Jones, but gently, i.e., $\neg m \vee (m \wedge g)$. We therefore obtain a P-sequence $\mathcal{S} = \langle \Phi, \prec \rangle$ where $\Phi = \{\neg m, (\neg m \vee (m \wedge g))\}$ and \prec is such that $(\neg m \vee (m \wedge g)) \prec \neg m$.

Now, for any set of states S and valuation \mathcal{I} interpreting atoms m and g , Definition 2 gives us an induced betterness relation $\preceq_{\mathcal{S}}$. For instance, consider a state s satisfying formula $m \wedge g$ and a state t satisfying formula $\neg m$. We have that $s \preceq_{\mathcal{S}} t$ since for all properties φ in Φ (e.g., $\neg m \vee (m \wedge g)$) if s satisfies φ , so does




Figure 2. Gentle Murder as an ordered partition of a state space in clusters of equally preferred states, going from more to less ideal from left to right

t. Consequently, the relation orders states in three disjoint clusters, shown in Figure 2: the most ideal states satisfy $\neg m$, strictly worse but not worst are the states satisfying $m \wedge g$ and strictly worst are the states with $m \wedge \neg g$, that is, the states that do not satisfy any of the formulae mentioned in the P-sequence.

Using dyadic obligations (Formula 1), we can therefore represent the scenario by the set of formulae $\{\mathbf{O}(\neg m \mid \top), \mathbf{O}(m \mid g)\}$.

Example 2 (The Chisholm scenario). Here is our second example, reported in the formulation proposed by Åqvist (1967):

1. It ought to be that Smith refrains from robbing Jones.
2. Smith robs Jones.
3. If Smith robs Jones, he ought to be punished for robbery.
4. It ought to be that if Smith refrains from robbing Jones he is not punished for robbery.

Again, three items (1, 3 and 4) specify a priority sequence: $\neg r \vee (r \wedge p) \prec (\neg r \wedge \neg p)$ (where r stands for “Smith robs Jones” and p for “Smith is punished”). It is most ideal that Smith refrains from robbing Jones while Smith is not punished; and that the most ideal states under the assumption that Smith robs Jones are states in which Smith is punished.

Using dyadic obligations, we can therefore represent statements 1, 3 and 4 of the scenario by the set of formulae $\{\mathbf{O}(\neg r \mid \top), \mathbf{O}(p \mid r), \mathbf{O}(\neg p \mid \neg r)\}$. The main difference with Example 1 is the factual statement “Smith robs Jones” (item 2), whose informational point will be our cue for the novel dynamic analysis of the scenario in section 5.

It is worth making a methodological point about our formalization of the above deontic scenarios. Our use of P-graphs emphasizes a non-inferential aspect of deontic logic. Obligations and permissions are handled as usual given their induced normative ordering. But equally important are the criteria themselves that generate the normative order on which our judgements of obligation are based. This criterial structure supports many specific deontic inferences, and hence it may be considered

a more permanent part of what drives an agent's behaviour. In other words, our formalization should not just aim for extracting obligation and permission operators in a text, but also for cues as to the normative priorities – and what we should maintain generally, is not just a record of the reasoning: the text also helps us maintain a record of the relevant normative priorities.

4.2 CTDs and priority structures

The literature on deontic logic has used CTDs as benchmarks to evaluate logics of conditional obligation with respect to the sort of reasoning they support (see, for instance, Prakken and Sergot, 1996; Carmo and Jones, 2001). We think of P-graphs as the natural formalization of a CTD: some norms are given, and obligations are computed going 'down the line'. What is the difference here with standard semantic models? A mere betterness order on states has, so to speak, forgotten its origins – whereas now we have these still available as reasons for the ordering, and relevant exception zones.

Theorem 1 backs the above intuition by showing that conditional obligations (as in Formula 1) can be captured by simple subsumptions dictated by the structure of a priority graph. We now illustrate the theorem by elaborating on Examples 1 and 2.

Example 3 (Gentle murder, continued). Recall the P-sequence of the gentle murder $\mathcal{S}: (\neg m \vee (m \wedge g)) \prec \neg m$. Now Theorem 1 guarantees that, on models where the ideality relation is generated by this P-sequence, the CTD obligation $\mathbf{O}(g \mid m)$ is equivalent to the subsumption one obtains by instantiating Formula 6 through the following sequence of equivalent formulae:

$$\begin{aligned} & [\mathbf{U}] \left(\left(\bigvee_{\langle \varphi_1, \dots, \varphi_n \rangle \in \{S\}} \bigwedge_{1 \leq i \leq n} (\langle \mathbf{U} \rangle (\varphi_i \wedge m) \rightarrow (\varphi_i \wedge m)) \right) \rightarrow g \right) \\ & [\mathbf{U}] (((\langle \mathbf{U} \rangle ((\neg m \wedge m) \rightarrow (\neg m \wedge m))) \wedge (\langle \mathbf{U} \rangle ((m \wedge g) \wedge m) \rightarrow ((m \wedge g) \wedge m))) \rightarrow g) \\ & [\mathbf{U}] ((m \wedge g) \rightarrow g) \end{aligned}$$

In other words, on the preorders generated by the P-sequence \mathcal{S} , the statement $\mathbf{O}(g \mid m)$ is equivalent to a subsumption statement expressing that sub-ideal states (i.e., states where $m \wedge g$ is the case) are states where g is the case. On those models, the first statement is valid if and only if the second is.⁸

⁸ The reader might feel puzzled by the fact that $[\mathbf{U}]((m \wedge g) \rightarrow g)$ is logically valid ($(m \wedge g) \rightarrow g$ is a propositional validity) while $\mathbf{O}(g \mid m)$ is clearly not. Thus it may be worth stressing another reading of Theorem 1, as illustrated by this example. The fact that $[\mathbf{U}]((m \wedge g) \rightarrow g)$ is logically valid implies that preorders generated by the gentle murder P-sequence necessarily satisfy the truth conditions of statement $\mathbf{O}(g \mid m)$, and vice versa.

Example 4 (Chisholm scenario, continued). Similarly, the conditional obligations of Example 2 $\mathbf{O}(p \mid r)$ and $\mathbf{O}(\neg p \mid \neg r)$ are equivalent, on the preorders generated by the P-sequence of the Chisholm scenario to, respectively, $[\mathbf{U}](r \wedge p \rightarrow p)$ and $[\mathbf{U}](\neg p \wedge \neg r \rightarrow \neg p)$.

We are not claiming a totally new analysis of CTDs. Ideal versus ‘subideal’ zones occur in Jones and Pörn (1985) and Tan and van der Torre (1996), and even the classic Hansson (1969, section XIV) says this:

The problem of conditional obligation is what happens if somebody nevertheless performs a forbidden act. Ideal worlds are excluded. But, it may be the case that among the still achievable worlds, some are better than others. There should then be an obligation to make the best out of the sad circumstances.

Chains of properties are used for the representation of CTDs in van Fraassen (1973), Åqvist (1997)⁹ and Governatori and Rotolo (2005)¹⁰. Our contribution is doing this in a general format (P-graphs) that has proven to work in other areas, such as plausibility-based belief (Liu, 2008), as a prolific abstraction.

4.3 The Kanger-Anderson reduction and P-sequences

We now connect our analysis to a proposal from the founding period of deontic logic. Anderson (1957) and Kanger (1971) reduced deontic unconditional \mathbf{O} -formulae to alethic modal \Box -formulae with a constant for violation \mathbf{V} or ideality \mathbf{I} :

$$\mathbf{O}(\varphi \mid \top) := \Box(\neg\varphi \rightarrow \mathbf{V}) \quad (7)$$

$$\mathbf{O}(\varphi \mid \top) := \Box(\mathbf{I} \rightarrow \varphi). \quad (8)$$

This reductionist view has been criticized in the deontic logic literature by observing that it cannot accommodate a satisfactory representation of CTDs. We will show briefly how P-graphs, properly specialized, offer a natural extension to Anderson’s and Kanger’s proposals that does deal with CTDs along the lines they advocated.

9 Åqvist (1997) uses a special variant of P-sequences $\langle q_1, \dots, q_n \rangle$ where each q_i and q_j – called there *systematic frame constants* – are pairwise disjoint. So q_1 states are strictly better than q_2 and so on. Note that his is a special case of the sort of graphs we mentioned in Remark 1, and it is equivalent to our P-sequences. The interested reader is referred to Theorem 4 in the Appendix.

10 Governatori and Rotolo (2005) propose an approach to CTDs where a Gentzen calculus is developed for handling formulae of the type $\varphi_1 \otimes \dots \otimes \varphi_n$ with \otimes being a connective representing a reparatory obligation in a CTD structure: φ_1 ought to be the case but if φ_1 is violated then φ_2 ought to be the case and so on up to φ_n . Unlike this proof-theoretic approach, our approach is geared towards semantics and aims at connecting such CTD structures to modal logics interpreted on orders.

Definition 5 (KA-sequences). Let $\{l_1, \dots, l_n\} \subseteq \mathbf{P}$. A Kanger-Anderson sequence (“KA-sequence”) for $\mathcal{L}(\mathbf{P})$ is a sequence defined as follows:

$$\left\langle \bigvee_{1 \leq j \leq i} l_j \right\rangle_{1 \leq i \leq n}$$

So, KA-sequences are tuples $\langle l_1, l_1 \vee l_2, \dots, l_1 \vee \dots \vee l_n \rangle$ which are built by using ideality atoms to construct n layers spanning from the most to the least ideal.

The above Theorem 1 specializes to KA-sequences yielding the following interesting corollary, which highlights a generalization of the Kanger-Anderson reduction to CTD reasoning:¹¹

Corollary 1 (Obligations from better to worse). Let \mathcal{G} be a KA-sequence. For any model $\mathcal{M}_{\mathcal{G}}$, state s , and $1 \leq i < n$ it holds that:

$$\mathcal{M}_{\mathcal{G}}, s \models \mathbf{O}(\varphi \mid \top) \Leftrightarrow \mathcal{M}_{\mathcal{G}}, s \models [\mathbf{U}](l_1 \rightarrow \varphi) \quad (9)$$

$$\mathcal{M}_{\mathcal{G}}, s \models \mathbf{O}(\varphi \mid l_i) \Leftrightarrow \mathcal{M}_{\mathcal{G}}, s \models [\mathbf{U}](l_i \rightarrow \varphi) \quad (10)$$

$$\mathcal{M}_{\mathcal{G}}, s \models \mathbf{O}\left(\varphi \mid \bigvee_{1 \leq j \leq i} l_j\right) \Leftrightarrow \mathcal{M}_{\mathcal{G}}, s \models [\mathbf{U}]\left(\bigvee_{1 \leq j \leq i+1} l_j \rightarrow \varphi\right) \quad (11)$$

$$\mathcal{M}_{\mathcal{G}}, s \models \mathbf{O}\left(\varphi \mid \neg \bigvee_{1 \leq j \leq n} l_j\right) \Leftrightarrow \mathcal{M}_{\mathcal{G}}, s \models [\mathbf{U}]\left(\neg \bigvee_{1 \leq j \leq n} l_j \rightarrow \varphi\right) \quad (12)$$

Formula 9 says that an unconditional obligation $\mathbf{O}(\varphi \mid \top)$ is what the most ideal states dictate. The corollary shows how obligations change as we move from most to least ideal circumstances. In most ideal states, where l_1 holds, what ought to be the case is what already is the case (Formula 10). Formula 11 states that, if the i^{th} element has been violated, what ought to be is what follows from the $(i+1)^{\text{th}}$ element in the sequence. And in the least ideal states, where l_n is false, what ought to be the case is again what is already the case (Formula 12).

The latter point deserves some words of comment. KA-sequences (and P-sequences in general) are finite. This means that they can prescribe responses only to a finite number of sub-ideal classes of situations. Once even the last sub-ideal obligation has been violated, that is, l_n is false, ‘all is lost’ in the sense that the P-sequence is unable to prescribe a further response to this last violation. This is a

¹¹ A variant of Corollary 1 occurs in the preference logic of Åqvist (1997), with special propositions $\langle q_1, \dots, q_n \rangle$. See n. 13.

natural and, we argue, desirable feature as no system of norms can possibly handle infinite sequences of violations of increasing severity.

Remark 5 (Normative conflicts). This section has focused on CTDs and P-sequences. P-sequences represent ‘coherent’ systems of norms where no normative conflict arises. However, P-graphs can easily represent normative conflicts by generating preorders that are not total. Consider for instance the simple graph $\langle \{p, \neg p\}, \emptyset \rangle$ with two inconsistent formulae (p and $\neg p$) and an empty priority relation. By Definition 2 this P-graph generates a preorder consisting of two disconnected clusters of p and $\neg p$ -states. This represents a full-blown normative conflict whereby only validities (e.g., $p \vee \neg p$) are satisfied by the most ideal states (the union of the two clusters), failing to provide the agent with reasons for deciding between p and $\neg p$. We will return to P-graphs in section 6.

5. Second Application: Information Dynamics in Deontic Settings

So far we have proposed a more richly structured priority-based model for deontic notions under static circumstances. But deontic reasoning is crucially also about changes. To deal with this, we use some basic methods from dynamic-epistemic logic (*DEL*), a current framework for dealing with actions that change agents’ information about possible worlds, or their evaluation of these worlds.

5.1 Logical information dynamics

We just state some basics; much more information can be found in van Benthem (2011). Epistemic logic describes what agents know on the basis of their current semantic information. Public announcement logic (*PAL*), the simplest case of *DEL*, combines epistemic logic with one dynamic event, namely, the *announcement* of new ‘hard information’ expressed in some proposition φ . The corresponding action $!\varphi$ transforms a current epistemic model (\mathcal{M}, s) into its submodel $(\mathcal{M}|\varphi, s)$ where all worlds that did not satisfy φ have been eliminated. This reflects the basic intuition of information gain, both in science and in common sense, as shrinking one’s current epistemic range of uncertainty. The typical new dynamic formula $[!\varphi]\psi$ of this system says that “after announcing the true proposition φ , formula ψ holds”. Here is its semantics:

$$\mathcal{M}, s \models [!\varphi]\psi \Leftrightarrow \text{if } \mathcal{M}, s \models \varphi \text{ then } \mathcal{M}|\varphi, s \models \psi. \quad (13)$$

PAL has been used as a pilot example for the analysis of a variety of epistemic changes capturing the many subtleties involved in, e.g., belief revision (van Benthem, 2004), but also of deontic changes (van Benthem and Liu, 2007). Agents must constantly cope with changes in information because they learn more.

In particular, deontic dynamics under hard information extends our modal axiom system for betterness reasoning in section 3 with further axioms for the dynamic modalities $[!\varphi]$. A typical reduction axiom of the latter kind is this:

$$[!\varphi][\preceq]\psi \leftrightarrow (\varphi \rightarrow [\preceq][!\varphi]\psi) \quad (14)$$

That is, ψ holds in all better states after the announcement of φ if and only if either φ is not the case or all better states are such that, after the announcement of φ , they satisfy ψ .¹² Given the definition of conditional obligation by Formula 5, one can then derive the following reduction axiom:

$$[!\varphi]\mathbf{O}(\psi \mid \chi) \leftrightarrow (\varphi \rightarrow \mathbf{O}([!\varphi]\psi \mid (\varphi \wedge [!\varphi]\chi))) \quad (15)$$

As a special case for unconditional obligation, setting χ to the always true proposition \top , we get the reduction axiom:

$$[!\varphi]\mathbf{O}\psi \leftrightarrow (\varphi \rightarrow \mathbf{O}([!\varphi]\psi \mid \varphi)) \quad (16)$$

Factual formulas Keeping track of the dynamic modalities in this way is crucial here since complex modal statements may change their truth values when a model changes. However, our subsequent illustrations will usually be about purely propositional *factual formulas*, whose truth values cannot change under update of models, and then some simplifications will arise.

We will now show how this simple dynamic logic allows us to model some basic deontic phenomena involved with information change. Dynamic-epistemic methods can also model more properly normative changes in how agents evaluate situations or worlds, but we postpone these till later on.

5.2 Information dynamics in conditional obligations

We have seen that obligations are typically conditional, so changes in circumstances determine changes in what ought to be the case. Semantically, this means that maximally ideal states change under different circumstances, while syntactically, this means that properties in the priority structure that are incompatible with current circumstances can be disregarded. This is often called *factual detachment* in deontic logic: conditional obligations remain stable, but what changes is what follows from them under different circumstances. As a result, obligations are subjected to what is sometimes referred to as ‘factual defeasibility’ (van der Torre and Tan, 1997): new information about the world

¹² It is not difficult to see that Formula 14 is valid on preorders. Proofs can be found in van Benthem (2011) and Liu (2011a).

may change our obligations. We show now how our framework naturally captures these phenomena.

We start by revisiting an earlier scenario:

Example 5 (The Chisholm scenario: a dynamic perspective). The Chisholm scenario (Example 2) consisted of three normative statements:

- It ought to be that Smith refrains from robbing Jones.
- If Smith robs Jones, he ought to be punished for robbery.
- It ought to be that if Smith refrains from robbing Jones he is not punished for robbery.

plus a factual one:

- Smith robs Jones.

Now, intuitively, there is a difference in function here. The normative statements seem global guides to behaviour, but the scenario suggests a dynamic reading of the factual statement. One way of doing so is by taking the robbery to be an event that takes place, changing the world.¹³ But in this article, we rather want to model deontic deliberation, where relevant facts become known, and enter the reasoning. In this line, the acquisition of the factual statement is a dynamic-epistemic event where new information becomes available and it becomes then settled that “Smith robs Jones”. This triggers precise normative consequences, namely that Smith ought to be punished. Thus, a conditional obligation is ‘in suspended animation’ until we get the hard information settling that its antecedent obtains. Then the model changes, and the conditional obligation becomes an absolute one. The Chisholm’s scenario is an excellent example of this phenomenon, and it can be formalized using logical dynamics as follows:

$$(\mathbf{O}(\neg r \mid \top) \wedge \mathbf{O}(p \mid r)) \rightarrow [!r]\mathbf{O}(p \mid \top). \quad (17)$$

Intuitively: if it ought to be the case that $\neg r$, and that p if r , then, after it is settled that r is the case, it ought to be the case that p . The formula is valid in all betterness **S4** models derived by the priority sequence of the Chisholm scenario.

Formula 17 is a special case of the following validity of our logic, that holds for factual formulae φ and ψ :¹⁴

$$\mathbf{O}(\varphi \mid \psi) \rightarrow [!\psi]([U]\psi \wedge \mathbf{O}(\varphi \mid \top)) \quad (18)$$

If we interpret the universal modality $[U]$ as an epistemic operator, as not uncommon in epistemic logic, Formula 18 formalizes the following intuitive principle:

¹³ Modelling physical events would take us to temporal deontic logics (van Eck, 1982; Horty, 2001).

¹⁴ The restriction to factual formulas is a useful simplification, but it is not essential: cf. Liu (2011a).

“If it ought to be the case that ϕ under condition that ψ then, if it is announced that ψ is the case, it is known that ψ and it ought to be (unconditionally) the case that ϕ ”. This suggests a natural form of conditional obligations taken in an epistemic sense: I should do something if I *know* the antecedent to be the case. A number of telling cases of such epistemic notions behind deontic logic were proposed in Pacuit *et al.* (2006).

Turning from concrete scenarios to general reasoning, our dynamic logic also illuminates deontic discussions about putative inference principles. The so-called *factual detachment* in deontic logic has the form

$$(\mathbf{O}(\phi \mid \psi) \wedge \psi) \rightarrow \mathbf{O}(\phi \mid \top). \quad (19)$$

Hansson’s semantics clearly does not validate this principle, which is therefore unreasonable for a maximality-based view of obligations.¹⁵ However, the following is valid, which highlights a completely different view on how conditions in conditional obligations can be formalized to capture a detachment principle:

$$\mathbf{O}(\phi \mid \psi) \rightarrow [!\psi]\mathbf{O}(\phi \mid \top) \quad (20)$$

for ψ, ϕ factual. So, compare Formula (20) with Formula (19). To allow for factual detachment under a Hansson-type dyadic obligation, a condition ψ is not just taken to be the case, but it is taken to be *settled* or *known*. Finally, a static principle that does come close to Formula (20) is:

$$(\mathbf{O}(\phi \mid \psi) \wedge [\mathbf{U}]\psi) \rightarrow \mathbf{O}(\phi \mid \top) \quad (21)$$

where $[\mathbf{U}]$ is again interpreted as an epistemic operator. We leave its derivation in our logic to the reader.

It must be stressed that this dynamic take on factual detachment and the Chisholm’s scenario is a novel contribution to the formalization of the principles underpinning reasoning with CTDs. Our dynamic logic is a style of deontic reasoning mixing effects of informational events with unpacking of duties. This, we have shown, can extend in a useful manner the repertoire of principles deontic logic is traditionally about.

Remark 6 (Non-monotonicity of conditional obligations). In the context of the above discussion, it is fitting to say a few words about one of the key critiques that

¹⁵ A simple countermodel is $\mathcal{M} = \langle \{s, s'\}, \preceq, \mathcal{I} \rangle$ with $s \preceq s'$ and $\mathcal{I}(p) = \mathcal{I}(q) = s$. Then $\mathcal{M}, s \models q \wedge \mathbf{O}(p \mid q)$ but $\mathcal{M}, s \not\models \mathbf{O}(p \mid \top)$, as the maximal state s' falsifies p . Detachment makes sense in special conditional or epistemic settings that require assumptions of “centring” for worlds around the current one.

has been moved in the deontic logic literature against the formalization of conditional obligations *à la* Hansson. Conditionals $\mathbf{O}(\varphi \mid \psi)$ are clearly non-monotonic. From our epistemic standpoint, this means that if $\mathbf{O}(\varphi \mid \psi)$, by learning $\psi \wedge \chi$ one cannot conclude that $\mathbf{O}(\varphi \mid \top)$, i.e.: $\mathbf{O}(\varphi \mid \psi) \rightarrow [!(\psi \wedge \chi)]\mathbf{O}\varphi$ is not valid. A neatly formulated critique to this aspect of our conditionals can be found in Horty (1997), where it is argued that although monotonicity is not a desirable principle for conditional obligation, some limited form of monotonicity should still be available to draw ‘safe’ monotonic conclusions in the logic, e.g., via some non-standard consequence relation.

The key point is made by using these two statements: 1) You ought to put your napkin on your lap; 2) If you are served asparagus, you ought to eat it with your fingers. In our logic this would be represented by $\mathbf{O}(n \mid \top)$ and $\mathbf{O}(f \mid a)$. Now the criticism of Horty (1997) is that the logic does not support the inference to $\mathbf{O}(n \mid a)$, i.e., if you are served asparagus, you ought to put your napkin on your lap. This is correct; however, it must be observed that to enable the inference it suffices to assume that there exist maximally ideal states where you are served asparagus or, in other words, it is not incompatible with ideality – it is *permitted* (cf. Remark 3) – that you are served asparagus. We argue this is a rather innocuous and intuitive assumption (and left unspecified in Horty, 1997). By $\mathbf{O}(n \mid \top)$ and $\mathbf{O}(f \mid a)$, this assumption allows you to infer that, if you learn that you are served asparagus, then it ought to be the case you put your napkin on your lap:

$$(\mathbf{O}(\chi \mid \top) \wedge \mathbf{O}(\varphi \mid \psi) \wedge \neg \mathbf{O}(\neg \psi \mid \top)) \rightarrow [!\psi]\mathbf{O}(\chi \mid \top)$$

We hope this brief discussion further strengthens our claim that conditional obligations *à la* Hansson are a solid ground on which to base the investigations proposed in this article and a versatile logical set up allowing natural and insightful interfaces with other logics such as dynamic-epistemic logic.

5.3 Information and priority

Up to now, our discussion of information dynamics was about conditional obligations at the base level of deontic ideality relations. But information dynamics also shows at the level of priority graphs. To conclude this section, we define an operation on P-graphs that matches the $[!\varphi]$ relational update modalities.

Definition 6 (P-graph restriction). Let $\mathcal{G} = \langle \Phi, \prec \rangle$ be a P-graph, and ψ a formula. The restriction of \mathcal{G} by ψ is the graph $\mathcal{G}^\psi = \langle \Phi^\psi, \prec^\psi \rangle$ where:

- $\Phi^\psi = \{ \varphi \wedge \psi \mid \varphi \in \Phi \};$
- $\prec^\psi = \{ (\varphi \wedge \psi, \varphi' \wedge \psi) \mid \varphi \prec \varphi' \}.$

The restriction of a P-graph \mathcal{G} by ψ simply intersects the elements of the original graph with ψ and keeps the original order:

Theorem 2 (Harmony of P-graph restriction). *The following diagram commutes for all P-graphs \mathcal{G} , propositional formula φ and valuation \mathcal{I} :*

$$\begin{array}{ccc}
 \mathcal{G} & \xrightarrow{\quad} & \mathcal{G}^\psi \\
 \text{sub} \downarrow & & \downarrow \text{sub} \\
 \langle S, \preceq_{\mathcal{G}}, \mathcal{I} \rangle & \xrightarrow{!_{\psi}} & \langle \llbracket \psi \rrbracket, \preceq_{\mathcal{G}^\psi}, \mathcal{I}|_{\psi} \rangle
 \end{array}$$

Proof. Let $s \preceq_{\mathcal{G}} s'$ and $s, s' \in \llbracket \psi \rrbracket$. By Definition 2, for all $\varphi \in \Phi$ if $s \in \llbracket \varphi \rrbracket$ then $s' \in \llbracket \varphi \rrbracket$. Hence, if $s \in \llbracket \varphi \wedge \psi \rrbracket$ then $s' \in \llbracket \varphi \wedge \psi \rrbracket$, so $s \preceq_{\mathcal{G}^\psi} s'$. The other direction is similar. \square

In other words, the order $\preceq_{\mathcal{G}}$ obtained by update of φ is the same order obtained by graph restriction, and the information dynamics at our two levels lives in harmony. As a further illustration, here is an instance of Theorem 1 in this setting of information dynamics:

$$\mathcal{M}_{\mathcal{G}}|_{\psi}, s \models \mathbf{O}_{\preceq}(\varphi | \xi) \Leftrightarrow \mathcal{M}_{\mathcal{G}^\psi}, s \models \mathbf{O}_{\mathcal{G}}(\varphi | \psi).$$

This relates conditional obligation based on a maximality statement about $\preceq_{\mathcal{G}}$ after an announcement ψ (i.e., in model $\mathcal{M}_{\mathcal{G}}|_{\psi}$) with the conditional obligation based on a subsumption statement ($\llbracket \mathbf{U} \rrbracket(\xi(\psi) \rightarrow \varphi)$) about the restricted priority graph $\mathcal{G}|_{\psi}$.¹⁶

This gives us a first example of how the two-level perspective we have developed in section 3 and applied to static scenarios in section 4 has a natural dynamic extension. The next section will push this line further analysing the sort of dynamics that arises from genuinely normative updates.

6. Third Application: Deontic Dynamics Proper

Deontically relevant events are of many kinds. Some of them are purely informational, as we have seen already. Others are about changes in evaluation of worlds, whether at base level, or at the level of priority structure. We will look at some examples of the latter kind, and then show how to model them at various levels that stay in harmony.

¹⁶ A more radical view germane to this article would make informational events syntactic in graphs, in the spirit of the evidence dynamics of van Benthem and Pacuit (2011).

6.1 Generating priorities in concrete scenarios

To get some examples, we add a dynamic twist to our earlier treatment of CTDs. This time, consider how their normative component arises:

Example 6 (Postfixing norms). Classic deontic scenarios come with a normative priority order, but the latter has usually been created. For instance, start the earlier Gentle Murder scenario with the P-sequence

$$\langle \neg m \rangle$$

By Definition 2, this generates a total preorder with all $\neg m$ states above all m states: “It is obligatory under the law that Smith not murder Jones”. Suppose this is the given deontic situation. Now, a lawgiver, or a person with moral authority comes in, and introduces the sub-ideal obligation to kill gently: “it is obligatory that, if Smith murders Jones, Smith murders Jones gently”? This can be done by postfixing the original sequence with the property $\neg m \vee g$:

$$\langle \neg m, \neg m \vee g \rangle$$

to obtain the sequence encountered in Example 1.

Placing the new moral priority last leaves the original prohibition on murder intact. There may also be cases where we want to do the opposite.

Example 7 (Prefixing norms). Recall the P-sequence of the gentle murder case:

$$\langle \neg m, \neg m \vee g \rangle$$

Now we want to introduce a stronger norm than just “It is obligatory under the law that Smith not murder Jones”, like “It is obligatory under the law that Smith not murder Jones and that Smith not be aggressive against Jones”. Also we do not want to introduce any further change in the priority ordering. This can be achieved syntactically by prefixing the P-sequence with the property $\neg m \wedge \neg a$, where a stands for “Smith is aggressive against Jones”:

$$\langle (\neg m \wedge \neg a), \neg m, \neg m \vee g \rangle$$

These examples suggest the following definition of general pre- and postfixing operations on P-graphs:

Definition 7 (Prefixing and postfixing in P-graphs). Let $\mathcal{G} = \langle \Phi, \prec \rangle$ be a P-graph, and ϕ a propositional formula:

- the prefixing of \mathcal{G} by ϕ yields the graph $\phi; \mathcal{G}$ where a new maximal element $\phi \wedge \bigwedge \max(\mathcal{G})$ is added to \mathcal{G} , consisting of the conjunction of ϕ with the conjunction of the maximal elements of \mathcal{G} ;
- the postfixing of \mathcal{G} by ϕ yields the graph $\mathcal{G}; \phi$ where a new minimal element $\phi \vee \bigvee \min(\mathcal{G})$ is added to \mathcal{G} , consisting of the disjunction of ϕ with the disjunction of the minimal elements of \mathcal{G} .




Figure 3. Hasse diagrams illustrating the pre- and post-fixing by φ of the P-graph on the right-hand side of Figure 1

These operations are illustrated in Figure 3. Observe that the constraints defining P-graphs – namely the logical dependencies of the elements of the graph – are respected by the definition, and that the above examples are special cases of the definition with $\varphi := g$ (Example 6) and, respectively, $\varphi := \neg a$ (Example 7).

6.2 Logic of base level betterness change

Modifying priority structure has concrete effects on betterness order of states induced in the manner of section 2. These base-level changes can also be studied directly, by means of relation transformers in dynamic-epistemic style. Here is a typical example that has been used widely in the literature, from belief revision to learning theory. It may be viewed as a ‘strong command’ in favour of realizing some proposition φ :

Definition 8 (Radical upgrade). Let $\mathcal{M} = \langle S, \preceq, \mathcal{I} \rangle$ be a model and φ be a propositional formula. A radical upgrade $\uparrow \varphi$ yields model $\mathcal{M}_{\uparrow \varphi} = \langle S, \preceq_{\uparrow \varphi}, \mathcal{I} \rangle$ where $\preceq_{\uparrow \varphi}$ is equal to:

$$\{(s, t) \in \preceq \mid s, t \in \llbracket \varphi \rrbracket\} \cup \{(s, t) \in \preceq \mid s, t \notin \llbracket \varphi \rrbracket\} \cup \{(s, t) \in S^2 \mid s \notin \llbracket \varphi \rrbracket \text{ and } t \in \llbracket \varphi \rrbracket\}$$

In other words, a radical upgrade $\uparrow \varphi$ changes the current order \preceq to that of a new model $\mathcal{M}_{\uparrow \varphi}$ where all φ -states become better than all $\neg\varphi$ -states, while, within those two zones, the old ordering remains intact.¹⁷ As a simple example of radical upgrade, consider a model consisting of two equally ideal states: s satisfying $\neg p$ and t satisfying p . Now a radical upgrade $\uparrow p$ modifies the order by placing t on top of s . In the resulting model formula $\mathbf{O}(p \mid \top)$ is valid, while it fails in the original model (as s is a maximal state which satisfies $\neg p$). Example 9 below offers a more elaborated illustration.

It is easy to see that radical upgrade preserves preorders and is thus well-defined for our semantics. Many other operators are definable in this format, including weaker commands or just ‘suggestions’: cf. Yamada (2008). This variety in normative force is deontically relevant, since agents are exposed to many normative cues, from direct orders to mere hints. It will not be our main topic here, however, and radical upgrade will do for making most of our points.

Remark 7 (Modal syntax for radical upgrade). The operation of radical upgrade does not change the domain of models, like the earlier public announcements of hard information: it changes the normative order. Still, reasoning with such actions falls within the scope of dynamic-epistemic logic, and Liu (2011a) shows how they support reduction axioms for absolute and conditional obligations. We merely display an example of axiom for a modal operator of radical upgrade (cf. Formula (14)):

$$\begin{aligned} [\uparrow \varphi][\preceq]\psi &\leftrightarrow (\varphi \wedge [\preceq](\varphi \rightarrow [\uparrow \varphi]\psi)) \\ &\vee (\neg\varphi \wedge [\preceq](\neg\varphi \rightarrow [\uparrow \varphi]\psi)) \wedge [\mathbf{U}](\varphi \rightarrow [\uparrow \varphi]\psi) \end{aligned} \quad (22)$$

Similar reductions for obligation operators can be obtained from this axiom and Formula 5. As with public announcement logics, this allows for much finer analysis of deontic reasoning, now looking also at how one should act in the presence of a much richer variety of normative events that can change one’s obligations.

6.3 Harmony

We have seen that deontic dynamics can be located both at the level of P-graphs and at the level of their underlying states. Both represent natural ways of thinking about deontic changes. What is the relation between the two levels? Like for the static and information dynamics cases, we obtain an harmony theorem:

17 A more compact but efficacious formulation of $\preceq^{\uparrow \varphi}$ can be given by the following regular expression: $\preceq^{\uparrow \varphi} = (? \varphi; \preceq; ? \varphi) \cup (? \neg \varphi; \preceq; ? \neg \varphi) \cup (? \neg \varphi; S^2; ? \varphi)$.

Theorem 3 (Harmony of P-graph pre-postfixing). *The following diagram commutes for all P-graphs \mathcal{G} , propositional formulae φ and valuations \mathcal{I} :*

$$\begin{array}{ccc}
 \mathcal{G} & \xrightarrow{\star\varphi} & \mathcal{G} \star \varphi \\
 \text{sub} \downarrow & & \downarrow \text{sub} \\
 \langle S, \preceq_{\mathcal{G}}, \mathcal{I} \rangle & \xrightarrow{\uparrow f_{\star}(\varphi)} & \langle S, \preceq_{\mathcal{G} \star \varphi}, \mathcal{I} \rangle
 \end{array}$$

where $\mathcal{G} \star \varphi$ denotes either the pre-fixing φ ; \mathcal{G} or the post-fixing \mathcal{G} ; φ of \mathcal{G} by φ and $f_{\star}(\varphi)$ denotes accordingly $\varphi \wedge \bigwedge \max(\mathcal{G})$ or $\varphi \vee \bigvee \min(\mathcal{G})$.

Proof (Sketch). Consider the case of post-fixing \mathcal{G} ; φ . We have two cases: 1) $\llbracket \varphi \vee \bigvee \min(\mathcal{G}) \rrbracket = \llbracket \bigvee \min(\mathcal{G}) \rrbracket$, in which case clearly $\preceq_{\mathcal{G}} = \preceq_{\mathcal{G} \star \varphi}$ and $\preceq_{\mathcal{G} \star \varphi} = \preceq_{\uparrow \varphi}$. 2) $\llbracket \varphi \vee \bigvee \min(\mathcal{G}) \rrbracket \supset \llbracket \bigvee \min(\mathcal{G}) \rrbracket$. By Definitions 2 and 7, the total preorder $\preceq_{\mathcal{G} \star \varphi}$ consists therefore of the same equivalence classes of $\preceq_{\mathcal{G}}$ except for splitting the class of states satisfying $\neg \bigvee \min(\mathcal{G})$ into two classes: consisting of $\varphi \wedge \neg \bigvee \min(\mathcal{G})$ -states and one (the bottom one) of $\neg \varphi \wedge \neg \bigvee \min(\mathcal{G})$ -states. Now consider $\preceq_{\uparrow(\varphi \vee \bigvee \min(\mathcal{G}))}$. Since $\varphi \vee \bigvee \min(\mathcal{G})$ is weaker than all the properties in \mathcal{G} , it will not affect the ordering $\preceq_{\mathcal{G}}$ except for splitting in two equivalence classes the class of worst elements (i.e., the set of $\neg \bigvee \min(\mathcal{G})$ -states). These two classes are: one, the $\neg \bigvee \min(\mathcal{G})$ -states which satisfy φ ; and two, the $\neg \bigvee \min(\mathcal{G})$ -states which do not satisfy φ . The case for pre-fixing is similar. \square

A concrete illustration of the theorem can be found in Example 9 below. Thus, we have seen how both informational and normative events can be represented in our priority framework, while also showing how the effects of these are in harmony with a natural base-level dynamics of betterness change on modal state models. As we have seen in a few examples, this offers a much richer view of what deontic scenarios actually involve, and how their normative structures are constructed and modified.

6.4 Case study: strong permission and CTDs

We conclude with one more concrete example of how a dynamic priority setting suggests new takes on old problems in deontic logic. The distinction between weak and strong permission dates back to von Wright (1963):

An act will be said to be permitted in the weak sense if it is not forbidden; and it will be said to be permitted in the strong sense if it is not forbidden but subject to norm. [...] Weak permission is not an independent norm-character. Weak permissions are not prescriptions or norms at all. Strong permission only is a norm-character. (p. 86)

Thus, the weak permission “it is permitted that φ ” amounts to mere absence of the prohibition “it is forbidden that φ ”, which is definable as $\neg \mathbf{O} \neg \varphi$. But – and that is

the quote's claim – this is not the case for strong permission. How can we do justice to this distinction?¹⁸ Defining strong permission is listed in Hansen *et al.* (2007) as one of ten major 'philosophical problems' in deontic logic.

Strong permission does not seem to be one single deontic act. On the just noted epistemic analogy, it could be viewed as opening up a new relevant possibility that had not been considered previously. It may also be viewed as allowing *actions* rather than endorsing possibilities, making comparisons with propositional "may" less immediate. We have nothing to say about these senses, even though they look congenial to a dynamic perspective. But the Von Wright quote in terms of "subject to norm" also highlights another sense of strong permission as related to normative behaviour, more in line with legal theory:

Telling me what I am permitted to do provides no guide to conduct unless the permission is taken as an exception to a norm of obligation [. . .]. Norms of permission have the normative function only of indicating, within some system, what are the exceptions from the norms of the obligation of the system. (Ross, 1968, p. 120)

Viewing strong permission as exception-making raises the issue of what "exception" means. On a radical view, it restricts the action of some earlier norm; on a less radical view, it introduces a new subnorm:

Example 8 (Killing in self-defence). Let us start again with the gentle murder scenario, slightly rephrased:

It is obligatory that Smith not kill Jones.

Like in the examples at the beginning of this section, we want to change this norm, now by a permission stating that:

Smith is allowed to kill Jones, provided he does that in self-defence.

We abbreviate "killing" with k and "killing in self-defence" with d .

There are two ways of interpreting the effects of such a strong permission:

1. The permission gives up the validity of $\mathbf{O}(\neg k \mid \top)$ for that of the weaker obligation $\mathbf{O}(k \rightarrow d \mid \top)$. Syntactically, this removes $\neg k$ from the P-sequence and substitutes $k \rightarrow d$. In this case, a strong permission repeals an earlier obligation, and then introduces a weaker one widening the field of permissibility (Hilpinen, 1980). This is close to an act of "derogation" in the law:

The difference between weak and strong permission becomes clear when thinking about the function of permissive norms. [. . .] A permissive norm is necessary when we have to repeal a preceding imperative norm or to derogate to it. That is to abolish a part of it [. . .]. (Bobbio, 1980, pp. 891–892)¹⁹

18 Strong permission does not stand alone, and it has analogues in other areas, such as strong epistemic "may" highlighting an epistemic possibility. We will not go into all natural language meanings of this term.

19 Quoted in the paper on permission and obligation (Boella and van der Torre, 2003).

2. The strong permission does not modify $\mathbf{O}(\neg k \mid \top)$, but introduces a CTD stating that, in case Smith kills Jones, he should do that in self-defence: $\mathbf{O}(d \mid k)$. Example 9 will show how this can be analysed as an instance of Theorem 3 where the CTD is introduced either by means of a radical upgrade on the ideality ordering or as an operation on the P-sequence.

Both options have been object of investigations in the deontic logic literature (e.g., Brown, 2000, and Aucher *et al.*, 2009, for the first option and Makinson and van der Torre, 2003, and Stolpe, 2010, for the second). Here we will illustrate the second one in further detail:

Example 9 (Two-level dynamics of strong permission). Following Example 8 consider a P-sequence $\langle \neg k \rangle$, whose derived model \mathcal{M} validates $\mathbf{O}(\neg k \mid \top)$. Enacting a strong permission to kill if acting in self-defence can be modelled by introducing a modified CTD that makes killing satisfying some specified condition (here, self-defence) better than killing when those conditions are violated: $\neg k \vee (k \wedge d)$, i.e., $k \rightarrow d$. What we obtain then is an instance of the earlier harmony theorem in its post-fixing format, for any valuation \mathcal{I} :

$$\begin{array}{ccc}
 \langle \neg k \rangle & \xrightarrow{\quad} & \langle \neg k \rangle; d \\
 \text{sub} \downarrow & & \text{sub} \downarrow \\
 \langle S, \preceq_{\langle \neg k \rangle}, \mathcal{I} \rangle & \xrightarrow{\uparrow (\neg k \vee d)} & \langle S, \preceq_{\langle \neg k, k \rightarrow d \rangle}, \mathcal{I} \rangle
 \end{array}$$

So, introducing a strong permission to kill in self-defence as a CTD that refines the previous unconditional obligation not to kill ($\mathbf{O}(\neg k \mid \top)$) can be viewed equivalently as:

- an operation of radical upgrade $\uparrow (k \rightarrow d)$ on the underlying betterness ordering which upgrades d states among the k -states and leaves $\neg k$ -states on top (recall Definition 8);
- an operation of postfixing of the original P-sequence by d , thereby generating the new P-sequence $\langle \neg k, k \rightarrow d \rangle$.

In other words, strong permission in our exception sense may be modelled by inserting a predicate in a P-graph at some specified position (in this case the first, or the last position).

Our treatment by no means exhausts all deontic views on strong permission,²⁰ but it does emphasize a link with CTDs and their dynamics which, to the best of our knowledge, had not yet been discussed. Our claim here is that the process of incremental specification of a CTD sequence via betterness update, be it syntactic or semantic, can legitimately be viewed as the enactment of strong permissions. Such permissions are refinements of existing obligations or, to say it otherwise, exceptions to existing obligations that do not reject such obligations altogether, but rather specify conditions under which a violation of such obligations is tolerable. Notice also that, against the intuitions that might be dictated by the natural language formulation of the notion, this interpretation of strong permission has actually more to do with **O**-statements of obligation (albeit of a CTD type) rather than with $\neg\mathbf{O}\neg$ -statements of (weak) permission.

6.5 Generalizing deontic dynamics: graph composition

We conclude this section with an observation which, we believe, should set the stage for the future development of the set of tools presented in this section. We have studied how formulae can update P-graphs and their associated preorders. However, a formula φ is nothing but a special P-graph $\mathcal{G}' = \langle \varphi \rangle$ consisting of one single property. So what we have seen up till now is how to update a given graph \mathcal{G} by the special graph $\langle \varphi \rangle$. The natural question arises then of how two graphs \mathcal{G} and \mathcal{G}' can be combined in general and which preorders they induce.

This is not just of technical interest, but it is interesting from the point of view of deontic logic as an abstract setting from which to study how norms – viewed as P-graphs – can be combined to form new norms. Reverting to our running example, here is a concrete illustration of what we have in mind.

Example 10 (Quick murder). Let us assume now there are two normative sources. According to the first one:

1. It is obligatory under the law that Smith not murder Jones.
2. It is obligatory that, if Smith murders Jones, Smith murders Jones gently.

According to the second one:

1. It is obligatory under the law that Smith not murder Jones.
2. It is obligatory that, if Smith murders Jones, Smith murders Jones *quickly*.

We want to merge the two CTDs and obtain a P-graph which contains the normative information of both components. Assuming the alphabet $\{m, g, q\}$ with the obvious

20 For instance, Bulygin (1986), Boella and van der Torre (2008) and Makinson and van der Torre (2003) pursue a view where strong permissions set boundaries to any prohibitions that normative authorities could possibly enact.




Figure 4. Merging of P-sequences in Example 10. The merging is depicted both in the form of parallel composition of P-sequences (top) and as intersection of total preorders in the form of spheres (bottom)

intuitive interpretation, we can model this scenario by means of two P-sequences \mathcal{S}_g with $\neg m \prec \neg m \vee (m \wedge g) (= m \rightarrow g)$, and \mathcal{S}_q with $\neg m \prec \neg m \vee (m \wedge q) (= m \rightarrow q)$. Formally, the graph operation we have in mind consists of taking the disjoint union of \mathcal{S}_g and \mathcal{S}_q – also called *parallel composition* (Liu, 2011a) – which we denote by symbol “ $||$ ”. The resulting preorder should be such that the best states are $\neg m$ -states and the sub-ideal states are split into two incomparable classes, the class of $m \wedge g$ -states and the class of $m \wedge q$ -states (see Figure 4).²¹ This is nothing but the intersection of the betterness (total) preorders of the two P-sequences.

The example instantiates a general result, from Andreka *et al.* (2002), relating the parallel composition of P-graphs as disjoint union and the intersection of their derived preorders:

Fact 3 (Harmony of parallel composition of P-graphs) Let $\mathcal{G} = \langle G, \prec \rangle$ and $\mathcal{G}' = \langle G', \prec' \rangle$ be two P-graphs. The following diagram commutes:

$$\begin{array}{ccc}
 \mathcal{G} & \xrightarrow{||\mathcal{G}'} & \mathcal{G}||\mathcal{G}' \\
 \text{sub} \downarrow & & \downarrow \text{sub} \\
 \langle S, \preceq_{\mathcal{G}} \rangle & \xrightarrow{\cap \preceq_{\mathcal{G}'}} & \langle S, \preceq_{\mathcal{G}||\mathcal{G}'} \rangle
 \end{array}$$

21 The incomparability requirement here models the fact that the resulting ordering contains some element of conflict (cf. Remark 5) deriving from the unrelatedness of the two normative sources.

Proof. The proof is given by the following equivalences obtained by iterated application of Definition 2, for any valuation:

$$\begin{aligned}
 s \preceq_{\mathcal{G} \parallel \mathcal{G}'} s' &\Leftrightarrow \forall \varphi \in G \cup G' : s \in \llbracket \varphi \rrbracket \Rightarrow s' \in \llbracket \varphi \rrbracket \\
 &\Leftrightarrow \forall \varphi \in G : s \in \llbracket \varphi \rrbracket \Rightarrow s' \in \llbracket \varphi \rrbracket \text{ and } \forall \varphi \in G' : s \in \llbracket \varphi \rrbracket \Rightarrow s' \in \llbracket \varphi \rrbracket \\
 &\Leftrightarrow s \preceq_G s' \text{ and } s \preceq_{G'} s'.
 \end{aligned}$$

This completes the proof. \square

Remark 8 (Generalizing pre- and postfixing operations). The earlier operations of pre- and post-fixing of formulae can also be generalized by an operation of sequential composition of two priority graphs $(\mathcal{G}; \mathcal{G}')$. Defining this is a bit more tricky in our setting, but we can build on results from [2]. There, $\mathcal{G}; \mathcal{G}'$ is simply defined by putting all nodes in \mathcal{G}' behind all those of \mathcal{G} without taking care of whether the logical entailment between the nodes is preserved.²² At the level of betterness relations, [2] establishes that the following diagram commutes

$$\begin{array}{ccc}
 \mathcal{G} & \xrightarrow{\quad ; \mathcal{G}' \quad} & \mathcal{G}; \mathcal{G}' \\
 \text{lex} \downarrow & & \downarrow \text{lex} \\
 \langle S, \preceq_{\mathcal{G}} \rangle & \xrightarrow{* \preceq'_{\mathcal{G}}} & \langle S, \preceq_{\mathcal{G}; \mathcal{G}'} \rangle
 \end{array}$$

where $\preceq_{\mathcal{G}} * \preceq'_{\mathcal{G}} = (\preceq_{\mathcal{G}} \cap \preceq_{\mathcal{G}'}) \cup \prec_{\mathcal{G}}$. Notice that in this case, the ordering is derived from the P-graph through Definition 4.

In this last section we have looked at natural composition operations on priority graphs generalizing the specific examples discussed in the article. As argued in Andreka *et al.* (2002), this abstract perspective on priority structures is a felicitous mathematical choice when analysing issues of merging and combining preferences. Our observations, however, are clearly just the beginning of a more general account of operations that construct and modify normative systems viewed as priority graphs.

7. Conclusions

In this article, we have shown how deontic scenarios can be mined for more structure than just deontic inferences. Equally crucial is normative structure,

²² The techniques presented in the Appendix allow us to effectively transform $\mathcal{G}; \mathcal{G}'$ defined in this way to our earlier inclusion format.

represented in priority graphs, and the dynamics of informational and deontic events that change current obligations. We have shown how this view can be implemented by merging ideas from graph representations of criteria for preference with dynamic epistemic logics of informational events. The result is a framework for representing obligations that fits with current trends in other areas of logic of agency, while also throwing fresh light on old issues in the deontic logic literature. Moreover, adopting our framework brings many relevant new phenomena into the scope of deontic logic, such as norm change and general calculi of normative code.

While our examples will have shown the flavour of the style of analysis that we have proposed, many topics remain for further research. We list a few here:

- Analysing the rich linguistic repertoire of commands and suggestions uttered by agents with different deontic roles, thus connecting to *speech act theories* (cf. Yamada, 2008).
- Extending our analysis from obligations in terms of propositions to obligations among *actions* as the primary carriers of moral qualifications (see Czelakowski, 1997, on action vs. deontology). Related to this, we would like to extend our analysis from the *deliberative stance* in this article, based on handling propositions about a fixed current world that we get to know better, to the *action stance* of a changing world where agents have to choose actions, which is closer to first-person experienced agency (cf. van der Meyden, 1996, about deontic transitions).
- Enriching the information dynamics in our treatment to beliefs and plausibility orders. Priority order in this case becomes like ‘entrenchment’, and one could have dual priority graphs for plausibility and betterness. One can then extend our analysis to games as a richer setting of preference entangled with belief, and strategic interaction (cf. Kooi and Tamminga, 2008; Sun and Liu, 2011; Grossi and Turrini, 2012).
- Adding social agents and their obligations, with the appropriate priority graphs for social knowledge and action, linking to the original motivation of priority graphs in Girard (2008) as models for social choice. We can further connect our dynamics of local informational or deontic events to long term deontic phenomena in agency over time (cf. Fagin *et al.*, 1995, and van Benthem, Gerbrandy *et al.*, 2009, on temporal ‘protocols’, and Horty, 2001, on deontic STIT logic).
- Exploring detailed legal argumentation as a natural test for the richer deontic modelling apparatus proposed here and developing our graph calculus to deal with a richer repertoire of natural operations of norm merging and construction of moral codes.

These extensions would tie deontic logic firmly to current logical studies of agency, social choice and games, increasing its range and impact.

Acknowledgements

The authors would like to thank the anonymous reviewers of *Theoria* for the many helpful remarks that allowed us to considerably improve the presentation of our work. The article has also greatly benefited from invaluable discussions with Jörg Hansen on the preference-based tradition in deontic logic, and from many helpful comments from Marek Sergot and the participants of DEON'10 and DEON'12. Fenrong Liu was supported by Project 2012WHYX003 of Tsinghua University.

Appendix: Equivalence of P-graph formats

The article has used a particular format for P-graphs based on logical entailment (Definition 1) and a particular way of constructing preorders from them (Definition 2). Those are not the only ways in which priority structures can be defined and used to specify preorders (Remark 1 has introduced one based on lexicographic ordering).

This technical appendix, which recapitulates results from van Benthem and Grossi (2011), is devoted to establishing the equivalence of different variants of P-graphs. Concretely, it shows how each P-graph (even without the logical entailment constraint) can be put in two syntactically different, but semantically equivalent formats.

Before starting, we fix some auxiliary notation. Given a P-graph \mathcal{G} , we denote with $\uparrow_{\mathcal{G}} \varphi = \{\psi \in \Phi \mid \varphi \prec \psi \text{ or } \psi = \varphi\}$ the *upset* of φ in \mathcal{G} . The set of all upsets of \mathcal{G} is denoted $\uparrow \mathcal{G}$. Given a preorder \preceq we denote by \sim its equivalence part, by $|s|_{\sim}$ the equivalence class of s and by $|S|_{\sim}$ the set of equivalence classes in S .

A.1 Exclusive normal form

We introduce the first type of normal form:²³

Definition 9 (Exclusive normal form for P-graphs). Let $\mathcal{G} = \langle \Phi, \prec \rangle$ be a P-graph. The exclusive normal form of \mathcal{G} is a P-graph $\mathcal{G}_{\text{ex}} = \langle \Phi_{\text{ex}}, \prec_{\text{ex}} \rangle$ such that:

- $\Phi_{\text{ex}} = 2^{\Phi}$. Each element $\Psi \in \Phi_{\text{ex}}$ has to be read as a finite conjunction:

$$\bigwedge \Psi \wedge \bigwedge \neg(\Phi - \Psi)$$

²³ This is the format we alluded to in Remark 1 and which, in a special limited form, had been proposed in Åqvist (1997) (see also n. 13).

that is, the conjunction of all properties in Ψ and all the negations of the properties not in Ψ .

- The relation \prec_{ex} is defined as follows:

$$\Psi \prec_{\text{ex}} \Psi' \Leftrightarrow \exists \varphi \in \Phi : [\varphi \in \Psi' \text{ and } \varphi \notin \Psi \\ \text{and } \forall \varphi' : [\varphi' \notin \Psi' \text{ and } \varphi \in \Psi \Rightarrow \varphi' \prec \varphi]].$$

An example of the exclusive normal form of a graph is given in Figure 5. It should be clear by the construction of the exclusive normal forms that these graphs consist of logically disjoint elements. Notice also that the number of elements in these graphs is bound by the cardinality of Φ being equal to $2^{|\Phi|}$.

Recall the lexicographic order-derivation rule lex introduced in Remark 1. We can now prove a simple normal form theorem guaranteeing that every P-graph has an equivalent exclusive normal form.

Lemma 1 (Adequacy of exclusive normal forms). *For each graph \mathcal{G} and valuation $\mathcal{I} : \preceq_{\mathcal{G}}^{\text{lex}} = \preceq_{\mathcal{G}_{\text{ex}}}^{\text{lex}}$.*




Figure 5. Hasse diagrams of a general P-graph (left) and its exclusive normal form (right)

Proof. We prove the claim by showing that, for any valuation \mathcal{I} :

- (i) $|S|_{\sim_{\mathcal{G}}^{\text{lex}}} = |S|_{\sim_{\mathcal{G}_{\text{ex}}}^{\text{lex}}}$, i.e., the two preorders give rise to the same equivalence classes;
- (ii) $\prec_{\mathcal{G}}^{\text{lex}} = \prec_{\mathcal{G}_{\text{ex}}}^{\text{lex}}$, i.e., the strict part of the two preorders is the same.

As to (i) it suffices to observe that all equivalence classes $|s|_{\sim_{\mathcal{G}}^{\text{lex}}}$ yielded by $\preceq_{\mathcal{G}}^{\text{lex}}$ must be the truth-sets – under \mathcal{I} – of some conjunctions of length $|\Phi|$ of formulae of the form $\bigwedge \Psi \wedge \bigwedge \neg(\Phi \setminus \Psi)$ with $\Psi \subseteq \Phi$. But these are precisely the disjoint elements of the normal form. As to (ii), it is proven by the following series of equivalent statements:

$$s \prec_{\mathcal{G}}^{\text{lex}} s'$$

$$\exists \varphi \in \Phi : s \notin \llbracket \varphi \rrbracket \text{ and } s' \in \llbracket \varphi \rrbracket \text{ and } \forall \varphi' : [s \in \llbracket \varphi' \rrbracket \text{ and } s' \notin \llbracket \varphi' \rrbracket \Rightarrow \varphi' \prec \varphi]$$

$$\forall \Psi, \Psi' \in \Psi_{\text{ex}} : \left[\text{if } \llbracket \Psi \rrbracket = |s|_{\sim_{\mathcal{G}}^{\text{lex}}} \text{ and } \llbracket \Psi' \rrbracket = |s'|_{\sim_{\mathcal{G}}^{\text{lex}}} \text{ then } \Psi \prec_{\text{ex}} \Psi' \right]$$

$$\exists \Xi \in \Phi_{\text{ex}} : [s \notin \llbracket \Xi \rrbracket \text{ and } s' \in \llbracket \Xi \rrbracket \text{ and } \forall \Xi' : [s \in \llbracket \Xi' \rrbracket \text{ and } s' \notin \llbracket \Xi' \rrbracket \Rightarrow \Xi' \prec_{\text{ex}} \Xi]]$$

$$s \prec_{\mathcal{G}_{\text{ex}}}^{\text{lex}} s'$$

where, to simplify notation, by Ψ, Ξ we denote the finite conjunction of the elements of Ψ, Ξ and of the negations of the members of its complement. The first equivalence holds by the definition of lex . The second one holds by the definition of exclusive normal form and the fact that equivalence classes are definable as Boolean compounds of elements of the graph. The third one holds by the fact that those compounds are logically disjoint. \square

A.2 Inclusive normal form

Definition 10 (Inclusive normal forms). Let $\mathcal{G} = \langle \Phi, \prec \rangle$ be a P -graph with logically disjoint elements. The inclusive normal form of \mathcal{G} is a P -graph $\mathcal{G}_{\text{in}} = \langle \Phi_{\text{in}}, \prec_{\text{in}} \rangle$ such that:

- $\Phi_{\text{in}} = \uparrow \mathcal{G}$, that is, the set of upsets of \mathcal{G} . To simplify notation, each element $\Psi \in \Phi_{\text{in}}$ will be read also as the finite disjunction: $\bigvee \Psi$.
- The relation \prec_{in} is defined as follows:

$$\Psi \prec_{\text{in}} \Psi' \Leftrightarrow \Psi' \subset \Psi.$$

An example of inclusive normal form is given in Figure 6. Notice that this form is defined only for graphs with logically disjoint elements.



Figure 6. Hasse diagrams of a P -graph with logically disjoint elements (left) and its inclusive normal form (right)

Like for the case of exclusive normal forms, we obtain a lemma guaranteeing that every exclusive P -graph has an equivalent inclusive normal form.

Lemma 2 (Adequacy of inclusive normal forms). *For each graph \mathcal{G} and valuation \mathcal{I} : $\preceq_{\mathcal{G}}^{\text{lex}} = \preceq_{\mathcal{G}_{in}}^{\text{sub}}$.*

Proof (Sketch). It is easy to see that $\preceq_{\mathcal{G}}^{\text{lex}}$ and $\preceq_{\mathcal{G}_{in}}^{\text{sub}}$ generate the same equivalence classes. As to the strict part of the order, assume $s \prec_{\mathcal{G}}^{\text{lex}} s'$. We then have by the definition of lex and the fact that \mathcal{G} contains disjoint elements: $\exists \varphi \in \Phi : (s \notin \llbracket \varphi \rrbracket \text{ and } s' \in \llbracket \varphi \rrbracket \text{ and } \varphi' \prec \varphi)$. It follows that $\uparrow_{\mathcal{G}} \varphi \subset \uparrow_{\mathcal{G}} \varphi'$. But $\uparrow_{\mathcal{G}} \varphi$ and $\uparrow_{\mathcal{G}} \varphi'$ belong to Φ_{in} by definition, from which we conclude $s \prec_{\mathcal{G}_{in}}^{\text{sub}} s'$. \square


So any graph of disjoint components has an equivalent inclusive normal form. What about other graphs not necessarily consisting of disjoint components? In that case an inclusive normal form can be obtained from the exclusive normal form of the graph. To make it clear, the inclusive normal form of the graph given in Figure 5 (left) is not the graph given in Figure 6 (right), but a graph obtained from Figure 5 (right) which consists of all the upsets of the latter, ordered by set inclusion.

A.3 Graph equivalences

We can now pull together Lemmata 1 and 2 into one characterization theorem showing that the class of P -graphs defines, by lexicographic derivation, the same class of preorders that can be derived via the exclusive normal forms of those graphs by lexicographic derivation, or by subsumption-based derivation from the inclusive normal form of the exclusive normal form of the graph.

Theorem 4 (Equivalence of classes of graphs). *For any P -graph \mathcal{G} and valuation \mathcal{I} : $\preceq_{\mathcal{G}}^{\text{lex}} = \preceq_{\mathcal{G}_{in}}^{\text{lex}} = \preceq_{\mathcal{G}_{in}}^{\text{sub}}$.*

Put it in a different way, the result can be illustrated by the following commutative diagram:



The preorder \preceq obtained from a general graph \mathcal{G} via lex can alternatively be obtained by first extracting the exclusive normal form of \mathcal{G} and then applying lex , or by extracting the inclusive normal form of (the exclusive normal form of) \mathcal{G} and then applying sub .

References

- ANDERSON, A. R. (1957) "The Formal Analysis of Normative Concepts." *American Sociological Review* 22: 9–17.
- ANDREKA, H., RYAN, M. and SCHOBBS, P. (2002) "Operators and Laws for Combining Preference Relations." *Journal of Logic and Computation* 12: 13–53.
- ÅQVIST, L. (1967) "Good Samaritans, Contrary-to-Duty Imperatives and Epistemic Obligations." *Nous* 1: 361–379.
- ÅQVIST, L. (1997) "Systematic Frame Constants in Defeasible Deontic Logic." In D. Nute (ed.), *Defeasible Deontic Logic*, Synthese Library, pp. 59–78. Dordrecht, Boston, London: Kluwer.
- AUCHER, G., GROSSI, D., HERZIG, A. and LORINI, E. (2009) "Dynamic Context Logic." In X. He, J. Horty and E. Pacuit (eds), *Proceedings of LORI 2009*, Vol. 5834 of LNAI, pp. 15–26. Springer.
- BOBBIO, N. (1980) "Norma." In *Enciclopedia Einaudi*, Vol. 9, pp. 876–907. Turin: Einaudi.
- BOELLA, G. and VAN DER TORRE, L. (2003) "Permissions and Obligations in Hierarchical Normative Systems." In *Proceedings of ICAIL'03*, pp. 109–118. New York: ACM Press.
- BOELLA, G. and VAN DER TORRE, L. (2008) "Institutions with a Hierarchy of Authorities in Distributed Dynamic Environments." *Journal of Artificial Intelligence and Law* 16: 53–71.
- BOUTILIER, C. (1994) "Conditional Logics of Normality: A Modal Approach." *Artificial Intelligence* 68: 87–154.
- BROWN, M. (2000) "Conditional Obligation and Positive Permission for Agents in Time." *Nordic Journal of Philosophical Logic* 5(2): 83–112.
- BULYGIN, E. (1986) "Permissive Norms and Normative Systems." In A. Martino and F. Socci Natali (eds), *Automated Analysis of Legal Texts*, pp. 211–218. Amsterdam: Springer.
- CARMO, J. and JONES, A. J. I. (2001) "Deontic Logic and Contrary-to-Duties." In D. Gabbay (ed.), *Handbook of Philosophical Logic* (2nd edn), Vol. 4, pp. 287–363. Dordrecht, Boston, London: Kluwer.
- CZELAKOWSKI, J. (1997) "Action and Deontology." In E. Ejerhed and S. Lindström (eds), *Logic, Action and Cognition-Essays in Philosophical Logic*, pp. 47–87. Dordrecht, Boston, London: Kluwer.
- FAGIN, R., HALPERN, J. Y., MOSES, Y. and VARDI, M. Y. (1995) *Reasoning about Knowledge*. Cambridge, MA: MIT Press.

- FORRESTER, J. (1984) "Gentle Murder, or the Adverbial Samaritan." *Journal of Philosophy* 81: 193–197.
- GIRARD, P. (2008) "Modal Logic for Belief and Preference Change." PhD thesis, University of Amsterdam.
- GOBLE, L. (2000) "Multiplex Semantics for Deontic Logic." *Nordic Journal of Philosophical Logic* 5(2): 113–134.
- GOVERNATORI, G. and ROTOLO, A. (2005) "Logic of Violations: A Gentzen System for Reasoning with Contrary-to-Duty Obligations." *Australasian Journal of Logic* 3.
- GROSSI, D. and TURRINI, P. (2012) "Short-Sight in Extensive Games." In V. Conitzer and M. Winikof (eds), *Proceedings of the 11th International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pp. 805–812. IFAAMAS.
- HANSEN, J. (2001) "Sets, Sentences, and Some Logics about Imperatives." *Fundamenta Informaticae* 48: 205–226.
- HANSEN, J. (2005) "Conflicting Imperatives and Dyadic Deontic Logic." *Journal of Applied Logic* 3: 484–511.
- HANSEN, J., PIGOZZI, G. and VAN DER TORRE, L. (2007) "Ten Philosophical Problems in Deontic Logic." In G. Boella, L. van der Torre and H. Verhagen (eds), *Normative Multi-agent Systems*, no. 07122 in *DROPS Proceedings. Internationales Begegnungs- und Forschungszentrum fuer Informatik (IBFI)*, Schloss Dagstuhl, Germany.
- HANSSON, B. (1969) "An Analysis of Some Deontic Logics." *Nous* 3: 373–398.
- HILPINEN, R. (1980) "On Normative Change." In E. Morscher and R. Stranzinger (eds), *Ethics. Foundations, Problems, and Applications. Proceedings of the 5th International Wittgenstein Symposium*, pp. 155–164. Vienna: Hölder-Pichler-Tempsky.
- HORTY, J. (1997) "Nonmonotonic Foundations of Deontic Logic." In D. Nute (ed.), *Defeasible Deontic Logic*, pp. 17–44. Dordrecht, Boston, London: Kluwer.
- HORTY, J. (2001) *Agency and Deontic Logic*. Oxford University Press.
- JONES, A. J. I. and PÖRN, I. (1985) "Ideality, Sub-Ideality and Deontic Logic." *Synthese* 65: 275–290.
- KANGER, S. (1971) "New Foundations for Ethical Theory." In R. Hilpinen (ed.), *Deontic Logic: Introductory and Systematic Readings*, pp. 36–58. Dordrecht: Reidel Publishing Company.
- KOOL, B. and TAMMINGA, A. (2008) "Conditional Obligations in Strategic Situations." In M. Singh G. Boella, G. Pigozzi and H. Verhagen (eds), *3rd International Workshop on Normative Multiagent Systems*, pp. 188–200. Dordrecht, Boston, London: Kluwer.
- LEWIS, D. (1974) "Semantic Analysis for Dyadic Deontic Logic." In S. Stenlund (ed.), *Logical Theory and Semantical Analysis*, pp. 1–14. Dordrecht: Reidel Publishing Company.
- LIU, F. (2008) "Changing for the Better. Preference Dynamics and Agent Diversity." PhD thesis, University of Amsterdam.
- LIU, F. (2011a) *Reasoning about Preference Dynamics*, Vol. 354 of Synthese Library. Heidelberg: Springer.
- LIU, F. (2011b) "A Two-Level Perspective on Preference." *Journal of Philosophical Logic* 40(3): 421–439.
- MAKINSON, D. (1993) "Five Faces of Minimality." *Studia Logica* 52: 339–379.
- MAKINSON, D. and VAN DER TORRE, L. (2003) "Permission from an Input/Output Perspective." *Journal of Philosophical Logic* 32: 391–416.
- MOORE, G. E. (1903) *Principia Ethica*. Cambridge University Press.
- NAYAK, A., NELSON, P. and POLANSKY, H. (1996) "Belief Change as Change in Epistemic Entrenchment." *Synthese* 109: 143–174.

- PACUIT, E., PARIKH, R. and COGAN, E. (2006) "The Logic of Knowledge Based on Obligation." *Synthese* 149: 311–341.
- PRAKKEN, H. and SERGOT, M. (1996) "Contrary-to-Duty Obligations." *Studia Logica* 57: 91–115.
- PRAKKEN, H. and SERGOT, M. (1997) "Dyadic Deontic Logic and Contrary-to-Duty Obligations." In D. Nute (ed.), *Defeasible Deontic Logic*, Vol. 263 of Synthese Library Series, pp. 223–262. Dordrecht, Boston, London: Kluwer.
- ROSS, A. (1968) *Directives and Norms*. London: Routledge & Kegan Paul.
- SPOHN, W. (1975) "An Analysis of Hansson's Dyadic Deontic Logic." *Journal of Philosophical Logic* 4: 237–252.
- STOLPE, A. (2010) "A Theory of Permission Based on the Notion of Derogation." *Journal of Applied Logic* 8: 97–113.
- SUN, X. and LIU, F. (2011) "Consequentialist Deontic Logic for Decisions and Games." In 2011 *LIRa Yearbook*, pp. 374–397. ILLC.
- TAN, Y.-H. and VAN DER TORRE, L. (1996) "How to Combine Ordering and Minimizing in a Deontic Logic Based on Preferences." In M. Brown and J. Carmo (eds), *Deontic Logic, Agency and Normative Systems. DEON '96: Third International Workshop on Deontic Logic in Computer Science*, pp. 216–232. Heidelberg: Springer.
- VAN BENTHEM, J. (2004) "Dynamic Logic for Belief Revision." *Journal of Applied Non-Classical Logic* 14(2).
- VAN BENTHEM, J. (2011) *Logical Dynamics of Information and Interaction*. Cambridge: Cambridge University Press.
- VAN BENTHEM, J., GERBRANDY, J., HOSHI, T. and PACUIT, E. (2009) "Merging Frameworks for Interaction." *Journal of Philosophical Logic* 38(5): 491–526.
- VAN BENTHEM, J., GIRARD, P. and ROY, O. (2009) "Everything Else Being Equal: A Modal Logic for Ceteris Paribus Preferences." *Journal of Philosophical Logic* 38: 83–125.
- VAN BENTHEM, J. and GROSSI, D. (2001) "Normal Forms for Priority Graphs." Technical Report PP-2011-02, Institute for Logic, Language and Computation.
- VAN BENTHEM, J. and LIU, F. (2007) "Dynamic Logic of Preference Upgrade." *Journal of Applied Non-Classical Logic* 17: 157–182.
- VAN BENTHEM, J. and PACUIT, E. (2011) "Dynamic Logics of Evidence-Based Beliefs." *Studia Logica* 99(1): 61–92.
- VAN DER MEYDEN, R. (1996) "The Dynamic Logic of Permission." *Journal of Logic and Computation* 6: 465–479.
- VAN DER TORRE, L. (1997) "Reasoning about Obligations: Defeasibility in Preference-Based Deontic Logic." PhD thesis, Erasmus University Rotterdam.
- VAN DER TORRE, L. and TAN, Y.-H. (1997) "The Many Faces of Defeasibility in Defeasible Deontic Logic: Essays in Nonmonotonic Normative Reasoning." In D. Nute (ed.), *Defeasible Deontic Logic*, Vol. 263 of Synthese Library, pp. 79–121. Dordrecht, Boston, London: Kluwer.
- VAN ECK, J. (1982) "A System of Temporally Relative Modal and Deontic Predicate Logic and Its Philosophical Applications." *Logique et Analyse* 99: 249–290.
- VAN FRAASSEN, B. (1973) "Values and the Heart's Command." *Journal of Philosophy* 70(1): 5–19.
- VON WRIGHT, G.H. (1963) *Norm and Action. A Logical Inquiry*. London: Routledge.
- YAMADA, T. (2008) "Logical Dynamics of Some Speech Acts that Affect Obligations and Preferences." *Synthese* 165(2): 295–315.