Deontic Reasoning on the Basis of Consistency Considerations

Christian Straßer* Aleks Knoks† Joke Meheus*

*Ghent University †University of Maryland

February 10, 2014

Abstract

Deontic conflicts pose an important challenge to deontic logicians. The standard account —standard deontic logic, **SDL**— is not apt for addressing this challenge since it trivializes conflicts. Two main stratagems for gaining conflict-tolerance have been proposed: to weaken **SDL** in various ways, and to contextualize the reign of **SDL** to consistent subsets of the premise set. The latter began with the work of van Fraassen and has been further developed by Horty. In this paper we characterize this second approach in general terms. We also study three basic ways to contextualize **SDL** and supplement each of these with a dynamic proof theory in the framework of adaptive logics.

Keywords: deontic conflicts; maximal consistent subsets; adaptive logics; nonmonotonic logic; conflict-tolerance;

1 Introduction

One of the major challenges for the deontic logic community is the handling of deontic conflicts. A deontic conflict occurs in a situation in which an agent faces conflicting norms. For instance, our agent may have made two promises: one – to finish and send off a review this evening, and another – to take out his spouse for a romantic dinner. However, it is impossible to fulfill the corresponding obligations to keep both promises.

Let us write $\mathsf{O}A$ for an obligation to bring about A. Each of the following principles may be considered a good candidate for a fundamental principle of deontic reasoning (see [11]):

Agg $\mathsf{O}A$ and $\mathsf{O}B$ imply $\mathsf{O}(A \wedge B)$

Inh If A classically implies B then $\mathsf{O}A$ implies $\mathsf{O}B$

D $\mathsf{O}A$ implies $\neg\mathsf{O}\neg A$

And, indeed, the three inference types are valid in standard deontic logic (henceforth – **SDL**).[1] However, with these three principles we run into a serious com-

[1] We present the full axiomatization of **SDL** in Section 2.

plication when facing a deontic conflict – the so-called *deontic explosion*. That is to say, from any deontic conflict it follows that everything is obligatory (see [12] for a careful analysis). There are two general ways this problem can be (and has been) addressed: we can either soften the strict standards of **SDL** by abandoning or weakening one of the above principles, or we can contextualize the reign of **SDL** to parts of the premise set where trouble does not lurk.

As for the first option, various systems have been proposed: in [10] **Agg** is given up, in [11] **Inh** is restricted, in [28] **Agg** and **Inh** are rejected in two respective phases, and in [10, 9, 6, 5] a paraconsistent negation is used instead of the classical one. However, with plenitude comes choice, and, in absence of guiding meta-principles, the choice may feel somewhat arbitrary. Moreover, compared to **SDL** the new systems are rather weak: many intuitively non-problematic inferences are not anymore valid. A frequently proposed solution is to strengthen these systems non-monotonically (see, e.g., [13, 25, 19, 9, 5, 6]). In addition to the strict rules provided by the weakened version of **SDL** also defeasible reasoning steps are allowed that compensate for the previous weakening.

In this paper we take the second route. Our standard of deontic reasoning remains full **SDL**. However, we restrict or *contextualize* its reign to consistent fragments of the given premise set. This approach is not entirely new. Both the system that van Fraassen put forward in his [29] and the consequence relations studied in Horty's [14, 15, 16] qualify as its instances.[2] Let us now characterize van Fraassen's system for exemplary purposes (the notational conventions are adapted to this paper). In Section 2 we give a more abstract and general characterization of the approach based on contextualizing, as well as some more examples of concrete consequence relations based on this idea.

Suppose we are presented with a set of obligations[3] $\Gamma^{\mathsf{O}} = \{\mathsf{O}A_1, \ldots, \mathsf{O}A_n\}$ where A_i $(1 \leq i \leq n)$ are propositions without normative content (hence, without occurrences of O). $\mathsf{O}B$ is a normative consequence of Γ^{O} if and only if $\mathsf{O}B$ follows by means of **SDL** from some consistent subset of Γ^{O}. Now, it can easily be established that $\mathsf{O}B$ follows by means of **SDL** from Θ^{O} (where the propositions in Θ are without normative content) if and only if B follows from Θ in classical propositional logic. Hence, we can alternatively and equivalently stipulate: $\mathsf{O}B$ is a normative consequence of Γ^{O} if and only if B follows classically from some consistent subset of Γ.[4]

[2] The reader who is interested in a philosophical motivation for this second approach is referred to Horty's more recent [17]. The general idea is that instead of basing a theory of deontic reasoning on general principles as **Agg** or **Inh** it is conceived as an instance of a more general theory of *reasoning in terms of reasons*. The obligations (or norms, more generally) are seen as reasons for acting, and they, rather than some abstract deontic principles, are taken to be fundamental for deontic logic. It should be added that Horty takes great care in linking this abstract picture to ethical conceptions.

[3] We will consider permissions in a follow-up paper. A brief discussion can be found in Section 8.

[4] The characterization of van Fraassen's system as a contextualization of **SDL** is only one way to look at it. Neither van Fraassen, nor Horty characterize it as such. Moreover, in contrast to the syntactic approach in terms of consistent subsets (that serve as contexts) presented in this paper, van Fraassen's approach is semantic. He looks at classes of models

Let us give a small illustration. Suppose Anne's mother tells her to wash the dishes and not to do the homework ($d \wedge \neg h$), while her father orders her to do the opposite —not to wash the dishes and to do the homework ($\neg d \wedge h$)—, and that, in addition, Anne's grandmother requests that she brings her the newspaper (n). According to the characterization of the normative consequences just introduced, Anne has, for instance, to bring the newspaper (n), to wash the dishes (d), but also not to wash the dishes ($\neg d$). On the right hand side of Figure 1 we see the valid contexts $\mathcal{C}_1$–$\mathcal{C}_5$. Every formula that follows from one of these contexts is obligatory in van Fraassen's system.

In this paper we present dynamic proof theories for consequence relations based on contextualizing **SDL**. Let us now give the gist of the approach. (All the technical details are specified in Sections 4–6.) Suppose we are interested in a proof theory that is adequate with respect to the consequence relation of the system just defined. We took the set of contexts $\mathfrak{C}$ to which **SDL** is applied to consist of all consistent subsets of the premise set Γ^{O}. Now, in a dynamic proof, a proof line is equipped with a column in which we keep track of the context that is used to derive the respective formula of the given line. A retraction mechanism marks lines whose associated context is not in $\mathfrak{C}$. Take our example with Anne. This is schematically illustrated in Figure 1.[5] Our consistent contexts are $\mathcal{C}_1$–$\mathcal{C}_5$. Thus, lines l_1–l_6 are derived from consistent contexts. However, line l_7 is derived from a context that is not consistent and hence the respective line will be marked and accordingly considered revoked.

$\vdots$	$\vdots$	$\vdots$	$\vdots$	
l_1	$\mathsf{O}(d \wedge \neg h)$	$\ldots$	$\mathcal{C}_2$	
l_2	$\mathsf{O}n$	$\ldots$	$\mathcal{C}_1$	
l_3	$\mathsf{O}d$	Inh; l_1	$\mathcal{C}_2$	
l_4	$\mathsf{O}(n \wedge d)$	Agg; l_2, l_3	$\mathcal{C}_4$	
l_5	$\mathsf{O}(\neg d \wedge h)$	$\ldots$	$\mathcal{C}_3$	
l_6	$\mathsf{O}h$	Inh; l_5	$\mathcal{C}_3$	
l_7	$\mathsf{O}(d \wedge h)$	Agg; l_3, l_6	$\mathcal{C}_2 \cup \mathcal{C}_3$	*invalid context* $\longrightarrow$ *retraction*

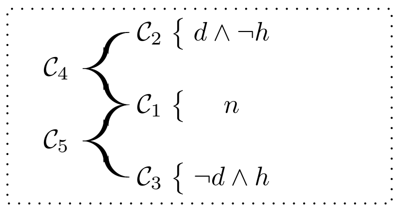

Figure 1

We will spell out these proofs in terms of adaptive logics in standard format. An advantage is that their meta-theory is well-investigated [3]. Adaptive logics in the standard format consist of three parameters: the so-called lower limit logic, a set of so-called abnormalities, and an adaptive strategy. All the adap-

with maximal *score*, where the score of a model consists of all the formulas in the premise set that it satisfies. See the discussion in [16], especially, Fact 2 (p. 598), where the semantic approach is shown to be equivalent to the syntactic approach.

[5] Figure 1 is slightly simplified for didactic purposes. As will be explained in Section 3, the fourth column of an adaptive proof carries the so-called condition of a line. It is not identical to a context (as in the simplified schematic representation of Figure 1), but is uniquely and in a very straightforward way associated with a context.

tive logics presented here are distinguished merely by the third parameter: the strategy. They are all based on **SDL** as the lower-limit logic and use the same abnormalities. We will see that the three standard strategies give rise to three intuitive consequence relations two of which have already been described in the deontic logic literature. Moreover, varying the parameters gives rise to other consequence relations based on the idea of contextualizing **SDL**. We shortly discuss them in Section 8.

The paper is structured as follows. In Section 2 we give a generic account of the idea of contextualizing **SDL** and subsequently introduce some concrete basic consequence relations: $\mid\!\sim_C, \mid\!\sim_D$, and $\mid\!\sim_U$. In Section 3 we introduce the three components that make up the adaptive logics that are presented in this paper. In Sections 4 (resp. 5, resp. 6) we introduce the dynamic proof theory for $\mid\!\sim_C$ (resp. $\mid\!\sim_U$, resp. $\mid\!\sim_D$). In Section 7 we have a brief look at the semantics of our logics. Finally, in Section 8 we compare our systems to some other approaches from the literature, indicate how other consequence relations based on the idea of contextualizing **SDL** can be characterized by adaptive logics, and suggest some directions for future research.

2 Contextualizing SDL

Our formal language. Where $\mathcal{A}$ is a set of propositional letters, we let $\mathcal{W}_{\text{pro}}$ stand for the $\langle \wedge, \vee, \supset, \neg \rangle$-closure of $\mathcal{A}$. We denote members of $\mathcal{A}$ by lower case letters $a, b, c, \ldots$. We write $\vdash$ for the consequence relation of classical propositional logic (**CL**) based on the language $\mathcal{W}_{\text{pro}}$.

Since we are in a deontic setting, we must be able to represent obligations. To this end we introduce two operators: $\circ$ and O. As we will see, the core idea behind the contextualizing approach can be realized in several ways, thus, giving rise to different consequence relations. In all of these, we will use $\circ$ to indicate something like explicitly stated norms, norms which we consider but do not yet have a commitment to, norms as they are when we start reasoning about them.[6] O, on the other hand, represents norms that we accept as consequences of our deliberation process. The different consequence relations differ in their treatment of O. Hence, the exact meaning of O depends on the formal system and the application context in which it is used. We will give examples below.

Since we do not allow for nested occurrences of deontic operators, our full set of well-formed formulas $\mathcal{W}$ is formed by taking the $\langle \wedge, \vee, \supset, \neg \rangle$-closure of $\mathcal{W}_{\text{pro}} \cup \mathcal{W}_{\circ} \cup \mathcal{W}_{\mathsf{O}}$, where $\mathcal{W}_{\circ} = \{\circ A \mid A \in \mathcal{W}_{\text{pro}}\}$ and $\mathcal{W}_{\mathsf{O}} = \{\mathsf{O}A \mid A \in \mathcal{W}_{\text{pro}}\}$. Lastly, given some $\Gamma \subseteq \mathcal{W}_{\text{pro}}$, we use Γ° to refer to the set $\{\circ A \mid A \in \Gamma\}$.

Standard Deontic Logic (SDL). One standard way to obtain **SDL** is by means of:

$$\text{If } \vdash A \text{ then } \vdash \mathsf{O}A \qquad \text{(NEC)}$$

[6]For some additional discussion and motivation of $\circ$ see Section 8.1.

$$\vdash \mathsf{O}(A \supset B) \supset (\mathsf{O}A \supset \mathsf{O}B) \quad \text{(K)}$$

$$\vdash \mathsf{O}A \supset \neg\mathsf{O}\neg A \quad \text{(D)}$$

Hence, O is a normal **KD**-modality. Note that in our setting $\circ$ is basically a 'dummy'-operator: $\circ A$ is treated just like a propositional atom. For instance, we do not have weakening for such formulas (e.g., $\circ A$ does not imply $\circ(A \vee B)$), nor replacement of equivalents (from $\vdash A \equiv B$ and $\circ A$ it does not follow that $\circ B$).

Semantically, **SDL** is characterized by a standard Kripke-frame for **KD**. An **SDL**-model M is a tuple $\langle W, R, v, a\rangle$ where W is a set of points (called worlds), R is an accessibility relation, $v : W \times (\mathcal{A} \cup \mathcal{W}_\circ) \to \{0, 1\}$ assigns truth values to atoms and formulas preceded by $\circ$ at each world,[7] and $a \in W$ is the actual world. Usually, in **SDL** R is a serial relation between worlds, but, since in this paper we do not allow for nested occurrences of O, we can simplify the setup by letting R be a non-empty subset of W containing all the worlds accessible from the actual world a. Similarly, since we do not allow for nested occurrences of $\circ$, we need not consider $v(w, \circ A)$ where $w \neq a$. Hence, we let $v : (W \times \mathcal{A}) \cup (\{a\} \times \mathcal{W}_\circ) \to \{0, 1\}$. We let conventionally $v_\circ(A) =_{\text{df}} v(a, \circ A)$.

Truth at a world $w \in W$ is then defined as follows:

- $M, w \models p$ iff $v(w, p) = 1$ where $p \in \mathcal{A}$
- $M, w \models \neg A$ iff $M, w \not\models A$ where $A \in \mathcal{W}$
- $M, w \models A \wedge B$ iff $M, w \models A$ and $M, w \models B$ where $A, B \in \mathcal{W}$
- similarly for the other Boolean connectives
- $M, a \models \mathsf{O}A$ iff for all $w \in R$, $M, w \models A$, where $A \in \mathcal{W}_{\text{pro}}$
- $M, a \models \circ A$ iff $v_\circ(A) = 1$, where $A \in \mathcal{W}_{\text{pro}}$

Contextualizing SDL. **SDL** can be contextualized in various different ways. From an abstract point of view the idea is as follows (it is illustrated in Figure 2). Suppose our premise set is Γ°, where $\Gamma \subseteq \mathcal{W}_{\text{pro}}$.

a. First, we define a set of *contexts* $\mathfrak{C} = \Gamma_1^{\mathsf{O}}, \ldots, \Gamma_m^{\mathsf{O}}$ where $\Gamma_i \subseteq \Gamma$ for each $1 \leq i \leq m$.

b. Second, **SDL** is applied to each context delivering $Cn_{\mathbf{SDL}}\left(\Gamma_i^{\mathsf{O}}\right)$ (where $1 \leq i \leq m$).

c. Third, we define the consequence set in one of two ways: either *credulously* by taking the union $\bigcup_{i=1}^m Cn_{\mathbf{SDL}}\left(\Gamma_i^{\mathsf{O}}\right)$, or *skeptically* by taking the intersection $\bigcap_{i=1}^m Cn_{\mathbf{SDL}}\left(\Gamma_i^{\mathsf{O}}\right)$.

If we focus on premise sets that consist of obligations only, there is even no need to consider **SDL** since

Fact 2.1. *$\mathsf{O}A_1, \ldots, \mathsf{O}A_n \vdash_{\mathbf{SDL}} \mathsf{O}B$ iff $A_1, \ldots, A_n \vdash_{\mathbf{CL}} B$ where A_i $(1 \leq i \leq n)$ has no occurrences of O.*

[7] Recall that formulas of the type $\circ A$ are treated like atoms. Of course, instead of using a unary 'dummy'-operator $\circ$ we could have as well used an enriched set of atoms: $\mathcal{A} \cup \{A_\circ \mid A \in \mathcal{W}_{\text{pro}}\}$.

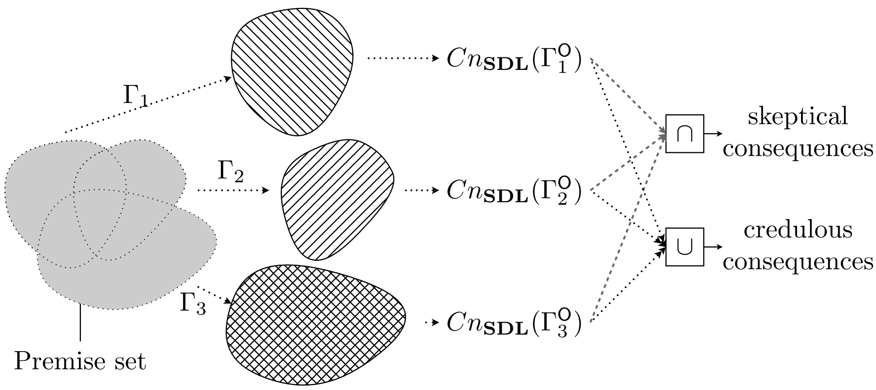

Figure 2

The situation is different if more expressiveness is required, and we want, for instance, also to deal with permissions. (Being a more involving topic, permissions are beyond the scope of this paper, but see Section 8.)

Steps (a)–(c) are summarized in the following definition. (i) is the credulous approach, and (ii) is the skeptical:

Definition 1. Where $\mathfrak{C}(\Gamma) \subseteq \wp(\Gamma)$:

(i) $\Gamma^\circ \mathrel{|\!\!\sim}^{\cap}_{\mathfrak{C}} \mathsf{O}A$ iff for all $\Theta \in \mathfrak{C}(\Gamma)$, $\Theta \vdash A$

(ii) $\Gamma^\circ \mathrel{|\!\!\sim}^{\cup}_{\mathfrak{C}} \mathsf{O}A$ iff for some $\Theta \in \mathfrak{C}(\Gamma)$, $\Theta \vdash A$

Here are some examples.

1. Letting $\mathfrak{C}$ consist of all consistent subsets of Γ and opting for a credulous approach we get exactly the approach we have already discussed in Section 1. Formally,

Definition 2. $\mathrel{|\!\!\sim}_C =_{\mathrm{df}} \mathrel{|\!\!\sim}^{\cup}_{\mathsf{CS}}$, where $\mathsf{CS}(\Gamma)$ is the set of **CL**-consistent subsets of Γ.

The fact that every consistent subset of Γ is contained in one of its maximally consistent subsets and the monotonicity of **CL** immediately imply:

Fact 2.2. $\mathrel{|\!\!\sim}_C = \mathrel{|\!\!\sim}^{\cup}_{\mathsf{MCS}}$, *where* $\mathsf{MCS}(\Gamma) =_{df} \max_{\subset}(\mathsf{CS}(\Gamma))$ *is the set of maximally consistent subsets of* Γ.

2. Another option is to let the contexts $\mathfrak{C}$ be the set of all maximally consistent subsets of Γ and to use the skeptical approach.

Definition 3. $\mathrel{|\!\!\sim}_D =_{\mathrm{df}} \mathrel{|\!\!\sim}^{\cap}_{\mathsf{MCS}}$

3. Let the only context in $\mathfrak{C}$ be the set of all obligations in Γ^{O} that are not involved in deontic conflicts. Evidently, the credulous and skeptical approach coincide in this case. Formally,

Definition 4. $\mathrel{|\!\!\sim}_U =_{\rm df} \mathrel{|\!\!\sim}^{\cap}_{\mathsf{IB}}$, where $\mathsf{ib}(\Gamma) =_{\rm df} \bigcap_{\Lambda \in \mathsf{MCS}(\Gamma)} \Lambda$ and $\mathsf{IB}(\Gamma) =_{\rm df} \{\mathsf{ib}(\Gamma)\}$.[8]

4. Similarly, we can let $\mathfrak{C}$ consist of all consistent subsets of Γ with maximal cardinality.

Definition 5. $\mathrel{|\!\!\sim}^{\#}_D =_{\rm df} \mathrel{|\!\!\sim}^{\cap}_{\mathsf{MCS}_\#}$, where $\mathsf{MCS}_\#(\Gamma) = \max_{<_\#}(\mathsf{CS}(\Gamma))$ and $\Theta <_\# \Theta'$ iff $\mathsf{card}(\Theta) < \mathsf{card}(\Theta')$.

In this paper we focus on the consequence relations $\mathrel{|\!\!\sim}_C, \mathrel{|\!\!\sim}_D$ and $\mathrel{|\!\!\sim}_U$. Let us, therefore, comment more on these, as well as some of their possible application contexts.

The first relation — the 'C-consequence' — is equivalent to the original proposal of van Fraassen, and it also appears in Horty's [14, 15] and [16]. Van Fraassen's motivation was to have a logic which, in contrast to **SDL**, would allow for normative dilemmas. And, indeed, with the C-consequence we get an abundance of conflicts among the derived norms. Thus, also the name: the 'C' stands for 'Conflict'. Let us illustrate this approach with a simple example.

Recall the situation we considered in the introduction: Anne's mother tells her to wash the dishes and not do the homework — $d \wedge \neg h$ —, while her father orders her to do the opposite — $\neg d \wedge h$—, and, in addition, Anne's grandmother requests that she brings her a newspaper: n. The explicitly given obligations in this situation are represented by the set $\Gamma_1^\circ = \{\circ(d \wedge \neg h), \circ(\neg d \wedge h), \circ n\}$. Now, $\Gamma_1 = \{d \wedge \neg h, \neg d \wedge h, n\}$ has two MCSs: $\{d \wedge \neg h, n\}$ and $\{\neg d \wedge h, n\}$, and, as a consequence, we have both $\Gamma_1^\circ \mathrel{|\!\!\sim}_C \mathsf{O}d$ and $\Gamma_1^\circ \mathrel{|\!\!\sim}_C \mathsf{O}\neg d$. Thus, our poor Anne is both obliged to do the dishes and obliged not to do the dishes. Although this result may strike us as odd, it is exactly what we should get, given the rationale behind the C-approach: anything that follows from a consistent set of obligations is to be considered a derived obligation.

Apparently, this approach is not good for contexts such as deriving obligations that would form the basis of actions (e.g., Anne cannot both do and not do the dishes), but it is useful for other contexts. For instance, it can be used when we are interested in mapping out or analyzing all the obligations that we can 'consistently' derive and to discover inconsistencies — a job that is not at all a trivial matter when this set is large or obligations themselves are complex. The knowledge gained may, for instance, be useful for a subsequent revision of the normative code(s) in question. On apprehending that she is both obliged and forbidden to do the dishes, Anne might confront her parents with this apparent problem and ask them to sort it out.

'D' in 'D-consequence' stands for 'disjunctive', and our toy example will show why. Notice that for $\mathsf{O}A$ to be derivable from Γ_1°, A has to be implied by all $\Lambda \in \mathsf{MCS}(\Gamma_1)$. This means that, unlike the case of $\mathrel{|\!\!\sim}_C$, we get neither $\Gamma_1^\circ \mathrel{|\!\!\sim}_D \mathsf{O}(d \wedge \neg h)$ nor $\Gamma_1^\circ \mathrel{|\!\!\sim}_D \mathsf{O}(\neg d \wedge h)$, but we get the corresponding disjunction, i.e., $\Gamma_1^\circ \mathrel{|\!\!\sim}_D \mathsf{O}((d \wedge \neg h) \vee (\neg d \wedge h))$. Now, this also happens to hold in general. Whenever Γ° contains $\circ A$ and $\circ B$ such that $\vdash A \supset \neg B$ and both A and B

[8] Evidently, since $\{\mathsf{ib}(\Gamma)\}$ is a singleton we could have equivalently defined: $\mathrel{|\!\!\sim}_U =_{\rm df} \mathrel{|\!\!\sim}^{\cup}_{\mathsf{IB}}$.

are consistent, the C-consequence gives us both $\mathsf{O}A$ and $\mathsf{O}B$, while the D-consequence allows only for a disjunctive obligation — $\mathsf{O}(A \vee B)$ (given that there is no consistent $\Theta \subseteq \Gamma$ for which $\Theta \vdash \neg A \wedge \neg B$). Notice that such an obligation still constraints the agent's actions: s/he has a choice between doing A or B, but is committed to doing at least one. Since the D-approach always gives a consistent set of O-obligations (as can be easily seen), one of its possible contexts of application are situations in which we want to derive obligations that would form the basis for action. The D-consequences were first introduced in [15], and the reader can find an attempt to link them to ethical conceptions in [16, pp. 569–71].

The 'U' in $\mid\!\sim_U$ stands for 'un-conflicted'. The U-consequence is even more cautious than the previous one. Intuitively, it does not take into account those of the explicitly given obligations that give rise to conflicts. The set $\mathsf{ib}(\Gamma)$ — '*ib*' standing for 'innocent bystander', a notion that was introduced already in [21], — reflects this idea since it consists of exactly those formulas of Γ that in no way contribute to its inconsistency. This can also be seen when characterizing $\mathsf{ib}(\Gamma)$ in an alternative and equivalent way by means of minimally inconsistent subsets:

Definition 6. A subset Γ' of Γ is said to be *minimally inconsistent* iff it is inconsistent and no subset Γ'' of Γ' is inconsistent. We write $\mathsf{MIS}\,(\Gamma)$ for the set of all minimally inconsistent subsets of Γ.

As has been pointed out in [7],

Fact 2.3. $A \in \mathsf{ib}(\Gamma)$ *iff* A *is not in any minimal inconsistent subset of* Γ. *In signs:* $\mathsf{ib}(\Gamma) = \Gamma \setminus \bigcup_{\Lambda \in \mathsf{MIS}(\Gamma)} \Lambda$.

Let us illustrate this by an example. Let $\Gamma_2^\circ = \{\circ p, \circ(\neg p \wedge q), \circ r\}$. Now, the corresponding Γ_2 is inconsistent because of formulas p and $\neg p \wedge q$, but r is 'innocent', and we, therefore, get $\Gamma_2^\circ \mid\!\sim_U \mathsf{O}r$. Notice also that we can derive $\mathsf{O}(p \vee (\neg p \wedge r))$ on the D-approach, but we do not get $\Gamma_2^\circ \mid\!\sim_U \mathsf{O}(p \vee (\neg p \wedge r))$.

The U-consequence provides intuitive results in application contexts where the risk of error is high. Let us narrow the interpretation of $\circ A$ to 'It is explicitly stated that A should be done' and that of $\mathsf{O}A$ to 'A should be done'. Now, consider a context in which a first expert system recommends that you should invest 70% of your money into some stocks S while another expert system recommends that you should rather invest 75% of your money into stocks S'. In this case you may want to be careful and not to derive that you should either invest 75% into S or 75% into S'. After all, the disagreement of the two expert systems may itself indicate a high degree of risk.[9]

[9]To the best of our knowledge, Neither $\mid\!\sim_U$, nor $\mid\!\sim_D^{\#}$ have been discussed in the context of deontic logics yet. It should be noted that general versions of all four relations — i.e., not restricted to deontic settings — have been considered in the literature. The ones corresponding to the C- and the D-consequences have appeared under various names: they have been called 'weak' and 'inevitable' in [21], 'existential' and 'universal' in [7], and 'weak' and 'strong' in [1]. The more general version of $\mid\!\sim_U$ is referred to as 'free' and that of $\mid\!\sim_D^{\#}$ – as 'cardinality-based' in [7].

In our opinion, all three of the relations defined give rise to reasonable consequences — albeit, on different readings of $\circ$ and O —, and are therefore worthy of further study. We will supply each of the three relations with a proof theory. Before turning to this task, let us notice that the three consequence relations are related as follows:

Proposition 2.1. $\Gamma^{\circ} \mid\!\sim_U \mathsf{O}A$ *implies* $\Gamma^{\circ} \mid\!\sim_D \mathsf{O}A$ *implies* $\Gamma^{\circ} \mid\!\sim_C \mathsf{O}A$.

The following examples may help the reader to get a better grip on $\mid\!\sim_C$, $\mid\!\sim_D$, and $\mid\!\sim_U$.

Example 1. Let $\Gamma_3 = \{a \wedge b, \neg a \wedge b\}$,

- $\Gamma_3^{\circ} \mid\!\sim_C \mathsf{O}a, \mathsf{O}\neg a, \mathsf{O}b$
- $\Gamma_3^{\circ} \mid\!\sim_D \mathsf{O}b$ while $\Gamma_3^{\circ} \not\mid\!\sim_D \mathsf{O}a$ and $\Gamma_3^{\circ} \not\mid\!\sim_D \mathsf{O}\neg a$
- $\Gamma_3^{\circ} \not\mid\!\sim_U \mathsf{O}b$; $\Gamma_3^{\circ} \not\mid\!\sim_U \mathsf{O}a$ and $\Gamma_3^{\circ} \not\mid\!\sim_U \mathsf{O}\neg a$

Example 2. Let $\Gamma_4 = \{a \wedge b, \neg a, c\}$

- $\Gamma_4^{\circ} \mid\!\sim_C \mathsf{O}a, \mathsf{O}\neg a, \mathsf{O}b, \mathsf{O}c$
- $\Gamma_4^{\circ} \mid\!\sim_D \mathsf{O}c$; $\Gamma_4^{\circ} \not\mid\!\sim_D \mathsf{O}b$; and $\Gamma_4^{\circ} \mid\!\sim_D \mathsf{O}((a \wedge b) \vee \neg a)$
- $\Gamma_4^{\circ} \mid\!\sim_U \mathsf{O}c$; $\Gamma_4^{\circ} \not\mid\!\sim_U \mathsf{O}b$; and $\Gamma_4^{\circ} \not\mid\!\sim_U \mathsf{O}((a \wedge b) \vee \neg a)$

3 Adaptive Logics in the Standard Format

From the next section on we develop adaptive proof theories for the three consequence relations $\mid\!\sim_C, \mid\!\sim_D$, and $\mid\!\sim_U$. We will present three adaptive logics. Since all of them are in the so-called standard format, we will immediately be equipped with a rich meta-theory (see, e.g., [4]). Adaptive logics in the standard format are characterized by triples (**LLL**, Ω, $\mathbf{x}$) where:

1. **LLL** is the so-called *Lower Limit Logic.* It forms the nonmonotonic core-logic: while all the axioms and inference rules of **LLL** are available in the adaptive logic, the **LLL** is strengthened by means of allowing for additional defeasible inferences (see below). We will use the lower limit logic **SDL** based on the language $\mathcal{W}$.

2. Ω is the set of abnormalities which are characterized by a logical form, in our case: $\circ A \wedge \neg \mathsf{O}A$ where $A \in \mathcal{W}_{\text{pro}}$. Let, thus,

$$\Omega_{\circ} =_{\text{df}} \{\circ A \wedge \neg \mathsf{O}A \mid A \in \mathcal{W}_{\text{pro}}\}$$

   In order to avoid notational clutter we will often use the abbreviations: $\lightning A =_{\text{df}} \circ A \wedge \neg \mathsf{O}A$ and $\lightning \Delta =_{\text{df}} \{\lightning A \mid A \in \Delta\}$ where $\Delta \subseteq \mathcal{W}_{\text{pro}}$ and $A \in \mathcal{W}_{\text{pro}}$.

   It seems intuitive to regard formulas of this form as abnormal: although A is an explicitly stated obligation ($\circ A$), it is nevertheless not the case that A is a derived obligation ($\neg \mathsf{O}A$). The core idea of adaptive logics is to regard abnormalities as false, unless there is a reason to suppose that they are true, or, in other words, to interpret abnormalities as false 'as much as possible'. Since $\circ A \wedge \neg \mathsf{O}A$ is equivalent to $\neg(\circ A \supset \mathsf{O}A)$ this means that we want to verify $\circ A \supset \mathsf{O}A$ 'as much as possible'. In other words, we want to derive $\mathsf{O}A$ from $\circ A$ 'as much as possible'.

3. The phrase 'as much as possible' is disambiguated by means of the adaptive strategy, indicated by $\mathbf{x}$. In this paper, we will work with three different strategies: *normal selections* ($\mathbf{x} = \mathbf{n}$), *reliability* ($\mathbf{x} = \mathbf{r}$), and *minimal abnormality* ($\mathbf{x} = \mathbf{m}$).[10]

In what follows we present three adaptive logics all based on the set of abnormalities $\Omega_\circ$ and the same lower limit logic **SDL**. The only difference concerns the strategy. We begin with normal selections in Section 4 which gives us a characterization of the C-consequence, then move to reliability in Section 5 which gives a characterization of the U-consequence, and, lastly, we discuss minimal abnormality in Section 6 characterizing the D-consequence. In accordance with the main thrust of our paper, the presentation of adaptive logics focuses on the proof theory and, hence, on the syntactic consequence relation. In a nutshell, we will obtain the following representation theorem (which is proven in the respective sections 4.2, 5.2, and 6.2):

Theorem 3.1. *Where* $\Gamma \subseteq \mathcal{W}_{\text{pro}}$*,*

(i) $\Gamma^\circ \mid\!\sim_C \mathsf{O}A$ *iff* $\Gamma^\circ \vdash_{\mathbf{SDL^n}} \mathsf{O}A$*,*

(ii) $\Gamma^\circ \mid\!\sim_U \mathsf{O}A$ *iff,* $\Gamma^\circ \vdash_{\mathbf{SDL^r}} \mathsf{O}A$

(iii) $\Gamma^\circ \mid\!\sim_D \mathsf{O}A$ *iff* $\Gamma^\circ \vdash_{\mathbf{SDL^m}} \mathsf{O}A$*,*

4 Normal Selections and C-Consequences

The first adaptive logic we introduce uses the normal selections strategy. Let, thus, $\mathbf{SDL^n}$ be characterized by the triple $\langle \mathbf{SDL}, \Omega_\circ, \mathbf{n} \rangle$. We start by explaining the proof format, which is common to all three systems to be presented.

4.1 The adaptive proof theory for C-consequences

Lines in adaptive proofs consist of four elements: a line number, a formula, a justification (calling upon previous line numbers and a rule, see below), and a condition which is a finite set of abnormalities. The general idea is this: the formula is derived on the assumption that no abnormality occurring in the condition is true. Inferences in proofs are made by means of three generic rules.

[10] Although the standard format only comprises reliability and minimal abnormality, it is shown that normal selections is under a modal translation reducible to either of the two (see [24]). The need to disambiguate the meaning of 'as much as possible' is due to the fact that sometimes only disjunctions of abnormalities are derivable, but neither of the disjuncts are derivable on their own. The reader will find examples below.

The first rule is PREM. It allows introducing premises on the empty condition.

$$\text{If } A \text{ is a premise:} \qquad \begin{array}{cccc} \vdots & \vdots & \vdots & \vdots \\ \hline l & A & \text{PREM} & \emptyset \end{array}$$

Next, we have the unconditional rule RU. Where $l_{n+1} > l_n > \ldots > l_1$,

$$\text{If } A_1, \ldots, A_n \vdash_{\mathbf{SDL}} B: \qquad \begin{array}{cccc} l_1 & A_1 & \ldots & \Delta_1 \\ \vdots & \vdots & \vdots & \vdots \\ l_n & A_n & \ldots & \Delta_n \\ \hline l_{n+1} & B & l_1, \ldots, l_n; \text{RU} & \bigcup_{i=1}^{n} \Delta_i \end{array}$$

Because of these two rules all the derivative power of our lower limit logic **SDL** is available in adaptive proofs: whenever A is derivable from $A_1, \ldots, A_n$ in **SDL** it can also be derived from $A_1, \ldots, A_n$ in the adaptive logic. Notice that the conditions of the lines at which $A_1, \ldots, A_n$ have been derived are carried forward and merged.

The most interesting rule is the conditional rule RC. It allows for defeasible, conditional derivations. Due to their central place in the adaptive proof- and meta-theory it is practical to use a designated notation for disjunctions of abnormalities: $\mathsf{Dab}(\Delta) =_{\text{df}} \bigvee_{A \in \Delta} A$ where $\Delta \subseteq \Omega_\circ$. RC allows to derive A on the condition Δ from $A_1, \ldots, A_n$ iff $A \vee \mathsf{Dab}(\Delta)$ is derivable from $A_1, \ldots, A_n$ in the lower limit logic. The intuitive meaning is that A is derived on the assumption that neither of the abnormalities in Δ is true. Evidently, we have $\circ A \vdash_{\mathbf{SDL}} \mathsf{O}A \vee \lightning A$. This enables us to derive $\mathsf{O}A$ from $\circ A$ on the condition $\{\lightning A\}$. Similarly, we have for instance $\circ B, \mathsf{O}A \vdash_{\mathbf{SDL}} \mathsf{O}(A \wedge B) \vee \lightning B$. So, suppose we have already derived $\mathsf{O}A$ on the condition $\{\lightning A\}$ and we also have $\circ B$, then we can derive $\mathsf{O}(A \wedge B)$ from $\mathsf{O}A$ and $\circ B$ by adding $\lightning B$ to the condition $\{\lightning A\}$ resulting in $\{\lightning A, \lightning B\}$. In this way we keep track of the context, i.e., of the explicitly stated obligations (preceded by '$\circ$') that we use to derive O-obligations, such as $\{\circ A, \circ B\}$ in this case. More generally: a context Δ° is tracked as $\lightning \Delta$ in the condition column of an adaptive proof.

Let us now state the generic rule RC (where $l_1 < \ldots < l_{n+1}$):

$$\text{If } A_1, \ldots, A_n \vdash_{\mathbf{SDL}} B \vee \mathsf{Dab}(\Theta): \qquad \begin{array}{cccc} l_1 & A_1 & \ldots & \Delta_1 \\ \vdots & \vdots & \vdots & \vdots \\ l_n & A_n & \ldots & \Delta_n \\ \hline l_{n+1} & B & l_1, \ldots, l_n; \text{RC} & \bigcup_{i=1}^{m} \Delta_i \cup \Theta \end{array}$$

Let us make use of the example from Section 1 to illustrate these rules of inference. Recall that the orders given to Anne were $\Gamma_1^\circ = \{\circ(d \wedge \neg h), \circ(\neg d \wedge h), \circ n\}$. We begin our $\mathbf{SDL^n}$-proof by introducing all the premises.

1	$\circ(d \wedge \neg h)$	PREM	$\emptyset$
2	$\circ(\neg d \wedge h)$	PREM	$\emptyset$
3	$\circ n$	PREM	$\emptyset$

At this point we apply RC to the first three lines. The proof is extended as follows:

4	$\mathsf{O}(d \wedge \neg h)$	1; RC	$\{\lightning(d \wedge \neg h)\}$
5	$\mathsf{O}(\neg d \wedge h)$	2; RC	$\{\lightning(\neg d \wedge h)\}$
6	$\mathsf{O}n$	3; RC	$\{\lightning n\}$

Take, for instance, line 4. Note that $\circ(d \wedge \neg h) \vdash_{\mathbf{SDL}} \mathsf{O}(d \wedge \neg h) \vee \lightning(d \wedge \neg h)$. Hence, we can 'move $\lightning(d \wedge \neg h)$ to the condition column' and conditionally derive $\mathsf{O}(d \wedge h)$.

Let us derive some more formulas using the unconditional rule. Notice how the conditions of the lines to which RU is applied are carried over (and, in the case of lines 9–11, also merged).

7	$\mathsf{O}d$	4; RU	$\{\lightning(d \wedge \neg h)\}$
8	$\mathsf{O}\neg d$	5; RU	$\{\lightning(\neg d \wedge h)\}$
9	$\mathsf{O}(\neg d \wedge n)$	6, 8; RU	$\{\lightning(\neg d \wedge h), \lightning n\}$
10	$\mathsf{O}d \wedge \mathsf{O}\neg d$	7,8; RU	$\{\lightning(d \wedge \neg h), \lightning(\neg d \wedge h)\}$
11	$\mathsf{O}(d \wedge \neg d)$	10; RU	$\{\lightning(d \wedge \neg h), \lightning(\neg d \wedge h)\}$

Recall the intuition behind the C-consequence: $\mathsf{O}A$ is a consequence iff it is derived from a consistent context, i.e., a consistent subset of the premise set. We have already mentioned that the conditions in our dynamic proofs keep track of the contexts. Take a look at the conditions of lines 7–9. These represent consistent contexts. The situation is different for the conditions of lines 10 and 11. Obviously, $(d \wedge \neg h)$ and $(\neg d \wedge h)$ are not classically consistent (or, equivalently, $\mathsf{O}(d \wedge \neg h)$ and $\mathsf{O}(\neg d \wedge h)$ are not **SDL**-consistent). Thus, we expect that the formulas derived at these lines are retracted from the proof so that they are not counted as consequences.

In adaptive logics lines are revoked by means of the so-called *marking*. Notably, definitions of marking are what differentiates our three proof theories. Before we specify the marking for normal selection it is instructive to make the following observation. There is a close relation between, on the one hand, the question whether the context presented by a condition Δ is consistent, and, on the other hand, the question whether $\mathsf{Dab}(\Delta)$ is derivable on the empty condition in a proof. As we will prove in the next section, the answers to the two question coincide:

Corollary 4.1. *Where $\Delta \subseteq \Gamma \subseteq \mathcal{W}_{\text{pro}}$: Δ is* **CL***-inconsistent (resp. Δ^{O} is* **SDL***-inconsistent) iff $\mathsf{Dab}(\lightning\Delta)$ is derivable on the empty condition in a dynamic proof from Γ°.*

And indeed, since $\vdash_{\mathbf{SDL}} \neg\mathsf{O}(d \wedge \neg h) \vee \neg\mathsf{O}(\neg d \wedge h)$ we get: $\{\circ(d \wedge \neg h), \circ(\neg d \wedge h)\} \vdash_{\mathbf{SDL}} \lightning(d \wedge \neg h) \vee \lightning(\neg d \wedge h)$:

12	$\lightning(d \wedge \neg h) \vee \lightning(\neg d \wedge h)$	1,2; RU	$\emptyset$

Hence, at line 12 we derive a disjunction of two abnormal formulas. This disjunction means that we cannot hold to the assumption that formulas $\lightning(d \wedge \neg h)$ and $\lightning(\neg d \wedge h)$ are both false.[11]

Definition 7 (Marking for Normal Selections)**.** A line l is marked at stage s of a proof iff, where Δ is its condition, $\mathsf{Dab}(\Delta)$ has been derived on the empty condition at stage s.[12]

According to our definition, lines 10 and 11 are marked at stage 12. Let us restate the last few lines of the proof so that we see Definition 7 at work.

7	$\mathsf{O}d$	4; RU	$\{\lightning(d \wedge \neg h)\}$
8	$\mathsf{O}\neg d$	5; RU	$\{\lightning(\neg d \wedge h)\}$
9	$\mathsf{O}(\neg d \wedge n)$	6, 8; RC	$\{\lightning(\neg d \wedge h), \circ n \wedge \neg\mathsf{O}n\}$
✓10	$\mathsf{O}d \wedge \mathsf{O}\neg d$	7,8; RU	$\{\lightning(d \wedge \neg h), \lightning(\neg d \wedge h)\}$
✓11	$\mathsf{O}(d \wedge \neg d)$	10; RU	$\{\lightning(d \wedge \neg h), \lightning(\neg d \wedge h)\}$
12	$\lightning(d \wedge \neg h) \vee \lightning(\neg d \wedge h)$	1,2; RU	$\emptyset$

We see that lines whose conditions represent inconsistent contexts are marked, just as expected. Corollary 4.1 and Definition 7 warrant this in general.

If a line is marked at a stage s, the formula occurring as its second element is not considered derived at stage s. In view of the fact that lines that are not marked at some stage may become marked at a later stage, we need a stable — *stage-independent* — notion of derivability in order to define a consequence relation. Thus, we have the following definition.

Definition 8. A formula A is *finally derived* in a proof from Γ at a stage s iff (i) A occurs as the second element on some unmarked line l at stage s, and (ii) line l does not get marked in any further extension of the proof.

The adaptive consequences of some given set Γ are exactly the formulas that can be finally derived.

Definition 9. $\Gamma \vdash_{\mathbf{SDL^n}} A$ iff A is finally derivable in an $\mathbf{SDL^n}$-proof from Γ.

The following representation theorem establishes a link between the marking definition and final derivability. Instead of Dab-formulas derived at a given stage on the empty condition, we now consider Dab-formulas that are **SDL**-consequences. Instead of considering formulas derived on conditions in a dynamic proof we consider formulas derivable in disjunction with Dab-formulas. It is easy to see that

[11]Note that we cannot, in general, know in advance whether the assumptions we make are justified or not. This can only be discovered as the proof proceeds. The reader might not immediately see this because of the simplicity of our toy example, where, indeed, it is clear from the very beginning that there is something wrong with the orders of Anne's parents. This, however, does not hold in general. As sets of obligations grow large and obligations themselves become complex, we have to draw inferences in order to analyze the given obligations before we actually see which disjunctions of abnormalities are derivable.

[12]A stage of a proof is a list of proof lines such that all the lines referred to in any of the justifications of lines in the list are also contained in the list.

Fact 4.1 (see [3], Lemma 1)**.** *A is derivable in a dynamic proof from* Γ *on the condition* Δ *iff* $\Gamma \vdash_{\mathbf{SDL}} A \vee \mathsf{Dab}(\Delta)$.

Given these correlations, the following representation theorem for adaptive consequences mirrors derivability at a stage in view of the adaptive marking.[13]

Theorem 4.1. $\Gamma \vdash_{\mathbf{SDL^n}} A$ *iff there is a finite* $\Delta \subseteq \Omega$ *such that* $\Gamma \vdash_{\mathbf{SDL}} A \vee \mathsf{Dab}(\Delta)$ *and* $\Gamma \nvdash_{\mathbf{SDL}} \mathsf{Dab}(\Delta)$.

In our example we thus have both $\Gamma_1^\circ \vdash_{\mathbf{SDL^n}} \mathsf{O}d$ and $\Gamma_1^\circ \vdash_{\mathbf{SDL^n}} \mathsf{O}\neg d$, since we have $\Gamma_1^\circ \vdash_{\mathbf{SDL}} \mathsf{O}d \vee \lightning(d \wedge \neg h)$ and $\Gamma_1^\circ \vdash_{\mathbf{SDL}} \mathsf{O}\neg d \vee \lightning(\neg d \wedge h)$, while $\Gamma_1^\circ \nvdash_{\mathbf{SDL}} \lightning(d \wedge \neg h)$ and $\Gamma_1^\circ \nvdash_{\mathbf{SDL}} \lightning(\neg d \wedge h)$.

Admittedly, some readers might find it puzzling (and, possibly, counter-intuitive) that, as soon as we discovered that either $\lightning(d \wedge \neg h)$ or $\lightning(\neg d \wedge h)$, we marked the line at which the conjunction $\mathsf{O}d \wedge \mathsf{O}\neg d$ appears, but did nothing about lines 7 and 8 where the two conjuncts appear isolated —, especially, since the derivation of each conjunct depends on one of the 'unsafe' disjuncts. Two notes are in order here.

First, recall that we had $\Gamma_1^\circ \mid\!\sim_C \mathsf{O}d$ and $\Gamma_1^\circ \mid\!\sim_C \mathsf{O}\neg d$ and that $\mathbf{SDL^n}$ equips $\mid\!\sim_C$ with a proof theory. Given that the C-consequence allows one to derive contradictory obligations, but not their conjunction, it should not be very surprising that in an $\mathbf{SDL^n}$-proof the two obligations are finally derivable, while their conjunction gets marked. Second and more importantly, we have to recall that the framework we are developing in this section (just as the C-consequence) is meant to derive an obligation $\mathsf{O}A$ iff A follows from some consistent subset of the explicitly stated obligations. Indeed, $\mathsf{O}d$ then follows in view of $\circ(d \wedge \neg h)$ and $\mathsf{O}\neg d$ follows in view of $\circ(\neg d \wedge h)$. In other words, on one internally consistent interpretation of Γ_1°, Anne is obliged to do the dishes and not to do her homework. On another one – she is obliged not to do the dishes while she is obliged to do her homework. But on no consistent interpretation is she obliged to do the dishes and also not to do them. Thus, we do not get $\mathsf{O}(d \wedge \neg d)$. In fact, this is something that our proof itself has taught us. For what the Dab-formula on line 12 says is that we cannot consistently suppose that both $\neg\lightning(d \wedge \neg h)$ and $\neg\lightning(\neg d \wedge h)$ hold, which is exactly the working assumption at lines 10 and 11.

This concludes the exposition of the proof theory of $\mathbf{SDL^n}$, and now it only remains to link it to the C-consequence relation.

4.2 Meta-Theoretic Results

First, note that we only need to consider those derivations in which only the elements of Γ are featured in the conditions:

Lemma 4.1. *Where* $\Delta \subseteq \Gamma \subseteq \mathcal{W}_{\text{pro}}$*:*

(i) $\Gamma^\circ \vdash_{\mathbf{SDL}} \mathsf{O}A \vee \mathsf{Dab}(\lightning\Delta)$ *iff* $\Gamma^\circ \vdash_{\mathbf{SDL}} \mathsf{O}A \vee \mathsf{Dab}(\lightning(\Delta \cap \Gamma))$*;*

(ii) $\Gamma^\circ \vdash_{\mathbf{SDL}} \mathsf{Dab}(\lightning\Delta)$ *iff* $\Gamma^\circ \vdash_{\mathbf{SDL}} \mathsf{Dab}(\lightning(\Delta \cap \Gamma))$.

[13] This is an instance of a meta-theorem of adaptive logics: see [22, Thm. 2.8.2].

Proof. Ad (i): ($\Leftarrow$) is trivial. ($\Rightarrow$) Suppose $\Gamma^\circ \nvdash_{\mathbf{SDL}} \mathsf{O}A \vee \mathsf{Dab}(\lightning(\Delta \cap \Gamma))$. Hence, there is a model $M = \langle W, R, v, a\rangle$ of Γ° such that $M \not\models \mathsf{O}A$ and for all $B \in \Delta \cap \Gamma$, $M \not\models \lightning B$ and thus $M \models \mathsf{O}B$. Let $M' = \langle W, R, v', a\rangle$ where $v'(w, C) = v(w, C)$ for all $C \in \mathcal{A}$ and all $w \in W$, and $v'_\circ = \mathbf{1}_\Gamma$ (the characteristic function of Γ, i.e., $C \mapsto 1$ if $C \in \Gamma$, else $C \mapsto 0$). Obviously, M' is a model of Γ° and $M' \not\models \mathsf{O}A \vee \mathsf{Dab}(\lightning\Delta)$.

Ad (ii): Let in (i) $A = B \wedge \neg B$ for some arbitrary $B \in \mathcal{W}_{\text{pro}}$. Then $\Gamma^\circ \vdash_{\mathbf{SDL}} \mathsf{Dab}(\lightning\Delta)$, iff [by **SDL**-properties], $\Gamma^\circ \vdash_{\mathbf{SDL}} \mathsf{O}A \vee \mathsf{Dab}(\lightning\Delta)$, iff [by (i)] $\Gamma^\circ \vdash_{\mathbf{SDL}} \mathsf{O}A \vee \mathsf{Dab}(\lightning(\Delta \cap \Gamma))$, iff [by **SDL**-properties], $\Gamma^\circ \vdash_{\mathbf{SDL}} \mathsf{Dab}(\lightning(\Delta \cap \Gamma))$. □

The next theorem shows that $\mathsf{O}A$ is derivable on a condition $\lightning\Delta$ iff Δ implies A classically:

Theorem 4.2. *Where $\Delta \subseteq \Gamma \subseteq \mathcal{W}_{\text{pro}}$: $\Gamma^\circ \vdash_{\mathbf{SDL}} \mathsf{O}A \vee \mathsf{Dab}(\lightning\Delta)$ iff $\Delta \vdash A$.*

Proof. ($\Rightarrow$) Suppose $\Delta \nvdash A$. Hence $\Delta \cup \{\neg A\}$ is **CL**-consistent and there is a maximal **CL**-consistent extension $\Lambda \subseteq \mathcal{W}_{\text{pro}}$ of $\Delta \cup \{\neg A\}$. Let $M = \langle \{a\}, \{a\}, v, a\rangle$ where $v : (a, C) \mapsto \begin{cases} 1 & \text{if } C \in \Lambda \cap \mathcal{A} \\ 0 & \text{else} \end{cases}$ and $v_\circ = \mathbf{1}_\Gamma$. Obviously, M is a model of Γ°, $M \not\models \mathsf{O}A$ and $M \not\models \lightning B$ for all $B \in \Delta$ (since $M \models \mathsf{O}B$). Hence, $\Gamma^\circ \nvdash_{\mathbf{SDL}} \mathsf{O}A \vee \mathsf{Dab}(\lightning\Delta)$.

($\Leftarrow$) Suppose $\Delta \vdash A$. Let M be a model of Γ° (note that for all $\Gamma \subseteq \mathcal{W}_{\text{pro}}$, Γ° is $\mathbf{SDL}_\circ$-consistent). We have to show that $M \models \mathsf{O}A \vee \mathsf{Dab}(\lightning\Delta)$ and hence that $M \models \mathsf{O}A$ or $M \models \mathsf{Dab}(\lightning\Delta)$. Suppose that $M \not\models \mathsf{Dab}(\lightning\Delta)$. Hence, $M \models \bigwedge_{B \in \Delta} \neg\lightning B$. Since, $\Delta \subseteq \Gamma$, $M \models \bigwedge_{B \in \Delta} \mathsf{O}B$, and hence (by **SDL** properties) $M \models \mathsf{O}A$. □

Thus, we also have (to see this just let $A = B \wedge \neg B$ in Theorem 4.2):

Corollary 4.2. *Where $\Delta \subseteq \Gamma \subseteq \mathcal{W}_{\text{pro}}$: $\Gamma^\circ \vdash_{\mathbf{SDL}} \mathsf{Dab}(\lightning\Delta)$ iff Δ is* **CL**-*inconsistent.*

Note that Fact 4.1 implies that $\mathsf{Dab}(\Delta)$ is derivable on the empty condition from Γ iff $\Gamma \vdash_{\mathbf{SDL}} \mathsf{Dab}(\Delta)$. Together with Corollary 4.2 this implies Corollary 4.1.

An immediate consequence is the following theorem that associates minimal Dab-consequences with minimally inconsistent sets (see also Table 1).

Definition 10. $\mathsf{Dab}(\Delta)$ is a *minimal* Dab-*consequence of* Γ iff $\Gamma \vdash_{\mathbf{SDL}} \mathsf{Dab}(\Delta)$ and for all $\Theta \subset \Delta$, $\Gamma \nvdash_{\mathbf{SDL}} \mathsf{Dab}(\Theta)$.

Theorem 4.3. *Where $\Delta \subseteq \Gamma \subseteq \mathcal{W}_{\text{pro}}$: $\mathsf{Dab}(\lightning\Delta)$ is a minimal* Dab-*consequence from Γ° iff $\Delta \in \mathsf{MIS}(\Gamma)$.*

Proof. $\Delta \in \mathsf{MIS}(\Gamma)$, iff, Δ is **CL**-inconsistent and for all $\Delta' \subset \Delta$, Δ' is not **CL**-inconsistent, iff [by Corollary 4.2], $\Gamma^\circ \vdash_{\mathbf{SDL}} \mathsf{Dab}(\lightning\Delta)$ and for all $\Delta' \subset \Delta$, $\Gamma^\circ \nvdash_{\mathbf{SDL}} \mathsf{Dab}(\lightning\Delta')$. □

The idea behind C-consequences is that they are derivable from a consistent subset of Γ. Translated to adaptive proofs this means that a C-consequence A is expected to be derivable on an assumption that is compatible with some consistent subset of Γ: i.e., it is expected to be derivable on a condition $\not\natural\Delta$ such that there is a $\Gamma' \in \mathsf{MCS}(\Gamma)$ for which $\Delta \subseteq \Gamma'$. Corollary 4.4 below shows exactly this. In order to see this, first note that:

Lemma 4.2. *Where $\Gamma \subseteq \mathcal{W}_{\text{pro}}$: Λ is an inconsistent subset of Γ iff $\Lambda \subseteq \Gamma$ and for all $\Theta \in \mathsf{MCS}(\Gamma)$, $\Lambda \setminus \Theta \neq \emptyset$.*

Proof. ($\Leftarrow$) Suppose $\Lambda \subseteq \Gamma$ is consistent. Hence, there is a $\Theta \in \mathsf{MCS}(\Gamma)$ such that $\Lambda \subseteq \Theta$ and thus $\Lambda \setminus \Theta = \emptyset$. ($\Rightarrow$) Suppose there is a $\Theta \in \mathsf{MCS}(\Gamma)$ for which $\Lambda \setminus \Theta = \emptyset$ and hence $\Lambda \subseteq \Theta$. Hence, Λ is a consistent subset of Γ. □

By Lemma 4.2 and Theorem 4.3 we get:

Corollary 4.3. *Where $\Delta \subseteq \Gamma \subseteq \mathcal{W}_{\text{pro}}$:*

(i) $\Gamma^{\circ} \vdash_{\mathbf{SDL}} \mathsf{Dab}(\not\natural\Delta)$ *iff for all $\Theta \in \mathsf{MCS}(\Gamma)$, $\Delta \setminus \Theta \neq \emptyset$; and equivalently*

(ii) $\Gamma^{\circ} \nvdash_{\mathbf{SDL}} \mathsf{Dab}(\not\natural\Delta)$ *iff there is a $\Theta \in \mathsf{MCS}(\Gamma)$ such that $\Delta \subseteq \Theta$.*

By Corollary 4.3 and Theorem 4.1 we immediately get the expected:

Corollary 4.4. *Where $\Gamma \subseteq \mathcal{W}_{\text{pro}}$: $\Gamma^{\circ} \vdash_{\mathbf{SDL^n}} \mathsf{O}A$ iff there is a $\Delta \in \mathsf{MCS}(\Gamma)$ and a $\Delta' \subseteq \Delta$ for which $\Gamma^{\circ} \vdash_{\mathbf{SDL}} \mathsf{O}A \vee \mathsf{Dab}(\not\natural\Delta')$.*

In view of Corollary 4.4 it is easy to prove our main theorem of this section.

Theorem 4.4. *Where $\Gamma \subseteq \mathcal{W}_{\text{pro}}$: $\Gamma^{\circ} \vdash_{\mathbf{SDL^n}} \mathsf{O}A$ iff $\Gamma \mathrel{|\!\!\sim_C} A$.*

Proof. $\Gamma^{\circ} \vdash_{\mathbf{SDL^n}} \mathsf{O}A$, iff [by Corollary 4.4], there is a $\Delta \in \mathsf{MCS}(\Gamma)$ and a finite $\Delta' \subseteq \Delta$ such that $\Gamma^{\circ} \vdash_{\mathbf{SDL}} \mathsf{O}A \vee \mathsf{Dab}(\not\natural\Delta')$, iff [by Theorem 4.2], there is a $\Delta \in \mathsf{MCS}(\Gamma)$ and a finite $\Delta' \subseteq \Delta$ such that $\Delta' \vdash A$, iff [by the compactness of **CL**], there is a $\Delta \in \mathsf{MCS}(\Gamma)$ such that $\Delta \vdash A$, iff [by the definition of $\mathrel{|\!\!\sim_C}$], $\Gamma \mathrel{|\!\!\sim_C} A$. □

5 Reliability and U-consequences

In this section we introduce our second adaptive logic that relies on the so-called reliability strategy. $\mathbf{SDL^r}$ is characterized by the triple $\langle \mathbf{SDL},\ \Omega_{\circ},\ \mathbf{r} \rangle$.

5.1 The adaptive proof theory for U-consequences

Recall that the U-consequences are the ones that follow from the unique context consisting of the 'innocent bystanders' of the premise set, i.e., they follow from those obligations that are not involved in conflicts. Hence, we expect that lines are marked which are derived on conditions that contain obligations that are not 'innocent' in this sense. In order to see how this is achieved, let us look again at the premise set Γ_1°.

1	$\circ n$	PREM	$\emptyset$
2	$\mathsf{O}n$	1; RC	$\{\lightning n\}$
3	$\circ(d \wedge \neg h)$	PREM	$\emptyset$
4	$\circ(\neg d \wedge h)$	PREM	$\emptyset$
✓5	$\mathsf{O}(d \wedge \neg h)$	3; RC	$\{\lightning(d \wedge \neg h)\}$
✓6	$\mathsf{O}d \wedge \mathsf{O}n$	2,5; RU	$\{\lightning(d \wedge \neg h), \lightning n\}$
7	$\lightning(d \wedge \neg h) \vee \lightning(\neg d \wedge h)$	3,4; RU	$\emptyset$

At lines 1, 3, and 4 we introduce premises, at lines 2 and 5 we have conditional derivations in accordance with RC, and, finally, at lines 6 and 7 we have applications of the unconditional rule RU. We have seen all this. What has changed is the way marking works (recall that it is exactly what makes the **SDL**-based proof theories differ from each other).

The idea is now as follows. Recall that some $A \in \Gamma$ is not an innocent bystander of Γ iff A is involved in a minimally inconsistent subset of Γ (see Fact 2.3). Hence, we expect to mark a line iff its condition contains some $\lightning A$ such that $A \in \Delta$ where $\Delta \in \mathsf{MIS}(\Gamma)$. By Theorem 4.3 this means that we expect to mark a line iff its conditions contains some $\lightning A$ which is part of a minimal Dab-formula. Since the marking is applied stage-wise we need a stage-dependent notion of minimal Dab-formulas:

Definition 11. A disjunction of abnormalities $\mathsf{Dab}(\Delta)$ is a *minimal* Dab*-formula at stage s* in a proof from Γ iff, at stage s $\mathsf{Dab}(\Delta)$ is derived on the condition $\emptyset$ and if $\mathsf{Dab}(\Theta)$ with $\Theta \subseteq \Delta$ is derived at stage s on the empty condition, then $\Theta = \Delta$.

It is obvious that the Dab-formula occurring at line 7 is minimal. Whenever several Dab-formulas are derived in the same proof, all of their disjuncts together are deemed to be *unreliable*:

Definition 12. Where $\mathsf{Dab}(\Delta_1), \mathsf{Dab}(\Delta_2), \ldots$ are the minimal Dab-formulas derived at stage s of a proof from Γ, the abnormalities in $U_s(\Gamma) = \Delta_1 \cup \Delta_2 \cup \ldots$ are called *unreliable at stage s*.

With these definitions in place, we can state the marking definition for reliability:

Definition 13 (Marking for Reliability)**.** A line l with condition Δ is marked at stage s of a proof from Γ iff, $\Delta \cap U_s(\Gamma) \neq \emptyset$.

Note the difference from $\mathbf{SDL^n}$. There, for some line l to get marked, all the members of its conditions had to be derived in the form of a disjunction, showing, explicitly, that at least one of them has to hold. Here, on the other hand, we mark l as soon as we get to know that one of them is involved in a minimal disjunction of abnormalities. In fact, this is just what happens in the proof above: when we discover at stage 7 that $\lightning(d \wedge \neg h)$ is unreliable, we conclude that lines 5 and 6 are based on unreliable assumptions and mark them.

In $\mathbf{SDL^n}$-proofs marking could only come, but in $\mathbf{SDL^r}$-proofs (and in $\mathbf{SDL^m}$-proofs, see Section 6) we also have cases of 'unmarking'. This is due to the fact that previously minimal Dab-formulas cease to be minimal when some of their sub-disjunctions are derived (they are 'shortened'). This double dynamics is taken into account in the following definition —in fact, a generalization of Definition 8— which is used for both strategies: reliability and minimal abnormality.

Definition 14. A formula A is *finally derived* in a proof from Γ at a stage s iff, (i) there is a non-marked line l with formula A at stage s, (ii) if l is marked in any extension of the proof, the proof can be further extended in a way such that l is unmarked again.

Again, the adaptive consequences of Γ are the formulas that are finally derivable (but now with respect to Definition 14):

Definition 15. $\Gamma \vdash_{\mathbf{SDL^r}} A$ iff A is finally derivable in an $\mathbf{SDL^r}$-proof from Γ.

This consequence relation can be characterized by means of **SDL**, which will give us an interesting link to the U-consequences. First, we define a stage-independent notion of unreliable formulas:

Definition 16. Where $\mathsf{Dab}(\Delta_1), \mathsf{Dab}(\Delta_2), \ldots$ are the minimal Dab-consequences of Γ, the set of *unreliable formulas* is given by $U(\Gamma) =_{\mathrm{df}} \Delta_1 \cup \Delta_2 \cup \ldots$

The $\mathbf{SDL^r}$-consequences from Γ can be characterized by means of $U(\Gamma)$ and **SDL** as follows: where $\Theta^{\neg} =_{\mathrm{df}} \{\neg A \mid A \in \Theta\}$,[14]

Theorem 5.1. $\Gamma \vdash_{\mathbf{SDL^r}} A$ *iff* $\Gamma \cup (\Omega_\circ \setminus U(\Gamma))^{\neg} \vdash_{\mathbf{SDL}} A$.

5.2 Meta-Theoretic Results

In the remainder of this section we relate $\mathbf{SDL^r}$ to the U-consequences. The idea was that A is a U-consequence from Γ iff we can derive A from the innocent bystanders of Γ, $\mathsf{ib}(\Gamma)$. Translated in the context of adaptive proofs, this means that we expect a U-consequence to be derivable on an assumption that is compatible with $\mathsf{ib}(\Gamma)$: i.e., on a condition $\lightning\Delta$ such that $\Delta \subseteq \mathsf{ib}(\Gamma)$. And, indeed, we can immediately rephrase Theorem 5.1 as Theorem 5.2 below. First, however, note that by Theorem 4.3:

Corollary 5.1. *Where* $\Gamma \subseteq \mathcal{W}_{\mathrm{pro}}$*:* $U(\Gamma^\circ) = \bigcup_{\Delta \in \mathsf{MIS}(\Gamma)} \lightning\Delta$.

Theorem 5.2. *Where* $\Gamma \subseteq \mathcal{W}_{\mathrm{pro}}$*:* $\Gamma^\circ \vdash_{\mathbf{SDL^r}} \mathsf{O}A$ *iff there is a* $\Delta \subseteq \mathsf{ib}(\Gamma)$ *such that* $\Gamma^\circ \vdash_{\mathbf{SDL}} \mathsf{O}A \vee \mathsf{Dab}(\lightning\Delta)$.

Proof. $\Gamma^\circ \vdash_{\mathbf{SDL^r}} \mathsf{O}A$, iff [by Theorem 5.1], $\Gamma^\circ \cup (\Omega \setminus U(\Gamma^\circ))^{\neg} \vdash_{\mathbf{SDL}} \mathsf{O}A$, iff [by classical properties (compactness, deduction and resolution theorem)], there is a finite $\lightning\Delta \subseteq \Omega \setminus U(\Gamma^\circ)$ such that $\Gamma^\circ \vdash_{\mathbf{SDL}} \mathsf{O}A \vee \mathsf{Dab}(\lightning\Delta)$, iff [by Corollary

[14]This is an instance of a meta-theorem of the standard format: see [4, Thm. 6].

5.1 and Lemma 4.1], there is a finite $\Delta \subseteq \Gamma \setminus \bigcup_{\Theta \in \mathsf{MIS}(\Gamma)} \Theta$ such that $\Gamma^\circ \vdash_{\mathbf{SDL}} \mathsf{O}A \vee \mathsf{Dab}(\lightning\Delta)$, iff [by Fact 2.3], there is a finite $\Delta \subseteq \mathsf{ib}(\Gamma)$ such that $\Gamma^\circ \vdash_{\mathbf{SDL}} \mathsf{O}A \vee \mathsf{Dab}(\lightning\Delta)$. □

With Theorem 4.2 we gain a representation theorem:

Theorem 5.3. *Where* $\Gamma \subseteq \mathcal{W}_{\text{pro}}$*:* $\Gamma^\circ \vdash_{\mathbf{SDL^r}} \mathsf{O}A$ *iff* $\Gamma \mid\!\sim_U A$.

Proof. $\Gamma^\circ \vdash_{\mathbf{SDL^r}} \mathsf{O}A$, iff [by Theorem 5.2], there is a $\Delta \subseteq \mathsf{ib}(\Gamma)$ such that $\Gamma^\circ \vdash_{\mathbf{SDL}} \mathsf{O}A \vee \mathsf{Dab}(\lightning\Delta)$, iff [by Theorem 4.2], there is a $\Delta \subseteq \mathsf{ib}(\Gamma)$ such that $\Delta \vdash A$, iff [by the compactness of **CL**], $\mathsf{ib}(\Gamma) \vdash A$, iff [by the definition of $\mid\!\sim_U$], $\Gamma \mid\!\sim_U A$. □

6 Minimal Abnormality and D-consequences

In this section we present a proof theory for the D-consequence relation. $\mathbf{SDL^m}$ is characterized by the triple $\langle \mathbf{SDL}, \Omega_\circ, \mathbf{m} \rangle$.

6.1 The adaptive proof theory for D-consequences

The marking definition for $\mathbf{SDL^m}$ is more engaged in comparison to the previous two strategies. It relies on the idea of *minimal choice sets*. Here we explain how it works, while in Section 6.2 we show how it captures the idea behind the D-consequence relation.

Suppose that we have an $\mathbf{SDL^m}$-proof from some set Γ, and that the following minimal Dab-formulas occur in this proof at stage s: $\mathsf{Dab}(\Delta_1), \ldots, \mathsf{Dab}(\Delta_n)$. Consider the set of all choice sets of $\{\Delta_1, \ldots, \Delta_n\}$, that is to say, all the sets that contain a member of every Δ_i. The minimal choice sets are those that are not proper subsets of other choice sets. We let $\Phi_s(\Gamma)$ stand for the set of minimal choice sets of $\{\Delta_1, \ldots, \Delta_n\}$. We can now state the definition for marking:

Definition 17 (Marking for Minimal Abnormality)**.** A line on which A is derived on the condition Δ is marked at stage s iff

(i) there is no $\phi \in \Phi_s(\Gamma)$ such that $\phi \cap \Delta = \emptyset$, or,

(ii) for some $\phi \in \Phi_s(\Gamma)$, there is no line at stage s on which A is derived on a condition Θ for which $\phi \cap \Theta = \emptyset$.

This becomes more transparent when 'reversed', i.e., if we consider when lines are *not marked*: Where A is derived on the condition Δ on line l, this line l is not marked at a stage s iff (a) there is a minimal choice set $\phi \in \Phi_s(\Gamma)$ for which $\phi \cap \Delta = \emptyset$, and (b) for every other choice set $\phi' \in \Phi_s(\Gamma)$ there is some line l' (possibly l itself) at which A is derived on a condition Θ such that $\phi' \cap \Theta = \emptyset$.

Let us have a look at a concrete $\mathbf{SDL^m}$-proof to see this version of marking at work. Again, we reason about orders issued to Anne, albeit her grandmother's order is left out of consideration.

1	$\circ(d \wedge \neg h)$	PREM	$\emptyset$
2	$\circ(\neg d \wedge h)$	PREM	$\emptyset$
✓3	$\mathsf{O}(d \wedge \neg h)$	1; RC	$\{\lightning(d \wedge \neg h)\}$
✓4	$\mathsf{O}d$	3; RU	$\{\lightning(d \wedge \neg h)\}$
✓5	$\mathsf{O}(d \vee h)$	4; RU	$\{\lightning(d \wedge \neg h)\}$
6	$\lightning(d \wedge \neg h) \vee \lightning(\neg d \wedge h)$	1,2; RU	$\emptyset$

As in the previous cases, we have done some defeasible reasoning on the basis of the premises —lines 3, 4, and 5— and afterwards we discovered that something may not be right with the previous steps since we derived the Dab-formula at line 6. Note that at the given stage $\lightning(d \wedge \neg h) \vee \lightning(\neg d \wedge h)$ is the only minimal Dab-formula occurring in this proof. Thus, the set of minimal choice sets $\Phi_6(\Gamma_1^\circ)$ of $\{\{\lightning(d \wedge \neg h), \lightning(\neg d \wedge h)\}\}$ contains exactly two singleton sets: $\phi_1 = \{\lightning(d \wedge \neg h)\}$ and $\phi_2 = \{\lightning(\neg d \wedge h)\}$. Let us now see why lines 3–5 are marked. If they were to be unmarked, two conditions would have to hold. First, for each line there would have to be some minimal choice set which does not share a single element with the respective conditions of these lines. It is immediate to see that this holds, for ϕ_2 is such a set. Second, for each other minimal choice set that intersects with the respective conditions — such as ϕ_1 — the respective formulas have to be derived without relying on any of its elements. This is evidently not the case for ϕ_1. Now let us continue our proof a bit further and see how this changes.

✓3	$\mathsf{O}(d \wedge \neg h)$	1; RC	$\{\lightning(d \wedge \neg h)\}$
✓4	$\mathsf{O}d$	3; RU	$\{\lightning(d \wedge \neg h)\}$
5	$\mathsf{O}(d \vee h)$	4; RU	$\{\lightning(d \wedge \neg h)\}$
6	$\lightning(d \wedge \neg h) \vee \lightning(\neg d \wedge h)$	1,2; RU	$\emptyset$
✓7	$\mathsf{O}(\neg d \wedge h)$	2; RC	$\{\lightning(\neg d \wedge h)\}$
✓8	$\mathsf{O}h$	7; RU	$\{\lightning(\neg d \wedge h)\}$
9	$\mathsf{O}(d \vee h)$	8; RU	$\{\lightning(\neg d \wedge h)\}$

In view of the minimal Dab-formula at line 6, as soon as lines 7 and 8 are added, they get marked. Note that there is nothing wrong with adding new lines to the proof on the basis of others that are already marked. Now, at stage 9 something noteworthy happens: we have again derived $\mathsf{O}(d \vee h)$, but this time on a different condition. This makes a difference for marking, for now we have both: (1) the minimal choice set ϕ_1 which does not share any elements with the condition of line 9 — point (a) above —, and (2) another line at which $\mathsf{O}(d \vee h)$ is derived on a different condition at line 5 that does not share any elements with the second minimal choice set ϕ_2. Given that here $\Phi_9(\Gamma_1^\circ)$ contains only the two elements ϕ_1 and ϕ_2, (2) is enough to satisfy point (b). What is more, now the situation has also changed for line 5 which, therefore, gets unmarked. The reader is welcome to verify that both line 5 and line 9 will never get marked in any extension of the proof, and are, thus, finally derived (Definition 14).

It is worth pausing here to say a few words about the intuition behind the minimal abnormality strategy. In short, it is already contained in its name.

When a minimal Dab-formula is derived in a proof, it is clear that at least one of its disjuncts $\delta_1, \ldots, \delta_n$ must be true. Where reliability strategy considers all of these disjuncts unreliable —allowing that all of them may be true—, minimal abnormality assumes that only a minimal number of them are true. These minimal sets of abnormalities that are true are determined by the minimal choice sets in view of all the minimal Dab-formulas derived at some stage of the proof.

It still remains to define derivability for $\mathbf{SDL^m}$. Recall that we already specified when a formula is counted as finally derived in Definition 14.

Definition 18. $\Gamma \vdash_{\mathbf{SDL^m}} A$ iff A is finally derivable in an $\mathbf{SDL^m}$-proof from Γ.

The following representation theorem establishes a link between the marking definition and final derivability.

Definition 19. Where $\Sigma(\Gamma)$ is the set of all $\Delta \subseteq \Omega$ such that $\mathsf{Dab}(\Delta)$ is a minimal Dab-consequence from Γ, $\Phi(\Gamma)$ is the set of minimal choice sets over $\Sigma(\Gamma)$.

Theorem 6.1. $\Gamma \vdash_{\mathbf{SDL^m}} A$ *iff for each* $\phi \in \Phi(\Gamma)$ *there is a* $\Delta \subseteq \Omega_\circ \setminus \phi$ *such that* $\Gamma \vdash_{\mathbf{SDL}} A \vee \mathsf{Dab}(\Delta)$.[15]

6.2 Meta-Theoretic Results

The minimal choice sets in $\Phi(\Gamma^\circ)$ correspond exactly to the complements of the maximal consistent subsets of Γ. In order to see that we first show:

Lemma 6.1. *Where* $\Gamma \subseteq \mathcal{W}_{\text{pro}}$*:* ϕ *is a minimal choice set of* $\mathsf{MIS}(\Gamma)$ *iff* $\Gamma \setminus \phi \in \mathsf{MCS}(\Gamma)$.

Proof. ($\Rightarrow$) Suppose $\Gamma \setminus \phi \notin \mathsf{MCS}(\Gamma)$. If $\Gamma \setminus \phi$ is inconsistent then evidently ϕ is not a (minimal) choice set of $\mathsf{MIS}(\Gamma)$ since $\phi \cap (\Gamma \setminus \phi) = \emptyset$. If $\Gamma \setminus \phi$ is consistent then there is a non-empty $\phi' \subseteq \phi$ such that $(\Gamma \setminus \phi) \cup \phi' \in \mathsf{MCS}(\Gamma)$. By Lemma 4.2, for all $\Delta \in \mathsf{MIS}(\Gamma)$, $\Delta \setminus ((\Gamma \setminus \phi) \cup \phi') = (\Delta \cap \phi) \setminus \phi' \neq \emptyset$. Hence, $\phi \setminus \phi'$ is a choice set of $\mathsf{MIS}(\Gamma)$ and since $\phi' \neq \emptyset$, ϕ is not a minimal choice set of $\mathsf{MIS}(\Gamma)$.

($\Leftarrow$) Suppose $\Gamma \setminus \phi \in \mathsf{MCS}(\Gamma)$. By Lemma 4.2, for all $\Delta \in \mathsf{MIS}(\Gamma)$, $\Delta \setminus (\Gamma \setminus \phi) = \Delta \cap \phi \neq \emptyset$. Hence, ϕ is a choice set of $\mathsf{MIS}(\Gamma)$. Where $\emptyset \subset \phi' \subseteq \phi$, $(\Gamma \setminus \phi) \cup \phi'$ is inconsistent by our supposition. Since $(\phi \setminus \phi') \cap ((\Gamma \setminus \phi) \cup \phi') = \emptyset$, $\phi \setminus \phi'$ is not a choice set of $\mathsf{MIS}(\Gamma)$. □

Now we can show the announced result (see also Table 1):

Theorem 6.2. *Where* $\Gamma \subseteq \mathcal{W}_{\text{pro}}$*:* $\lightning\phi \in \Phi(\Gamma^\circ)$ *iff* $\Gamma \setminus \phi \in \mathsf{MCS}(\Gamma)$.

Proof. By Theorem 4.3, $\Phi(\Gamma^\circ)$ is the set of minimal choice sets over $\{\lightning\Delta \mid \Delta \in \mathsf{MIS}(\Gamma)\}$. The rest follows immediately by Lemma 6.1. □

In view of Theorem 6.1, Theorem 6.2 and Lemma 4.1 we get:

[15] This is an instance of a meta-theorem of the standard format: see [4, Thm. 8].

adaptive logic	*MCS-approach*
minimal Dab-consequence of Γ°: $\mathsf{Dab}(\lightning\Delta)$	minimal inconsistent sets: $\Delta \in \mathsf{MIS}(\Gamma)$
reliable abnormalities: $\lightning A \in \lightning\Gamma \setminus U(\Gamma^\circ)$	innocent bystanders: $A \in \mathsf{ib}(\Gamma)$
maximal choice sets: $\lightning\phi \in \Phi(\Gamma^\circ)$	maximal consistent sets: $\Gamma \setminus \phi \in \mathsf{MCS}(\Gamma)$

Table 1: Correspondences

Corollary 6.1. *Where $\Gamma \subseteq \mathcal{W}_{\text{pro}}$: $\Gamma^\circ \vdash_{\mathbf{SDL^m}} \mathsf{O}A$ iff for each $\phi \in \mathsf{MCS}(\Gamma)$ there is a $\Delta \subseteq \phi$ such that $\Gamma^\circ \vdash_{\mathbf{SDL}} \mathsf{O}A \vee \mathsf{Dab}(\lightning\Delta)$.*

Together with Theorem 4.2 this implies the following representation theorem:

Theorem 6.3. *Where $\Gamma \subseteq \mathcal{W}_{\text{pro}}$: $\Gamma^\circ \vdash_{\mathbf{SDL^m}} \mathsf{O}A$ iff $\Gamma \mathrel{|\!\!\sim}_D A$.*

Proof. $\Gamma^\circ \vdash_{\mathbf{SDL^m}} \mathsf{O}A$, iff [by Corollary 6.1], for each $\phi \in \mathsf{MCS}(\Gamma)$ there is a finite $\Delta \subseteq \phi$ such that $\Gamma^\circ \vdash_{\mathbf{SDL}} \mathsf{O}A \vee \mathsf{Dab}(\lightning\Delta)$, iff [by Theorem 4.2], for each $\phi \in \mathsf{MCS}(\Gamma)$ there is a finite $\Delta \subseteq \phi$ such that $\Delta \vdash A$, iff [by the compactness of **CL**], for each $\phi \in \mathsf{MCS}(\Gamma)$, $\phi \vdash A$, iff [by the definition of $\mathrel{|\!\!\sim}_D$], $\Gamma \mathrel{|\!\!\sim}_D A$. □

7 Semantics

In this section we define semantics corresponding to the three systems developed above: $\mathbf{SDL^n}$, $\mathbf{SDL^r}$, and $\mathbf{SDL^m}$. Since in this paper we focus on proof theory, our treatment of the semantics is brief.

In general, semantics for adaptive logics in the standard format are obtained by selecting a subset of those **LLL**-models (in our case — **SDL**-models) that verify some given Γ. On the intuitive level, we restrict our attention to exactly those models of the lower limit logic that are the 'most normal' in view of the given adaptive strategy. It is very important to note that this selection is made on the basis of some set of premises Γ. Thus, it is not accurate to say that some **LLL**-model is an adaptive model *per se*. Rather we should say that it is an adaptive model of the given premise set.

The selection of the 'most normal' models is carried out by an appeal to the notion of the *abnormal part* of a model. It is defined as follows. Let M be an arbitrary model of **SDL**. Its abnormal part — $Ab(M)$ — is the set of all and only those abnormalities that it verifies, i.e., $Ab(M) = \{A \in \Omega_\circ \mid M \models A\}$. The semantics of *minimal abnormality* is the simplest of the three, and, hence, is a good point to start with.

Definition 20. An **SDL**-model M of Γ is *minimally abnormal* iff there is no **SDL**-model M' of Γ such that $Ab(M') \subset Ab(M)$.

Definition 21. $\Gamma \Vdash_{\mathbf{SDL^m}} A$ iff A is verified by all minimally abnormal models of Γ.

Surprising as it might be, these definitions do not appeal to the minimal choice sets that were so important for the definition of marking. Nevertheless, the latter are tightly connected to abnormal parts of the models picked out by Definition 20. More concretely, in [1] it was proven that for every $\phi \in \Phi(\Gamma)$ there is some minimally abnormal model M of Γ such that $Ab(M) = \phi$, and *vice versa*, whenever M is a minimally abnormal model of Γ, $\mathrm{Ab}(M) \in \Phi(\Gamma)$. There are, usually, more minimally abnormal models than choice sets, and a natural question arises: which formulas are made true by all minimally abnormal **LLL**-models of Γ that share the same abnormal part? As it turns out, after we formalize the intuition behind this question, we arrive at the semantics for *normal selections*.

Definition 22. A set $\mathcal{M}$ of **SDL**-models of Γ is a *normal selection of* Γ iff for some $\phi \in \Phi(\Gamma)$, $\mathcal{M} = \{M \text{ is a } \mathbf{SDL}\text{-model of } \Gamma \mid Ab(M) = \phi\}$.

In view of the fact that the abnormal parts of minimally abnormal models of Γ constitute exactly $\Phi(\Gamma)$, we can equivalently define: $\mathcal{M}$ is a normal selection of Γ iff there is a minimally model M of Γ such that $\mathcal{M} = \{M' \text{ is an } \mathbf{SDL}\text{-model of } \Gamma \mid \mathrm{Ab}(M') = \mathrm{Ab}(M)\}$.

Definition 23. $\Gamma \Vdash_{\mathbf{SDL^n}} A$ iff A is verified by every model of some normal selection $\mathcal{M}$ of Γ.

Now it only remains to introduce the semantics of $\mathbf{SDL^r}$. Recall that the marking for *reliability* was carried out by a reference to the set of *unreliable* formulas — $U(\Gamma)$. In the corresponding semantics we also appeal to it.

Definition 24. An **SDL**-model M of Γ is *reliable* iff $Ab(M) \subseteq U(\Gamma)$.

Definition 25. $\Gamma \Vdash_{\mathbf{SDL^r}} A$ iff A is verified by all reliable models of Γ.

This concludes the exposition of the semantics. In view of the previous discussion and the fact that all three of our logics are in the standard format we immediately get the following result, which also serves as a summary of Sections 4–7.

Theorem 7.1. *Where* $\Gamma \subseteq \mathcal{W}_{\mathrm{pro}}$,

(i) $\Gamma^\circ \mid\!\sim_C \mathsf{O}A$ *iff* $\Gamma^\circ \vdash_{\mathbf{SDL^n}} \mathsf{O}A$ *iff* $\Gamma^\circ \Vdash_{\mathbf{SDL^n}} \mathsf{O}A$,

(ii) $\Gamma^\circ \mid\!\sim_U \mathsf{O}A$ *iff* $\Gamma^\circ \vdash_{\mathbf{SDL^r}} \mathsf{O}A$ *iff* $\Gamma^\circ \Vdash_{\mathbf{SDL^r}} \mathsf{O}A$,

(iii) $\Gamma^\circ \mid\!\sim_D \mathsf{O}A$ *iff,* $\Gamma^\circ \vdash_{\mathbf{SDL^m}} \mathsf{O}A$ *iff* $\Gamma^\circ \Vdash_{\mathbf{SDL^m}} \mathsf{O}A$.

8 Discussion

In this section we make two remarks on premise sets (Section 8.1) and compare ours systems to other approaches from the literature (Section 8.2). We then discuss some directions for future research (Section 8.3). In particular, we indicate how other consequence relations based on the idea of contextualizing **SDL** can be characterized by adaptive logics and consider various ways the systems defined above can be extended.

8.1 Two Remarks Concerning Premises

So far we have mainly focused on the setup for which van Fraassen's original account and for that matter also Horty's systems are defined, namely, the situation in which our premise set is simply a set of norms Γ°. It is important to notice that our adaptive logics can deal with more complex premise sets as well.

We can, for instance, deal with disjunctions of obligations, such as $\circ A \vee \circ B$. This may be relevant in situations in which our information about the given norms is incomplete: e.g., Anne may know that either her mother asked her to wash the dishes, or she asked her to bring out the garbage ($\circ d \vee \circ g$). Later Anne may learn that she was not asked to do the dishes $\neg \circ d$, from which Anne can infer that her mother asked her to bring out the garbage $\circ g$. Similarly, we can enforce constraints as part of our premises: e.g., we can add $\neg \mathsf{O}(p \wedge q)$ to the premise set in order to make sure that $\mathsf{O}(p \wedge q)$ will not be derived in the reasoning process. In this way we can represent physical constraints. Similarly, in some modeling situations we might be interested in adding $\mathsf{O}p$ to the premise set. Recall our reading of $\mathsf{O}p$: the obligation to p is a result of our deliberation process. So far, we have only considered reasons for accepting $\mathsf{O}p$ that are internal, i.e., provided by the underlying logic (e.g., in $\mathbf{SDL^m}$ due to the fact that p follows from all MCSs of Γ). Of course, the reasoning agent may also have external reasons for accepting $\mathsf{O}p$. She may, for instance, face a deontic system which contains both $\circ p$ and $\circ \neg p$, but, for some (external) reasons, may have a preference for $\circ p$. The easiest, albeit somewhat coarse, way of modeling this preference and making sure that $\mathsf{O}p$ is accepted as a consequence is by adding it to the premise set. In the end, in (dogmatic) political discourse it is not uncommon for all the reasoning process to be structured around a pre-established conclusion. We discuss a more sophisticated way of modeling preferences in Section 8.3.

Our second remarks concerns the motivations behind the distinction between $\circ$ and O. We have already pointed out in Section 2 that the two operators have different meanings. Still some readers may feel that an application of Occam's Razor is in order, unless there are some technical reasons that necessitate the usage of $\circ$. The technical reason here is the reflexivity of the consequence relation of adaptive logics in the standard format: $\Gamma \vdash_{\mathbf{AL}} A$ for all $A \in \Gamma$. Were we to use exclusively O and, hence, Γ^{O} instead of Γ°, we would get $\mathsf{O}A$ for all $A \in \Gamma$ as a consequence. This is obviously inadequate for $\mid\!\sim_D$ and $\mid\!\sim_U$: e.g., $\{\circ p, \circ \neg p\} \not\mid\!\sim_D \mathsf{O}p$. Also $\mid\!\sim_C$ does not support a reflexive treatment, since we have, for instance, $\{\circ(p \wedge \neg p)\} \not\mid\!\sim_C \mathsf{O}(p \wedge \neg p)$. Now, when facing a choice between either giving up reflexivity or adding a second normative operator, it seems more natural to go for the latter. For, first, a distinction between explicitly given norms and norms that are accepted as a consequence of our deliberation process is conceptually transparent, and, second, reflexivity seems to be warranted under our reading of $\circ$. It is obvious that an explicitly given norm does not cease to be explicitly given after we have scrutinized it (the same concerns the 'prima facie' reading). Also note that, even if such a $\circ A$ is not turned into $\mathsf{O}A$ at the end of our reasoning process, it can still have a normative force which may be witnessed by feelings of guilt or regret if A is not realized (cf. [30]).

8.2 Comparisons

In [11] Lou Goble distinguished three basic ways of devising conflict-tolerant deontic logics. All three are based on the idea of weakening **SDL**. According to the first approach, we invalidate or restrict aggregation [10], according to the second, we invalidate or restrict inheritance [10, 11], and, according to the third, we go paraconsistent and weaken our negation (at least in the scope of the deontic modalities, [10, 9, 5]). The basic thrust behind these approaches is ampliative. They equip us with basic principles of normative reasoning which are rather weak so that we have to supplement them with further, more contextual reasoning in order to get a more encompassing and realistic model. For instance, inheritance is only applicable to $\mathsf{O}A$ in Goble's **DPM**-systems if we have $\neg\mathsf{O}\neg A$, i.e., in case A is non-conflicted: e.g., $\mathsf{O}(A \wedge B) \nvdash_{\mathbf{DPM}} \mathsf{O}A$ while $\{\mathsf{O}(A \wedge B), \neg\mathsf{O}\neg(A \wedge B)\} \vdash_{\mathbf{DPM}} \mathsf{O}A$. These systems have been nonmonotonically strengthened by means of adaptive logics in such a way that the given normative code is interpreted as 'non-conflicting as possible' (see e.g., [13, 25, 19, 9, 5, 6]). Hence, in the adaptive strengthening we can derive $\mathsf{O}A$ from the premise set $\{\mathsf{O}(A \wedge B)\}$.

Some of these adaptive systems are quite powerful. For instance, many of the systems presented in [13] allow for the following inferences, neither of which is possible in the systems presented above (however, see Section 8.3):

$$\{\mathsf{O}(t \wedge r), \mathsf{O}(\neg t \wedge v)\} \vdash \mathsf{O}r \wedge \mathsf{O}v \quad (1)$$

$$\{\mathsf{O}(t \wedge r), \mathsf{O}(\neg t \wedge v)\} \vdash \mathsf{O}(r \wedge v) \quad (2)$$

$$\{\mathsf{O}(t \wedge (f \vee s)), \mathsf{O}(\neg t \wedge \neg f)\} \vdash \mathsf{O}s \quad (3)$$

Goble criticized the adaptive approach based on lower-limit logics that are weakenings of **SDL** since they treat normative conflicts as 'abnormal' while, according to Goble, "if one regards such conflicts as real, as commonplace, there is no reason to consider them abnormal, or as presumptively false" [13, p. 27]. This more philosophical problem has repercussions for the adaptive consequences. For instance, in none of the proposed adaptive systems based on weakenings of **SDL** do we get: $\{\mathsf{O}(f \vee s), \mathsf{O}\neg f, \mathsf{O}\neg s\} \vdash \mathsf{O}s$, although according to Goble this inference "should stand" (ibid.). Note that (a) in the systems presented in this paper deontic conflicts are not treated as abnormalities, rather an abnormality is an explicitly stated obligation which is not part of the consequence set, and (b) the inference $\{\circ(f \vee s), \circ\neg f, \circ\neg s\} \mathrel{|\!\sim}_C \mathsf{O}s$ holds in the C-account.

Horty has proposed a way to model the reasoning leading to the consequence relations $\mathrel{|\!\sim}_C$ and $\mathrel{|\!\sim}_D$ using Reiter's default logic. Since his system is developed for conditional obligations, let us restrict our attention to the non-conditional case — which is in line with the setting of this paper. Given some Γ° we consider the default theory $\langle \mathcal{F}, \mathcal{D} \rangle$ where the set of facts $\mathcal{F}$ is empty and the set of defaults is defined by $\mathcal{D}_\Gamma = \{(\top : A \,/\, A) \mid A \in \Gamma\}$. We reason on the basis of our default theory by building so-called *extensions*. Here is how it works. A default $(A : B \,/\, C)$ is interpreted as follows: if A is the case, we can conclude C unless the set of formulas derived so far is inconsistent with B. An extension of a default theory $\mathcal{D}$ is a set of formulas $\mathcal{E} \subseteq \mathcal{W}_{\text{pro}}$ that is a fixed point ($\mathcal{E} = Cn_{\mathbf{CL}}(\mathcal{E})$) and that satisfies: for all $(A : B \,/\, C) \in \mathcal{D}$, if $A \in \mathcal{E}$ and $\neg B \notin \mathcal{E}$ then $C \in \mathcal{E}$. In our simplified scenario the latter condition reads: for all $(\top : A \,/\, A) \in \mathcal{D}_\Gamma$, if $\neg A \notin \mathcal{E}$ then $A \in \mathcal{E}$. The credulous consequences are obtained by the union of all the obtained extensions, the skeptical consequences are obtained by their intersection.

Looking at our adaptive proofs we can see interesting connections. First note that the obtained extensions correspond to the classical closure of the maximal consistent subsets of Γ. In view of this, Fact 4.1, and Theorem 4.2, it is easy to see that some $\mathsf{O}A$ is derivable on a condition $\not\sharp\Delta$ where $\Delta \subseteq \Gamma$ and Δ is **CL**-consistent iff $A \in Cn_{\mathbf{CL}}(\mathcal{E})$ for some extension $\mathcal{E}$ of $\mathcal{D}_\Gamma$. Looking back at Section 4.2, the bridge is even stronger for $\mathbf{SDL^n}$: recall that a line with a condition $\not\sharp\Delta$ will get marked according to the normal selections strategy iff Δ is classically inconsistent. In view of Corollary 4.4 we have: $\mathsf{O}A$ is finally derivable from Γ° iff it is derivable on a condition $\Delta \subseteq \mathcal{E}$ for some extension $\mathcal{E}$ of $\mathcal{D}_\Gamma$ and $A \in Cn_{\mathbf{CL}}(\mathcal{E})$. In this sense one can interpret proofs in $\mathbf{SDL^n}$ in terms of 'building extensions'.

There are also some advantages in terms of transparency that comes with our dynamic proofs. The first concerns the compactness of the modeling. This works as follows. In the dynamic proof we can easily make a reasoning step on the basis of previous inferences that are based on different conditions. All we have to do is simply merge the conditions of the previous inferences. The marking mechanism takes care of the cases in which the newly obtained conditions are 'too strong'. Altogether the reasoning process is unified in one dynamic proof. In the default approach we build extensions one-by-one, and, thus, the reasoning process is modeled in a more fragmented manner.

Another interesting aspect is that we can use one and the same proof as a basis for all three consequence relations. All we need is to alter the adaptive strategy, i.e., the marking definition, in order to model the reasoning that is adequate with respect to the three different consequence relations.

We should also note that the idea to use a dummy operator (like $\circ$) was first utilized in the adaptive systems presented in [22] and [20]. There the idea was to apply $\circ A \supset A$ as much as possible: in [22] the focus was propositional logic, while in [20] it was generalized to the predicative case. In the latter paper it was shown that the standard adaptive strategies (reliability, minimal abnormality, and normal selections) exactly characterize predicative versions of the consequence relations of Rescher and Manor, [21].[16]

8.3 Some directions for future research

So far, our adaptive characterizations concerned only three consequence relations that are based on the idea of contextualizing **SDL**. We begin this section by briefly describing how other consequence relations could be characterized by adaptive logics and, hence, — equipped with dynamic proof theories.

Recently a general framework has been proposed for adaptive logics that are, for instance, able to handle quantitative considerations and priorities (see [22, Chapter 5]). The semantic characterization and the dynamic proof theory of any adaptive logic that fits this format is similar to, respectively, the semantics and the proof theory of minimal abnormality. Let us explain the basic idea in semantic terms (for the proof theory the reader is referred to [22]). Given a premise set Γ, we again select specific models out of the lower limit logic models of Γ that are deemed 'normal enough' in view of their abnormal part. The latter is determined by means of the so-called selection function Λ. For instance, for minimal abnormality Λ simply selects the models of Γ whose abnormal part is minimal. However, the framework also allows for selection functions that do not correspond to a strategy in the standard format. For instance, it can be used to devise adaptive logics characterizing quantitative variants of the consequences studied in this paper. Consider $\mid\!\sim^{\#}_{D}$ which was defined in Section 2. In this case the selection function Λ selects all **SDL**-models of Γ° for which the cardinality of their respective abnormal parts is minimal.

Furthermore, this framework also allows us to introduce priorities. We can use different $\circ_i$-s (where $i \geq 1$) to model the priority of the source of an obligation (with 1 being the highest priority). Let us give two examples for consequence relations based on the idea of contextualizing **SDL** that are informed by priorities. Where $\Gamma^{\star}$ is a set of obligations of the form $\circ_i A$ let $\Gamma_i = \{A \mid \circ_i A \in \Gamma^{\star}\}$ (where $i \geq 1$) and $\Gamma = \{A \mid \circ_i A \in \Gamma^{\star} \text{ where } i \geq 1\}$. Analogously to Definition 1, we need to specify a function $\mathfrak{C}(\Gamma)$ that generates contexts for a given $\Gamma^{\star}$. Given these contexts we can then define: $\Gamma^{\star} \mid\!\sim^{\cap}_{\mathfrak{C}} \mathsf{O}A$ iff for all $\Theta \in \mathfrak{C}(\Gamma)$, $\Theta \vdash A$; and $\Gamma^{\star} \mid\!\sim^{\cup}_{\mathfrak{C}} \mathsf{O}A$ iff for some $\Theta \in \mathfrak{C}(\Gamma)$, $\Theta \vdash A$. Here are two examples:

1. We let $\mathfrak{C}(\Gamma)$ be $\max_{\sqsubset^{\Gamma^{\star}}_{\mathsf{lex}}}(\mathsf{MCS}(\Gamma))$ where $\varphi \sqsubset^{\Gamma^{\star}}_{\mathsf{lex}} \psi$ iff there is an $i \in \mathbb{N}$ such that (i) for all $j < i$, $\varphi \cap \Gamma_j = \psi \cap \Gamma_j$ and (ii) $\varphi \cap \Gamma_i \subset \psi \cap \Gamma_i$.[17] Let e.g.,

[16] These were not the first adaptive systems that give characterizations of the consequence relations defined by Rescher and Manor: in [1] they were characterized by means of inconsistency-adaptive logics and in [2] strengthenings were proposed on the basis of the modal logic **S5**. It should be noted that the ALs presented in these papers are not in the standard format.

[17] In [27] the so-called lexicographic adaptive logics (see also [26]) were used to model this (lexicographic) preference order in a prioritization of the deontic logic from [19]. Lexicographic adaptive logics fall within the generic framework of [22, Chapter 5]. Outside of the context of

$\Gamma^\star = \{\circ_1 a, \circ_2 b, \circ_3 \neg b\}$. We have two MCSs, $\phi = \{a, b\}$ and $\psi = \{a, \neg b\}$. Since $\psi \sqsubset^{\Gamma^\star}_{\mathsf{lex}} \phi$ we get $\mathsf{O}a$ and $\mathsf{O}b$ in this account.

2. In another approach we could associate obligations in $\Gamma^\star$ that are not selected in a given context with a penalty that is relative to the priority of the obligation. E.g., we could use the penalty function $\mu_{\Gamma^\star}(\phi) = \sum_{A \in \Gamma \setminus \phi} \sum_{\circ_i A \in \Gamma_A} \frac{1}{i}$, where $\Gamma_A = \{\circ_i A \mid \circ_i A \in \Gamma^\star\}$, and define $\varphi \sqsubset_{\Gamma^\star} \psi$ iff $\mu_{\Gamma^\star}(\varphi) < \mu_{\Gamma^\star}(\psi)$.[18] Now, we let $\mathfrak{C}(\Gamma)$ be $\min_{\sqsubset_{\Gamma^\star}}(\mathsf{MCS}(\Gamma))$.

   For instance, $\Gamma^\star = \{\circ_2 a, \circ_2 b, \circ_3 \neg a\}$ has two MCSs: $\phi_1 = \{a, b\}$ and $\phi_2 = \{\neg a, b\}$ where $\phi_1 \sqsubset_{\Gamma^\star} \phi_2$ since $\mu_{\Gamma^\star}(\phi_1) = \frac{1}{3} < \mu_{\Gamma^\star}(\phi_2) = \frac{1}{2}$. This is as expected, since a is an obligation stemming from a prioritized source compared to the source of $\neg a$. However, if we add $\circ_4 \neg a$ to $\Gamma^\star$ resulting in $\Gamma'^\star$ the situation is different. Now, $\phi_2 \sqsubset_{\Gamma'^\star} \phi_1$, since $\mu_{\Gamma'^\star}(\phi_2) = \frac{1}{2} < \mu_{\Gamma'^\star}(\phi_1) = \frac{1}{3} + \frac{1}{4} = \frac{1}{2} + \frac{1}{12}$. Here, two obligations concerning $\neg a$ that stem from 'weaker' sources overpower the obligation a from the stronger source.

The corresponding selection functions for adaptive logics in the generalized format of [22, Chapter 5] are fairly straightforward (see [22, Chapter 5.8.3] for a selection function for (1) and [22, Chapter 5.8.4] for an example that is similar to (2)).

In this paper we presented a framework for non-conditional obligations. Extending it to the conditional setting is another direction for future research. Two approaches from the literature are especially interesting since they deal with conflicts on the basis of consistency considerations similar to the approach we presented for non-conditional obligations: (i) Input-Output logics (with constraints) [18] and (ii) Horty's approach in terms of defaults that we discussed above [14, 15, 16]. An adaptive characterization of the former is carried out in [23] while a characterization of the latter is a topic for future research.

In a follow-up paper that is in preparation we demonstrate how to overcome a well-known shortcoming of almost any approach based on consistency considerations, namely, the fact that conjunctions (conjunctive obligations, in our case) are treated as indivisible wholes. Recall the inferences in (1)–(3). Presently they do not go through in any of our systems because the premises are treated as indivisible wholes. Note that the only maximal consistent subsets of $\{\circ(t \wedge r), \circ(\neg t \wedge v)\}$ are $\{\circ(t \wedge r)\}$ and $\{\circ(\neg t \wedge v)\}$. We will show how to integrate an approach into our dynamic proof theories in which premises are analyzed or, intuitively speaking, broken into their constituent parts. In the case of $\{\circ(t \wedge r), \circ(\neg t \wedge v)\}$ it will give us the following maximally consistent contexts: $\{t, r, v\}$ and $\{\neg t, r, v\}$. Now, it should be apparent that this is just a step away from having, say, $\mathsf{O}r$ as a consequence on any of the strengthened version of our standard systems.

---

deontic logic this preference order has been used in Brewka's 'preferred subtheories' approach [8].

[18] We suppose here that Γ is finite.

In this follow-up research we will also extend the above consequence relations (as well as the corresponding adaptive logics) in a way that allows us to model permissive norms. Notice that with the introduction of permissions we also have to deal with a new kind of normative conflicts. Let us use another variation on our hackneyed example as an illustration. Suppose Anne's father requests her to wash the dishes, while her mother *allows* her not to do so (since, for instance, it is her birthday). This shows that normative conflicts can occur not only between obligations, but also between, for instance, an obligation and a permissive norm. It is true that here, unlike the other examples we have considered, Anne has a practical option in which neither of the norms is violated (i.e., doing the dishes). However, this situation is still a genuine conflict, for Anne is confronted with a practical choice between two conflicting options both of which can be normatively justified, and, in the absence of any further higher-order deontic principles, none of the two should be given priority.

References

[1] Diderik Batens, *Towards the unification of inconsistency handling mechanisms*, Logic and Logical Philosophy **8** (2000), 5–31.

[2] ______, *A strengthening of the Rescher–Manor consequence relations*, Logique at Analyse **183–184** (2003), 289–313, Appeared 2005.

[3] ______, *A universal logic approach to adaptive logics*, Logica Universalis **1** (2007), 221–242.

[4] ______, *A universal logic approach to adaptive logics*, Logica Universalis **1:1** (2007), 221–242.

[5] Mathieu Beirlaen and Christian Straßer, *Two adaptive logics of norm-propositions*, Journal of Applied Logic **11** (2013), no. 2, 147–168.

[6] Mathieu Beirlaen, Christian Straßer, and Joke Meheus, *An inconsistency-adaptive deontic logic for normative conflicts*, Journal of Philosophical Logic **2** (2013), no. 42, 285–315.

[7] Salem Benferhat, Didier Dubois, and Henri Prade, *Some syntactic approaches to the handling of inconsistent knowledge bases: A comparative study part 1: The flat case*, Studia Logica **58:1** (1997), 17–45.

[8] Gerhard Brewka, *Preferred subtheories: An extended logical framework for default reasoning.*, IJCAI, vol. 89, 1989, pp. 1043–1048.

[9] Newton C.A. da Costa and Walter Carnielli, *On paraconsistent deontic logic*, Philosophia **16** (1986), 293–305.

[10] Lou Goble, *Preference semantics for deontic logic. Part I: Simple models*, Logique at Analyse **183–184** (2003), 383–418.

[11] ———, *A logic for deontic dilemmas*, Journal of Applied Logic **3** (2005), 461–483.

[12] ———, *Normative conflicts and the logic of 'ought'*, Noûs **43** (2009), no. 3, 450–489.

[13] ———, *Deontic logic (adapted) for normative conflicts*, Logic Journal of IGPL (2013), online first.

[14] John F. Horty, *Moral dilemmas and nonmonotonic logic*, Journal of Philosophical Logic **23** (1994), 17–44.

[15] ———, *Nonmonotonic foundations for deontic logic*, Defeasible Deontic Logic (Donald Nute, ed.), Kluwer, Dordrecht, 1997, pp. 17–44.

[16] ———, *Reasoning with moral conflicts*, Noûs **37:4** (2003), 557–605.

[17] John. F. Horty, *Reasons as defaults*, Oxford University Press, 2012.

[18] David Makinson and Leendert Van Der Torre, *Constraints for Input/Output logics*, Journal of Philosophical Logic **30** (2001), no. 2, 155–185.

[19] Joke Meheus, Mathieu Beirlaen, and Frederik Van De Putte, *Avoiding deontic explosion by contextually restricting aggregation*, DEON (10th International Conference on Deontic Logic in Computer Science) (Guido Governatori and Giovanni Sartor, eds.), Lecture Notes in Artificial Intelligence, vol. 6181, Springer, 2010, pp. 148–165.

[20] Joke Meheus, Christian Straßer, and Peter Verdée, *Which style of reasoning to choose in the face of conflicting information?*, Journal of Logic and Computation (201x), Forthcoming.

[21] Nicholas Rescher and Ruth Manor, *On inference from inconsistent premisses*, Theory and Decision **1:2** (1970), 179–217.

[22] Christian Straßer, *Adaptive logic and defeasible reasoning. applications in argumentation, normative reasoning and default reasoning.*, Trends in Logic, vol. 38, Springer, 2013.

[23] Christian Straßer, Mathieu Beirlaen, and Frederik Van De Putte, *Dynamic proof theories for input/output logic*, Under review.

[24] Christian Strasser and Joke Meheus, *Simple representations of some adaptive consequence relations in a Kripke frame*, Forthcoming.

[25] Christian Straßer, Joke Meheus, and Mathieu Beirlaen, *Tolerating deontic conflicts by adaptively restricting inheritance*, Logique at Analyse (2012), no. 219, 477–506, Improved version in [22], Chap. 10.

[26] Frederik Van De Putte and Christian Straßer, *Extending the standard format of adaptive logics to the prioritized case*, Logique at Analyse **55** (2012), no. 220, 601–641.

[27] ———, *A logic for prioritized normative reasoning*, Journal of Logic and Computation **X** (2012), xx–xx, Published online first.

[28] Leendert van der Torre and Yao Hua Tan, *Two-phase deontic logic*, Logique at Analyse (2000), no. 171-172, 411–456.

[29] Bas C. van Fraassen, *Values and the heart's command*, The Journal of Philosophy **70:1** (1973), 5–19.

[30] Bernard Williams, *Symposium: Ethical consistency*, Proceedings of the Aristotelian Society, Supplementary Volumes **39** (1965), 103–138.