Violation Contexts and Deontic Independence

Leendert van der Torre

Department of Artificial Intelligence, Vrije Universiteit
De Boelelaan 1081a, 1081 HV Amsterdam, the Netherlands
torre@cs.vu.nl
http://www.cs.vu.nl/~torre

Abstract. In this paper we discuss the role of context and independence in normative reasoning. First, deontic operators – obligations, prohibitions, permissions – referring to the ideal context may conflict with operators referring to a violation (or contrary-to-duty) context. Second, deontic independence is a powerful concept to derive deontic operators from such operators of other violation contexts. These two concepts are used to determine how to proceed once a norm has been violated, a key issue of deontic logic applications in computer science. We also show how violation contexts and deontic independence can be used to give a new analysis of several notorious paradoxes of deontic logic.

1 Introduction

Deontic logic is a modal logic in which $\bigcirc p$ is read as 'p ought to be (done),' Fp as 'p is forbidden to be (done)' and Pp as 'p is permitted to be (done).' Deontic logic has traditionally been used by philosophers to analyze the structure of the normative use of language. In the eighties deontic logic had a revival, when it was discovered by computer scientists that this logic can be used for the formal specification and verification of a wide variety of topics in computer science (for an overview and further references see [42]). For example, deontic logic can be used to formally specify soft constraints in planning and scheduling problems as norms. The advantage is that norm violations do not create an inconsistency in the formal specification, in contrast to violations of hard constraints. Recently new interest in deontic logic has arisen, because in the multi agent system discipline *social* norms and laws are used to improve coordination and cooperation [10, 37], in qualitative decision theory norms are used to represent the agent's goals [22, 7, 17], and in computer security norms are used to analyze and specify security issues for electronic commerce [12].

In this paper we study the violability of norms, and in particular how to proceed once a norm has been violated. Clearly this issue is of great practical relevance, because in most applications norms are violated frequently. In the fine print of a contract it is therefore usually stipulated what has to be done if a term in the contract is violated. For example, if the delivery time is over due the responsible agent might be obliged to pay the extra transport and warehousing costs that result from the delay. If the violation is not too serious, or was not

intended by the violating party, the contracting parties usually do not want to consider this as a breach of contracts, but simply as a disruption in the execution of the contract that has to be repaired.

The crucial problem of violability is that a violation changes the context of normative reasoning. New obligations may arise as a result of the violation, and old obligations may no longer be in force. Consequently, obligations are context-dependent. The so-called *violation context* of an obligation distinguishes between ideal and varying sub-ideal contexts to formalize contrary-to-duty (hereafter: CTD) reasoning as it occurs in Forrester's and Chisholm's notorious deontic paradoxes [9, 13, 23, 24, 36]. In these paradoxes obligations referring to the ideal context conflict with obligations referring to the sub-ideal context. If the violation context is left implicit, then the paradoxes result in a counterintuitive inconsistency. This type of context should be distinguished from epistemic or world-based contexts, as is explained later in detail, because it refers to an implicit set of norms called the *normative system*. In this paper we are interested in the relation between violation contexts and independence. They can be used to analyze the following example of Prakken and Sergot [24], that motivates much of their study in contextual deontic reasoning.

Example 1 (Cottage housing). Consider the following sentences.

(1) There must be no dog.
(2) If there is a dog, then there must be a sign.
(3) There must be no sign.

Prakken and Sergot argue that (2) and (3) are inconsistent, because the CTD obligation (2) regulating the violation of (1) does not respect the primary obligation (3). CTD obligations are intended to regulate norm violation, but they cannot just ignore all other norms. In the absence of prioritisation we may conclude that on the one hand there must be a sign if there is a dog, but on the other hand there must not. They continue to consider the following variant.[1]

(1a) There must be no dog.
(2a) If there is a dog, it must be a poodle.
(3a) There must be no poodle.

The first two sentences are a version of Forrester's gentle murderer [13] ('Smith should not murder Jones,' but 'if he murders him, then he should do it gently'), because poodles are dogs. Intuitively (1a) and (2a) should therefore be consistent. Moreover, Prakken and Sergot suggest that (2a) and (3a) are inconsistent in precisely the same way that (2) and (3) are. (2a) regulates the violation of (1a)

[1] Since Forrester's paradox is better analyzed when different types of killing are replaced by different colors [14], Prakken and Sergot also consider the following alphabetic variant: (1b) The door must be painted red; (2b) If the door is not painted red, it must be left unpainted; (3b) The door must not be left unpainted. Again they want to say that (1b)-(3b) are inconsistent, even though (1b)-(2b) are consistent, and (1b) implies (3b).

but it does not respect another primary obligation, (3a). The new feature of the poodle example, of course, is that (3a) is implied by (1a). At first sight, it seems that we have to abandon consequential closure, because this is the only way to make (1a)-(3a) logically stronger than (1a)-(2a) (such that we can add constraints that make only the latter set inconsistent). However, if a normgiver forbids having dogs he surely also implicitly forbids having any particular kind of dog. Consequential closure should not be disregarded when determining what is obligatory.

Prakken and Sergot conclude that 'the difference between what appears consistent and what appears inconsistent seems to depend critically on what is stated explicitly.' Moreover, they sketch the development of an entailment relation $\Gamma \models A$ in which designated *explicit obligations* in premises Γ will be given special status.

In this paper we argue that the crucial distinction between (1a)-(2a) and (1a)-(3a) is an implicit independence assumption. In this paper we show that the three obligations are inconsistent only if the independence assumption is added. Prakken and Sergot's conclusion simply means that the explicit manner of representation suggests that (3a) depends on (1a) if it is derived from (1a), whereas it is independent of (1a) when it is given explicitly.

This paper is organized as follows. In Section 2 we discuss the relation between deontic and contextual reasoning and we introduce a simple contextual deontic logic. In Section 3 we discuss deontic independence and we analyze the cottage housing regulations. Finally in Section 4 we show how independence assumptions can be derived from the explicit manner of representation, which also shows the distinction between deontic independence and world-based factual independence.

2 Violation Contexts

Deontic reasoning is context dependent in different though related ways.

The norm context of a descriptive obligation (a normative proposition) refers to a set of prescriptive obligations (the norms or the normative system) [39]. This reflects the distinction, much debated in the philosophical literature, between descriptive obligations, that have a truth value, and prescriptive obligations, that do not [1, 41, 18, 39, 35, 38, 37]. The norm context is explicitly represented in labeled deontic logic [31, 18, 32], see also [19, 30]. A descriptive obligation $\bigcirc\alpha_L$ contains a label L that represents the norm context in which the obligation for α is derived, and it thus explains *why* the obligation is derived.

The violation context of a descriptive obligation refers to a set of violations of prescriptive obligations, where the violation of a norm 'α should be (done) if β is (done)' is $\neg\alpha \wedge \beta$, i.e. the negation of the material implication $\neg(\beta \rightarrow \alpha)$. It distinguishes between ideal and varying sub-ideal contexts. Obligations that refer to the ideal context are called ideal (or primary) obligations, and

obligations that refer to the sub-ideal context are called contrary-to-duty (or secondary) obligations.

The reasoning context of an obligation refers to the circumstances which are considered. Within a reasoning context some fact or action may be optimal and therefore obligatory, whereas it is not ideal in the most general reasoning context. In other words, an obligation may refer to a sub-ideal context if the reasoning context excludes the ideal context. The reasoning context may extend the actual evidence in order to reason about hypothetical violations. Two reasoning contexts in temporal deontic reasoning are the context of deliberation and the context of justification, used by Thomason [28] to distinguish a (deliberative) ought that implies 'practical-temporal can' from a (judgmental) ought that does not imply 'practical-temporal can.'

The following example illustrates the distinction between the different kinds of contexts.

Example 2 (Cottage housing, continued). Consider only the first two obligations, i.e. 'there must be no dog ' and 'if there is a dog, it must be a poodle,' together with the fact that there is a pitbull. There are three violation contexts: the ideal context there is no dog, the sub-ideal context there is a dog and the sub-sub-ideal context there is a dog that is not a poodle. There is an obligation *'there must be no dog'* referring to the ideal context, and there is an obligation *'there must be a poodle'* referring to the sub-ideal context. Obviously, both obligations are violated by the pitbull. Selling the dog as well as exchanging it for a poodle are improvements of the present state of affairs. However, in the general reasoning context its owner should sell it; only in the reasoning context in which it is not sold she should exchange it. Whether the dog must be sold or exchanged depends on the context referred to, because different reasoning contexts make different violation contexts optimal.

The two derived obligations *'there must be no dog'* and *'there must be a poodle'* are contradictory. In a logic that does not distinguish between contexts, like so-called Standard Deontic Logic,[2] the facts together with the two sentences of the regulations are inconsistent. Two extensions of Standard Deontic Logic, developed to analyze contrary-to-duty reasoning, introduce either temporal references or an explicit context. In the first approach the context is implicitly represented by facts that refer to the same moment in time as the obligatory fact, and the condition refers to an earlier moment in time. To formalize the regulations without introducing the background knowledge that poodles and pitbulls are dogs, Prakken and Sergot represent the poodle by $p \wedge d$ and the pitbull by $i \wedge d$, where d stands for dog.

Explicit time. We write $\beta \Rightarrow \bigcirc\alpha$ for 'if β then $\bigcirc\alpha$,' where $\Rightarrow$ is a suitably defined conditional. If the reasoning context is the ideal context then dogs

[2] Standard Deontic Logic is a normal modal system that is closed under the inference rules modus ponens $\frac{p, p\rightarrow q}{q}$ and necessitation $\frac{\vdash p}{\vdash \bigcirc p}$ and contains the propositional tautologies and the axioms $K : \bigcirc(\alpha \rightarrow \beta) \rightarrow (\bigcirc\alpha \rightarrow \bigcirc\beta)$ and $D : \neg(\bigcirc\alpha \wedge \bigcirc\neg\alpha)$.

are forbidden ($\bigcirc\neg d_t$) and if the reasoning context is the sub-ideal context then there must be a poodle ($d_t \Rightarrow \bigcirc(p_t \wedge d_t)$). These obligations hold regardless of the actual evidence. If there is a pitbull then we still have $i_{t-1} \wedge d_{t-1} \Rightarrow \bigcirc\neg d_t$ and $i_{t-1} \wedge d_{t-1} \wedge d_t \Rightarrow \bigcirc(p_t \wedge d_t)$.

Condition-context. We write $\bigcirc_\gamma(\alpha|\beta)$ for 'α ought to be (done) if β is (done) in the context γ,' and we represent the reasoning context γ by a propositional sentence. Again, if the reasoning context is the ideal context then dogs are forbidden ($\bigcirc\neg d$) and if the reasoning context is the sub-ideal context then there must be a poodle $\bigcirc_d(p \wedge d)$. Moreover, if there is a pitbull then the obligations are both still in force $\bigcirc(\neg d|i \wedge d)$ and $\bigcirc_d(p \wedge d|i \wedge d)$. Contextual obligations can formalize hypothetical reasoning, because the latter two obligations can be read as 'I have a pitbull and I should sell it, but if I will not sell it (hypothesis) then I should at least exchange it for a poodle.'

In this paper we are interested in the second type of logic. We use the phrase 'contextual deontic logic' (hereafter: CDL) to refer to the logic that contains formulas of the type $\bigcirc_\gamma(\alpha\,|\,\beta)$, $F_\gamma(\alpha\,|\,\beta)$ and $P_\gamma(\alpha\,|\,\beta)$. If we have $\bigcirc\alpha$ without $\bigcirc(\alpha\,|\,\beta)$, then we say that β represents exceptional circumstances in which the obligation for α is defeated. A priori there is an obligation for α, but de facto there is not. Conversely, if we have $\bigcirc\alpha$ without $\bigcirc_\beta\alpha$ then it does *not* mean that the obligation for α is defeated by exceptional circumstances β. It means that if we assume β then we refer to another violation context that no longer prefers α.

We keep the semantics as simple as possible. The dynamic preference semantics proposed below generalizes Hansson's preference semantics [15], that is also used by Boutilier [7] in his qualitative decision theory. The basic idea is that, when evaluating the truth condition of the obligation for α, the condition β determines the relevant preference ordering, and the context γ determines the relevant part of the preference ordering. Now, the obligation for α is true if α is true in the preferred states of this sub-ordering.

Definition 1 (CDL). *Let $M = \langle W, \leq_\alpha, V\rangle$ with W a set of worlds, $\leq_\alpha$ a set (for each propositional formula) of transitive, reflexive and totally connected accessibility relations, and V the valuation function. We have $M \models \bigcirc_\gamma(\alpha\,|\,\beta)$ iff there exists a γ world w such that for all worlds $w' \leq_\beta w$ we have $M, w' \models \gamma \rightarrow \alpha$ (all preferred γ worlds of the β ordering satisfy α). Prohibition and (weak) permission are defined by $F_\gamma(\alpha|\beta) =_{def} \bigcirc_\gamma(\neg\alpha|\beta)$ and $P_\gamma(\alpha|\beta) =_{def} \neg\bigcirc_\gamma(\neg\alpha|\beta)$. We write $\bigcirc_\gamma\alpha$ for $\bigcirc_\gamma(\alpha|\top)$, $\bigcirc(\alpha|\beta)$ for $\bigcirc_\top(\alpha|\beta)$, $\bigcirc\alpha$ for $\bigcirc_\top(\alpha|\top)$. Similarly for F and P.*

In the semantics the violation contexts are represented implicitly by the preference relation $\leq$. It could be represented more explicitly in several different ways. For example, we can explicitly represent the violations of the violation contexts, or we can explicitly represent the equivalence classes of the relation $\{(w, w') \mid w \leq w' \text{ and } w' \leq w\}$ by a set of constants, like e.g. Åqvist's optimality classes [2]. Moreover, we can explicitly represent the norm context or the prescriptive obligations in the semantics, and deduce the preference relation, violation contexts or optimality classes from them. Since in this paper we are

interested in deontic independence we keep our semantics as simple as possible and we do not further consider this kind of semantic extensions.

In contextual deontic logic we cannot represent the problematic conditions that occur later than the consequents, as the obligation in Chisholm's notorious paradox 'if the man goes to the assistance, then he should tell them before that he will come' [9]. We thus 'solve' the backwards version of the paradox by restricting the language, see [34] for a discussion. The following example adapted from Pearl [22] further illustrates the dynamic preference semantics.

Example 3 (Switch [22]). We will test the assertability of the following dialogue:

Robot 1: It is too dark here.
Robot 2: Then you ought to push the switch up.
Robot 1: The switch is already up.
Robot 2: Then you ought to push the switch down.

The challenge would be to explain the reversal of the "ought" statement in response to the new observation "The switch is already up." A CDL model can, in contrast to a single preference ordering of static preference semantics, represent that after it is known that the switch is up but still it is dark, the preference is reversed.

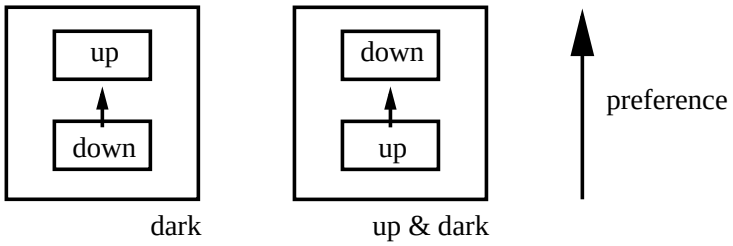

Fig. 1. Switch

CDL does not impose any restrictions on the condition. Such restrictions can be added depending on the application of the logic. For example, for reasoning by cases one could add the following disjunction rule.

$$\bigcirc_\gamma(\alpha|\beta_1) \wedge \bigcirc_\gamma(\alpha|\beta_2) \rightarrow \bigcirc_\gamma(\alpha|\beta_1 \vee \beta_2)$$

3 Deontic Independence

In this paper we are interested in the relation between contextual deontic reasoning and independence, and in this section we consider explicitly given independence assumptions. Qualitative independence has been studied to deal with the irrelevance problem of default reasoning: 'red birds fly' $r \wedge b \Rightarrow f$ can be derived

from 'birds fly' $b \Rightarrow f$ if red 'r' is independent of flying 'f' [11]. In contextual deontic reasoning there are two types of qualitative independence which can be used to deal with irrelevance problems. *Factual independence* is independence with respect to the condition, and *deontic independence* is independence with respect to the context. Thus, β' is factually independent with respect to the obligation for α if β in the context γ if and only if we have

$$FI : \bigcirc_\gamma(\alpha|\beta) \leftrightarrow \bigcirc_\gamma(\alpha|\beta \wedge \beta')$$

Following [11], we call $\bigcirc_\gamma(\alpha \mid \beta)$ the *a priori* obligation, and $\bigcirc_\gamma(\alpha \mid \beta \wedge \beta')$ the *de facto* obligation, and we discriminate between nine different situations (a priori and de facto α is either obliged, forbidden, or it and its negation are both permitted). The following continuation of Example 2 shows that factual independence plays a similar role in defeasible deontic reasoning as independence in default reasoning.

Example 4 (Cottage housing, continued). The CDL prohibition 'dogs are forbidden' Fd is factually independent of the existence of a pitbull, and from this independence relation we can derive the prohibition 'if there is a pitbull, then dogs are forbidden' $F(d|i \wedge d)$. Moreover, the obligation is factually dependent on the needs of a blind owner, from which we can derive the permission *'if the owner is blind, then it is permitted that there is a dog'* $P(d|b)$.

It is well-known that we have to distinguish between decision variables and parameters or events (controllable and uncontrollable variables), because it does not make sense to oblige someone to see to it that a parameter has a desirable value [7, 29, 17]. In the example we take blindness as a parameter, because we do not want to be able to infer that it is obligatory or forbidden that the owner is blind.

Moreover, γ' is deontically independent with respect to the obligation for α if β in the context γ if and only if we have[3]

$$DI : \bigcirc_\gamma(\alpha|\beta) \leftrightarrow \bigcirc_{\gamma\wedge\gamma'}(\alpha|\beta)$$

Factual and deontic independence are very powerful concepts. We can derive obligations from obligations and permissions from permissions in case of independence, and we can derive obligations from permissions and permissions from obligations in case of dependence. The following example illustrates how deontic independence can be used to analyze the Cottage Housing Regulations in Example 1.

Example 5 (Cottage housing, continued). Consider the two sets of CDL formulas $S_1 = \{Fd, \bigcirc_d(p \wedge d)\}$ and $S_2 = \{Fd, \bigcirc_d(p \wedge d), F(p \wedge d)\}$. Without the independence relations, the two sets are equivalent. A typical CDL model consists of three equivalence classes (of $\{(w, w') \mid w \leq w' \text{ and } w' \leq w\}$), in order

[3] Prakken and Sergot [23] call the formula $O_{A\wedge B}C \rightarrow \bigcirc_A C$ up, and the formula $O_A C \rightarrow \bigcirc_{A\wedge B} C$ is called down.

of preference the $\neg d$ worlds, the $p \wedge d$ worlds and the $\neg p \wedge d$ worlds. However, in S_1 the prohibition of poodles deontically depends on dog, whereas they are deontically independent in S_2. Hence, S_2 also contains the following formula, and S_1 its negation.

$$DI : F(p \wedge d) \leftrightarrow F_d(p \wedge d)$$

With the deontic dependence and independence assumption we have respectively $S_1 \models P_d(p \wedge d)$ and $S_2 \models F_d(p \wedge d)$. The latter obligation is inconsistent with $\bigcirc_d(p \wedge d)$, and S_2 is therefore inconsistent. We thus explained Prakken and Sergot's problem in Example 1. Moreover, the deontic independence relation is suggested by the explicit manner of representation, because $F(p \wedge d)$ is deontically independent of d only if it is given explicitly (in S_2).

Deontic independence can also be used to analyze the Reykjavik Scenario [5, 29]. The example is taken from a discussion by Makinson [18]. The explicit obligations are that neither Reagan nor Gorbatchov should be told the secret, if Reagan is told the secret then Gorbatchov should be told, and if Gorbatchov is told then Reagan should be told. Intuitively, under the condition that Gorbatchov is told, we would like to be able to conclude that Reagan should be told, without also deriving its negation. On the other hand, if the first item is broken down into two parts, saying separately that Reagan should not be told and that Gorbatchov should not be told, so that there are four explicit obligations, then the conclusion seems intuitively more ambiguous, depending on whether promulgations are prioritized according to, say, specificity of their antecedents. In the absence of prioritisation we seem authorized to conclude both that Reagan should be told and that he should not; with prioritisation, only the former.

Example 6 (Reykjavik Scenario). Consider the two sets of CDL formulas $S_1 = \{\bigcirc(\neg r \wedge \neg g), \bigcirc_r g, \bigcirc_g r\}$ and $S_2 = \{\bigcirc(\neg r) \wedge \bigcirc(\neg g), \bigcirc_r g, \bigcirc_g r\}$. Without independence relations S_1 and S_2 are logically equivalent, and a typical CDL model consists of three equivalence classes, containing in order of preference the $\neg r \wedge \neg g$ worlds, the $r \wedge g$ worlds and the $r \leftrightarrow \neg g$ worlds. Moreover, let in S_1 the obligation $\bigcirc \neg r$ be deontically dependent on $\neg g$, and let them in S_2 be independent (as suggested by the explicit manner of representation). That is, S_2 additionally contains the following formula, and S_1 its negation.

$$DI : \bigcirc \neg r \leftrightarrow \bigcirc_g \neg r$$

With the additional formulas we have $S_1 \models P_g \neg r$ and $S_2 \models \bigcirc_g \neg r$. The set S_2 is inconsistent, because the premise $\bigcirc_g r$ and the derived $\bigcirc_g \neg r$ are contradictory. Only with the independence assumption the sets are inconsistent.

Summarizing, an obligation is deontically independent if it refers to the same violation context, and factually independent if it is not overridden. Deontic independence and factual independence are thus quite different. One aspect of this difference is that different types of independence relations can be inferred from the explicit representation, as discussed in the following section.

4 Manner of Representation

Independence relations can either be given explicitly, or inferred from the explicit manner of presentation of the obligations. In this section we discuss the second option. For simplicity we assume that the premise set only contains contextual obligations, i.e. no disjunctions, negations, permissions, prohibitions or facts. We study the explicitly represented premises in combination with derived obligations, i.e. deontic operators without independence assumptions.

We start to propose a formal approach to study the explicit manner of representation. It is based on the introduction of an additional deontic operator $①_\gamma(\alpha|\beta)$, to be read as '$\bigcirc_\gamma(\alpha|\beta)$ is explicitly represented,' i.e. we have $\bigcirc_\gamma(\alpha|\beta)$ as well as independence relations suggested by the fact that it is explicitly represented. Note that this approach is quite different from the explicit obligations suggested by Prakken and Sergot, because in their approach logical properties of explicit operators are not studied. Proof-theoretically, the formal approach leads to a two-phase proof theory [27]. We therefore replace the operator $\bigcirc$ by the more suggestive ②.

Definition 2 (The two phase approach). *Premises are phase-1 operators like* $①_\gamma(\alpha \mid \beta)$*, conclusions are phase-2 operators like* $②_\gamma(\alpha \mid \beta)$*, and the two phases are related via the axiom* $①_\gamma(\alpha|\beta) \to ②_\gamma(\alpha|\beta)$.

We first discuss deontic independence and then factual independence. The following independency notion says that a premise is deontically independent of any fact, unless it violates 'ought implies can.'

Definition 3 (Deontic Independence). *If* $\bigcirc_\gamma(\alpha \mid \beta)$ *is a premise, then we have that it is deontically independent of* γ' *iff* $\alpha \wedge \gamma \wedge \gamma'$ *is consistent.*

In other words, if we define $M \models \Diamond(\alpha)$ as 'M contains an α world,' and $Indep(②_\gamma(\alpha|\beta), \gamma')$ as '$②_\gamma(\alpha|\beta)$ is deontically independent of γ',' then we have:

$$M \models ①_\gamma(\alpha|\beta) \Leftrightarrow M \models ②_\gamma(\alpha|\beta) \text{ and } \forall(\gamma') : \Diamond(\alpha \wedge \gamma \wedge \gamma') \to Indep(②_\gamma(\alpha|\beta), \gamma')$$

This definition can be simplified to the following one.

$$M \models ①_\gamma(\alpha|\beta) \Leftrightarrow \forall(\gamma') : \Diamond(\alpha \wedge \gamma \wedge \gamma') \to M \models ②_{\gamma \wedge \gamma'}(\alpha|\beta)$$

Several properties of this new operator are that it does not have weakening of the consequent, but it has strengthening of the context and a kind of cut rule.

$$\models ①_\gamma(\alpha|\beta) \wedge \Diamond(\alpha \wedge \gamma \wedge \gamma') \to ①_{\gamma \wedge \gamma'}(\alpha|\beta)$$

$$\models ①_\gamma(\alpha|\beta) \wedge ①_{\gamma \wedge \alpha}(\alpha'|\beta) \wedge \Diamond(\alpha \wedge \alpha' \wedge \gamma) \to ①_\gamma(\alpha \wedge \alpha'|\beta)$$

The following example illustrates why we cannot use a similar strong independence notion for factual independence, well-known from dyadic deontic logic [40, 33].

Example 7 (Window [40]). The diamond 'the window should be closed if it rains' $①(c \mid r)$ and 'the window should be open if the sun shines' $①(\neg c|s)$ should not derive the contradictory $②(c|r \wedge s)$ and $②(\neg c|r \wedge s)$.

Obviously, the factual independence relations implied by the explicit representation are much weaker than the deontic independence relations due to possible conflicts. As a consequence of these conflicts, in more specific circumstances there are only *prima facie* obligations [25, 3, 21, 35]. For example, there are only prima facie obligations to close and to open the window when it rains and the sun shines. In this paper we can only sketch the logic of prima facie obligations. The crucial distinction is that the connected ordering is replaced by a possibly non-connected ordering. There is a prima facie obligation for α if there is an equivalence class of preferred worlds that satisfies α.

Definition 4 (CDL with prima facie). *Let $M = \langle W, \leq_\alpha, V\rangle$ with $\leq_\alpha$ a set (for each propositional formula) of transitive, reflexive but possibly not totally connected accessibility relations. We have $M \models \bigcirc_\gamma(\alpha \mid \beta)$ iff for all γ worlds w there is a γ world $w' \leq_\beta w$ such that for all worlds $w'' \leq_\beta w'$ we have $M, w'' \models \gamma \rightarrow \alpha$ (the preferred γ worlds of the β ordering satisfy α). We have $M \models \bigcirc^{pf}_\gamma(\alpha|\beta)$ iff there exists a γ world w such that for all worlds $w' \leq_\beta w$ we have $M, w' \models \gamma \rightarrow \alpha$ (all worlds of an equivalence class of preferred γ worlds of the β ordering satisfy α). The obligation $\bigcirc_\gamma(\alpha|\beta)$ weakly factually depends on β' iff we have*

$$\bigcirc^{pf}_\gamma(\alpha|\beta) \leftrightarrow \bigcirc^{pf}_\gamma(\alpha|\beta \wedge \beta')$$

Definition 5 (Factual Independence). *If $\bigcirc_\gamma(\alpha \mid \beta)$ is a premise, then we have that it is weakly factual independent of β' iff $\alpha \wedge \beta \wedge \beta'$ is consistent.*

Additional machinery has to be introduced in the logic to formalize conflict resolution mechanisms based on the non-monotonic overriding of obligations.

5 Further Research

In this paper we argued that deontic reasoning should be considered as context-dependent reasoning, and we showed how independence assumptions can be used in this perspective. There are many ways to continue this line of research.

First, the preference-based semantics has to be weakened to invalidate the counterintuitive axioms characteristic of the preference-based approach, such as for example:

$$\bigcirc_\gamma \alpha \leftrightarrow \bigcirc_\gamma(\alpha \wedge \gamma)$$

A counterintuitive consequence of this theorem is that two independent obligations $①p$ and $①q$ derive not only $②_{\neg(p\wedge q)}p$ and $②_{\neg(p\wedge q)}q$ but also the contradictory $②_{\neg(p\wedge q)}(p \wedge \neg q)$ and $②_{\neg(p\wedge q)}(\neg p \wedge q)$ [33]. This illustrates that the strong independence assumption is usually too strong for the preference-based contextual obligations discussed in this paper. For example, in [19, 30] a logic has been proposed that does not have the counterintuitive axioms but that still

validates the intuitive derivation of 'if a then $\bigcirc y$' from 'if a then $\bigcirc x$' and 'if $a \wedge x$ then $\bigcirc y$.'

Second, the here developed concepts can be used to analyze and classify other recently proposed deontic logics. For example, we already mentioned that Åqvist's optimality classes [2] can be interpreted as contextual parameters. Moreover, Jones and Pörn's bimodal logic [16] and more recent proposals of Carmo and Jones [8] are also based on different degrees of ideality and the idea of some facts being settled (like our context of reasoning).

Third, the repair approach to conditional knowledge bases [6] has been proposed as an alternative to independence assumptions. The main idea is to take advantage of the main features of nonmonotonic inference with respect to rule addition. Contrary to classical logic, it is legitimate to expect that undesirable conclusions can be blocked by adding new rules (and thus not only by deleting them). Moreover, it is also conjectured that reasoning on basis of a ranking is expressive and flexible enough to capture any 'reasonable' expected behavior of a normative base. It is an open problem whether this repair approach can be used to formalize contrary-to-duty reasoning too.

Fourth, the relation between the notion of context introduced in this paper and the recent ideas of formalizing context in the AI literature (such as for example McCarthy's formalism based on an ist-operator) has to be considered. Moreover, the relation between deontic dependence and dependence in other areas has to be investigated. Gärdenfors studied independence inspired by the axiom of probability theory that for disjoint sets A and B we have that the probability of $A \cup B$ is the sum of the probabilities of A and B, i.e. $Pr(A \cup B) = Pr(A) + Pr(B)$ if $A \cap B = \emptyset$. Moreover, independence relations between sets of variables is an important concept to reduce complexity in probability and utility functions in respectively Bayesian networks (reasoning about uncertainty and default reasoning, and in epistemic reasoning) and decision-theoretic planning (reasoning about goals and desires), such that for example for independent x and y we have $p(x, y) = p(x)p(y)$ and $u(x, y) = u(x) + u(y)$. In normative reasoning, such independence relations between variables could be used for example to derive for independent p and q the permission $P(p \wedge q)$ from Pp and Pq.

Fifth, a technical problem for future research is the interference between factual and deontic independence. We have $\bigcirc\alpha$ without $\bigcirc(\alpha \mid \beta)$ if β expresses exceptional circumstances, and we have $\bigcirc\alpha$ without $\bigcirc_\gamma\alpha$ if γ expresses a violation. The obligation is factually dependent if it is overridden, and deontically dependent if it is violated. The two different types of independence are related to different types of defeasibility, that have been observed in [29, 31, 32]. Another topic for further research is the relation between contextual deontic reasoning and defeasible [36] and temporal reasoning [28].

Sixth, a practical problem for future research is how the user can efficiently state independence relations. One way is to derive independence relations from the explicit representation. Another one is the graphical representation of independence by independence networks [4].

6 Conclusion

An important question for deontic logic applications is how to proceed once a norm has been violated. In this paper we studied violation contexts and deontic independence in contextual deontic logic. We showed how factual and deontic independence can be used to deal with different types of irrelevance problems in deontic logic. In contrast to the use of a relevant implication [20], and relevant deontic logic [26], the use of independence relations enables the derivation of deontic operators from such operators of other violation contexts. It is therefore a powerful tool to deal with contrary-to-duty reasoning.

References

[1] C.E. Alchourrón and Bulygin. The expressive conception of norms. In R. Hilpinen, editor, *New Studies in Deontic Logic: Norms, Actions and the Foundations of Ethics*, pages 95–124. D. Reidel, 1981.

[2] L. Åqvist. Systematic frame constants in defeasible deontic logic. In D. Nute, editor, *Defeasible Deontic Logic*, pages 59–77. Kluwer, 1997.

[3] N. Asher and D. Bonevac. Prima facie obligation. *Studia Logica*, 57:19–45, 1996.

[4] F. Bacchus and A.J. Grove. Utility independence in a qualitative decision theory. In *Proceedings of KR'96*, pages 542–552, 1996.

[5] M. Belzer. A logic of deliberation. In *Proceedings of the AAAI'86*, pages 38–43, 1986.

[6] S. Benferhat, D. Dubois, and H. Prade. Practical handling of exception-tainted rules and independence information in possibilistic logic. *Applied Intelligence*, 9:101–127, 1998.

[7] C. Boutilier. Toward a logic for qualitative decision theory. In *Proceedings of the KR'94*, pages 75–86, 1994.

[8] J. Carmo and A.J.I. Jones. A new approach to contrary-to-duty obligations. In D. Nute, editor, *Defeasible Deontic Logic*, pages 317–344. Kluwer, 1997.

[9] R.M. Chisholm. Contrary-to-duty imperatives and deontic logic. *Analysis*, 24:33–36, 1963.

[10] R. Conte and R. Falcone. ICMAS'96: Norms, obligations, and conventions. *AI Magazine*, 18,4:145–147, 1997.

[11] D. Dubois, L. Farinas del Cerro, A. Herzig, and H. Prade. Qualitative relevance and independence: a roadmap. In *Proceedings of the IJCAI'97*, pages 62–67, 1997.

[12] B.S. Firozabadi and L.W.N. van der Torre. Towards a formal analysis of control systems. In *Proceedings of the ECAI'98*, pages 317–318, 1998.

[13] J.W. Forrester. Gentle murder, or the adverbial Samaritan. *Journal of Philosophy*, 81:193–197, 1984.

[14] L. Goble. Murder most gentle: the paradox deepens. *Philosophical Studies*, 64:217–227, 1991.

[15] B. Hansson. An analysis of some deontic logics. In R. Hilpinen, editor, *Deontic Logic: Introductory and Systematic Readings*, pages 121–147. D. Reidel Publishing Company, Dordrecht, Holland, 1971.

[16] A.J.I. Jones and I. Pörn. Ideality, sub-ideality and deontic logic. *Synthese*, 65:275–290, 1985.

[17] J. Lang. Conditional desires and utilities - an alternative approach to qualitative decision theory. In *Proceedings of the ECAI'96*, pages 318–322, 1996.

[18] D. Makinson. On a fundamental problem of deontic logic. In P. McNamara and H. Prakken, editors, *Norms, Logics and Information Systems. New Studies on Deontic Logic and Computer Science*, pages 29–54. IOS Press, 1999.

[19] D. Makinson and L. van der Torre. *The logic of reusable propositional output.* 1999. Submitted.

[20] L.T. McCarty. Modalities over actions: 1. model theory. In *Proceedings of the KR'94*, pages 437–448, 1994.

[21] M. Morreau. *Prima Facie* and seeming duties. *Studia Logica*, 57:47–71, 1996.

[22] J. Pearl. From conditional oughts to qualitative decision theory. In *Proceedings of the UAI'93*, pages 12–20, 1993.

[23] H. Prakken and M.J. Sergot. Contrary-to-duty obligations. *Studia Logica*, 57:91–115, 1996.

[24] H. Prakken and M.J. Sergot. Dyadic deontic logic and contrary-to-duty obligations. In D. Nute, editor, *Defeasible Deontic Logic*, pages 223–262. Kluwer, 1997.

[25] D. Ross. *The Right and the Good.* Oxford University Press, 1930.

[26] W. Stelzner. Relevant deontic logic. *Journal of Philosophical Logic*, 21:193–216, 1992.

[27] Y. Tan and L. van der Torre. How to combine ordering and minimizing in a deontic logic based on preferences. In *Deontic Logic, Agency and Normative Systems. Proceedings of the ΔEON'96*, Workshops in Computing, pages 216–232. Springer, 1996.

[28] R. Thomason. Deontic logic as founded on tense logic. In R. Hilpinen, editor, *New Studies in Deontic Logic: Norms, Actions and the Foundations of Ethics*, pages 165–176. D. Reidel, 1981.

[29] L. van der Torre. Violated obligations in a defeasible deontic logic. In *Proceedings of the ECAI'94*, pages 371–375, 1994.

[30] L. van der Torre. The logic of reusable propositional output with the fulfilment constraint. In *Labelled Deduction*, Applied Logic Series. Kluwer, 1999.

[31] L. van der Torre and Y. Tan. Cancelling and overshadowing: two types of defeasibility in defeasible deontic logic. In *Proceedings of the IJCAI'95*, pages 1525–1532, 1995.

[32] L. van der Torre and Y. Tan. The many faces of defeasibility in defeasible deontic logic. In D. Nute, editor, *Defeasible Deontic Logic*, pages 79–121. Kluwer, 1997.

[33] L. van der Torre and Y. Tan. Prohairetic Deontic Logic (PDL). In *Logics in Artificial Intelligence*, LNAI 1489, pages 77–91. Springer, 1998.

[34] L. van der Torre and Y. Tan. The temporal analysis of Chisholm's paradox. In *Proceedings of the AAAI'98*, pages 650–655, 1998.

[35] L. van der Torre and Y. Tan. An update semantics for prima facie obligations. In *Proceedings of the Thirteenth European Conference on Artificial Intelligence (ECAI'98)*, pages 38–42, 1998.

[36] L. van der Torre and Y. Tan. Contextual deontic logic: violation contexts and factual defeasibility. In M. Cavalcanti, editor, *Formal Aspects in Context*, Applied Logic Series. Kluwer, 1999.

[37] L. van der Torre and Y. Tan. Rights, duties and commitments between agents. In *Proceedings of the IJCAI'99*, 1999.

[38] L. van der Torre and Y. Tan. An update semantics for defeasible obligations. In *Proceedings of the UAI'99*, 1999.

[39] L. van der Torre and Y. Tan. An update semantics for deontic reasoning. In P. McNamara and H. Prakken, editors, *Norms, Logics and Information Systems. New Studies on Deontic Logic and Computer Science*, pages 73–90. IOS Press, 1999.

[40] G.H. von Wright. A new system of deontic logic. In R. Hilpinen, editor, *Deontic Logic: Introductory and Systematic Readings*, pages 105–120. D. Reidel Publishing Company, Dordrecht, Holland, 1971.

[41] G.H. von Wright. Deontic logic: as I see it. In P. McNamara and H. Prakken, editors, *Norms, Logics and Information Systems. New Studies on Deontic Logic and Computer Science*, pages 15–25. IOS Press, 1999.

[42] R.J. Wieringa and J.-J.Ch. Meyer. Applications of deontic logic in computer science: A concise overview. In J.-J. Meyer and R. Wieringa, editors, *Deontic Logic in Computer Science*, pages 17–40. John Wiley & Sons, Chichester, England, 1993.