A Unifying Framework for Reasoning about Normative Conflicts

Mathieu Beirlaen*

Abstract

First, two context-dependent desiderata are presented for devising calculi of deontic logic that can consistently accommodate normative conflicts. Conflict-tolerant deontic logics (CTDLs) can be evaluated by their treatment of the trade-off between these desiderata. Next, it is argued that CTDLs defined within the standard format for adaptive logics are particularly good at overcoming this trade-off.

1 Normative conflicts

One of the many challenges in the field of deontic logic concerns the consistent accommodation of normative conflicts by formal calculi. Intuitively, a *normative conflict* occurs whenever we find ourselves in a situation in which our normative directives are inconsistent or not uniquely action-guiding in the sense that we are permitted or even obliged to do something that is forbidden. We may for instance be permitted to break a promise in view of a more binding obligation to rescue someone in need.

Normative conflicts also occur in e.g. jurisprudence (cfr. Alchourrón & Bulygin, 1971) and theoretical ethics. As a classic example of a normative conflict in the latter context, consider one of the many formulations of the so-called trolley problem: a trolley is headed toward five people walking on the track, and its conductor has fainted. An

*I am very grateful to the organizers and participants of the *Logica 2011* conference. I also wish to thank Christian Straßer for many valuable comments and suggestions.

agent is standing next to a switch, which she can throw, that will turn the trolley onto a parallel side track, thereby preventing it from killing the five people. However, there is a man standing on the side track with his back turned. If the agent throws the switch, she will kill this man. Should the agent throw the switch, or shouldn't she interfere?

2 Accommodating normative conflicts

Let the operators 'O' and 'P' represent obligations, resp. permissions. Using these operators, we can distinguish between various types of normative conflicts (Beirlaen, Straßer, & Meheus, n.d.):

- OO-conflicts between two or more obligations, e.g. $\mathsf{O}p \wedge \mathsf{O}\neg p$,
- OP-conflicts between an obligation and a permission, e.g. $\mathsf{O}p \wedge \mathsf{P}\neg p$,
- contradictory obligations [permissions]: conflicts between an obligation [permission] and its negation, e.g. $\mathsf{O}p \wedge \neg\mathsf{O}p$, $\mathsf{P}q \wedge \neg\mathsf{P}q$,

In turn, these various types can be combined into more complex conflicts, e.g. $\mathsf{O}p \wedge \mathsf{O}q \wedge \mathsf{O}\neg(p \wedge q)$, $\mathsf{O}(p \wedge \neg q) \wedge (\mathsf{O}\neg p \vee \mathsf{P}q)$, or $\mathsf{P}(p \wedge q) \wedge (\mathsf{P}p \supset \mathsf{O}\neg q)$. Depending on the properties of the logic in question, one or more types of normative conflicts may imply or be equivalent to a conflict of another type. For instance, in systems in which – for any well-formed formula A – $\mathsf{P}\neg A \equiv \neg\mathsf{O}A$, an OP-conflict $\mathsf{O}p \wedge \mathsf{P}\neg p$ is equivalent to the pair of contradictory obligations $\mathsf{O}p \wedge \neg\mathsf{O}p$.

In Standard Deontic Logic (**SDL**)[1], normative conflicts cause explosion: **SDL** cannot consistently accommodate OO-conflicts, OP-conflicts, or contradictory obligations [permissions]. The modal logician can try to overcome this limitation in one of two ways.

A first way of trying to make deontic logic more conflict-tolerant consists of enriching the language of **SDL**. This can be done by adding sub- and/or superscripts to the deontic operators for indicating the authorities, normative standards and/or interest groups in view of which the conflicting norms hold (e.g. Kooi & Tamminga, 2008). For

[1]**SDL** is obtained by adding to the basic normal modal logic **K** the axiom schema $\Box A \supset \Diamond A$ (and by subsequently replacing instances of the alethic modal operator $\Box$ [$\Diamond$] by the deontic operator O [P]). This logic is also known as **KD** or simply **D**.

instance, in the trolley example above we could use different superscripts for indicating that the obligation to throw the switch holds in view of the 'utilitarian' conviction that saving five lives is preferable to losing one life, whereas the obligation not to throw the switch holds in view of the 'deontological' conviction that refrainment is preferable to actively and consciously killing someone.

Another way to enrich the language of **SDL** is to introduce a preference ordering on our obligations and permissions in order to resolve conflicts between norms of different hierarchies in the order (e.g. Hansson, 2001). Doing this would allow us to model situations in which more binding obligations or permissions override less binding ones. Yet another extension of **SDL** consists in making its deontic operators dyadic in order to properly express under which conditions our obligations and permissions hold true (for a good oversight, see Åqvist, 2002).

These extensions are very successful in increasing the expressive power of **SDL**. Furthermore, they effectively allow us to consistently model conflicts between norms of different hierarchies, norms issued by different authorities, norms arising from different normative standards, norms that hold in different circumstances, etc. However, they fall short of tolerating so-called *symmetrical* normative conflicts (Sinnott-Armstrong, 1988). These are normative conflicts of the same preference, arising from one and the same authority and normative standard, that hold in view of one and the same interest group in the same circumstances.

Imagine, for instance, a situation where two identical twins are drowning and the situation is such that some agent, say Ann, can save either one of them, but she cannot save both (because, for instance, Ann is not a very good swimmer and there is not enough time to save both). Morally, Ann ought to save the first twin ($\mathsf{O}t_1$) and she ought to save the second twin ($\mathsf{O}t_2$). However, it is impossible for her to save both twins ($\neg\Diamond(t_1 \wedge t_2)$) (Gowans, 1987, p. 192). Following Lou Goble (2005), we can represent Ann's dilemma in purely deontic terms by means of the formulas $\mathsf{O}t_1 \wedge \mathsf{O}\neg t_1$ and $\mathsf{O}t_2 \wedge \mathsf{O}\neg t_2$.[2] Given our assumptions, these OO-conflicts are perfectly symmetrical and will still cause explosion in any of the proposed enrichments of **SDL**.

[2]Goble arrives at this representation via the principle $\neg\Diamond(A \wedge \neg B) \supset (\mathsf{O}A \supset \mathsf{O}B)$, by means of which the obligations $\mathsf{O}\neg t_1$ and $\mathsf{O}\neg t_2$ follow from $\mathsf{O}t_1, \mathsf{O}t_2$, and $\neg\Diamond(t_1 \wedge t_2)$.

A second way of allowing for the consistent possibility of normative conflicts in deontic logic consists in weakening **SDL**. Instead of making **SDL** more expressive, this approach proceeds by rejecting or restricting certain **SDL**-valid inferences so as to avoid explosion in view of normative conflicts. The aim is not to extend the language of **SDL**, but rather to avoid absurdities from following from a normative conflict. Where NC is an OO- or OP-conflict or a pair of contradictory obligations or permissions, and where $\mathbf{L}$ is some logic, we also want to avoid things like (Goble, 2005):

$$\mathsf{NC} \vdash_{\mathbf{L}} \mathsf{O}B \tag{DEX-1}$$

$$\mathsf{NC} \vdash_{\mathbf{L}} \mathsf{P}B \supset \mathsf{O}B \tag{DEX-2}$$

$$\mathsf{NC} \vdash_{\mathbf{L}} \mathsf{O}B \vee \mathsf{O}\neg B \tag{DEX-3}$$

Likewise for premise sets containing more complex or combined normative conflicts. By weakening **SDL** to a system that invalidates principles like (DEX-1), (DEX-2), and (DEX-3), we can arrive at a logic in which we can sensibly accommodate even the symmetrical cases discussed above. For this reason and for the reason that we might simply be constrained by a not very expressive formal language, I focus on this latter approach from now on.
Where $i \in \{1, 2, 3\}$, the variants of (DEX-i) where NC is a OO-conflict (resp. OP-conflict, pair of contradictory norms) are in the remainder denoted by (OO-DEX-i) (resp. (OP-DEX-i), ($\perp$-DEX-i)). Henceforth, logics devised in order to consistently accommodate one or more types of normative conflicts are called *conflict-tolerant deontic logics* (CTDLs).

3 Deontic pluralism

The aim of this paper is not to present and defend one particular conflict-tolerant deontic logic that allows for the consistent possibility of *all* types of normative conflicts. I believe instead that the adequacy of a given CTDL is a context-dependent matter: both its rules of inference and its degree of conflict-tolerance depend on the concrete application of the logic. Let me illustrate this claim by means of three examples, each of which is situated in a different 'deontic' context.

(1) As a first example, consider a *moral* context. In discussions on moral dilemmas, philosophers have typically focussed on conflicting obligations. Moral dilemmas are conceived as situations in which an agent ought to adopt each of two or more alternatives which are equally compelling from a moral point of view, and in which the agent cannot do both (or all) of the actions (e.g. Sinnott-Armstrong, 1988). In view of the discussion in Section 2, moral dilemmas can be formalized as OO-conflicts (cfr. footnote 2).
Since there is nothing particularly 'dilemmatic' about an OP-conflict (here, the agent can still safely fulfill all of her moral requirements, i.e. all of her obligations), then – assuming that a rational agent facing a moral dilemma is not facing any contradictory obligations or permissions – a CTDL that allows for the consistent possibility of OO-conflicts is sufficiently conflict-tolerant for dealing with moral dilemmas.
How, then should **SDL** be weakened in this context? One suggestion is to reject or restrict the aggregation schema (AND):

$$(\mathsf{O}A \wedge \mathsf{O}B) \supset \mathsf{O}(A \wedge B) \tag{AND}$$

In the moral context, (AND) was disputed (amongst others) by Bernard Williams, who argued that an agent facing an OO-conflict thinks that she should fulfill each of the conflicting obligations, but does not think that she should fulfill both (Williams, 1965).

(2) Next, consider the context of *normative systems*. In this context, formulas of the form $\mathsf{O}A$ [$\mathsf{P}A$] are read as "there exists a norm to the effect that A is mandatory [permitted]". Even though jurists have created a number of principles in order to resolve legal conflicts, normative systems often contain irresolvable conflicts between norms. These conflicts can be formalized as OO- or OP-conflicts, e.g. Alchourrón (1969); Alchourrón and Bulygin (1971).
In the context of normative systems, formulas of the form $\mathsf{O}A$ or $\mathsf{P}A$ abbreviate statements *about* norms. For instance, a formula $\neg\mathsf{O}p$ [$\neg\mathsf{P}p$] denotes the absence of a norm to the effect that p is mandatory [permitted].[3] Whereas a normative system may very well contain both a norm to the effect that p is mandatory as well as a norm to the effect that $\neg p$ is mandatory or permitted, it is less clear how such

[3] Formulas of the form "PA" are interpreted here as *strong* or *positive* permissions, in accordance with their interpretation in (Alchourrón, 1969).

a system could both contain and not contain a norm to the effect that p is mandatory or permitted.[4] Thus it is reasonable to construct a logic of normative systems that takes into account the possibility of OO- and OP-conflicts, but not the possibility of contradictory norms. Due to the possibility of OP-conflicts and the specific interpretation of the deontic operators in this context, a concrete CTDL for normative systems should invalidate the interdefinability schema (DfP) (Wright, 1963):

$$\mathsf{P}A \equiv \neg\mathsf{O}\neg A \tag{DfP}$$

(3) As a third and final illustration, consider the logic of commands. In this setting, the O- and P-operators are interdefinable, i.e. (DfP) is valid in this context (see e.g. (Wright, 1963)). Since it is possible for a (confused) authority to assert that p is obligatory, and also that $\neg p$ is obligatory or permitted, OO- and OP-conflicts should be tolerated. Moreover, we should allow for contradictory obligations and permissions in this setting, since a formula $\mathsf{O}A \wedge \mathsf{P}\neg A$ is equivalent to $\mathsf{O}A \wedge \neg\mathsf{O}A$ and $\neg\mathsf{P}\neg A \wedge \mathsf{P}\neg A$ in view of (DfP). In Beirlaen et al. (n.d.), a CTDL is presented that is *fully* conflict-tolerant in the sense that it invalidates all of (DEX-1), (DEX-2), and (DEX-3) for all types of normative conflicts presented above. This logic weakens **SDL** by turning its classical negation into a paraconsistent one, thus invalidating the *Ex Contradictione Quodlibet* schema:

$$(A \wedge \neg A) \supset B \tag{ECQ}$$

One need not agree with all the details in illustrations (1)-(3) in order to be convinced by the main argument, namely that different normative contexts require different CTDLs. From the illustrations, it is also clear that the degree of conflict-tolerance of a given CTDL, i.e. the variety of types of normative conflicts that the CTDL should consistently allow for, is also context-dependent.

Given these insights, we can formulate a first of two desiderata for CTDLs:

Desideratum 1 Given the normative context to which it is applied, a CTDL should be sufficiently conflict-tolerant.

[4]Exceptions can be made, for instance, when one of two parties argues that system S *does* contain a norm to the effect that A is permitted, whereas the other argues that S *doesn't* contain such a norm. However, such a context is different from the one discussed here.

More formally, this desideratum boils down to the demand that, depending on the types of conflicts we want our logic **L** to be able to accommodate, **L** should invalidate principles like (DEX-1)-(DEX-3) for these types of conflicts. If, for instance, we want to devise a logic **L** for dealing with moral dilemmas as in illustration (1), then – according to Desideratum 1 – **L** must invalidate (at least) (OO-DEX-1), (OO-DEX-2), and (OO-DEX-3).

4 CTDLs and inferential power

By weakening **SDL** in order to make it more conflict-tolerant, we run the risk of losing inferences that are intuitively valid. Suppose, for instance, that we want to devise a CTDL for dealing with moral dilemmas, as in illustration (1) in Section 3. Suppose further that we follow Williams' suggestion and reject the (AND)-schema. Then consider the following example from (Horty, 1994, p. 39): suppose that some agent should either fight in the army or perform alternative service to his country ($\mathsf{O}(f \vee s)$). As a pacifist, this agent is opposed to warfare, so he ought not to fight in the army ($\mathsf{O}\neg f$). Then he can consistently satisfy all of his obligations by performing alternative service to his country. However, although $\mathsf{O}s$ is **SDL**-derivable from $\mathsf{O}(f \vee s)$ and $\mathsf{O}\neg f$, we can no longer make this inference if we give up (AND). This led Horty to the conclusion that "apparently, what is needed is some degree of agglomeration [aggregation], but not too much; and the problem of formulating a principle allowing for exactly the right amount of agglomeration [aggregation] raises delicate issues that have generally been ignored in the literature" (Horty, 2003, p. 580).
Horty's argument boils down to the need for validating at least *some* instances of the Deontic Disjunctive Syllogism schema:

$$(\mathsf{O}(A \vee B) \wedge \mathsf{O}\neg A) \supset \mathsf{O}B \qquad \text{(DDS)}$$

(DDS) is lost in its entirety if we reject all instances of (AND). In general, Horty's example illustrates the need for the following (second) desideratum for CTDLs:

Desideratum 2 Given the normative context, a CTDL should be strong enough to account for all intuitively valid inferences.

The main problem then in devising adequate CTDLs consists in finding the right equilibrium between Desideratum 1 and Desideratum

2: we want a logic that is sufficiently *weak* in order to meet the first, and sufficiently *strong* to meet the second desideratum.

5 The adaptive logics framework

In the remainder of this paper, I defend the 'adaptive' approach for devising and evaluating CTDLs. Logics developed within the adaptive logics framework are well-suited for modeling non-monotonic reasoning. Like most human inferencing, our normative reasoning is non-monotonic. We may, for instance, withdraw a conclusion $\mathsf{O}p$ drawn from two premises $\mathsf{O}(p \vee q)$ and $\mathsf{O}\neg q$ in view of the new information $\mathsf{O}\neg p$. In order for a logic $\mathbf{L}$ to model such reasoning processes, it needs to be non-monotonic: $\mathsf{O}(p \vee q), \mathsf{O}\neg q \vdash_{\mathbf{L}} \mathsf{O}p$, yet $\mathsf{O}(p \vee q), \mathsf{O}\neg q, \mathsf{O}\neg p \nvdash_{\mathbf{L}} \mathsf{O}p$.

Most adaptive CTDLs (ACTDLs) devised so far are defined within the *standard format* for adaptive logics from (Batens, 2007). An ACTDL in this standard framework is characterized as a triple, consisting of:

1. A lower limit logic **LLL**: a reflexive, transitive, monotonic and compact CTDL that contains classical logic and has a characteristic semantics.
2. A set of abnormalities Ω: a set of **LLL**-contingent formulas, characterized by a logical form $\mathcal{F}$; or a union of such sets.
3. An adaptive strategy: Reliability or Minimal Abnormality.

The lower limit logic is the stable part of the adaptive logic (AL); anything that follows from the premises by **LLL** will never be revoked. In the case of CTDLs, we want this 'base' logic to be sufficiently conflict-tolerant given its context of application, i.e. the lower limit logic of an ACTDL has to meet Desideratum 1.
Typically, an AL enables one to derive, for most premise sets, some additional consequences on top of those that are **LLL**-derivable. These supplementary consequences are obtained by interpreting a premise set "as normally as possible", or, equivalently, by supposing abnormalities to be false "as much as possible". For sensible ACTDLs, the set Ω is defined in such a way that, for every type of normative conflict that we want to be able to accommodate, a member of Ω is **LLL**-derivable from a normative conflict. A concrete illustration of this idea follows in Section 6.

The formal disambiguation of the phrases "as normally as possible" and "as much as possible" in the previous paragraph is relative to the adaptive strategy used. The two strategies currently defined within the standard framework are Reliability and Minimal Abnormality. Generally, the Minimal Abnormality strategy allows one to derive some extra consequences on top of those derivable by means of the Reliability strategy.
ALs defined within the standard format are well-behaved syntactically and semantically. Syntactically, adaptive proofs extend a Fitch-style proof theory (Fitch, 1952), which is illustrated informally in Section 6. Semantically, ALs have a Shoham-style preferential model semantics (Shoham, 1987). A detailed account of the semantics and proof theory for ALs in the standard format can be found in Batens (2007).

6 Illustration: the logic $\mathbf{DP^r}$

The logic $\mathbf{DP^r}$ from (Beirlaen et al., n.d.) aims to meet the demands described in illustration (3) of Section 3. Its lower limit logic $\mathbf{DP}$ is a fully conflict-tolerant paraconsistent CTDL:[5] it invalidates (DEX-1)-(DEX-3) for any type of normative conflict. Its set of abnormalities is the set of well-formed formulas of the form $A \wedge \neg A$, where '$\neg$' is a paraconsistent negation connective, and where A is either an atomic proposition or a formula of the form $\mathsf{O}B$, where B is a well-formed formula of classical propositional logic. In view of the construction of $\mathbf{DP}$, an abnormality is derivable from any normative conflict: $\mathsf{O}A \wedge \mathsf{O}\neg A \vdash_{\mathbf{DP}} \mathsf{O}\neg A \wedge \neg\mathsf{O}\neg A, \mathsf{O}A \wedge \mathsf{P}\neg A \vdash_{\mathbf{DP}} \mathsf{O}A \wedge \neg\mathsf{O}A$, and $\mathsf{P}A \wedge \neg\mathsf{P}A \vdash_{\mathbf{DP}} \mathsf{O}\neg A \wedge \neg\mathsf{O}\neg A$. The strategy employed by $\mathbf{DP^r}$ is Reliability (hence the superscript r).
Due to its invalidation of all instances of (DEX-1)-(DEX3) for all types of normative conflicts, $\mathbf{DP}$ meets Desideratum 1. However, it does not meet Desideratum 2. It is, for instance, too weak in order to account for Horty's example from Section 4: $\mathsf{O}(f \vee s), \mathsf{O}\neg f \nvdash_{\mathbf{DP}} \mathsf{O}s$. This problem is solved by $\mathbf{DP^r}$. Let $\Gamma = \{\mathsf{O}(f \vee s), \mathsf{O}\neg f\}$. We start a $\mathbf{DP^r}$-proof from Γ by introducing the premises. This can be done in an adaptive proof via the premise introduction rule PREM:

1	$\mathsf{O}(f \vee s)$	PREM	$\emptyset$

[5] A logic is *paraconsistent* if it invalidates the (ECQ) schema.

2	$\mathsf{O}\neg f$	PREM	$\emptyset$

The last column in the proof is called the *condition*. For lines at which premises are introduced, the condition is always empty. Due to the paraconsistency of "$\neg$", the Disjunctive Syllogism rule is **DP**-invalid. So is its deontic variant (DDS), which would allow us to derive $\mathsf{O}s$ from $\mathsf{O}(f \vee s)$ and $\mathsf{O}\neg f$. However, the weaker formula $\mathsf{O}s \vee (\mathsf{O}\neg f \wedge \neg\mathsf{O}\neg f)$ *is* **DP**-derivable from Γ. We can add this formula to the proof by using the *unconditional rule* RU. This rule allows us to derive – for any adaptive logic in the standard format – all formulas that follow from the premises by the lower limit logic. In the condition column of a line at which RU is applied, we find the union of the conditions found at the lines used in this derivation. The formula at line 3 below is derived on the empty condition because it relies on the formulas at lines 1 and 2, the conditions of which are empty too.

3	$\mathsf{O}s \vee (\mathsf{O}\neg f \wedge \neg\mathsf{O}\neg f)$	1,2; RU	$\emptyset$

Note that the second disjunct of the formula derived at line 3 is a member of Ω. At a line in an adaptive proof, we can move to the condition those members of Ω which were derived in disjunction with some other formula. This is realized by an application of the *conditional rule* RC:

4	$\mathsf{O}s$	3; RC	$\{\mathsf{O}\neg f \wedge \neg\mathsf{O}\neg f\}$

Informally, moving an abnormality to the condition column corresponds to making the assumption that this abnormality is false. Thus, at line 4 we have derived the formula $\mathsf{O}s$ on the assumption that $\mathsf{O}\neg f \wedge \neg\mathsf{O}\neg f$ is false.

If it would later turn out that $\mathsf{O}\neg f \wedge \neg\mathsf{O}\neg f$ were true after all, then the lines in the proof at which we assumed this formula to be false would become *marked* in the proof. In fact, for the Reliability strategy lines become marked as soon as an element of their condition occurs in a disjunction of abnormalities that follows from the premise set by means of the lower limit logic.

Since no such disjunction containing the formula $\mathsf{O}\neg f \wedge \neg\mathsf{O}\neg f$ is **DP**-derivable from Γ, line 4 remains unmarked in any extension of this proof. By the criterion for final derivability in an adaptive proof, it

follows that $\mathsf{O}s$ is $\mathbf{DP^r}$-derivable from Γ: $\Gamma \vdash_{\mathbf{DP^r}} \mathsf{O}s$.[6]
As opposed to its lower limit logic, $\mathbf{DP^r}$ meets Desideratum 2. In fact, for all **SDL**-consistent premise sets Γ, $\Gamma \vdash_{\mathbf{DP^r}} A$ iff $\Gamma \vdash_{\mathbf{SDL}} A$. In adaptive terminology: **SDL** is the *upper limit logic* of $\mathbf{DP^r}$.
For reasons of space, I cannot here provide more details about the workings of $\mathbf{DP^r}$. For a full formal account of the proof theory of $\mathbf{DP^r}$ (including its marking definition and the definition of final derivability in a $\mathbf{DP^r}$-proof), for its semantics, and for more illustrations for this logic, see (Beirlaen et al., n.d.).

7 Conclusion and outlook

The logic $\mathbf{DP^r}$ is but one of a large family of adaptive CTDLs. Other adaptive CTDLs can be found in e.g. Goble (n.d.); Meheus, Beirlaen, and Van De Putte (2010); Straßer (2010); Straßer, Beirlaen, and Meheus (n.d.).
Assuming a pluralistic (context-dependent) attitude with respect to CTDLs, adaptive logics provide a unifying framework for studying various CTDLs in different normative contexts. The list of normative contexts and respective CTDLs hinted at in this paper is not meant to be exhaustive. Different options are available for the deontic logician in devising (A)CTDLs.
Of course, the adaptive logics framework itself only represents one of the many options open for devising CTDLs. Its advantages include a very intuitive treatment of the trade-off between Desiderata 1 and 2, and a well-behaved semantics and defeasible proof theory. Moreover, the framework itself is flexible in the sense that it allows for various extensions and enrichments in the sense discussed in Section 2 (for an example, see (Beirlaen & Straßer, 2011; Van De Putte & Straßer, n.d.)). Furthermore, ACTDLs defined within the standard format come with a meta-theory which automatically guarantees properties like smoothness, fixed point, soundness and completeness (see

[6] From the informal statement of the marking criterion for Reliability, it is clear that line 4 would become marked in a proof from $\Gamma \cup \{\neg\mathsf{O}\neg f\}$ as soon as the abnormality $\mathsf{O}\neg f \wedge \neg\mathsf{O}\neg f$ is derived in the proof (since $\Gamma \cup \{\neg\mathsf{O}\neg f\} \vdash_{\mathbf{DP}} \mathsf{O}\neg f \wedge \neg\mathsf{O}\neg f$, the latter formula is obtainable via RU on the empty condition). In view of the marking definition for Reliability, and the criterion for final derivability in a $\mathbf{DP^r}$-proof, it follows that $\Gamma \cup \{\neg\mathsf{O}\neg f\} \nvdash_{\mathbf{DP^r}} \mathsf{O}s$. This illustrates the non-monotonicity of $\mathbf{DP^r}$.

(Batens, 2007) for proofs and a more detailed list of meta-theoretical properties).
One drawback of this framework is that adaptive logics tend to be computationally complex. But then again, so is human (normative) reasoning (Batens, de Clercq, Verdée, & Meheus, n.d.).

References

Alchourrón, C. E. (1969). Logic of norms and logic of normative propositions. *Logique & Analyse*, *47*, 242–268.

Alchourrón, C. E., & Bulygin, E. (1971). *Normative Systems.* Springer-Verlag, Wien/New York.

Åqvist, L. (2002). Deontic logic. In D. Gabbay & F. Guenthner (Eds.), *Handbook of philosophical logic (2nd edition)* (Vol. 8, pp. 147–264). Kluwer Academic Publishers.

Batens, D. (2007). A universal logic approach to adaptive logics. *Logica Universalis*, *1*, 221-242.

Batens, D., de Clercq, K., Verdée, P., & Meheus, J. (n.d.). Yes fellows, most human reasoning is complex. *Synthese*(1), 113-131.

Beirlaen, M., & Straßer, C. (2011). A paraconsistent multi-agent framework for dealing with normative conflicts. In J. Leite, P. Torrini, T. Ågotnes, G. Boella, & L. van der Torre (Eds.), *Computational logic in multi-agent systems. Proceedings of the* 12^{th} *international workshop CLIMA XII* (Vol. 6814, p. 312-329). Springer-Verlag.

Beirlaen, M., Straßer, C., & Meheus, J. (n.d.). An inconsistency-adaptive deontic logic for normative conflicts. *Journal of Philosophical Logic*. (Forthcoming)

Fitch, F. B. (1952). *Symbolic Logic: an Introduction.* Ronald Press Co.

Goble, L. (n.d.). Notes on adaptive deontic logics for normative conflicts. *Unpublished paper*.

Goble, L. (2005). A logic for deontic dilemmas. *Journal of Applied Logic*, *3*, 461-483.

Gowans, C. W. (Ed.). (1987). *Moral dilemmas.* Oxford University Press.

Hansson, S. O. (2001). *The Structure of Values and Norms.* Cambridge University Press.

Horty, J. F. (1994). Moral dilemmas and nonmonotonic logic. *Journal of Philosophical Logic*, *23*(1), 35-66.

Horty, J. F. (2003). Reasoning with moral conflicts. *Noûs*, *37*, 557-605.

Kooi, B., & Tamminga, A. (2008). Moral conflicts between groups of agents. *Journal of Philosophical Logic*, *37*, 1–21.

Meheus, J., Beirlaen, M., & Van De Putte, F. (2010). Avoiding deontic explosion by contextually restricting aggregation. In G. Governatori & G. Sartor (Eds.), *Deon (10th international conference on deontic logic in computer science)* (Vol. 6181, pp. 148–165). Springer.

Shoham, Y. (1987). A semantical approach to nonmonotonic logics. In M. L. Ginsberg (Ed.), *Readings in nonmonotonic reasoning* (pp. 227–250). Morgan Kaufmann Publishers.

Sinnott-Armstrong, W. (1988). *Moral Dilemmas.* Basil Blackwell, Oxford/New York.

Straßer, C. (2010). An adaptive logic framework for conditional obligations and deontic dilemmas. *Logic and Logical Philosophy*, *19*(1-2), 95-128.

Straßer, C., Beirlaen, M., & Meheus, J. (n.d.). Tolerating deontic conflicts by adaptively restricting inheritance.
(Under review)

Van De Putte, F., & Straßer, C. (n.d.). A logic for prioritized normative reasoning.
(Under review)

Williams, B. (1965). Ethical consistency. *Proceedings of the Aristotelian Society (Supplementary Volumes)*, *39*, 103-124.

Wright, G. H. von. (1963). *Norm and Action. A Logical Enquiry.* Routledge and Kegan Paul, London.

Mathieu Beirlaen
Centre for Logic and Philosophy of Science, Ghent University
Blandijnberg 2
B-9000 Gent
Belgium
e-mail: `Mathieu.Beirlaen@UGent.be`