Normative Reasoning by Sequent-Based Argumentation

Christian Straßer [1] Ofer Arieli [2]

[1] Institute of Philosophy II, Ruhr University Bochum, Germany
[2] School of Computer Science, The Academic College of Tel-Aviv, Israel

Abstract

In this paper we present an argumentative approach to normative reasoning. Special attention is paid to deontic conflicts, contrary-to-duty and specificity cases, which are modeled by means of argumentative attacks. For this, we adopt a recently proposed framework for logical argumentation in which arguments are generated by a sequent calculus of a given base logic (see the papers of the second author in CLIMA'2013 and of the two authors in Volume 6 (No. 1) of Argument & Computation), and use standard deontic logic as our base logic. Argumentative attacks are realized by elimination rules that allow to discharge specific sequents. We demonstrate the usefulness of our approach by means of various well-known benchmark examples, and show that this approach is rich enough to capture a variety of paradigms for handling conflicting norms, such as reasoning with maximally consistent sets, prioritized norms, and deontic formalisms based on I/O logic.

1 Introduction

Normative reasoning is concerned with reasoning with and about notions such as obligations, permissions, etc. A paradigmatic instance is so-called *factual detachment*, which says that if φ holds and there is a commitment to ψ conditioned by φ, then there is a commitment to ψ. Another instance is *aggregation*: if there is an obligation to bring about φ and another obligation to bring about ψ, then there should be an obligation to bring about $\varphi \wedge \psi$. Allowing for unrestricted factual detachment or unrestricted aggregation is problematic in cases of normative conflicts [2, 25]. For instance, aggregating two conflicting obligations leads to an obligation that commits us to do the impossible. Other problematic cases are concerned with *specificity*: sometimes more specific obligations or permissions override more general ones. In such cases it is usually desired to block factual detachment from the overridden obligations or permissions.

Logical accounts of normative reasoning that are tolerant with respect to normative conflicts and/or specificity cases have been shown to be challenging, and have given rise to a variety of related approaches (for an overview see [25]).

In this paper we model normative reasoning by means of *logical argumentation*. Given a set of facts and a set of possibly conflicting and interdependent conditional obligations or permissions, we will show that this model helps us to identify conflict-free sets that are apt to guide the actions of a user. It will follow that the entailment relations that are obtained in our framework offer conflict-handling mechanisms for various types of conflicts, and as such they are adaptive to different application contexts.

Our starting point in modeling normative reasoning is concerned with Dung's well-known *abstract argumentation frameworks* [17]. These frameworks consist of a set of abstract objects (the 'arguments') and an attack relation between them. Their role is to serve as a tool to analyze and reason with arguments. Various procedures for selecting accepted arguments have been proposed, based on the dialectical relationships between the arguments. Usually, these methods avoid selecting arguments that conflict with each other and allow to respond to every possible attack on the argumentative stance with a counter-argument.

For formalizing normative reasoning we need to enhance abstract argumentation in order to model the structure of arguments. There are various ways of doing so, like those in [18] and [41]. In this paper, we settle for the representation in terms of *sequents* [3] (see also [5]). One advantage of this approach is that it immediately equips us with dynamic proof procedures in the style of adaptive logics [7, 47] that allow for automated reasoning [4]. Another advantage is that we can plug-in any Tarskian logic that comes with an adequate sequent calculus as a base logic that produces our arguments (as in [3, 5]).

In this paper (which is an extended version of [48]) we use SDL (standard deontic logic) as our base logic (see Section 2). In this context, the modality O is used to model obligations and permissions are modeled by P, defined by $\neg\mathsf{O}\neg$. Accordingly, arguments are sequents of the form $\Gamma \Rightarrow \phi$ (for some finite set of formulas Γ and a formula ϕ), which are derivable in a sequent calculus for SDL, extending Gentzen's calculus LK for classical logic [22]. Attacks between arguments are represented by *attack rules* that allow to derive elimination sequents of the form $\Gamma \not\Rightarrow \phi$, whose effect is the canceling or discharging of $\Gamma \Rightarrow \phi$ (see [3]).

The following example illustrates (still on the intuitive level) how the sequent-based argumentation framework described above models normative reasoning.

Example 1 Consider the following example by Horty [29]:

- When served a meal one should not eat with fingers.
- However, if the meal is asparagus one should eat with fingers.

These statements may be represented, respectively, by the formulas $m \supset \mathsf{O}\neg f$ and $(m \wedge a) \supset \mathsf{O}f$. Now, in case asparagus is indeed served $(m \wedge a)$ we expect to derive the (unconditional) obligation to eat with fingers ($\mathsf{O}f$) rather than to not eat with fingers ($\mathsf{O}\neg f$). This is a paradigmatic case of *specificity*: a more specific obligation cancels (or overrides) a less specific obligation. In our setting this will be handled by an attack rule advocating specificity (see Example 2 below), according to which the argument $\{m,\ a,\ (m \wedge a) \supset \mathsf{O}f\} \Rightarrow \mathsf{O}f$ attacks the argument $\{m, m \supset \mathsf{O}\neg f\} \Rightarrow \mathsf{O}\neg f$, and as a consequence $\mathsf{O}f$ will be inferable in this case while $\mathsf{O}\neg f$ will not.

The rest of this paper is divided to the following three parts:

- In the first part (Sections 2 and 3) we set up our framework. In Section 2 we define the base logic SDL and its sequent calculus $\mathcal{C}_{\mathsf{SDL}}$, which underly argumentation frameworks for normative reasoning. The latter are introduced in Section 3, where $\mathcal{C}_{\mathsf{SDL}}$ is augmented with a variety of attack rules for resolving conflicts among norms.

- In the second part of the paper (Sections 4–6) we illustrate reasoning in our framework by means of examples and relations to similar forms of reasoning. Section 4 is devoted to unconditional obligations. In this section we show, among others, that already by this kind of reasoning we are able to characterize within our framework the Rescher–Manor method for reasoning with maximally consistent subsets [43] (see also [52]), its strengthening by Benferhat, Dubois and Prade [9], and Brewka's approach to reasoning with preferred theories [11]. In Section 5 we allow also conditional obligations and demonstrate the usefulness of our approach by means of several well-known examples, such as Horty's conflict [29] discussed above, Forrester's Gentle Murderer scenario [20], Chisholm's paradox [16], and Caminada's deontic conflict [14]. In Section 6 we consider some enhancements of our framework which allow for more flexibility in the handling of relations between different forms of attacks.

- The last part of the paper (Sections 7 and 8) is devoted to further comparative studies and discussions. In Section 7 we relate our approach to formalisms that are based on Makinson and Van Der Torre's Input/Output logics [33]. In Section 8 we compare our approach to other deontic approaches (like Nute's defeasible logic [37, 38]) and conclude by providing an outlook.

2 The Base Logic SDL

The base logic that we shall use in this paper is SDL (standard deontic logic, i.e., the normal modal logic KD). The underlying language $\mathcal{L}_{\mathsf{SDL}}$ consists of a propositional constant $\bot$ (representing falsity), the standard operators for negation $\neg$, conjunction $\wedge$, disjunction $\vee$, and implication $\supset$, and the modal operator O representing obligations. Thus, for instance, the conditional obligation $\phi \supset \mathsf{O}\psi$ may be intuitively understood as "ϕ commits to bring about ψ". We denote the fragment of the language without modal operators by $\mathcal{L}$.

We shall denote formulas in $\mathcal{L}_{\mathsf{SDL}}$ by the lower Greek letter ψ, ϕ, and finite sets of formulas by the upper Greek letters Γ, Δ, Σ. Arbitrary (i.e., possibly infinite) sets of formulas in $\mathcal{L}_{\mathsf{SDL}}$ are denoted by the calligraphic uppercase letters $\mathcal{S}$ or $\mathcal{T}$. As usual, we incorporate the modality P for representing permissions, where $\mathsf{P}\psi$ is defined by $\neg\mathsf{O}\neg\psi$. Also, we shall abbreviate the formula $\bot \supset \bot$ by $\top$, abbreviate $(\phi \supset \psi) \wedge (\psi \supset \phi)$ by $\phi \equiv \psi$, write $\mathsf{O}\mathcal{S}$ for the set $\{\mathsf{O}\psi \mid \psi \in \mathcal{S}\}$, denote $\bigwedge\Gamma$ for the conjunction of the formulas in the (finite) set Γ, and write Γ_F for a (finite) set of modal-free (i.e., non-modal) formulas in the language $\mathcal{L}$.

Reasoning with SDL is done by $\mathcal{L}_{\mathsf{SDL}}$-*sequents* (or just sequents, for short), that is: expressions of the form $\Gamma \Rightarrow \psi$, where Γ is a finite set of $\mathcal{L}_{\mathsf{SDL}}$-formulas and $\Rightarrow$ is a symbol that does not appear in $\mathcal{L}_{\mathsf{SDL}}$. In what follows, we shall denote $\mathsf{Prem}(\Gamma \Rightarrow \psi) = \Gamma$.

Given a set $\mathcal{S}$ of formulas in $\mathcal{L}_{\mathsf{SDL}}$, we say that a formula ψ *follows* from $\mathcal{S}$ (in SDL), and denote this by $\mathcal{S} \vdash_{\mathsf{SDL}} \psi$, if there is a finite subset $\Gamma \subseteq \mathcal{S}$ such that the $\mathcal{L}_{\mathsf{SDL}}$-sequent $\Gamma \Rightarrow \psi$ is provable in the sequent calculus $\mathcal{C}_{\mathsf{SDL}}$ shown in Figure 1. Note that $\mathcal{C}_{\mathsf{SDL}}$ is the same as Gentzen's well-known sequent calculus LK for classical propositional logic, extended with rules for the modal operator O [50]. It is easy to verify that $\vdash_{\mathsf{SDL}}$ is a Tarskian consequence relation (that is, it is reflexive, monotonic and transitive).

Axioms: $\psi \Rightarrow \psi$

Structural Rules:

Weakening: $\dfrac{\Gamma \Rightarrow \Delta}{\Gamma, \Gamma' \Rightarrow \Delta, \Delta'}$ Cut: $\dfrac{\Gamma_1 \Rightarrow \Delta_1, \psi \quad \Gamma_2, \psi \Rightarrow \Delta_2}{\Gamma_1, \Gamma_2 \Rightarrow \Delta_1, \Delta_2}$

Logical Rules:

$[\wedge\Rightarrow]$ $\dfrac{\Gamma, \psi, \varphi \Rightarrow \Delta}{\Gamma, \psi \wedge \varphi \Rightarrow \Delta}$ $[\Rightarrow\wedge]$ $\dfrac{\Gamma \Rightarrow \Delta, \psi \quad \Gamma \Rightarrow \Delta, \varphi}{\Gamma \Rightarrow \Delta, \psi \wedge \varphi}$

$[\vee\Rightarrow]$ $\dfrac{\Gamma, \psi \Rightarrow \Delta \quad \Gamma, \varphi \Rightarrow \Delta}{\Gamma, \psi \vee \varphi \Rightarrow \Delta}$ $[\Rightarrow\vee]$ $\dfrac{\Gamma \Rightarrow \Delta, \psi, \varphi}{\Gamma \Rightarrow \Delta, \psi \vee \varphi}$

$[\supset\Rightarrow]$ $\dfrac{\Gamma \Rightarrow \psi, \Delta \quad \Gamma, \varphi \Rightarrow \Delta}{\Gamma, \psi \supset \varphi \Rightarrow \Delta}$ $[\Rightarrow\supset]$ $\dfrac{\Gamma, \psi \Rightarrow \varphi, \Delta}{\Gamma \Rightarrow \psi \supset \varphi, \Delta}$

$[\neg\Rightarrow]$ $\dfrac{\Gamma \Rightarrow \Delta, \psi}{\Gamma, \neg\psi \Rightarrow \Delta}$ $[\Rightarrow\neg]$ $\dfrac{\Gamma, \psi \Rightarrow \Delta}{\Gamma \Rightarrow \Delta, \neg\psi}$

$[\bot\Rightarrow]$ $\Gamma, \bot \Rightarrow \Delta$ $[\Rightarrow\neg\bot]$ $\Gamma \Rightarrow \Delta, \neg\bot$

KR: $\dfrac{\Gamma \Rightarrow \phi}{\mathsf{O}\Gamma \Rightarrow \mathsf{O}\phi}$ DR: $\dfrac{\Gamma \Rightarrow \phi}{\mathsf{O}\Gamma \Rightarrow \neg\mathsf{O}\neg\phi}$

Figure 1: The proof system $\mathcal{C}_{\mathsf{SDL}}$

Note 1 Instead of basing SDL on classical propositional logic and LK, our framework also allows for other variants, such as incorporating intuitionistic logic and Gentzen's LJ. This may be justified by the fact that, e.g., legal or medical systems often involve uncertainties, thus excluded middle is sometimes rejected in them, and proofs are required to be constructive. To keep things as simple as possible, we will proceed the discussion in terms of classical logic.

The logic SDL has the usual problems or 'paradoxes' that are associated with a material account of implication (such as $\phi \supset (\psi \supset \phi)$ which in terms of conditional obligations becomes $\mathsf{O}\phi \supset (\psi \supset \mathsf{O}\phi)$: given an obligation to bring about ϕ implies that under any condition ψ we have that very obligation). Furthermore, straightforward accounts of modeling conditional obligations with material implication is plagued with consistency problems for various types of conditional obligations (e.g., contrary-to-duty and specificity cases like those considered in the introduction). For instance, applying SDL to Forrester's example of the gentle murderer (see [20] and Example 8 below) results in triviality. The system developed in this paper may be seen as a sign of caution: although the 'paradoxes' of material implication are here to stay, we can give an account which solves the consistency problems and gives intuitive results for contrary-to-duty and specificity cases (as will be demonstrated below with various examples).

3 Argumentation for Normative Reasoning

In what has become the orthodox approach based on Dung's representation [17], formal argumentation is studied on the basis of so-called *argumentation frameworks*. In this section we consider argumentation frameworks for normative reasoning.

3.1 Normative Argumentation Frameworks

An argumentation framework in its most abstract form is a directed graph, where the nodes present (abstract) arguments and the arrows present argumentative attacks, as defined next.

Definition 1 *An* (abstract) argumentation framework [17] *is a pair* $\mathcal{AF} = \langle Args, Attack \rangle$*, where Args is an enumerable set of elements, called (abstract)* arguments*, and Attack is a binary relation between arguments whose instances are called* attacks*.*

When it comes to specific applications of formal argumentation it is often useful to provide an *instantiation* of (abstract) argumentation frameworks. Instantiations provide a specific account of the structure of arguments, and the concrete nature of argumentative attacks. There are various formal accounts available that provide frameworks for instantiating abstract argumentation, such as assumption-based argumentation [18], structured argumentation in the context of ASPIC [41], etc. Here we settle for a recently proposed account based on sequent-based calculi [3], which turns out to be a particularly useful framework for reasoning with argumentation systems (see also [5]).

The basic idea behind our instantiation is that, given a premise set $\mathcal{S}$ of formulas in $\mathcal{L}_{\mathsf{SDL}}$, arguments are sequents $\Gamma \Rightarrow \phi$ derivable in $\mathcal{C}_{\mathsf{SDL}}$ such that $\Gamma \subseteq \mathcal{S}$.

Definition 2 $\mathsf{Arg}(\mathcal{S})$ *is the set of sequents of the form* $\Gamma \Rightarrow \psi$ *for some finite* $\Gamma \subseteq \mathcal{S}$ *that are derivable in* $\mathcal{C}_{\mathsf{SDL}}$*.*

For specifying the attack relations we complement $\mathcal{C}_{\mathsf{SDL}}$ with *sequent elimination rules*. Unlike the inference (or, sequent introduction) rules of $\mathcal{C}_{\mathsf{SDL}}$, the conclusions of sequent elimination rules are of the form $\Gamma \not\Rightarrow \psi$, and their intuitive meaning is the discharging of the sequent $\Gamma \Rightarrow \psi$.

We use attacks to model normative conflicts as well as conflict resolution by rules such as specificity (e.g., 'lex specialis' in legal contexts). Normative conflicts occur in cases in which one can construct arguments for conflicting obligations (and permissions).

Example 2 Consider the following sequent elimination rule:

$$\mathsf{SpeReb}\ \frac{\Gamma,\ \phi \supset \psi \Rightarrow \psi \quad \Gamma \Rightarrow \phi \quad \Gamma' \Rightarrow \phi' \quad \phi \Rightarrow \phi' \quad \psi \Rightarrow \neg\psi' \quad \Gamma',\ \phi' \supset \psi' \Rightarrow \psi'}{\Gamma',\ \phi' \supset \psi' \not\Rightarrow \psi'}$$

The rule above (whose acronym stands for Specificity Rebuttal)[1] aims at formalizing the principle of specificity. It states that when two sequents $\Gamma,\ \phi \supset \psi \Rightarrow \psi$ and $\Gamma',\ \phi' \supset \psi' \Rightarrow \psi'$ are conflicting, the one which is more specific gets higher precedence, and so the other one is discharged. In more detail, the sequents making up the conditions of the elimination rule state the following: $\Gamma \Rightarrow \phi$ $[\Gamma' \Rightarrow \phi']$ expresses that the conditional $\phi \supset \psi$ $[\phi' \supset \psi']$ is 'triggered' in view of Γ $[\Gamma']$ (i.e., the antecedent ϕ $[\phi']$ follows from Γ $[\Gamma']$),[2] $\phi \Rightarrow \phi'$ expresses that ϕ is logically at least as strong as ϕ' (i.e., it is at least as *specific*), and $\psi \Rightarrow \neg\psi'$ expresses that the consequents of the two conditionals are conflicting.

In Example 1 for instance, SpeReb allows to discharge the sequent $m,\ m \supset \mathsf{O}\neg f \Rightarrow \mathsf{O}\neg f$ in light of the more specific sequent $m,\ a,\ (m \wedge a) \supset \mathsf{O}f \Rightarrow \mathsf{O}f$. In this case we say that the latter sequent *attacks* the former.

A variation of SpeReb that is in the form of an undercut specially devised for deontic cases is given below (where $\mathsf{NN'} \in \{\mathsf{OO}, \mathsf{OP}, \mathsf{PO}\}$):[3]

$$\mathsf{NN'SpeUcut}\ \frac{\begin{array}{c}\Gamma,\ \phi \supset \mathsf{N}\psi \\ \Rightarrow \neg(\phi' \supset \mathsf{N}'\psi')\end{array} \quad \Gamma \Rightarrow \phi \quad \phi \Rightarrow \phi' \quad \psi \Rightarrow \neg\psi' \quad \begin{array}{c}\Gamma',\ \phi' \supset \mathsf{N}'\psi' \\ \Rightarrow \psi''\end{array}}{\Gamma',\ \phi' \supset \mathsf{N}'\psi' \not\Rightarrow \psi''}$$

[1] In this paper the notion *rebuttal* is reserved for argumentative attacks that are based on incompatible conclusions of the involved arguments. In contrast, an *undercut* takes place if the conclusion of the attacking argument is incompatible with (parts of) the antecedent of the attacked argument. This is more in line with what Besnard and Hunter in [10] call undercuts and defeats than with Pollock's notion of undercut which is an attack on the application of a rule [40].

[2] The triggering condition $\Gamma \Rightarrow \phi$ is indeed essential for specificity attacks. Without it, the sequent $p, p \supset r, (p \wedge q) \supset r \Rightarrow r$ would attack the sequent $q, q \supset \neg r \Rightarrow \neg r$, but this would not reflect the rationale behind specificity attacks, because the consequent r of the attacking sequent is derived from the conditional $p \supset r$ and not from $(p \wedge q) \supset r$. However, only the latter is more specific than the conflicting conditional $q \supset \neg r$ in the attacked sequent.

[3] Note that a '$\mathsf{PPSpeUcut}$'-variant would not be sensible since permissions with incompatible content are not conflicting in any intuitive sense.

This rule is similar to SpeReb in the sense that it takes care of a conflict between two conditionals $\phi \supset \mathsf{N}\psi$ and $\phi' \supset \mathsf{N}'\psi'$, one of which ($\phi \supset \mathsf{N}\psi$) is at least as specific as the other ($\phi' \supset \mathsf{N}'\psi'$). In contrast to SpeReb we are now concerned with deontic conditionals and the argumentative attack is an undercut instead of a rebuttal (i.e., the consequent $\neg(\phi' \supset \mathsf{N}'\psi')$ of the attacking sequent negates one of the antecedents of the attacked sequent).

The above rules are useful for modeling different principles. For instance, the rule $\mathsf{POSpeUCut}$ (Permission–Obligation Specificity UnderCut) models permission as derogation [44]: a permission may suspend a more general obligation. These rules (as well as the ones introduced in what follows) also capture unconditional obligations and permissions such as $\mathsf{O}\psi$. For instance, when $\phi' = \top$ one may specify that $\phi, \phi \supset \mathsf{O}\psi \Rightarrow \neg\mathsf{O}\neg\psi$ $\mathsf{OOSpeUcut}$-attacks $\mathsf{O}\neg\psi \Rightarrow \psi''$.

Example 3 Consider the set $\mathcal{S} = \{a, b, a \supset \mathsf{O}c, b \supset \mathsf{O}\neg c\}$. One could imagine an argumentative context in which one proponent presents an argument for $\mathsf{O}c$ by proving $a, a \supset \mathsf{O}c \Rightarrow \mathsf{O}c$, while the opponent may rebut this argument for $\mathsf{O}\neg c$ by proving $b, b \supset \mathsf{O}\neg c \Rightarrow \mathsf{O}\neg c$. In the absence of an overriding principle such as specificity that resolves this conflict, one may need to introduce new sequent elimination rules like the following (where again $\mathsf{NN}' \in \{\mathsf{OO}, \mathsf{OP}, \mathsf{PO}\}$ and $\mathsf{N} \in \{\mathsf{O}, \mathsf{P}\}$):

$$\mathsf{NN'Reb} \quad \frac{\Gamma \Rightarrow \mathsf{N}\psi \quad \psi \Rightarrow \neg\psi' \quad \Gamma' \Rightarrow \mathsf{N}'\psi'}{\Gamma' \not\Rightarrow \mathsf{N}'\psi'}$$

$$\mathsf{NUcut} \quad \frac{\Gamma \Rightarrow \neg(\phi \supset \mathsf{N}\psi) \quad \Gamma', \phi \supset \mathsf{N}\psi \Rightarrow \psi'}{\Gamma', \phi \supset \mathsf{N}\psi \not\Rightarrow \psi'}$$

We will study these and some other elimination rules in the following sections.

The prerequisites of attack rules usually consist of three ingredients: the first sequent in the rule's prerequisites is the "attacking" sequent, the last sequent in the prerequisites is the "attacked" sequent, and the other prerequisites are the attack conditions. In this view, conclusions of sequent elimination rules are the eliminations of the attacked arguments.

Attacks between arguments in $\mathsf{Arg}(\mathcal{S})$ are triggered by applications of sequent elimination rules as defined next.

Definition 3 *Let R be a sequent elimination rule of the following form:*

$$\frac{\Gamma_1 \Rightarrow \phi_1 \ \ldots \ \Gamma_n \Rightarrow \phi_n}{\Gamma_n \not\Rightarrow \phi_n}$$

and let $\mathcal{R}$ be a set of such elimination rules.

- *A sequent $\mathfrak{s}$ R-*attacks *a sequent $\mathfrak{s}'$, if there is an $\mathcal{L}_{\mathsf{SDL}}$-substitution θ such that $\mathfrak{s} = \theta(\Gamma_1) \Rightarrow \theta(\phi_1)$, $\mathfrak{s}' = \theta(\Gamma_n) \Rightarrow \theta(\phi_n)$, and for every $1 < i < n$ $\theta(\Gamma_i) \Rightarrow \theta(\phi_i)$ is $\mathcal{C}_{\mathsf{SDL}}$-provable.*
- *We say that $\mathfrak{s}$ $\mathcal{R}$-*attacks *$\mathfrak{s}'$ if $\mathfrak{s}$ R-attacks $\mathfrak{s}'$ for some $R \in \mathcal{R}$.*

We are now ready to adjust Definition 1 to our setting.

Definition 4 *Let $\mathcal{S}$ be a set of $\mathcal{L}_{\mathsf{SDL}}$-formulas and let $\mathcal{R}$ be a set of elimination rules. A* normative argumentation framework *induced by $\mathcal{S}$ and $\mathcal{R}$ is the argumentation framework $\mathcal{AF}_{\mathcal{R}}(\mathcal{S}) = \langle \mathsf{Arg}(\mathcal{S}), Attack \rangle$, in which $(s, s') \in Attack$ iff s $\mathcal{R}$-attacks s'.*[4]

3.2 Normative Entailments

We are ready now to use (normative) argumentation frameworks for normative reasoning. As usual in the context of abstract argumentation, we do so by incorporating Dung's notion of extension [17], defined next.

Definition 5 *Let $\mathcal{AF} = \langle Args, Attack \rangle$ be an argumentation framework, and let $\mathcal{E} \subseteq Args$.*

- *We say that $\mathcal{E}$* attacks *an argument A if there is an argument $B \in \mathcal{E}$ that attacks A (i.e., $(B, A) \in Attack$). The set of arguments that are attacked by $\mathcal{E}$ is denoted $\mathcal{E}^+$.*
- *We say that $\mathcal{E}$* defends *A if $\mathcal{E}$ attacks every argument B that attacks A.*
- *The set $\mathcal{E}$ is called* conflict-free *if it does not attack any of its elements (i.e., $\mathcal{E}^+ \cap \mathcal{E} = \emptyset$), $\mathcal{E}$ is called* admissible *if it is conflict-free and defends all of its elements, and $\mathcal{E}$ is* complete *if it is admissible and contains all the arguments that it defends.*
- *The minimal complete subset of $Args$ is called the* grounded extension *of $\mathcal{AF}$,*[5] *and a maximal complete subset of $Args$ is called a* preferred extension *of $\mathcal{AF}$. A conflict-free set $\mathcal{E}$ is called a* stable extension *of $\mathcal{AF}$ if $\mathcal{E} \cup \mathcal{E}^+ = Args$.*[6]
- *We write $\mathsf{Adm}(\mathcal{AF})$ [respectively: $\mathsf{Cmp}(\mathcal{AF})$, $\mathsf{Prf}(\mathcal{AF})$, $\mathsf{Stb}(\mathcal{AF})$] for the set of all admissible [respectively: complete, preferred, stable] extensions of $\mathcal{AF}$ and $\mathsf{Grd}(\mathcal{AF})$ for the unique grounded extension of $\mathcal{AF}$.*

Next, we consider entailment relations that are induced by normative argumentation frameworks.

Definition 6 *Let $\mathcal{AF}_{\mathcal{R}}(\mathcal{S}) = \langle \mathsf{Arg}(\mathcal{S}), Attack \rangle$ be a normative argumentation framework.*

- *$\mathcal{S} \mathrel{{}_{\mathsf{gr}}{\mid\!\sim}_{\mathcal{R}}} \psi$ if there is an $A \in \mathsf{Grd}(\mathcal{AF}_{\mathcal{R}}(\mathcal{S}))$ such that $A = \Gamma \Rightarrow \psi$.*[7]

[4] It should be noted that the way normative argumentation frameworks are defined here is a bit different than the way they are defined in [48]. Here, arguments are sequents and so attacks are between sequents, while in [48] arguments are represented by their $\mathcal{C}_{\mathsf{SDL}}$-proofs. The present representation is simpler and is in line with the definition of sequent-based argumentation frameworks in [5].

[5] It is well-known that there is a unique grounded extension of $\mathcal{AF}$ [17, Theorem 25].

[6] It is well-known that a stable extension is always also preferred (see [17, Lemma 15]). The converse does not necessarily hold.

[7] Recall that by the definition of $\mathsf{Arg}(\mathcal{S})$, this implies that $\Gamma \subseteq \mathcal{S}$.

- $\mathcal{S}\ _{\mathsf{prf}}\!\mid\!\!\sim^{\cap}_{\mathcal{R}} \psi$ $[\mathcal{S}\ _{\mathsf{prf}}\!\mid\!\!\sim^{\cup}_{\mathcal{R}} \psi]$ *if for every [some]* $\mathcal{E} \in \mathsf{Prf}(\mathcal{AF}_{\mathcal{R}}(\mathcal{S}))$ *there is an* $A \in \mathcal{E}$ *such that* $A = \Gamma \Rightarrow \psi$.
- $\mathcal{S}\ _{\mathsf{prf}}\!\mid\!\!\sim^{\Cap}_{\mathcal{R}} \psi$ *if there is an* $A \in \bigcap \mathsf{Prf}(\mathcal{AF}_{\mathcal{R}}(\mathcal{S}))$ *where* $A = \Gamma \Rightarrow \psi$.
- $\mathcal{S}\ _{\mathsf{stb}}\!\mid\!\!\sim^{\cap}_{\mathcal{R}} \psi$ $[\mathcal{S}\ _{\mathsf{stb}}\!\mid\!\!\sim^{\cup}_{\mathcal{R}} \psi]$ *if for every [some]* $\mathcal{E} \in \mathsf{Stb}(\mathcal{AF}_{\mathcal{R}}(\mathcal{S}))$ *there is an* $A \in \mathcal{E}$ *such that* $A = \Gamma \Rightarrow \psi$.
- $\mathcal{S}\ _{\mathsf{stb}}\!\mid\!\!\sim^{\Cap}_{\mathcal{R}} \psi$ *if there is an* $A \in \bigcap \mathsf{Stb}(\mathcal{AF}_{\mathcal{R}}(\mathcal{S}))$ *where* $A = \Gamma \Rightarrow \psi$.

We will use the notation $\mid\!\sim$ *whenever a statement applies to each of the defined consequence relations.*

Note 2 Similar entailment relations may of-course be defined not only for complete, grounded, preferred and stable semantics, but also for other well-accepted semantics of abstract argumentation, such as the semi-stable semantics [13, 53], ideal semantics [19], etc.

4 Reasoning with Unconditional Obligations

We now exemplify reasoning with normative argumentation. Let us start with a simplified setting in which we have only non-conditional obligations.

Example 4 Suppose first that $\mathcal{S} = \{\mathsf{O}(a \wedge b), \mathsf{O}(\neg a \wedge b)\}$ and the attack rule is OUcut (defined in Example 3). We have for instance the following excerpt of the attack diagram:[8]

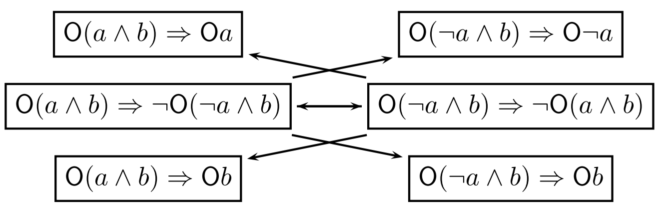

Figure 2: Example 4

In the diagram of Figure 2 one can identify two stable extensions: one extension consists of the formulas in the left column of the figure, and the other extension includes the formulas in the right column. Note that in both extensions there are arguments with $\mathsf{O}b$ on their right-hand sides. It follows that $\Sigma\ _{\mathsf{stb}}\!\mid\!\!\sim^{\cap}_{\mathsf{OUcut}}\ \mathsf{O}b$ while $\Sigma\ _{\mathsf{stb}}\!\not\mid\!\!\sim^{\Cap}_{\mathsf{OUcut}}\ \mathsf{O}b$.

The entailments obtained in the last example are not coincidental. As we show next, for premise sets $\mathcal{S}$ consisting of non-conditional obligations and the attack rule OUcut, some well-known consequence relations based on maximal consistent subsets are characterized within our framework.

[8] Recall the following 'unconditional' instance of OUcut where $\phi = \top$:

$$\frac{\Gamma \Rightarrow \neg\mathsf{O}\psi \quad \Gamma', \mathsf{O}\psi \Rightarrow \psi'}{\Gamma', \mathsf{O}\psi \not\Rightarrow \psi'}$$

Definition 7 *Let $\mathcal{S}$ be a set of $\mathcal{L}$-formulas.*[9] *We denote by* $\mathsf{MCS}(\mathcal{S})$ *the set of all the maximally consistent subsets of $\mathcal{S}$ (where maximality is taken with respect to the subset relation and consistency is with respect to classical logic).*

Proposition 1 *Let $\mathcal{S}$ be a set of formulas in $\mathcal{L}$. We denote by* $\mathrm{Cn}_{\mathsf{CPL}}(\mathcal{S})$ *the transitive closure of $\mathcal{S}$ with respect to classical logic. Then:*

1. $\mathsf{O}\mathcal{S} \;_{\mathsf{prf}}\!\!\mid\!\sim^{\Cap}_{\mathsf{OUcut}} \mathsf{O}\psi$ *iff* $\mathsf{O}\mathcal{S} \;_{\mathsf{stb}}\!\!\mid\!\sim^{\Cap}_{\mathsf{OUcut}} \mathsf{O}\psi$ *iff* $\psi \in \mathrm{Cn}_{\mathsf{CPL}}\,(\bigcap \mathsf{MCS}(\mathcal{S}))$.
2. $\mathsf{O}\mathcal{S} \;_{\mathsf{prf}}\!\!\mid\!\sim^{\cap}_{\mathsf{OUcut}} \mathsf{O}\psi$ *iff* $\mathsf{O}\mathcal{S} \;_{\mathsf{stb}}\!\!\mid\!\sim^{\cap}_{\mathsf{OUcut}} \mathsf{O}\psi$ *iff* $\psi \in \bigcap_{\mathcal{T} \in \mathsf{MCS}(\mathcal{S})} \mathrm{Cn}_{\mathsf{CPL}}(\mathcal{T})$
3. $\mathsf{O}\mathcal{S} \;_{\mathsf{prf}}\!\!\mid\!\sim^{\cup}_{\mathsf{OUcut}} \mathsf{O}\psi$ *iff* $\mathsf{O}\mathcal{S} \;_{\mathsf{stb}}\!\!\mid\!\sim^{\cup}_{\mathsf{OUcut}} \mathsf{O}\psi$ *iff* $\psi \in \bigcup_{\mathcal{T} \in \mathsf{MCS}(\mathcal{S})} \mathrm{Cn}_{\mathsf{CPL}}(\mathcal{T})$

Note 3 Inferences from the maximal consistent subsets of the premises, mentioned in Proposition 1, are known from [43] and [52]. Such inferences are also considered in many other contexts; See, e.g., [6, 9, 27, 32].

In order to establish the claim in Proposition 1 it is sufficient to prove the following lemma:

Lemma 1 *Let $\mathcal{S}$ be a (nonempty) set of $\mathcal{L}$-formulas and let $\mathcal{AF} = \mathcal{AF}_{\mathsf{OUcut}}(\mathsf{O}\mathcal{S})$. Then $\mathcal{E} \in \mathsf{Prf}(\mathcal{AF})$ iff $\mathcal{E} \in \mathsf{Stb}(\mathcal{AF})$ iff $\mathcal{E} = \mathsf{Arg}(\mathsf{O}\mathcal{T})$ for some $\mathcal{T} \in \mathsf{MCS}(\mathcal{S})$.*

Proof. Suppose that $\mathcal{T} \in \mathsf{MCS}(\mathcal{S})$ and let $\mathcal{E} = \mathsf{Arg}(\mathsf{O}\mathcal{T})$. Conflict-freeness follows by the consistency of $\mathcal{T}$ (which implies the consistency of $\mathsf{O}\mathcal{T}$): For every finite $\Delta \subseteq \mathcal{T}$ and $\phi \in \mathcal{T}$ we have that $\mathsf{O}\Delta \Rightarrow \neg\mathsf{O}\phi \notin \mathcal{E}$ since $\mathcal{T} \vdash \phi$ and so $\mathsf{O}\mathcal{T} \vdash \mathsf{O}\phi$ which implies $\mathsf{O}\mathcal{T} \nvdash \neg\mathsf{O}\phi$. For a finite $\Gamma \subseteq \mathcal{S}$, let $\mathsf{O}\Gamma \Rightarrow \phi \in \mathsf{Arg}(\mathsf{O}\mathcal{S}) \setminus \mathcal{E}$. Then $\Gamma \setminus \mathcal{T} \neq \emptyset$ and by the maximal consistency of $\mathcal{T}$, $\mathcal{T} \vdash \neg\psi$ for any $\psi \in \Gamma \setminus \mathcal{T}$. Thus, $\mathsf{O}\Delta \Rightarrow \mathsf{O}\neg\psi \in \mathcal{E}$ for some finite $\Delta \subseteq \mathcal{T}$ and in view of (DR) also $\mathsf{O}\Delta \Rightarrow \neg\mathsf{O}\psi \in \mathcal{E}$. Hence, $\mathcal{E}$ attacks all arguments in $\mathsf{Arg}(\mathsf{O}\mathcal{S}) \setminus \mathcal{E}$ which shows that $\mathcal{E}$ is stable and hence preferred.

For the converse, let $\mathcal{T}$ be a subset of $\mathcal{S}$ such that $\mathcal{T} \notin \mathsf{MCS}(\mathcal{S})$. If $\mathcal{T} \vdash \bot$, then by the compactness of classical propositional logic there is a finite $\Delta \subseteq \mathcal{T}$ for which $\Delta \vdash \bot$ and thus $\mathsf{O}\Delta \Rightarrow \bot \in \mathsf{Arg}(\mathsf{O}\mathcal{T})$. Hence, by [$\Rightarrow\supset$], $\mathsf{O}\Delta \setminus \{\mathsf{O}\phi\} \Rightarrow \neg\mathsf{O}\phi \in \mathsf{Arg}(\mathsf{O}\mathcal{T})$ for any $\phi \in \Delta$. Thus, $\mathsf{Arg}(\mathsf{O}\mathcal{T})$ is not conflict-free. If $\mathcal{T} \nvdash \bot$, then $\mathcal{T}$ is *a proper subset* of a set $\mathcal{T}'$ in $\mathsf{MCS}(\mathcal{S})$, and hence $\mathsf{Arg}(\mathsf{O}\mathcal{T}') \in \mathsf{Stb}(\mathsf{O}\mathcal{S})$. Since $\mathsf{Arg}(\mathsf{O}\mathcal{T}) \subsetneq \mathsf{Arg}(\mathsf{O}\mathcal{T}')$, we have that $\mathsf{Arg}(\mathsf{O}\mathcal{T})$ is neither preferred nor stable. □

Next, we extend our setting with preferences among obligations. The need of this has been recognized in e.g. [1, 28]. One way to introduce priorities among the obligations is to index the O-operator with a natural number, indicating the relative importance of the obligation. For instance, $\mathsf{O}^3\phi$ may indicate that 'there is an obligation of degree 3 to bring about ϕ'. Assuming that higher indices denote lower priorities, this can be reflected in our calculus as follows. Firstly, we rephrase the rules (KR) and (DR). For $k \geq \max(\{i_1, \ldots, i_n\})$, we define:

$$\frac{\phi_1, \ldots, \phi_n \Rightarrow \phi}{\mathsf{O}^{i_1}\phi_1, \ldots, \mathsf{O}^{i_n}\phi_n \Rightarrow \mathsf{O}^k\phi} \quad (\mathsf{KRp})$$

[9] Recall (Section 2) that this means that no modal operators appear in the formulas of $\mathcal{S}$.

$$\frac{\phi_1, \ldots, \phi_n \Rightarrow \phi}{\mathsf{O}^{i_1}\phi_1, \ldots, \mathsf{O}^{i_n}\phi_n \Rightarrow \neg\mathsf{O}^k\neg\phi} \quad (\mathsf{DRp})$$

Secondly, we adjust the undercut rule as follows. Let $k \geq \max(\{i_1, \ldots, i_n\})$ and Γ_F be a set of non-modal formulas in the language $\mathcal{L}$. We define:

$$\frac{\Gamma_F,\ \psi_1 \supset \mathsf{O}^{i_1}\phi_1,\ \ldots,\ \psi_n \supset \mathsf{O}^{i_n}\phi_n \Rightarrow \neg(\psi \supset \mathsf{O}^k\phi) \qquad \Gamma, \psi \supset \mathsf{O}^k\phi \Rightarrow \theta}{\Gamma, \psi \supset \mathsf{O}^k\phi \not\Rightarrow \theta} \quad (\mathsf{pOUcut})$$

The idea is that if on the basis of *only* using factual information Γ_F and (conditional) obligations of at most degrees k we can derive $\neg(\psi \supset \mathsf{O}^k\phi)$, then we can use this argument to undercut any argument making use of $\neg(\psi \supset \mathsf{O}^k\phi)$.[10]

Note 4 The restriction of Γ_F to non-modal formulas is essential to prevent cases in which unrelated obligations with indices $i \leq k$ are added to the antecedent of the attacker (e.g., an attack could simply be triggered by adding $\psi \supset \mathsf{O}^k\phi$ to the antecedent of any sequent with consequent $\neg(\psi \supset \mathsf{O}^k\phi)$).

By means of pOUcut ('prioritized Obligation Undercut') we can characterize another well-known consequence relation, namely (the propositional version of) Brewka's *preferred subtheories* [11]. For this, we first briefly review the main idea and the technical details behind Brewka's approach.

Prioritized theories consist of prioritized formulas of the form (ψ, i), where ψ is a formula in $\mathcal{L}$ and $i \geq 1$ is a numerical indication about the relative priority of ψ (Here, again, higher numbers correspond to lower priorities). Given a prioritized theory $\mathcal{S} = \{(\phi_j, i_j) \mid j \in J\}$ for an index set J, we denote by $\mathcal{S}_*$ the set $\{\phi \mid (\phi, i) \in \mathcal{S}\}$ and by $\mathcal{S}_O$ the set $\{\mathsf{O}^i\phi \mid (\phi, i) \in \mathcal{S}\}$. Now,

Definition 8 *Let $\mathcal{S} = \{(\phi_j, i_j) \mid j \in J\}$ be a prioritized theory, and let $\mathcal{S}_1, \mathcal{S}_2 \subseteq \mathcal{S}$. We denote:*

- $\mathsf{MCS}_{\text{lex}}(\mathcal{S}) = \max_{\prec_{\text{lex}}}(\mathsf{MCS}_\circ(\mathcal{S}))$ *where*
  - $\mathsf{MCS}_\circ(\mathcal{S}) = \{\mathcal{T} \subseteq \mathcal{S} \mid \mathcal{T}_* \in \mathsf{MCS}(\mathcal{S}_*)\}$ *and*
  - $\mathcal{S}_1 \prec_{\text{lex}} \mathcal{S}_2$ *iff there is a k such that:*

    $\{(\psi, l) \in \mathcal{S}_1 \mid l < k\} = \{(\psi, l) \in \mathcal{S}_2 \mid l < k\}$ *and*

    $\{(\psi, l) \in \mathcal{S}_1 \mid l = k\} \subsetneq \{(\psi, l) \in \mathcal{S}_2 \mid l = k\}$.
- $\mathrm{Cn}^k_{\text{pt}}(\mathcal{S}) = \mathrm{Cn}(\{\psi \mid (\psi, l) \in \bigcap \mathsf{MCS}_{\text{lex}}(\mathcal{S}),\ l \leq k\})$ *and* $\mathrm{Cn}_{\text{pt}}(\mathcal{S}) = \bigcup_{k \geq 1} \mathrm{Cn}^k_{\text{pt}}(\mathcal{S})$.

[10] Although our discussion of priorities in this section is limited to unconditional obligations (see also Proposition 2 below), the rule is phrased in a more general way covering also applications to conditional norms. Note that the following is an instance of the rule:

$$\frac{\Gamma_F, \mathsf{O}^{i_1}\phi_1, \ldots, \mathsf{O}^{i_n}\phi_n \Rightarrow \neg\mathsf{O}^k\phi \quad \Gamma, \mathsf{O}^k\phi \Rightarrow \theta}{\Gamma, \mathsf{O}^k\phi \not\Rightarrow \theta}.$$

Example 5 Let $\mathcal{S} = \{(d,1),(a,2),(\neg a,2),(\neg d,3)\}$. We have two members of $\mathsf{MCS}_{\rm lex}(\mathcal{S})$, namely $\mathcal{E}_1 = \{(a,2),(d,1)\}$ and $\mathcal{E}_2 = \{(\neg a,2),(d,1)\}$.[11] Hence, $d \in \mathrm{Cn}^2_{\rm pt}(\mathcal{S}) = \mathrm{Cn}_{\rm pt}(\mathcal{S})$.[12]

Reasoning with prioritized theories may now be characterized in our framework as follows.

Proposition 2 *Let $\mathcal{S} = \{(\phi_j, i_j) \mid j \in J\}$ be a prioritized theory. Then:*

1. $\mathcal{S}_O \;_{\mathsf{stb}}\!\mid\!\sim^{\Cap}_{\mathsf{pOUcut}} \mathsf{O}^k\psi$ *iff* $\mathcal{S}_O \;_{\mathsf{prf}}\!\mid\!\sim^{\Cap}_{\mathsf{pOUcut}} \mathsf{O}^k\psi$ *iff* $\psi \in \mathrm{Cn}^k_{\rm pt}(\bigcap \mathrm{MCS}_{\rm lex}(\mathcal{S}))$
2. $\mathcal{S}_O \;_{\mathsf{stb}}\!\mid\!\sim^{\cap}_{\mathsf{pOUcut}} \mathsf{O}^k\psi$ *iff* $\mathcal{S}_O \;_{\mathsf{prf}}\!\mid\!\sim^{\cap}_{\mathsf{pOUcut}} \mathsf{O}^k\psi$ *iff* $\psi \in \bigcap_{\mathcal{T} \in \mathsf{MCS}_{\rm lex}(\mathcal{S})} \mathrm{Cn}^k_{\rm pt}(\mathcal{T})$

It is sufficient to prove the following lemma:

Lemma 2 *Let $\mathcal{S} = \{(\phi_j, i_j) \mid j \in J\}$, be a (nonempty) prioritized theory and let $\mathcal{AF} = \mathcal{AF}_{\mathsf{pOUcut}}(\mathcal{S}_O)$. Then $\mathcal{E} \in \mathsf{Prf}(\mathcal{AF})$ iff $\mathcal{E} \in \mathsf{Stb}(\mathcal{AF})$ iff $\mathcal{E} = \mathsf{Arg}(\mathcal{T}_O)$ for a $\mathcal{T} \in \mathsf{MCS}_{\rm lex}(\mathcal{S})$*

Proof. Suppose that $\mathcal{T} \in \mathsf{MCS}_{\rm lex}(\mathcal{S})$. Conflict-freeness of $\mathsf{Arg}(\mathcal{T}_O)$ is trivial since $\mathcal{T}_*$ is consistent and hence $\Delta_O \not\Rightarrow \neg\mathsf{O}^k\psi$ for any $\psi \in \mathcal{T}_*$ and any finite $\Delta \subseteq \mathcal{T}$. Let $\Gamma_O \Rightarrow \phi \in \mathsf{Arg}(\mathcal{S}_O) \setminus \mathsf{Arg}(\mathcal{T}_O)$ (where $\Gamma \subseteq \mathcal{S}$). Hence, $\Gamma \setminus \mathcal{T} \neq \emptyset$. Let $(\psi, k) \in \Gamma \setminus \mathcal{T}$. Assume for a contradiction that there is no finite $\Delta \subseteq \mathcal{T}$ such that $\Delta_* \vdash \neg\psi$ and $\max(\{l \mid (\phi, l) \in \Delta\}) \leq k$. But then $\mathcal{T} \prec_{\rm lex} \mathcal{T}'$, where $\mathcal{T}' \in \mathsf{MCS}_*(\mathcal{S})$ is such that $\{(\phi, l) \in \mathcal{T} \mid l \leq k\} \cup \{(\psi, k)\} \subseteq \mathcal{T}'$. Note that such a $\mathcal{T}'$ exists since $\{\phi \mid (\phi, l) \in \mathcal{T},\ l \leq k\} \cup \{\psi\}$ is consistent. Altogether, this is a contradiction to the fact that $\mathcal{T} \in \mathsf{MCS}_{\rm lex}(\mathcal{S})$. Hence, there is a finite $\Delta \subseteq \mathcal{T}$ such that $\Delta_* \vdash \neg\psi$ and $\max(\{l \mid (\phi, l) \in \Delta\}) \leq k$. Clearly, $\Delta_O \Rightarrow \mathsf{O}^k\neg\psi \in \mathsf{Arg}(\Delta_O) \subseteq \mathsf{Arg}(\mathcal{T}_O)$. By (DRp), $\Delta_O \Rightarrow \neg\mathsf{O}^k\psi \in \mathsf{Arg}(\Delta_O) \subseteq \mathsf{Arg}(\mathcal{T}_O)$. This argument attacks $\Gamma_O \Rightarrow \phi$. It follows that $\mathsf{Arg}(\mathcal{T}_O)$ attacks any sequent (in $\mathsf{Arg}(\mathcal{S}_O)$) that is not in this set, thus $\mathsf{Arg}(\mathcal{T}_O)$ is a stable (and so also a preferred) extension of $\mathcal{AF}$. The other direction is shown in a similar way and is left to the reader. □

Let us take a look at the rebuttal rule OOReb. For this we go back to SDL without priorities. Note that, given a (finite and) inconsistent subset Γ of $\mathcal{S}$, one can produce just any conclusion and rebut any other sequent $\Gamma' \Rightarrow \mathsf{O}\phi$ with the sequent $\Gamma \Rightarrow \mathsf{O}\neg\phi$. This excessive behavior can be avoided, for instance, in the following way. We introduce an additional undercut-type of attack, called UCon ('Undercut Consistency') that allows to attack any sequent with an inconsistent premise set Γ by means of the sequent $\Rightarrow \neg\bigwedge\Gamma'$ for some $\Gamma' \subseteq \Gamma$:

$$\frac{\Rightarrow \neg\bigwedge\Gamma' \qquad \Gamma \Rightarrow \phi}{\Gamma \not\Rightarrow \phi} \quad \text{where } \Gamma' \subseteq \Gamma. \qquad (\mathsf{UCon})$$

[11] Note that for $\mathcal{E} = \{(\neg d,3),(a,2)\}$ we have $\mathcal{E} \notin \mathsf{MCS}_{\rm lex}(\mathcal{S})$ although $\mathcal{E}_* \in \mathsf{MCS}(\mathcal{S}_*)$.

[12] In [11] Brewka considers also a generalization of the system presented above, according to which instead of having a linear (numerical) order, inputs may be ordered by a strict partial order $\prec$. The idea is to 'linearize' these partial orders, that is, to consider strict total orders $\prec'$ of the input that respect $\prec$, and then to proceed on the basis of $\prec'$ as described above. We note that our approach is capable of handling this generalization as well, and shall elaborate on this in a future work.

To avoid situations in which sequents with inconsistent premises defend themselves by rebutting UCon-attackers, we require that rebuts can only affect arguments with non-empty premise sets. Thus, for instance, in the definition of $\mathsf{NN'Reb}$, one could add a side-condition that $\Gamma' \neq \emptyset$.

Using the attack rules OOReb (with the additional side condition mentioned above) and UCon, we can characterize a deontic variant of the entailment considered by Benferhat, Dubois and Prade in [9]:

Definition 9 *Consider the following strengthening (defined in [9]) of the entailment relations from Proposition 1: Given a set $\mathcal{S}$ of (modality-free) propositions in $\mathcal{L}$, we denote $\mathcal{S} \mid\!\sim_{\text{Argued}} \phi$ iff the following conditions are satisfied:*

1. *$\mathcal{T} \vdash \phi$ for some consistent subset $\mathcal{T}$ of $\mathcal{S}$.*
2. *There is no consistent subset $\mathcal{T}$ of $\mathcal{S}$ such that $\mathcal{T} \vdash \neg\phi$.*

The entailment $\mid\!\sim_{\text{Argued}}$ is characterizable in our framework as follows:

Proposition 3 *Let $\mathcal{S}$ be a set of (modality-free) formulas in $\mathcal{L}$ and ψ a formula in $\mathcal{L}$. Then $\mathcal{S} \mid\!\sim_{\text{Argued}} \psi$ iff $\mathsf{O}\mathcal{S} \;_{\mathsf{gr}}\!\mid\!\sim_{\mathsf{OOReb},\mathsf{UCon}} \mathsf{O}\psi$.*

Proof. Suppose that $\mathcal{S} \mid\!\sim_{\text{Argued}} \psi$. Hence, there is a $\mathcal{T} \in \mathsf{MCS}(\mathcal{S})$ such that $\mathcal{T} \vdash \psi$ and there is no $\mathcal{T}' \in \mathsf{MCS}(\mathcal{S})$ such that $\mathcal{T}' \vdash \neg\psi$. Thus (by the compactness of SDL), $A = \mathsf{O}\Delta \Rightarrow \mathsf{O}\psi \in \mathsf{Arg}(\mathsf{O}\mathcal{S})$ for some finite $\Delta \subseteq \mathcal{T}$. Suppose that $\mathsf{O}\mathcal{T}' \Rightarrow \mathsf{O}\neg\psi \in \mathsf{Arg}(\mathsf{O}\mathcal{S})$. Then $\mathcal{T}' \vdash \neg\psi$ and hence $\mathcal{T}' \vdash \bot$ for some finite $\mathcal{T}' \subseteq \mathcal{S}$. Then $\Rightarrow \neg\bigwedge \mathsf{O}\Delta'$ is derivable for some $\Delta' \subseteq \mathcal{T}'$. Clearly, $\Rightarrow \neg\bigwedge \mathsf{O}\Delta' \in \mathsf{Arg}(\mathsf{O}\mathcal{S}) \setminus \mathsf{Arg}(\mathsf{O}\mathcal{S})^{+}$. This means that any OOReb-attacker of A is counter-UCon-attacked by an argument in $\mathsf{Arg}(\mathsf{O}\mathcal{S})$ (which itself is not attacked), thus A is defended. It follows, then, that $A \in \mathsf{Grd}(\mathcal{AF}_{\mathsf{OOReb},\mathsf{UCon}}(\mathsf{O}\mathcal{S}))$.

Suppose now that $\mathcal{S} \not\mid\!\sim_{\text{Argued}} \psi$. This means that either there is no $\mathcal{T} \in \mathsf{MCS}(\mathcal{S})$ such that $\mathcal{T} \vdash \psi$, or there is a $\mathcal{T} \in \mathsf{MCS}(\mathcal{S})$ such that $\mathcal{T} \vdash \neg\psi$. In the first case the only sequents A with a consequence $\mathsf{O}\psi$ and $\mathsf{Prem}(A) \subseteq \mathsf{O}\mathcal{S}$ are those that $\Rightarrow \neg\bigwedge\Gamma$ where $\Gamma \subseteq \mathsf{Prem}(A)$, and hence all arguments for $\mathsf{O}\psi$ are not members of any admissible extension of $\mathcal{AF}_{\mathsf{OOReb},\mathsf{UCon}}(\mathsf{O}\mathcal{S})$. In the second case we can construct an admissible extension $\mathcal{E}$ such that $A \in \mathcal{E}^{+}$ for any $A = \Delta \Rightarrow \mathsf{O}\psi \in \mathsf{Arg}(\mathsf{O}\mathcal{S})$ by letting $\mathcal{E} = \mathsf{Arg}(\mathsf{O}\mathcal{T})$. We leave to the reader to verify that $\mathcal{E} \in \mathsf{Adm}(\mathcal{AF}_{\mathsf{OOReb},\mathsf{UCon}}(\mathsf{O}\mathcal{S}))$. Clearly, this also implies that $A \notin \mathsf{Grd}(\mathcal{AF}_{\mathsf{OOReb},\mathsf{UCon}}(\mathsf{O}\mathcal{S}))$ for any such $A \in \mathsf{Arg}(\mathsf{O}\mathcal{S})$. □

5 Reasoning with Conditional Obligations

We now turn to the general case, where conditional obligations are allowed. In this section we demonstrate our argumentative model for normative reasoning by means of various examples.

We start with the following two easily verified basic observations:

Observation 1 *Let $\mathcal{S}$ be a set of formulas in $\mathcal{L}_{\mathsf{SDL}}$. Then:*

1. *For any set of attack rules previously defined and any entailment $\mathrel{|\!\sim}$ induced from the corresponding normative argumentation framework by Definition 6, whenever $\mathcal{S}$ is* SDL*-consistent (i.e., $\mathcal{S} \nvdash_{\mathsf{SDL}} \bot$) then $\mathcal{S} \vdash_{\mathsf{SDL}} \psi$ iff $\mathcal{S} \mathrel{|\!\sim} \psi$. It is easy to verify that in this case all arguments in* $\mathsf{Arg}(\mathcal{S})$ *are selected since there are no argumentative attacks.*

2. *If* UCon *is part of the attack rules and the other attack rules are restricted in such a way that sequents with empty premise sets cannot be attacked, then: (i) $\mathcal{S} \mathrel{|\!\sim} \phi$ implies that ϕ is* SDL*-consistent (i.e., $\phi \nvdash_{\mathsf{SDL}} \bot$), thus (ii) $\mathrel{|\!\sim}$ is non-exploding, that is: for all $\mathcal{S}$, $\mathcal{S} \not\mathrel{|\!\sim} \bot$.*

Example 6 Let us recall Example 1, based on the following set of assertions:

$$\mathcal{S}_{\mathsf{asp}} = \{m,\ a,\ m \supset \mathsf{O}\neg f,\ (m \wedge a) \supset \mathsf{O}f\}.$$

Some arguments in $\mathsf{Arg}(\mathcal{S}_{\mathsf{asp}})$ are listed in Figure 3 (right). Figure 3 (left) shows an attack diagram where the only attack rule is $\mathsf{OOSpeUcut}$.

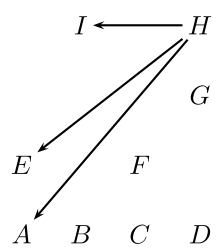

$A = m \supset \mathsf{O}\neg f \Rightarrow m \supset \mathsf{O}\neg f$
$B = m \Rightarrow m$
$C = a \Rightarrow a$
$D = (m \wedge a) \supset \mathsf{O}f \Rightarrow (m \wedge a) \supset \mathsf{O}f$
$E = m, m \supset \mathsf{O}\neg f \Rightarrow \mathsf{O}\neg f$
$F = m, a \Rightarrow m \wedge a$
$G = m, a, (m \wedge a) \supset \mathsf{O}f \Rightarrow \mathsf{O}f$
$H = m, a, (m \wedge a) \supset \mathsf{O}f \Rightarrow \neg(m \supset \mathsf{O}\neg f)$
$I = m, a, m \supset \mathsf{O}\neg f, (m \wedge a) \supset \mathsf{O}f \Rightarrow \mathsf{O}\bot$

Figure 3: (Part of) the normative argumentation framework of Example 6.

We observe that H $\mathsf{OOSpeUcut}$-attacks A, E, and I. It follows that, as expected, we have the following deductions:

- $\mathcal{S}_{\mathsf{asp}} \not\mathrel{|\!\sim} \mathsf{O}\neg f$. Indeed, one cannot derive $\mathsf{O}\neg f$ since the application of MP to $m \supset \mathsf{O}\neg f$ (depicted by argument E) gets attacked by H.[13]
- $\mathcal{S}_{\mathsf{asp}} \mathrel{|\!\sim} \mathsf{O}f$. Indeed, G is not $\mathsf{OOSpeUcut}$-attacked by an argument in $\mathsf{Arg}(\mathcal{S}_{\mathsf{asp}})$, thus it is part of every grounded, preferred, and stable extension of the underlying normative argumentation framework, and so its descendant follows from $\mathcal{S}_{\mathsf{asp}}$.[14]

The next example shows that we have to be careful when modeling conditional obligations with rebuttals.

[13] Note that $m \supset \mathsf{O}\neg f$ cannot be derived either due to the attack of H on A.

[14] It is important to note that G *is* $\mathsf{OOSpeUcut}$-attackable by SDL-derivable arguments, but none of them is in $\mathsf{Arg}(\mathcal{S}_{\mathsf{asp}})$. For instance, since the material implication allows for strengthening of antecedents ($\phi \supset \psi \Rightarrow (\phi \wedge \phi') \supset \psi$), we have that $m \supset \mathsf{O}\neg f \Rightarrow (m \wedge a) \supset \mathsf{O}\neg f$ is SDL-derivable, and so G is attacked by an argument with, say, the SDL-derivable sequent $m, m \supset \mathsf{O}\neg f, m, a, (m \wedge a) \supset \mathsf{O}\neg f \Rightarrow \neg((m \wedge a) \supset \mathsf{O}\neg f)$. Yet, since $m \wedge a \supset \mathsf{O}\neg f \notin \mathcal{S}_{\mathsf{asp}}$, this argument is not in $\mathsf{Arg}(\mathcal{S}_{\mathsf{asp}})$. We note, further, that the sequent $a, m, m \supset \mathsf{O}\neg f \Rightarrow \neg((m \wedge a) \supset \mathsf{O}f)$ is derivable, but it does not $\mathsf{OOSpeUcut}$-attack G and H, though it is attacked by H.

Example 7 In [14], Caminada gives the following example for a deontic conflict that is not resolved by a resolution principle such as specificity:

- snoring is a misbehavior ($s \supset m$),
- it is allowed to remove misbehaving people from the library ($m \supset \mathsf{P}r$),
- it is obliged not to remove a professor from the library ($p \supset \mathsf{O}\neg r$),
- people who misbehave are subject to a fine ($m \supset \mathsf{O}f$).

Now, suppose we have a snoring professor resulting in the following set:

$$\mathcal{S}_{\mathsf{pro}} = \{s,\ p,\ s \supset m,\ p \supset \mathsf{O}\neg r,\ m \supset \mathsf{P}r,\ m \supset \mathsf{O}f\}.$$

The sequents $A_1 = p,\ p \supset \mathsf{O}\neg r \Rightarrow \mathsf{O}\neg r$ and $A_2 = s,\ s \supset m,\ m \supset \mathsf{P}r \Rightarrow \mathsf{P}r$ are provable and rebut each other according to OPReb and POReb.

Caminada uses this example to illustrate what is sometimes considered a shortcoming of deductive approaches to defeasible reasoning [14, 42]. Given two conflicting inference steps described schematically by $\phi \rightsquigarrow \neg\theta$ and $\psi_1 \rightsquigarrow \psi_2 \rightsquigarrow \theta$. When "$\rightsquigarrow$" is contrapositive, we get $\phi \rightsquigarrow \neg\theta \rightsquigarrow \neg\psi_2$. This schematic representation applied to our example yields $p \rightsquigarrow \mathsf{O}\neg r \rightsquigarrow \neg\mathsf{P}r \rightsquigarrow \neg m$. Thus, the sequent $A_4 = p,\ p \supset \mathsf{O}\neg r,\ m \supset \mathsf{P}r \Rightarrow \neg m$ is provable and conflicting with $A_3 = s,\ s \supset m \Rightarrow m$. Caminada argues that this violates the aspiration to keep conflicts as local as possible and so deontic conditionals are not to be contrapositive.

In our case, although contraposition holds in SDL, we can 'undercut' A_4 by attack rules such as cNiC ('conditional Norms imply Can', where Γ_F is in the language $\mathcal{L}$ and $\mathsf{N}_i \in \{\mathsf{O}, \mathsf{P}\}$ $(1 \leq i \leq n)$)

$$\frac{\Gamma_F \Rightarrow \neg \bigwedge_{i=1}^{n}(\phi_i \supset \mathsf{N}_i\psi_i) \qquad \Gamma, \phi_1 \supset \mathsf{N}_1\psi_1, \ldots, \phi_n \supset \mathsf{N}_n\psi_n \Rightarrow \psi}{\Gamma, \phi_1 \supset \mathsf{N}_1\psi_1, \ldots, \phi_n \supset \mathsf{N}_n\psi_n \not\Rightarrow \psi} \quad (\mathsf{cNiC})$$

The idea is to undercut any sequent that features a deontic conflict in its antecedent. A deontic conflict is given by a set of obligations $\phi_i \supset \mathsf{N}_i\psi_i$ $(1 \leq i \leq n)$ that cannot be jointly realized in view of the given facts Γ_F. This is expressed by the attacking sequent.

In our example, in order to construct an argument for $\neg m$ we need the two conflicting conditional obligations $p \supset \mathsf{O}\neg r$ and $m \supset \mathsf{P}r$ as premises. The fact that they are both triggered and conflicting in view of the given factual information s, p and $s \supset m$ can be formally expressed by the derivable sequent $A_7 = s \wedge p \wedge (s \supset m) \Rightarrow \neg((m \supset \mathsf{P}r) \wedge (p \supset \mathsf{O}\neg r))$. Now, A_7 attacks all sequents that have $p \supset \mathsf{O}\neg r$ and $m \supset \mathsf{P}r$ as premises, such as A_4. Additionally, we also need to restrict rebuttals in such a way that no sequents are attacked with premise sets with only factual information (e.g., premise sets that do not contain any O or P).

Figure 4 depicts an excerpt of an attack diagram for $\mathcal{S}_{\mathsf{pro}}$ with the attack rules cNiC and NN'Reb (where $\mathsf{NN'} \in \{\mathsf{OO}, \mathsf{OP}, \mathsf{PO}\}$).

We get, e.g., that $\mathcal{S}_{\mathsf{pro}}\ {}_{\mathsf{gr}}{\mid\!\sim}_{\mathsf{cNiC,NN'Reb}}\ \mathsf{O}f$ and $\mathcal{S}_{\mathsf{pro}}\ {}_{\mathsf{gr}}{\mid\!\not\sim}_{\mathsf{cNiC,NN'Reb}}\ \mathsf{P}r$. If we use only cNiC, we get also $\mathcal{S}_{\mathsf{pro}}\ {}_{\mathsf{gr}}{\mid\!\sim}_{\mathsf{cNiC}}\ \mathsf{P}r$ and $\mathcal{S}_{\mathsf{pro}}\ {}_{\mathsf{gr}}{\mid\!\sim}_{\mathsf{cNiC}}\ \mathsf{O}\neg r$, but we are still

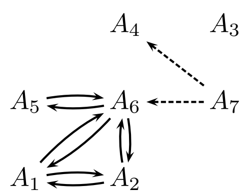

$$A_1 = p, p \supset \mathsf{O}\neg r \Rightarrow \mathsf{O}\neg r$$
$$A_2 = s, s \supset m, m \supset \mathsf{P}r \Rightarrow \mathsf{P}r$$
$$A_3 = s, s \supset m \Rightarrow m$$
$$A_4 = p, p \supset \mathsf{O}\neg r, m \supset \mathsf{P}r \Rightarrow \neg m$$
$$A_5 = s, s \supset m, m \supset \mathsf{O}f \Rightarrow \mathsf{O}f$$
$$A_6 = p, p \supset \mathsf{O}\neg r, m \supset \mathsf{P}r, s, s \supset m \Rightarrow \mathsf{O}\bot$$
$$A_7 = s, s \supset m, p \Rightarrow \neg((m \supset \mathsf{P}r) \wedge (p \supset \mathsf{O}\neg r))$$

Figure 4: A diagram for Example 7. Here, cNiC-attacks are denoted by dashed lines and NN'Reb-attacks are denoted by solid lines.

not able to accept arguments with both conflicting conditional obligations as premise (e.g., A_4 and A_6), since such arguments are cNiC-attacked by A_7.[15]

Example 8 Next, we take a look at contrary-to-duty (CTD) obligations. A paradigmatic example is Forrester's gentle murderer scenario [20]:

- Generally, one ought not to kill ($\top \supset \mathsf{O}\neg k$).
- However, upon killing, this should be done gently ($k \supset \mathsf{O}(k \wedge g)$).

This scenario may be represented as follows:

$$\mathcal{S}_{\mathsf{gm}} = \{k,\ \top \supset \mathsf{O}\neg k,\ k \supset \mathsf{O}(k \wedge g)\}.$$

We list various ways in which this example can be treated in our framework.

1. We can use specificity rules such as OOSpeUcut in Example 6 that do not only block detachment from overridden conditional obligations but do 'destroy' them in the sense that they do not appear in the consequence set.

Van Der Torre and Tan [51] distinguish CTD-obligations from cases of specificity. In the former the general obligations are not canceled or overridden but 'overshadowed': they still have normative force despite the fact that they are violated. By contrast, in cases of specificity the more general conditional obligations are canceled and thus deprived of normative force. There are various ways in which attack rules can be used to reflect this distinction.

2. We can make use of rules such as OOReb that preserve 'overshadowed' conditional CTD obligations despite the fact that detachment is blocked (see the upper right diagram of Figure 5, where OOReb (pointed arrows) and cNiC (solid arrow) are used). When $\mathcal{E} = \{A, B, C, G\}$, both $\mathcal{E} \cup \{D\}$ and $\mathcal{E} \cup \{E\}$, are preferred extensions, hence:

$$\mathcal{S}_{\mathsf{gm}}\ {}_{\mathsf{prf}}\!\mid\!\sim^{\cup}_{\mathsf{OOReb,cNiC}}\ \mathsf{O}\neg k \ \text{ and } \ \mathcal{S}_{\mathsf{gm}}\ {}_{\mathsf{prf}}\!\mid\!\sim^{\cup}_{\mathsf{OOReb,cNiC}}\ \mathsf{O}(k \wedge g).$$

[15] This is in line with Goble's [24, p. 27–28] analysis of an enriched version of Horty's Smith argument [32]: Given $\{\mathsf{O}\neg f,\ \mathsf{O}(f \vee s),\ \mathsf{O}\neg s\}$ (where f abbreviates fighting in the army and s is performing civil service) Goble advocates to let both $\mathsf{O}s$ and $\mathsf{O}\neg s$ be derivable without aggregating them to $\mathsf{O}(s \wedge \neg s)$.

$A = \top \supset \mathsf{O}\neg k \Rightarrow \top \supset \mathsf{O}\neg k$
$B = k \Rightarrow k$
$C = k \supset \mathsf{O}(k \wedge g) \Rightarrow k \supset \mathsf{O}(k \wedge g)$
$D = \top \supset \mathsf{O}\neg k \Rightarrow \mathsf{O}\neg k$
$E = k, k \supset \mathsf{O}(k \wedge g) \Rightarrow \mathsf{O}(k \wedge g)$
$F = k, \top \supset \mathsf{O}\neg k, k \supset \mathsf{O}(k \wedge g) \Rightarrow \mathsf{O}\bot$
$G = k \Rightarrow \neg((\top \supset \mathsf{O}\neg k) \wedge (k \supset \mathsf{O}(k \wedge g)))$

Figure 5: Two modelings of Forrester's gentle murderer

It follows that the tension created by overshadowing is expressed on the level of unconditional obligations: both the more specific and the overshadowed obligation are derivable in the credulous approach as unconditional obligations. In the skeptical approach we get:

$$\mathcal{S}_{\mathsf{gm}}\ {}_{\mathsf{prf}}\mid\!\sim^{\cap}_{\mathsf{OOReb,cNiC}}\ \mathsf{O}(\neg k \vee (k \wedge g)) \ \text{ and } \ \mathcal{S}_{\mathsf{gm}}\ {}_{\mathsf{prf}}\mid\!\sim^{\cap}_{\mathsf{OOReb,cNiC}}\ \mathsf{O}\neg k \vee \mathsf{O}(k \wedge g).$$

3. We can use uOiC ('unconditional Ought implies Can')

$$\frac{\Gamma \Rightarrow \phi \qquad \Delta \Rightarrow \mathsf{O}\neg\phi \qquad \Delta \Rightarrow \psi}{\Delta \not\Rightarrow \psi} \qquad (\mathsf{uOiC})$$

that blocks detachment from obligations that conflict with the given factual information. This is illustrated in the upper left-hand side of Figure 5. While $\mathcal{S}_{\mathsf{gm}}\ {}_{\mathsf{grd}}\mid\!\sim_{\mathsf{uOiC}}\ \mathsf{O}(k \wedge g)$, we also have $\mathcal{S}_{\mathsf{gm}}\ {}_{\mathsf{grd}}\not\mid\!\sim_{\mathsf{uOiC}}\ \mathsf{O}\neg k$. To block sequents that make use of contraposition to derive $\neg k$ (e.g., $\top \supset \mathsf{O}\neg k, k \supset \mathsf{O}(k \wedge g) \Rightarrow \neg k$) we may want to use in addition uNiC as motivated in Example 7.

4. Yet another option is to use a very liberal approach with cNiC only. This will block arguments with inconsistent premises such as in the sequent F of Figure 5, but otherwise allows, e.g., to derive both $\mathsf{O}\neg k$ and $\mathsf{O}(g \wedge k)$ even by the grounded approach:

$$\mathcal{S}_{\mathsf{gm}}\ {}_{\mathsf{grd}}\mid\!\sim_{\mathsf{cNiC}}\ \mathsf{O}\neg k \ \text{ and } \ \mathcal{S}_{\mathsf{gm}}\ {}_{\mathsf{grd}}\mid\!\sim_{\mathsf{cNiC}}\ \mathsf{O}(k \wedge g).$$

Example 9 Consider the paradigmatic CTD-case below, known as the Chisholm paradox [16]:

- It ought to be that Jones visits his neighbors.
- It ought to be that if Jones goes, he tells them that he is coming.

- If Jones doesn't go, then he ought not to tell them that he is coming.
- Jones doesn't visit his neighbors.

In the modelling of this configuration, specific requirements have been posed. First, the logical model should not trivialize the set. Second, the formal representation of the four sentences should be rendered logically independent. It is obvious that by modelling conditional obligations via $\phi \supset \mathsf{O}\psi$ we will fail to meet the second requirement since with material implication we have $\neg\phi \supset (\phi \supset \psi)$ and hence we get $\neg g \supset (g \supset \mathsf{O}t)$ (where g is going to the neighbors and t is telling them). Since $\{\neg g, g \supset \mathsf{O}t, \neg g \supset \mathsf{O}\neg t, \mathsf{O}g\}$ is SDL-consistent, argumentation frameworks based on this set and based on the previously discussed attack rules are conflict-free. Hence, the first criterion is met.[16]

Example 10 Next we take a look at a simple conflict that is neither a specificity nor a CTD-case. Let

$$\mathcal{S}_{\mathsf{bi}} = \{a,\ b,\ a \supset \mathsf{O}(c \wedge d),\ b \supset \mathsf{O}(\neg c \wedge d)\}.$$

Figure 6 shows the situation for the attack rule OiReb ('Obligation $\mathsf{i}$ndirect Rebuttal')

$$\frac{\Gamma \Rightarrow \mathsf{O}\phi \qquad \psi \Rightarrow \neg\mathsf{O}\phi \qquad \Delta \Rightarrow \psi}{\Delta \not\Rightarrow \psi} \quad \text{where } \Delta \neq \emptyset \qquad (\mathsf{OiReb})$$

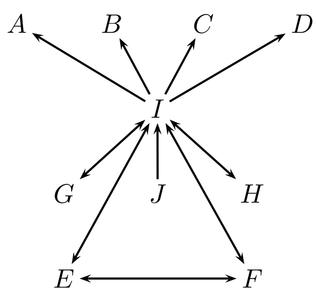

$A = a \supset \mathsf{O}(c \wedge d) \Rightarrow a \supset \mathsf{O}(c \wedge d)$
$B = a \Rightarrow a$
$C = b \Rightarrow b$
$D = b \supset \mathsf{O}(\neg c \wedge d) \Rightarrow b \supset \mathsf{O}(\neg c \wedge d)$
$E = a, a \supset \mathsf{O}(c \wedge d) \Rightarrow \mathsf{O}(c \wedge d)$
$F = b, b \supset \mathsf{O}(\neg c \wedge d) \Rightarrow \mathsf{O}(\neg c \wedge d)$
$G = a, a \supset \mathsf{O}(c \wedge d) \Rightarrow \mathsf{O}d$
$H = b, b \supset \mathsf{O}(\neg c \wedge d) \Rightarrow \mathsf{O}d$
$I = a, b, a \supset \mathsf{O}(c \wedge d), b \supset \mathsf{O}(\neg c \wedge d) \Rightarrow \mathsf{O}\bot$
$J = \ \Rightarrow \mathsf{O}\top$

Figure 6: A simple conflict (Example 10)

We have the following preferred extensions where $\mathcal{E} = \{A, B, C, D, G, H, J\}$: $\mathcal{E} \cup \{E\}$ and $\mathcal{E} \cup \{F\}$. Note that we have a *floating conclusion* here (i.e., a conclusion which is obtained from each of a set of otherwise conflicting arguments, see for instance [31]): $\mathcal{S}_{\mathsf{bi}}\ {}_{\mathsf{prf}}\!\mid\!\sim^{\cap}_{\mathsf{OiReb}} \mathsf{O}d$ and $\mathcal{S}_{\mathsf{bi}}\ {}_{\mathsf{prf}}\!\not\mid\!\sim^{\Cap}_{\mathsf{OiReb}} \mathsf{O}d$ since G and H are in every preferred extension.

Note that we have a similar situation when we use OUcut instead of OiReb: $\mathcal{S}_{\mathsf{bi}}\ {}_{\mathsf{prf}}\!\mid\!\sim^{\cap}_{\mathsf{OUcut}} \mathsf{O}d$. However, this time $\mathcal{S}_{\mathsf{bi}}\ {}_{\mathsf{prf}}\!\not\mid\!\sim^{\Cap}_{\mathsf{OUcut}} \mathsf{O}d$ since we get the following two preferred extensions of $\mathcal{AF}_{\mathsf{OUcut}}(\mathcal{S}_{\mathsf{bi}})$: $\mathsf{Arg}(\mathcal{S}_{\mathsf{bi}} \setminus \{a \supset \mathsf{O}(c \wedge d)\})$ and

[16] An alternative modelling of the second premise by $\mathsf{O}(g \supset t)$ is not appropriate here, since it would be ad-hoc to model some conditional obligations by $\phi \supset \mathsf{O}\psi$ and others by $\mathsf{O}(\phi \supset \psi)$ whenever we run into problems with logical dependency. Moreover, given $\mathsf{O}g$ and $\mathsf{O}(g \supset t)$ we would be able to derive $\mathsf{O}t$ although g is not derivable, i.e., although the conditional obligation is not triggered.

$\mathsf{Arg}(\mathcal{S}_{\mathsf{bi}} \setminus \{a \supset \mathsf{O}(\neg c \wedge d)\})$. This implies, for instance, that neither G nor H is in both extensions.

Another difference between ${}_{\mathsf{prf}}{\mid\!\sim}_{\mathsf{OUcut}}$ and ${}_{\mathsf{prf}}{\mid\!\sim}_{\mathsf{OiReb}}$ is that OiReb allows to derive the conditional obligations $a \supset \mathsf{O}(c \wedge d)$ and $b \supset \mathsf{O}(\neg c \wedge d)$ (sequents A and D are only attackable via sequents with inconsistent norms in their consequents and they are hence part of every preferred extension), while at the same time blocking detachment from them (sequents E and F are never part of the same preferred extension). In contrast, when OUcut is used, A is attacked by $a, b, b \supset \mathsf{O}(\neg c \wedge d) \Rightarrow \neg(a \supset \mathsf{O}(c \wedge d))$ and as a consequence neither $a \supset \mathsf{O}(c \wedge d)$ nor $b \supset \mathsf{O}(\neg c \wedge d)$ are derivable in ${}_{\mathsf{prf}}{\mid\!\sim}^{\cap}_{\mathsf{OUcut}}$ and ${}_{\mathsf{prf}}{\mid\!\sim}^{\Cap}_{\mathsf{OUcut}}$ but we only get $\mathcal{S}_{\mathsf{bi}}\ {}_{\mathsf{prf}}{\mid\!\sim}^{\cap}_{\mathsf{OUcut}}\ (a \supset \mathsf{O}(c \wedge d)) \vee (b \supset \mathsf{O}(\neg c \wedge d))$.

Example 11 Our next example illustrates a conflict between three obligations. Consider the following set:

$$\mathcal{S}_{\mathsf{tri}} = \{c,\ c \supset \mathsf{O}(a \vee b),\ c \supset \mathsf{O}(a \vee \neg b),\ c \supset \mathsf{O}\neg a\}.$$

Figure 7 presents an attack diagram where the set of attacks consists of OOReb (dotted arrows) and cNiC (dashed arrow).

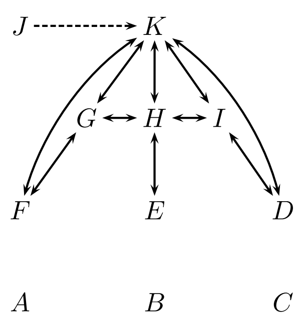

$A = c \supset \mathsf{O}(a \vee b) \Rightarrow c \supset \mathsf{O}(a \vee b)$
$B = c \supset \mathsf{O}(a \vee \neg b) \Rightarrow c \supset \mathsf{O}(a \vee \neg b)$
$C = c \supset \mathsf{O}\neg a \Rightarrow c \supset \mathsf{O}\neg a$
$D = c, c \supset \mathsf{O}(a \vee b) \Rightarrow \mathsf{O}(a \vee b)$
$E = c, c \supset \mathsf{O}(a \vee \neg b) \Rightarrow \mathsf{O}(a \vee \neg b)$
$F = c, c \supset \mathsf{O}\neg a \Rightarrow \mathsf{O}\neg a$
$G = c, c \supset \mathsf{O}(a \vee b), c \supset \mathsf{O}(a \vee \neg b) \Rightarrow \mathsf{O}a$
$H = c, c \supset \mathsf{O}(a \vee b), c \supset \mathsf{O}\neg a \Rightarrow \mathsf{O}(\neg a \wedge b)$
$I = c, c \supset \mathsf{O}(a \vee \neg b), c \supset \mathsf{O}\neg a \Rightarrow \mathsf{O}(\neg a \wedge \neg b)$
$J = c \Rightarrow \neg((c \supset \mathsf{O}(a \vee b)) \wedge (c \supset \mathsf{O}(a \vee \neg b)) \wedge (c \supset \mathsf{O}\neg a))$
$K = c, c \supset \mathsf{O}(a \vee b), c \supset \mathsf{O}(a \vee \neg b), c \supset \mathsf{O}\neg a \Rightarrow \mathsf{O}\bot$

Figure 7: A triple conflict (Example 11)

What is remarkable in this case is that no conflicts are triggered between D, E and F, since the triple-conflict is not reducible to a binary conflict that fits the attack rule OOReb, and so the set $\{A, B, C, D, E, F, J\}$ is a preferred extension in this case. This is undesirable since the unconditional obligations in the conclusions of D, E and F are inconsistent. This problem may be avoided by using OUcut (or OiReb) instead of OOReb. It can easily be seen that with these attack rules we have no preferred extension in which D, E and F are contained. Moreover, we get, for instance,

$$\mathcal{S}_{\mathsf{tri}}\ {}_{\mathsf{prf}}{\mid\!\sim}^{\cap}_{\mathsf{OUcut}}\ \mathsf{O}a \vee \mathsf{O}(\neg a \wedge b) \vee \mathsf{O}(\neg a \wedge \neg b)$$

which doesn't follow with OOReb. This example shows that elimination rules should be carefully chosen.[17]

[17] In Section 7 we will show that OUcut is rather well-behaved and can be used to give an argumentative account of a specific Input/Output logic.

Specificity, as discussed so far, is concerned with two conflicting conditional norms in the input set, one of which is more specific than the other. Sometimes, however, we may have to deal with situations where the two conditional norms that stand in a relationship of specificity are not contained in the input set but are derived from it. This has to be treated with caution though, since there are cases in which it is not intuitive to let the more specific derived conditional norm override the more general conditional norm. To see this, we take another look at Example 6. Since material implication allows for strengthening of the antecedent, we have $\vdash_{\mathsf{SDL}} (m \supset \mathsf{O}\neg f) \supset ((m \wedge a \wedge x) \supset \mathsf{O}\neg f)$ (where x is arbitrary). Intuitively, though, given the factual information $a, m, x \in \mathcal{S}$, the conditional obligation $(m \wedge a \wedge x) \supset \mathsf{O}\neg f$ should not override the conditional obligation $(m \wedge a) \supset \mathsf{O}f$.

Nevertheless, there are situations in which it makes sense to let a more specific derived conditional norm $\phi \supset \mathsf{N}\psi$ override a more general derived conditional norm $\lambda \supset \mathsf{N}'\delta$. This is the case, for instance, when both conditional norms are *equivalent* to a subset of conditional norms in the premise set $\mathcal{S}$, i.e., there are

- $\phi_i \supset \mathsf{N}_i\psi_i$ where $1 \leq i \leq n$ and $\mathsf{N}_i \in \{\mathsf{O}, \mathsf{P}\}$ and
- $\lambda_j \supset \mathsf{N}'_j\delta_j$ where $1 \leq j \leq m$ and $\mathsf{N}'_j \in \{\mathsf{O}, \mathsf{P}\}$

such that

- $\vdash_{\mathsf{SDL}} (\phi \supset \mathsf{N}\psi) \equiv \bigwedge_{i=1}^{n}(\phi_i \supset \mathsf{N}_i\psi_i)$ and
- $\vdash_{\mathsf{SDL}} (\lambda \supset \mathsf{N}'\delta) \equiv \bigwedge_{j=1}^{m}(\lambda_j \supset \mathsf{N}'_j\delta_j)$.

The following example demonstrates such a case.

Example 12 Consider the following set of formulas in $\mathcal{L}_{\mathsf{SDL}}$:

$$\mathcal{S}_{\mathsf{spei}} = \left\{ \begin{array}{llll} b \vee b', & a, & b \supset \mathsf{O}\neg c, & b' \supset \mathsf{O}\neg c, \\ b' \wedge a \supset \mathsf{O}c, & b \wedge a \supset \mathsf{O}d, & b' \wedge a \supset \mathsf{O}d & b \wedge a \supset \mathsf{O}c \end{array} \right\}.$$

First, we note that $\mathcal{S}_{\mathsf{spei}}$ is free of OOSpeUcut-attacks, because neither of the conditional obligations in $\mathcal{S}_{\mathsf{spei}}$ is triggered by the facts. However, we have an intuitive case of specificity on the level of the two derived conditional obligations $(b \vee b') \wedge a \supset \mathsf{O}(c \wedge d)$ and $(b \vee b') \supset \mathsf{O}\neg c$. To see this note that:

$$\vdash_{\mathsf{SDL}} [(b \supset \mathsf{O}\neg c) \wedge (b' \supset \mathsf{O}\neg c)] \equiv [(b \vee b') \supset \mathsf{O}\neg c]$$

$$\vdash_{\mathsf{SDL}} [(b \wedge a \supset \mathsf{O}c) \wedge (b' \wedge a \supset \mathsf{O}c) \wedge (b \wedge a \supset \mathsf{O}d) \wedge (b' \wedge a \supset \mathsf{O}d)] \\ \equiv [(b \vee b') \wedge a \supset \mathsf{O}(c \wedge d)].$$

It seems intuitive to derive here $\mathsf{O}(c \wedge d)$ rather than $\mathsf{O}\neg c$, since the antecedent of the derivable conditional obligation $(b \vee b') \wedge a \supset \mathsf{O}(c \wedge d)$ is more specific than the antecedent of the derivable conditional obligation $(b \vee b') \supset \mathsf{O}\neg c$.

Motivated by our discussion above, we propose a refinement NN'SpeIUcut ('Specificity for Indirect Undercut') of the specificity rule NN'SpeUcut.

Let $\mathsf{N}, \mathsf{N}', \mathsf{N}_i, \mathsf{N}'_j \in \{\mathsf{O}, \mathsf{P}\}$ for $1 \leq i \leq n$, $1 \leq j \leq m$. Then:[18]

[18]Due to the length of the rule it is presented in a subdivided form.

Attacker: $\Gamma,\ \phi_1 \supset \mathsf{N}_1\psi_1,\ \ldots,\ \phi_n \supset \mathsf{N}_n\psi_n \Rightarrow \neg(\lambda \supset \mathsf{N}\delta)$

Conditions:

$\Rightarrow (\phi \supset \mathsf{N}\psi) \equiv \bigwedge_{i=1}^{n}(\phi_i \supset \mathsf{N}_i\psi_i)$
$\Rightarrow (\lambda \supset \mathsf{O}\delta) \equiv \bigwedge_{j=1}^{m}(\lambda_j \supset \mathsf{N}'_j\delta_j)$
$\Rightarrow \Gamma \supset \phi$
$\Rightarrow \phi \supset \lambda$
$\Rightarrow \psi \supset \neg\delta$

Attacked: $\Gamma',\ \lambda_1 \supset \mathsf{N}'_1\delta_1,\ \ldots,\ \lambda_m \supset \mathsf{N}'_m\delta_m \Rightarrow \theta$

Note, first, that OOSpeUcut is a special case of the new rule above. Now, as can easily be seen, in Example 12 we have that

$$b \vee b', a, b \wedge a \supset \mathsf{O}c, b' \wedge a \supset \mathsf{O}c, b \wedge a \supset \mathsf{O}d, b' \wedge a \supset \mathsf{O}d \Rightarrow \neg((b \vee b') \supset \mathsf{O}\neg c)$$

OOSpeIUcut-attacks the sequent $b \vee b', b \supset \mathsf{O}\neg c, b' \supset \mathsf{O}\neg c \Rightarrow \mathsf{O}\neg c$.

6 Enhanced Types of Attacks

6.1 Prioritized Attacks

In Section 4 we have shown that prioritized information may be properly represented in our framework by augmenting the language with indices attached to obligations and revising some (inference and attack) rules accordingly (see Proposition 2). In this section we consider situations in which not only the raw data should be prioritized, but also it is necessary to establish priorities among the attack rules. As the next example shows, this may be necessary when one wants to handle both specificity cases and other types of conflicts that are not resolvable by specificity considerations.

Example 13 Consider an evolving system that integrates different norms and attack rules formed by different agents (or by any other sources of information). To illustrate some of the challenges of this process and ways to handle them in our setting we reexamine the following sets of assertions, considered previously in our examples.

- Let $\mathcal{S}_{\mathsf{asp}} = \{m,\ s,\ m \Rightarrow \mathsf{O}\neg f,\ (m \wedge s) \Rightarrow \mathsf{O}f\}$ be the set of assertions from Example 6 (we renamed the formula 'a' in that example by 's', to avoid confusion with another formula called 'a' that we shall use shortly). It is shown in that example that $\mathsf{O}f$ follows from $\mathcal{S}_{\mathsf{asp}}$ in a normative argumentation framework whose sole attack rule is OOSpeUcut.
- Let $\mathcal{S}_{\mathsf{bi}} = \{a,\ b,\ a \Rightarrow \mathsf{O}(c \wedge d),\ b \Rightarrow \mathsf{O}(\neg c \wedge d)\}$ be the set of assertions discussed in Example 10. Similar considerations to those in Example 4 show that in the normative argumentation framework that is based on $\mathcal{S}_{\mathsf{bi}}$ and the attack rule OUcut we get that $\mathcal{S}_{\mathsf{bi}}\ {}_{\mathsf{Stb}}\!\mid\!\sim^{\cap}_{\mathsf{OUcut}}\ \mathsf{O}d$.

Now, to combine these two normative frameworks, we consider the union $\mathcal{S}_{\mathsf{asp}} \cup \mathcal{S}_{\mathsf{bi}}$ together with the attack rules OOSpeUcut and OUcut. Note that this time $\mathsf{O}f$ does *not* follow from $\mathcal{S}_{\mathsf{asp}} \cup \mathcal{S}_{\mathsf{bi}}$ anymore, because the sequent

$$m,\ s,\ (m \wedge s) \supset \mathsf{O}f \Rightarrow \neg(m \supset \mathsf{O}\neg f)$$

that $\mathsf{OSpeUcut}$-attacks arguments like $m, m \supset \mathsf{O}\neg f \Rightarrow \mathsf{O}\neg f$ (and so allows in Example 6 to infer $\mathsf{O}f$ rather than $\mathsf{O}\neg f$) is now OUcut-attacked by the argument

$$m,\ s,\ m \supset \mathsf{O}\neg f \Rightarrow \neg((m \wedge s) \supset \mathsf{O}f).$$

One way to resolve the problem raised in Example 13 is to give precedence to $\mathsf{OOSpeUcut}$-attacks over OUcut-attacks. This may be formalized for the general case as follows:

Definition 10 *Let $\mathcal{AF}_{\mathcal{R}}(\mathcal{S}) = \langle \mathsf{Arg}(\mathcal{S}), Attack\rangle$ be a normative argumentation framework, $A, B \in \mathsf{Arg}(\mathcal{S})$, and $\preceq$ a partial order on $\mathcal{R}$.*

- *Let $R \in \mathcal{R}$. We say that an argument A $\langle R, \preceq\rangle$*-attacks *$B$, if A R-attacks B and there is no $R' \in \mathcal{R}$ such that $R \prec R'$ and B R'-attacks A.*
- *We say that A $\langle \mathcal{R}, \preceq\rangle$*-attacks *$B$ if there is some $R \in \mathcal{R}$ for which A $\langle R, \preceq\rangle$-attacks B. Accordingly, $Attack^{\preceq}$ is set of pairs (A, B) such that A $\langle \mathcal{R}, \preceq\rangle$-attacks B.*
- *The* rule-prioritized normative argumentation framework *that is induced by $\mathcal{AF}_{\mathcal{R}}(\mathcal{S})$ and $\preceq$, is the normative argumentation framework $\mathcal{AF}^{\preceq}_{\mathcal{R}}(\mathcal{S}) = \langle \mathsf{Arg}(\mathcal{S}), Attack^{\preceq}\rangle$.*

Reasoning with rule-prioritized normative frameworks is illustrated in Figure 8, where $\mathsf{OOUcut} \prec \mathsf{OOSpeUcut}$. Whereas the diagram on the left features all the OOUcut and $\mathsf{OOSpeUcut}$-attacks, the diagram on the right only shows the prioritized attacks. According to this diagram we have two preferred extensions: $\{A, D\}$ and $\{B, D\}$. And indeed, $\mathsf{O}f$ is derivable, as is expected.

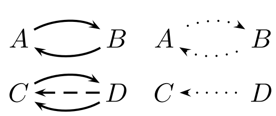

$$A = a, b, a \supset \mathsf{O}(c \wedge d) \Rightarrow \neg(b \supset \mathsf{O}(\neg c \wedge d))$$
$$B = a, b, b \supset \mathsf{O}(\neg c \wedge d) \Rightarrow \neg(a \supset \mathsf{O}(c \wedge d))$$
$$C = m, s, m \supset \mathsf{O}\neg f \Rightarrow \neg(m \wedge s \supset \mathsf{O}f)$$
$$D = m, s, m \wedge s \supset \mathsf{O}f \Rightarrow \neg(m \supset \mathsf{O}\neg f)$$

Figure 8: Left: Attack diagram for Example 13 (solid arrows are OOUcut-attacks, dashed arrows are $\mathsf{OOSpecUcut}$-attacks). Right: Attack diagram for the prioritized argumentation framework.

6.2 Restricted Counter-Attacks

Another refinement that may be needed when several attack rules are used is concerned with the notions of defense, that is: the way an argument A counter-attacks an argument B that attacks C (and so A defends C). To see this consider the next example.[19]

[19] Related issues are also discussed under the name *preemption*: A preempts the attack of B on C (see e.g. [37], or [30] in the context of inheritance networks).

Example 14 Let us consider a variant of Example 6. Suppose that beside the obligation not to eat with fingers we have the permission to do so in case asparagus is served, but it is considered impolite to eat asparagus with fingers if there is a guest who considers this rude. The enriched set of premises may look as follows:

$$\mathcal{S}_{\mathsf{rein}} = \{a,\ m,\ c,\ m \supset \mathsf{O}\neg f,\ (m \wedge a) \supset \mathsf{P}f,\ (m \wedge a \wedge c) \supset \mathsf{O}\neg f\}.$$

The situation is depicted in Figure 9, where the attack rules are $\mathsf{OPSpeUcut}$ (dotted arrows) and $\mathsf{POSpeUcut}$ (dashed arrows).

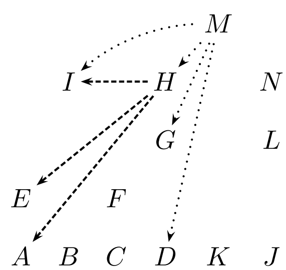

$$
\begin{aligned}
&D = (m \wedge a) \supset \mathsf{P}f \Rightarrow (m \wedge a) \supset \mathsf{P}f\\
&G = m, a, (m \wedge a) \supset \mathsf{P}f \Rightarrow \mathsf{P}f\\
&H = m, a, (m \wedge a) \supset \mathsf{P}f \Rightarrow \neg(m \supset \mathsf{O}\neg f)\\
&I = m, a, m \supset \mathsf{O}\neg f, (m \wedge a) \supset \mathsf{P}f \Rightarrow \mathsf{O}\bot\\
&J = c \Rightarrow c\\
&K = (m \wedge a \wedge c) \supset \mathsf{O}\neg f \Rightarrow (m \wedge a \wedge c) \supset \mathsf{O}\neg f\\
&L = m, a, c \Rightarrow m \wedge a \wedge c\\
&M = m, a, c, (m \wedge a \wedge c) \supset \mathsf{O}\neg f \Rightarrow \neg((m \wedge a) \supset \mathsf{P}f))\\
&N = m, a, c, (m \wedge a \wedge c) \supset \mathsf{O}\neg f \Rightarrow \mathsf{O}\neg f
\end{aligned}
$$

Figure 9: A normative argumentation framework for Example 14 (arguments A, B, C, E, F are as in Figure 3)

We have that $\mathcal{S}_{\mathsf{rein}} \mid\!\sim \mathsf{O}\neg f$ (as expected), since N is defended, while G is not. Note that arguments A and E are also defended, since their only attacker H is attacked by the defended M. In argumentation theory A and E are said to be *reinstated*.

While reinstatement seems reasonable in this example, it seems less reasonable if we alter the premise $(m \wedge a \wedge c) \supset \mathsf{O}\neg f$ to $(m \wedge a \wedge c) \supset \mathsf{P}\neg f$ and the premise $(m \wedge a) \supset \mathsf{P}f$ to $(m \wedge a) \supset \mathsf{O}f$. In that case the argument H in Figure 9 becomes $m, a, (m \wedge a) \supset \mathsf{O}f \Rightarrow \neg(m \supset \mathsf{O}\neg f)$, M becomes $m, a, c, (m \wedge a \wedge c) \supset \mathsf{P}\neg f \Rightarrow \neg((m \wedge a) \supset \mathsf{O}f)$, the attack of H on A and E becomes a $\mathsf{OOSpeUcut}$ and the attack of M on H becomes a $\mathsf{POSpeUcut}$. Note that in the most specific context we are only *allowed* to not eat with fingers, while in the least specific context we are *obliged* to do so. So, if we reinstate the least specific conditional norm via the attack of M on H, we at the same time introduce a stronger norm (the obligation $\mathsf{O}\neg f$) than what is expressed by the reinstating argument (the permission $\mathsf{P}\neg f$). Altogether, this indicates that an argument should not be defended from an $\mathsf{OOSpeUcut}$-attack by means of a $\mathsf{POSpeUcut}$-attack on its attacker.

The last example suggests that defense may be restricted to specific attack rules. This may be done as follows: Given a set $\mathcal{R}$ of attack rules, one may associate each $R \in \mathcal{R}$ with a set of rules $\mathcal{R}' \subseteq \mathcal{R}$ via a function $\pi : \mathcal{R} \to \wp(\mathcal{R})$. Intuitively speaking, $\pi(R)$ contains all the rules in $\mathcal{R}$ that are deemed sufficient to defend an argument which is R-attacked: A set $\mathcal{E}$ of arguments defends an argument A iff for every R-attack from an argument C on A, there is an argument $B \in \mathcal{E}$ that R'-attacks C for some $R' \in \pi(R)$.

Dung-style extensions may now be defined according to Definition 5, where the notion of defense in that definition is replaced by the refined notion of defense defined above.

Example 15 In Example 14, if we let $\mathsf{POSpeUcut} \notin \pi(\mathsf{OOSpeUcut})$, then M does not defend A (nor E) from H, although it still attacks H. While with reinstatement A, E and M would be in the grounded extension, now only M is. As before, H is not in the grounded extension since it is attacked by M.

7 Relations to I/O Logics

In this section we consider some relations to Input/Output logics [33] (I/O logics, for short). These logics represent standard approaches in conditional deontic logic, and many of them have simple and intuitive syntactic as well as semantic characterizations. The framework of I/O logics has been extended in order to deal with conflicts among conditionals such as in Contrary-to-Duty obligations [34] and for modeling permissions [44]. Moreover, there are links to other frameworks such as default logic (see [34]), logic programming [26], and adaptive logics [49]. In view of this it is interesting to note that I/O logics can also be related to our framework, as will be established below (see Proposition 4, Corollary 1, Proposition 6, and Corollary 2).

In the following we focus on premise sets $\mathcal{S}$ that consist of non-modal formulas (representing 'facts' or 'input') and formulas of the type $\phi \supset \mathsf{O}\psi$ (representing conditional obligations). For this, we consider the following sets:

- $\mathcal{S}_F$ – a set of non-modal propositional formulas,
- $\mathcal{S}_O$ – a set of pairs of non-modal formulas (ψ, ϕ) ('I/O-pairs'),
- $\mathcal{S}_O^* = \{\psi \supset \mathsf{O}\phi \mid (\psi, \phi) \in \mathcal{S}_O\}$.

The basic idea behind Input/Output logic is to treat I/O-pairs (ϕ, ψ) as the name suggests: if the input ϕ is derivable in classical logic from the input set $\mathcal{S}_F$, then generate the output ψ. This is realized in view of the following definition.

Definition 11 *Denote by* $\vdash_{\mathsf{CPL}}$ *the consequence relation of classical logic. We define:* $\mathsf{out}(\mathcal{S}_F, \mathcal{S}_O) = \{\psi \mid (\phi, \psi) \in \mathcal{S}_O,\ \mathcal{S}_F \vdash_{\mathsf{CPL}} \phi\}$.

This idea is still too simple to give rise to interesting output functions. For instance, in many applications,

- if we have (ϕ, ψ_1) and (ϕ, ψ_2) and $\mathcal{S}_F \vdash_{\mathsf{CPL}} \phi$ then we would like to get $\psi_1 \wedge \psi_2$ as output besides ψ_1 and ψ_2, or
- if we have (ϕ, ψ), $\phi \in \mathcal{S}_F$, and $\psi \vdash_{\mathsf{CPL}} \psi'$ then we would like to get ψ' as output.

More generally, we would expect that applying $\mathrm{Cn}_{\mathsf{CPL}}(\mathcal{S})$ to the output set $\mathcal{S}$ doesn't produce new outputs. Additionally, in many applications it seems desirable to get the output ψ from (ϕ_1, ψ) and (ϕ_2, ψ) with the input $\mathcal{S}_F = \{\phi_1 \vee \phi_2\}$. In order to fulfil these desiderata, we need a more refined mechanism. One option is given by the following definition:

Definition 12 *Given $\mathcal{S}_F$ and $\mathcal{S}_O$ as defined above, the set* out_2 *is defined as follows (where* $\mathrm{Cn}_{\mathsf{CPL}}(\bullet)$ *is defined as in Proposition 1):*

- *If $\mathcal{S}_F$ is classically consistent, then* $\mathsf{out}_2(\mathcal{S}_F, \mathcal{S}_O) = \bigcap\{\mathrm{Cn}_{\mathsf{CPL}}(\mathsf{out}(\mathcal{T}, \mathcal{S}_O)) \mid \mathcal{T}$ *is a* CPL*-maximal consistent set for which* $\mathcal{S}_F \subseteq \mathcal{T}\}$.
- *If $\mathcal{S}_F$ is classically inconsistent, then* $out_2(\mathcal{S}_F, \mathcal{S}_O) = \mathrm{Cn}_{\mathsf{CPL}}(\{\psi \mid (\phi, \psi) \in \mathcal{S}_O\})$.

The following is a corollary of Observation 4 in [33].[20] It offers an alternative modal characterization of out_2.

Lemma 3 *If $\mathcal{S}_F$ is classically consistent,* $\mathcal{S}_F \cup \mathcal{S}_O^* \vdash_{\mathsf{SDL}} \mathsf{O}\phi$ *iff* $\phi \in \mathsf{out}_2(\mathcal{S}_F, \mathcal{S}_O)$.

In order to deal with situations in which $\mathsf{out}(\mathcal{S}_F, \mathcal{S}_O)$ is inconsistent, Makinson and Van Der Torre [34] 'contextualize' their output-functions to maximal sets of conditionals that are consistent with $\mathcal{S}_F$, so-called maxfamilies:[21]

Definition 13 *Given $\mathcal{S}_F$ and $\mathcal{S}_O$ as defined above, we consider the following sets:*

- $\mathcal{T}_O \in \mathsf{maxfamily}_2(\mathcal{S}_F, \mathcal{S}_O)$ *iff* $\mathsf{out}_2(\mathcal{S}_F, \mathcal{T}_O)$ *is classically consistent and for all* $(\psi, \phi) \in \mathcal{S}_O \setminus \mathcal{T}_O$*, the set* $\mathsf{out}_2(\mathcal{S}_F, \mathcal{T}_O \cup \{(\psi, \phi)\})$ *is not classically consistent.*
- $\psi \in \mathsf{out}_2^{\cup}(\mathcal{S}_F, \mathcal{S}_O)$ *iff* $\psi \in \bigcup_{\mathcal{T}_O \in \mathsf{maxfamily}_2(\mathcal{S}_F, \mathcal{S}_O)} \mathsf{out}_2(\mathcal{S}_F, \mathcal{T}_O)$.
- $\psi \in \mathsf{out}_2^{\cap}(\mathcal{S}_F, \mathcal{S}_O)$ *iff* $\psi \in \bigcap_{\mathcal{T}_O \in \mathsf{maxfamily}_2(\mathcal{S}_F, \mathcal{S}_O)} \mathsf{out}_2(\mathcal{S}_F, \mathcal{T}_O)$.

We now show that in our argumentative approach the Input/Output logics in Definition 13 can be characterized by means of the attack rule OUcut.

Lemma 4 *If $\mathcal{S}_F$ is classically consistent and* $\mathcal{T}_O \in \mathsf{maxfamily}_2(\mathcal{S}_F, \mathcal{S}_O)$ *then* $\mathsf{Arg}(\mathcal{S}_F \cup \mathcal{T}_O^*)$ *is a stable extension of* $\mathcal{AF}_{\mathsf{OUcut}}(\mathcal{S}_F \cup \mathcal{S}_O^*)$.

Proof. Let $\mathcal{T}_O \in \mathsf{maxfamily}_2(\mathcal{S}_F, \mathcal{S}_O)$. By Lemma 3, $\mathcal{S}_F \cup \mathcal{T}_O^*$ is SDL-consistent and hence $\mathsf{Arg}(\mathcal{S}_F \cup \mathcal{T}_O^*)$ is conflict-free. Let $A \in \mathsf{Arg}(\mathcal{S}_F \cup \mathcal{S}_O^*) \setminus \mathsf{Arg}(\mathcal{S}_F \cup \mathcal{T}_O^*)$. This means that there is a $\psi \supset \mathsf{O}\phi \in \mathsf{Prem}(A) \cap (\mathcal{S}_O^* \setminus \mathcal{T}_O^*)$. Since $\mathsf{out}_2(\mathcal{S}_F, \mathcal{T}_O \cup \{(\psi, \phi)\})$ is classically inconsistent, we have by Lemma 3 that $\mathcal{S}_F \cup \mathcal{T}_O^* \cup \{\psi \supset \mathsf{O}\phi\}$ is SDL-inconsistent. Thus, there is a finite $\Gamma \subseteq \mathcal{S}_F \cup \mathcal{T}_O^*$ such that $\Gamma, \psi \supset \mathsf{O}\phi \Rightarrow \bot$ is $\mathcal{C}_{\mathsf{SDL}}$-provable. It follows that $B = \Gamma \Rightarrow \neg(\psi \supset \mathsf{O}\phi)$ is also $\mathcal{C}_{\mathsf{SDL}}$-provable, and so $B \in \mathsf{Arg}(\mathcal{S}_F \cup \mathcal{T}_O^*)$. Thus, B OUcut-attacks A.

[20] In [33] the correspondence is shown for all modal logics L for which $\vdash_{\mathsf{K}} \subseteq \vdash_{\mathsf{L}} \subseteq \vdash_{\mathsf{K45}}$.

[21] The approach in [34] is more general since it takes into account additional constraints beside the requirement of consistency.

We have thus shown that every argument not in $\mathsf{Arg}(\mathcal{S}_F \cup \mathcal{T}_O^*)$ is attacked by an argument in $\mathsf{Arg}(\mathcal{S}_F \cup \mathcal{T}_O^*)$, therefore $\mathsf{Arg}(\mathcal{S}_F \cup \mathcal{T}_O^*)$ is a stable extension of $\mathcal{AF}_{\mathsf{OUcut}}(\mathcal{S}_F \cup \mathcal{S}_O^*)$. $\square$

Lemma 5 *If $\mathcal{S}_F$ is classically consistent then for every conflict-free set of arguments $\mathcal{E}$ of the normative framework $\mathcal{AF}_{\mathsf{OUcut}}(\mathcal{S}_F \cup \mathcal{S}_O^*)$ there is a $\mathcal{T}_O \in \mathsf{maxfamily}_2(\mathcal{S}_F, \mathcal{S}_O)$ such that $\mathcal{E} \subseteq \mathsf{Arg}(\mathcal{S}_F \cup \mathcal{T}_O^*)$.*

Proof. Suppose that $\mathcal{E}$ is a set of arguments of $\mathcal{AF}_{\mathsf{OUcut}}(\mathcal{S}_F \cup \mathcal{S}_O^*)$ and there is no $\mathcal{T}_O \in \mathsf{maxfamily}_2(\mathcal{S}_F, \mathcal{S}_O)$ for which $\mathcal{E} \subseteq \mathsf{Arg}(\mathcal{S}_F \cup \mathcal{T}_O^*)$. Hence, there is no $\mathcal{T}_O \in \mathsf{maxfamily}_2(\mathcal{S}_F, \mathcal{S}_O)$ for which $\mathcal{T}_\mathcal{E} \subseteq \mathcal{T}_O$ where $\mathcal{T}_\mathcal{E} = \bigcup_{A \in \mathcal{E}} \{(\psi, \phi) \mid \psi \supset \mathsf{O}\phi \in \mathsf{Prem}(A)\}$. This means that $\mathsf{out}_2(\mathcal{S}_F, \mathcal{T}_\mathcal{E})$ is classically inconsistent. By Lemma 3, $\mathcal{S}_F \cup \mathcal{T}_\mathcal{E}^*$ is SDL-inconsistent. Hence, there are finite $\Gamma_F \subseteq \mathcal{S}_F$ and $\Gamma_O \subseteq \mathcal{T}_\mathcal{E}$ such that $A = \Gamma_F, \Gamma_O^* \Rightarrow \bot \in \mathsf{Arg}(\mathcal{S}_F \cup \mathcal{T}_\mathcal{E}^*)$. Since $\mathcal{S}_F$ is classically consistent, $\Gamma_O \neq \emptyset$. With Weakening and $[\Rightarrow\supset]$ we have an argument $B = \Gamma_F, \Gamma_O^* \setminus \{\psi \supset \mathsf{O}\phi\} \Rightarrow \neg(\psi \supset \mathsf{O}\phi) \in \mathsf{Arg}(\mathcal{S}_F \cup \mathcal{T}_\mathcal{E}^*)$, where $(\psi, \phi) \in \Gamma_O$. Since B OUcut-attacks A, $\mathsf{Arg}(\mathcal{S}_F \cup \mathcal{T}_\mathcal{E}^*)$ is not conflict-free. $\square$

By the last two lemmas, we immediately get:

Proposition 4 *If $\mathcal{S}_F$ is classically consistent, then:*

$$\begin{aligned}\mathsf{Stb}(\mathcal{AF}_{\mathsf{OUcut}}(\mathcal{S}_F \cup \mathcal{S}_O^*)) &= \mathsf{Prf}(\mathcal{AF}_{\mathsf{OUcut}}(\mathcal{S}_F \cup \mathcal{S}_O^*)) \\ &= \{\mathsf{Arg}(\mathcal{S}_F \cup \mathcal{T}_O^*) \mid \mathcal{T}_O \in \mathsf{maxfamily}_2(\mathcal{S}_F, \mathcal{S}_O)\}.\end{aligned}$$

Corollary 1 *If the only attack rule is OUcut, then for every $\lambda \in \{\cup, \cap\}$ it holds that $\psi \in \mathsf{out}_2^\lambda(\mathcal{S}_F, \mathcal{S}_O)$ iff $\mathcal{S}_F \cup \mathcal{S}_O^* \ {}_{\mathsf{prf}}\!\mid\!\sim^{\lambda}_{\mathsf{OUcut}} \mathsf{O}\psi$ iff $\mathcal{S}_F \cup \mathcal{S}_O^* \ {}_{\mathsf{stb}}\!\mid\!\sim^{\lambda}_{\mathsf{OUcut}} \mathsf{O}\psi$.*

Example 16 Let $\mathcal{S}_O = \{(p_1, q_1 \wedge q_2), (p_2, \neg q_1 \wedge q_2)\}$ and $\mathcal{S}_F = \{p_1, p_2\}$. Then:

$$\mathsf{maxfamily}_2(\mathcal{S}_F, \mathcal{S}_O) = \{\{(p_1, q_1 \wedge q_2)\}, \{(p_2, \neg q_1 \wedge q_2)\}\}.$$

Since $q_2 \in \mathsf{out}_2(\mathcal{S}_F, \{(p_1, q_1 \wedge q_2)\}) \cap \mathsf{out}_2(\mathcal{S}_F, \{(p_2, \neg q_1 \wedge q_2)\})$, we have that $q_2 \in \mathsf{out}_2^\cap(\mathcal{S}_F, \mathcal{S}_O)$.

In the normative argumentation framework $\mathcal{AF}_{\mathsf{OUcut}}(\mathcal{S}_F, \mathcal{S}_O^*)$ we have two preferred extensions:

1. An extension with e.g. the following arguments:
   $p_1, p_1 \supset \mathsf{O}(q_1 \wedge q_2) \Rightarrow \neg(p_2 \supset \mathsf{O}(\neg q_1 \wedge q_2))$,
   $p_1, p_1 \supset \mathsf{O}(q_1 \wedge q_2) \Rightarrow \mathsf{O}(q_1 \wedge q_2)$,
   $p_1, p_1 \supset \mathsf{O}(q_1 \wedge q_2) \Rightarrow \mathsf{O}q_2$.
2. An extension with e.g. the following arguments:
   $p_2, p_2 \supset \mathsf{O}(\neg q_1 \wedge q_2) \Rightarrow \neg(p_1 \supset \mathsf{O}(q_1 \wedge q_2))$,
   $p_2, p_2 \supset \mathsf{O}(\neg q_1 \wedge q_2) \Rightarrow \mathsf{O}(\neg q_1 \wedge q_2)$,
   $p_2, p_2 \supset \mathsf{O}(\neg q_1 \wedge q_2) \Rightarrow \mathsf{O}q_2$.

Thus, $\mathcal{S}_F \cup \mathcal{S}_O^* \ {}_{\mathsf{stb}}\!\mid\!\sim^{\cap}_{\mathsf{OUcut}} \mathsf{O}q_2$.

The result above provides us also a representational result for another I/O-logic based on out_4.

Definition 14 *Given $\mathcal{S}_F$ and $\mathcal{S}_O$ as defined above, the set* out_4 *is defined as follows:*

- *If $\mathcal{S}_F$ is consistent, then $\phi \in \mathsf{out}_4(\mathcal{S}_F, \mathcal{S}_O)$ iff $\phi \in Cn_{\mathsf{CPL}}(\mathsf{out}(\mathcal{T}, \mathcal{S}_O))$ for all maximal consistent set $\mathcal{T}$ for which $\mathcal{S}_F \subseteq \mathcal{T} \subseteq \mathsf{out}(\mathcal{T}, \mathcal{S}_O)$.*
- *If $\mathcal{S}_F$ is inconsistent, then $\mathsf{out}_4(\mathcal{S}_F, \mathcal{S}_O) = Cn_{\mathsf{CPL}}(\{\psi \mid (\phi, \psi) \in \mathcal{S}_O\})$.*

The definitions of $\mathsf{maxfamily}_4$, $\mathsf{out}_4^{\cap}$ and $\mathsf{out}_4^{\cup}$ are analogous to Definition 13 (where out_2 is replaced by out_4).

The main difference between out_4 and out_2 is that with out_4 we get the following transitivity property: $\phi \in \mathsf{out}_4(\{\psi\}, \{(\psi, \theta), (\psi \wedge \theta, \phi)\})$. This property does not hold for out_2: $\phi \notin \mathsf{out}_2(\{\psi\}, \{(\psi, \theta), (\psi \wedge \theta, \phi)\})$. Intuitively speaking, the generated output θ from ψ and (ψ, θ) is reused as input for $(\psi \wedge \theta, \phi)$ to generate ϕ. This is why Makinson and Van Der Torre call this approach also *reusable output*.

In [33] it has been shown that out_4 can be characterized as follows:[22]

Proposition 5 *$out_4(\mathcal{S}_F, \mathcal{S}_O) = out_2(\mathcal{S}_F \cup m(\mathcal{S}_O), \mathcal{S}_O)$ where $m(\mathcal{S}_O)$ is the materialization of $\mathcal{S}_O$ (i.e., each (ϕ, ψ) in $\mathcal{S}_O$ translates to $\phi \supset \psi$ in $m(\mathcal{S}_O)$)*

In view of this, we have the following results:

Proposition 6 *If $\mathcal{S}_F$ is classically consistent, then:*

$$\begin{aligned}\mathsf{Stb}(\mathcal{AF}_{\mathsf{OUcut}}(\mathcal{S}_F \cup \mathcal{S}_O^* \cup m(\mathcal{S}_O))) &= \mathsf{Prf}(\mathcal{AF}_{\mathsf{OUcut}}(\mathcal{S}_F \cup \mathcal{S}_O^* \cup m(\mathcal{S}_O))) \\ &= \{\mathsf{Arg}(\mathcal{S}_F \cup \mathcal{T}_O^*) \mid \mathcal{T}_O \in \mathsf{maxfamily}_4(\mathcal{S}_F, \mathcal{S}_O)\}.\end{aligned}$$

Proof. Immediate from Propositions 4 and 5. □

Corollary 2 *If the only attack rule is* OUcut*, then for every $\lambda \in \{\cup, \cap\}$ it holds that $\psi \in \mathsf{out}_4^{\lambda}(\mathcal{S}_F, \mathcal{S}_O)$ iff $\mathcal{S}_F \cup \mathcal{S}_O^* \cup m(\mathcal{S}_O) \,_{\mathsf{stb}}{\mid\!\sim}^{\lambda}_{\mathsf{OUcut}}\, \mathsf{O}\psi$ iff $\mathcal{S}_F \cup \mathcal{S}_O^* \cup m(\mathcal{S}_O) \,_{\mathsf{prf}}{\mid\!\sim}^{\lambda}_{\mathsf{OUcut}}\, \mathsf{O}\psi$.*

8 Discussion and Outlook

The idea to use argumentation in general and abstract argumentation in particular to model normative reasoning is not new. Two examples of this are described in [21] and [39]. The approach in [21] is based on bipolar abstract argumentation frameworks [15], in which beside an attack relation a support relation is used to express conditional obligations. In [39] Dung's framework is also enhanced by a support relation, this time signifying evidential support. Prolog-like predicates are used to encode argument schemes [54] of normative

[22] In [33, Observation 15] it has also been shown that $\phi \in out_4(\mathcal{S}_F, \mathcal{S}_O)$ iff $\mathcal{S}_F \cup \mathcal{S}_O^* \vdash_{\mathsf{LT}} O\phi$, where $\mathcal{S}_F$ is consistent and LT is the result of enriching any logic L for which $\vdash_{\mathsf{K}} \subseteq \vdash_{\mathsf{L}} \subseteq \vdash_{\mathsf{K45}}$ by (T) (i.e., $\vdash O\phi \supset \phi$). It is easy to see that if we enrich SDL with (T) in Corollary 1 we get a representational result for out_4.

reasoning and an algorithm is provided to translate them into an argumentation framework. One of the main differences in our approach is that we use a base logic (SDL) that generates all the given arguments (on the basis of a premise set). As a consequence an additional support relation is not needed since argumentative support is intrinsically modeled by considering arguments as sequents in SDL. A by-product of this is that our approach is more tightly linked to deontic logic.

The advantage of our approach is twofold:

1. It offers a formal explication of an argumentative account of defeasible reasoning in terms of considerations and counter-considerations (see, e.g., the work of Mercier and Sperber, [36], where it is argued for the adequacy of such a perspective for human reasoning). Moreover, the results in [4] (as well as some other results that are in preparation) indicate that dynamic proofs for sequent-based argumentation may be useful to automate normative reasoning as modeled in this paper. Combined together, these works may provide a procedural account of defeasible normative reasoning.
2. Our framework is expressive enough to distinguish and combine different types of defeasibility in terms of different types of argumentative attacks. For instance, we have shown that in the argumentative setting we have the resources to easily enhance the argumentation framework that characterized one of the I/O-logics with considerations concerning specificity.

Deontic logicians mainly agree that modeling conditional obligations on the basis of SDL and material implication is futile due to problems with CTD-obligations and specificity [2]. Therefore, more research interest has been directed towards bi-conditionals. Specificity considerations sometimes call for weakened principles of strengthening the antecedent which are still strong enough to support many intuitively valid inferences. For instance, the principle of Rational Monotonicity has been challenged in [23] and replaced by a weakened version which itself has been criticized in [45]. In contrast, SDL uses the standard implication of classical propositional logic to model conditional obligations and allows for full strengthening of the antecedent. As demonstrated e.g. in Section 5, undesired applications of the latter are avoided by means of argumentative attacks that are triggered e.g. in cases of specificity. As a consequence, our consequence relations are non-monotonic. Other non-monotonic accounts of normative reasoning are based on default logic [32], Input/Output logic [34], Nute's defeasible logic [37, 38], and adaptive logics [8, 24, 35, 46].

We have already seen that consequence relations based on considerations concerning maximally consistent subsets, both in the context of non-conditional obligations (see Propositions 1, 2, 3) and in the context of conditional obligations (see our discussion of I/O-logics and the related results in Corollaries 1 and 2) can be represented in our argumentation-based approach. This is not surprising since argumentative attacks and argumentation semantics can be seen as rather refined consistency maintenance devices.

Approaches based on maximal consistent subsets, as well as our approach are *corrective* in character. They apply a rather strong core logic to the input and localize consistent chunks of the consequence set (the 'extensions'). Adaptive deontic logics proceed differently: they take a rather weak core logic as the basis for *ampliative* defeasible reasoning in which the core logic is defeasibly (i.e., non-monotonically) strengthened by additional assumptions. For example, in [46] deontic detachment is applied on the assumption that the application does not create a deontic conflict. Ampliative reasoning has not been investigated for sequent-based argumentation yet. However, techniques from [18] could be used to supplement arguments (i.e., sequents) with assumptions.

Nute's defeasible logic has been applied for the modeling of normative reasoning [37]. There, strict and defeasible rules are used to build derivations. Both types of rules are domain specific information presented by $A \rightarrow B$ (A strictly implies B) or $A \rightsquigarrow B$ (A defeasibly implies B)[23] and derivations are basically (chains of) applications of detachment to these rules. Such derivations may be blocked in view of conflicts with other derivations, depending on the strength of the respective derivations. In case that defeasible rules are involved the strength of a derivation is also measured in view of a relation that allows to compare the strength of defeasible rules. One difference to our approach is that our inference rules are not domain specific but stem from the calculus of our core logic. One of the advantages of using a sequent-based calculus is that derivations allow also for the manipulation of the antecedents of rules. For instance, from $a,\ a \supset \mathsf{O}b \Rightarrow \mathsf{O}b$ and $c,\ c \supset \mathsf{O}b \Rightarrow \mathsf{O}b$ we get $a \vee c,\ a \supset \mathsf{O}b,\ c \supset \mathsf{O}b \Rightarrow \mathsf{O}b$, which means that $\mathcal{S} \mid\!\sim \mathsf{O}b$ (where $\mathcal{S} = \{a \vee c,\ a \supset \mathsf{O}b,\ c \supset \mathsf{O}b\}$) for any of the previously defined entailment relations (see Observation 1 from Section 5 and note that $\mathcal{S}$ is SDL-consistent). On the other hand, from $a \rightsquigarrow \mathsf{O}b$ and $c \rightsquigarrow \mathsf{O}b$ one cannot get $a \vee c \rightsquigarrow \mathsf{O}b$ in defeasible logic. As a consequence, $\mathsf{O}b$ is not derivable from the set of facts $\{a \vee c\}$ and the rules $\{a \rightsquigarrow \mathsf{O}b,\ c \rightsquigarrow \mathsf{O}b\}$. In Section 5 we have seen that this additional expressive power also allows to identify more involving cases of specificity (see Example 12) and to handle them with attack rules such as $\mathsf{NN'SpelUcut}$.

In our approach we do not distinguish between strict and defeasible rules. Nevertheless, depending on the attack-rules some sequents may not be attackable in our approach (like, e.g., modality-free sequents in Example 7). Unlike our approach, the approach in [37] does not allow for argumentative defense, i.e., a sequent that attacks another sequent and that is itself attacked by a third sequent doesn't loose its attacking power. In Section 6 we have seen how we can nevertheless avoid unwanted cases of reinstatement. Moreover, our approach can be used with the full variety of argumentative semantics that have been proposed for abstract argumentation (see Note 2). As in our approach, defeasible logic handles both specificity cases and other types of normative conflicts.

In future work we plan to investigate whether other nonmonotonic approaches and non truth-functional logics can be expressed in our framework. Also, we shall examine base logics that are obtained from SDL by removing

[23] Nute uses the double-arrow $\Rightarrow$ for defeasible rules. In order to avoid ambiguities with our notation for sequents we replaced it by $\rightsquigarrow$.

some of the inference rules in $\mathcal{C}_{\mathsf{SDL}}$, and so such logics may not have deterministic matrices. Furthermore, it would be interesting to investigate other resolution principles for conflicts besides the ones discussed in this paper. For instance, "lex posterior derogat legi priori" expressing that more recent laws override older ones. In order to model this we need to express temporal information and hence enhance our language. Also, priorities among norms were only discussed in passing in the current work and deserve a more thorough investigation. For instance, we shall investigate whether Horty's approach to prioritized normative reasoning based on default logic [32] can be characterized in our framework. Another subject that remains for future work is whether a model of permission as derogation based on techniques from belief revision [44] can be expressed in an argumentative setting. Finally, we plan to study rationality postulates like those in [12] for various of our systems. Some of our results, like the characterizations for Input/Output logic, already provide good indications for the well-behavior of some of our systems, such as those with the attack rule OUcut. In that respect, systems that incorporate the attack rules OUcut and OOSpeUcut (see Section 6.1) are of a special interest, as in view of the adequacy results in Section 7 they generalize Input/Output logic by handling specificity cases.

References

[1] C. E. Alchourrón and D. Makinson. Hierarchies of regulations and their logic. In *New studies in deontic logic*, pages 125–148. Springer, 1981.

[2] L. Åqvist. Deontic logic. In D. M. Gabbay and F. Guenthner, editors, *Handbook of Philosophical Logic*, volume 8, pages 147–264. Kluwer, 2002.

[3] O. Arieli. A sequent-based representation of logical argumentation. In *Proc. CLIMA'13*, LNCS 8143, pages 69–85. Springer, 2013.

[4] O. Arieli and C. Straßer. Dynamic derivations for sequent-based deductive argumentation. In *Proc. COMMA'14*, volume 266 of *Frontiers in Artificial Intelligence and Applications*, pages 89–100. IOS Press, 2014.

[5] O. Arieli and C. Straßer. Sequent-based logical argumentation. *Journal of Argument and Computation*, 6(1):73–99, 2015.

[6] C. Baral, S. Kraus, and J. Minker. Combining multiple knowledge bases. *IEEE Transactions on Knowledge and Data Engineering*, 3(2):208–220, 1991.

[7] D. Batens. A universal logic approach to adaptive logics. *Logica Universalis*, 1:221–242, 2007.

[8] M. Beirlaen, C. Straßer, and J. Meheus. An inconsistency-adaptive deontic logic for normative conflicts. *Journal of Philosophical Logic*, 2(42):285–315, 2013.

[9] S. Benferhat, D. Dubois, and H. Prade. Some syntactic approaches to the handling of inconsistent knowledge bases: A comparative study. Part 1: The flat case. *Studia Logica*, 58:17–45, 1997.

[10] Ph. Besnard and A. Hunter. A logic-based theory of deductive arguments. *Artificial Intelligence*, 128(1–2):203–235, 2001.

[11] G. Brewka. Preferred subtheories: An extended logical framework for default reasoning. In N. S. Sridharan, editor, *Proc. IJCAI'89*, pages 1043–1048. Morgan Kaufmann, 1989.

[12] M. Caminada and L. Amgoud. On the evaluation of argumentation formalisms. *Artificial Intelligence*, 171(5):286–310, 2007.

[13] M. W. A. Caminada. Semi-stable semantics. In Dunne. P. E. and TJ. M. Bench-Capon, editors, *Proc. COMMA'06*, volume 144 of *Frontiers in Artificial Intelligence and Applications*, pages 121–130. IOS Press, 2006.

[14] M. W. A. Caminada. On the issue of contraposition of defeasible rules. In *Proc. COMMA'08*, volume 172 of *Frontiers in Artificial Intelligence and Applications 266*, pages 109–115. IOS Press, 2008.

[15] C. Cayrol and M. C. Lagasquie-Schiex. On the acceptability of arguments in bipolar argumentation frameworks. In *Proc. ECSQARU'05*, LNCS 3571, pages 378–389. Springer, 2005.

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

[17] P. M. Dung. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. *Artifical Intelligence*, 77:321–358, 1995.

[18] P. M. Dung, R. A. Kowalski, and F. Toni. Assumption-based argumentation. *Argumentation in Artificial Intelligence*, pages 199–218, 2009.

[19] P. M. Dung, P. Mancarella, and F. Toni. Computing ideal sceptical argumentation. *Artificial Intelligence*, 171(10–15):642–674, 2007.

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

[21] D. Gabbay. Bipolar argumentation frames and contrary to duty obligations, preliminary report. In *Proc. CLIMA'12*, LNCS 7486, pages 1–24. Springer, 2012.

[22] G. Gentzen. Investigations into logical deduction, 1934. (In German. An English translation appears in 'The Collected Works of Gerhard Gentzen', edited by M. E. Szabo, North-Holland, 1969).

[23] L. Goble. A proposal for dealing with deontic dilemmas. In *Proc. DEON'04*, LNCS 3065, pages 74–113. Springer, 2004.

[24] L. Goble. Deontic logic (adapted) for normative conflicts. *Logic Journal of the IGPL*, 2013.

[25] L. Goble. Prima facie norms, normative conflicts, and dilemmas. In D. Gabbay, J. Horty, X. Parent, R. van der Meyden, and L. van der Torre, editors, *Handbook of Deontic Logic and Normative Systems*, pages 241–352. College Publications, 2013.

[26] R. Gonçalves and J. J. Alferes. An embedding of input-output logic in deontic logic programs. In *Deontic Logic in Computer Science*, pages 61–75. Springer, 2012.

[27] A. Grove. Two modellings for theory change. *Journal of Philosophical Logic*, 17:157–180, 1988.

[28] J. Hansen. Deontic logics for prioritized imperatives. *Artificial Intelligence and Law*, 14(1-2):1–34, 2006.

[29] J. F. Horty. Moral dilemmas and nonmonotonic logic. *Journal of Philosophical Logic*, 23:35–65, 1994.

[30] J. F. Horty. Some direct theories of nonmonotonic inheritance. In D. M. Gabbay, C. J. Hogger, and J. A. Robinson, editors, *Handbook of Logic in Artificial Intelligence and Logic Programming, Volume 3: Nonmonotonic Reasoning and Uncertain Reasoning*, pages 111–187. Oxford University Press, 1994.

[31] J. F. Horty. Skepticism and floating conclusions. *Artifical Intelligence*, 135(1-2):55–72, 2002.

[32] J. F. Horty. Reasoning with moral conflicts. *Noûs*, 37(4):557–605, 2003.

[33] D. Makinson and L. Van Der Torre. Input/Output logics. *Journal of Philosophical Logic*, 29(4):383–408, 2000.

[34] D. Makinson and L. Van Der Torre. Constraints for Input/Output logics. *Journal of Philosophical Logic*, 30(2):155–185, 2001.

[35] J. Meheus, M. Beirlaen, and F. Van De Putte. Avoiding deontic explosion by contextually restricting aggregation. In *Proc. DEON'10*, LNAI 6181, pages 148–165. Springer, 2010.

[36] H. Mercier and D. Sperber. Why do humans reason? Arguments for an argumentative theory. *Behavioral and Brain Sciences*, 34(2):57–74, 2011.

[37] D. Nute. Apparent obligation. In *Defeasible Deontic Logic*, pages 287–315. Springer, 1997.

[38] D. Nute. Defeasible logic. In *Web Knowledge Management and Ddecision Support*, pages 151–169. Springer, 2003.

[39] N. Oren, M. Luck, S. Miles, and T.J. Norman. An argumentation inspired heuristic for resolving normative conflict. In *Proc. AAMAS-08*, pages 41–56, 2008.

[40] J. Pollock. *Cognitive Carpentry. A Blueprint for How to Build a Person.* Bradford/MIT Press, 1995.

[41] H. Prakken. An abstract framework for argumentation with structured arguments. *Argument and Computation*, 1(2):93–124, 2011.

[42] H. Prakken. Some reflections on two current trends in formal argumentation. In *Logic Programs, Norms and Action*, LNCS 7360, pages 249–272. Springer, 2012.

[43] N. Rescher and R. Manor. On inference from inconsistent premises. *Theory and Decision*, 1:179–217, 1970.

[44] A. Stolpe. A theory of permission based on the notion of derogation. *Journal of Applied Logic*, 8(1):97–113, 2010.

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

[46] C. Straßer. A deontic logic framework allowing for factual detachment. *Journal of Applied Logic*, 9(1):61–80, 2010.

[47] C. Straßer. *Adaptive Logic and Defeasible Reasoning. Applications in Argumentation, Normative Reasoning and Default Reasoning.* Trends in Logic, Volume 38. Springer, 2014.

[48] C. Straßer and O. Arieli. Sequent-based argumentation for normative reasoning. In *Proc. DEON'14*, LNAI 8554, pages 224–240. Springer, 2014.

[49] C. Straßer, M. Beirlaen, and F. Van De Putte. Dynamic proof theories for input/output logic. Under review., 2014.

[50] S. Valentini. The sequent calculus for the modal logic D. *Unione Matematica Italiana. Bollettino. A. Serie*, 7(7):3, 1993.

[51] L. Van Der Torre and Y. H. Tan. Cancelling and overshadowing: Two types of defeasibility in defeasible deontic logic. In *Proc. IJCAI'95*, pages 1525–1532, 1995.

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

[53] B. Verheij. Two approaches to dialectical argumentation: admissible sets and argumentation stages. In *Proc. NAIC'96*, pages 357–368, 1996.

[54] D. Walton, C. Reed, and F. Macagno. *Argumentation Schemes.* Cambridge University Press, 2008.