Sequent-Based Argumentation for Normative Reasoning

Christian Straßer[1] and Ofer Arieli[2]

[1] Centre for Logic and Philosophy of Science, Ghent University, Belgium, christian.strasser@UGent.be

[2] School of Computer Science, The Academic College of Tel-Aviv, Israel, oarieli@mta.ac.il

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. These 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 [1]), and use standard deontic logic as our base logic. Argumentative attacks are realized by elimination rules that allow to discharge specific sequents. We demonstrate our system by means of various well-known benchmark examples.

1 Introduction

Normative reasoning concerns 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 conditional on φ, then there is a commitment to ψ. Another instance is aggregation: if there is a 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]. For instance, aggregating two conflicting obligations leads to an obligation that commits us to do the impossible. Other problematic cases concern specificity: sometimes more specific obligations or permissions override more general ones. In such cases we want to block factual detachment from the overridden obligations or permissions. Logical accounts of normative reasoning that is tolerant with respect to normative conflicts and/or specificity cases have been shown to be challenging. This has given rise to a variety of approaches (e.g., [3–8]).

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 demonstrate how this model helps us to identify conflict-free sets that are apt to guide the actions of a user. Furthermore, we will show how it offers an elegant tool to deal with specificity cases. It follows that the entailment relations that are obtained 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 [9]. 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 (e.g., [10, 11]). In this paper, we settle for the representation in terms of *sequents* [1]. One advantage of this approach is that it immediately equips us with dynamic proof procedures in the style of adaptive logics [12, 13] that allow for automated reasoning [14]. 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.

In this paper 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 (proofs of) derivable sequents $\Gamma \Rightarrow \phi$ (for some finite set of formulas Γ and a formula ψ) in a sequent calculus for SDL, based on Gentzen's LK proof system [15]. 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 uncharging of $\Gamma \Rightarrow \phi$ (see [1]).

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 [16]:

- When served a meal you ought to not eat with your fingers.
- However, if the meal is asparagus you ought to eat with your fingers.

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

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 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 shall denote formulas in $\mathcal{L}_{\mathsf{SDL}}$ by the lower Greek letter ψ, ϕ, and set of formulas by the upper Greek letters Γ, Δ, Σ. 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 $\top$, write $\mathsf{O}\Gamma$ for the set $\{\mathsf{O}\psi \mid \psi \in \Gamma\}$, and denote $\bigwedge\Gamma$ for the conjunction of the formulas in a finite set Γ.

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 $\mathcal{L}$-formulas and $\Rightarrow$ is a symbol that does not appear in $\mathcal{L}_{\mathsf{SDL}}$. We shall denote $\mathsf{Prem}(\Gamma \Rightarrow \psi) = \Gamma$.

Given a set Σ of formulas in $\mathcal{L}_{\mathsf{SDL}}$, we say that a formula ψ *follows* from Σ (in SDL), and denote this by $\Sigma \vdash_{\mathsf{SDL}} \psi$, if there is a subset $\Gamma \subseteq \Sigma$, 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. It is easy to verify that $\vdash_{\mathsf{SDL}}$ is a Tarskian consequence relation (that is, reflexive, monotonic and transitive).

Axioms: $\psi \Rightarrow \psi$

Structural Rules:

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

Logical Rules:

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

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

$$\text{MP:}\ \frac{}{\Gamma, \phi, \phi \supset \psi \Rightarrow \psi} \qquad [\Rightarrow\supset]\ \frac{\Gamma, \psi \Rightarrow \varphi, \Delta}{\Gamma \Rightarrow \psi \supset \varphi, \Delta}$$

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

$$\text{KR:}\ \frac{\Gamma \Rightarrow \phi}{\mathsf{O}\Gamma \Rightarrow \mathsf{O}\phi} \qquad \text{DR:}\ \frac{\Gamma \Rightarrow \phi}{\mathsf{O}\Gamma \Rightarrow \neg\mathsf{O}\neg\phi}$$

Fig. 1. The proof system $\mathcal{C}_{\mathsf{SDL}}$

Note 1. The proof system $\mathcal{C}_{\mathsf{SDL}}$ is equivalent to Gentzen's well-known sequent calculus LK for classical propositional logic, extended with rules for the modal operator O [17]. In particular, in $\mathcal{C}_{\mathsf{SDL}}$ the rule [MP] is primitive and the rule

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

is admissible (i.e., it is derivable from the rules of $\mathcal{C}_{\mathsf{SDL}}$), while in LK it is the other way around. We switched between [MP] and [$\supset\Rightarrow$] since this allows to simplify some of the formalities to be developed in the sequel (namely, it allows for a more straightforward formulation of some sequent elimination rules in Section 3).

Note 2. 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 $\mathsf{O}\phi \supset (\psi \supset \mathsf{O}\phi)$). 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). For instance, applying SDL to Forrester's example of the gentle murderer (see [18] and Example 6 below) results in triviality. These problems are usually taken to be the death sentence for an account of conditional obligations based on SDL and material implication. 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).

Note 3. Instead of basing SDL on classical propositional logic and LK, our framework also allows for other variants, such as 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.

3 Logical Argumentation for Normative Reasoning

In what has become the orthodox approach based on Dung's representation [9], formal argumentation is studied on the basis of so-called 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.

Definition 1. *An* (abstract) argumentation framework *is a pair $\langle Args, Attack\rangle$, where Args is an enumerable set of elements, called (abstract)* arguments*, and Attack is a 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 [10], ASPIC [11], etc. Here we settle for a recently proposed account based on sequent-based calculi [1].

The basic idea behind our instantiation is that arguments are $\mathcal{C}_{\mathsf{SDL}}$-proofs.

Definition 2. $\mathsf{Arg}(\Sigma)$ *is the set of* $\mathcal{C}_{\mathsf{SDL}}$*-proofs of sequents of the form* $\Gamma \Rightarrow \psi$ *for some* $\Gamma \subseteq \Sigma$.

For specifying the attack relation 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 we can construct arguments for conflicting obligations (and permissions).

Example 2. Consider the following sequent elimination rule:

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

This rule aims at formalizing the principle of specificity. It states that when two sequents $\Gamma' \Rightarrow \psi'$ and $\Gamma \Rightarrow \psi$ are conflicting, the one which is more specific gets higher precedence, and so the other one is discarded. Thus, in Example 1 for instance, SPEC allows to discharge the sequent $m,\, m \supset \mathsf{O}\neg f \Rightarrow \mathsf{O}\neg f$ in light of the more specific sequent $m \wedge a,\, (m \wedge a) \supset \mathsf{O}f \Rightarrow \mathsf{O}f$. We also say that the latter sequent *attacks* the former.

Some variations of SPEC are given below (where $\mathsf{NN'} \in \{\mathsf{OO}, \mathsf{OP}, \mathsf{PO}\}$):[3,4]

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

For instance, POSPEC models permission as derogation [19]: a permission may suspend a more general obligation.

Example 3. In order to illustrate a conflict for which there is no overriding principle such as specificity that resolves it, suppose we have two triggered conditional obligations that conflict: $\Sigma = \{a, b, a \supset \mathsf{O}c, b \supset \mathsf{O}\neg c\}$. One could imagine an argumentative context in which one proponent presents an argument for $\mathsf{O}c$ by proving $a, a \supset \mathsf{O}c \Rightarrow \mathsf{O}c$. 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 a unilateral context this may be considerations and counter-considerations of a single reasoner. Such argumentative attacks may be modeled by sequent elimination rules as (where $\mathsf{NN'} \in \{\mathsf{OO}, \mathsf{OP}, \mathsf{PO}\}$):

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

Some further sequent elimination rules for handling conflicting sequents are listed in Figure 2. We will not further discuss them here but we will come back to them in Section 4.

[3] In this and the following attack rules we also intend to capture unconditional obligations such as $\mathsf{O}\psi$. E.g., that $\phi, \phi \supset \mathsf{O}\psi \Rightarrow \mathsf{O}\psi$ OOSPEC-attacks $\mathsf{O}\neg\psi \Rightarrow \mathsf{O}\neg\psi$.

[4] Note that a 'PPSPEC'-variant would not be sensible since permissions with incompatible content do not conflict in any intuitive sense.

$$\text{CON}\quad \frac{\Rightarrow \neg\bigwedge\Gamma \qquad \Gamma,\Gamma' \Rightarrow \psi}{\Gamma,\Gamma' \not\Rightarrow \psi} \qquad\qquad \text{NIC}\quad \frac{\Gamma \Rightarrow \neg\phi \qquad \Gamma' \Rightarrow \mathsf{N}\phi}{\Gamma' \not\Rightarrow \mathsf{N}\phi}$$

$$\text{NN}'\text{CONFU}\quad \frac{\Gamma,\, \phi \supset \mathsf{N}\psi \Rightarrow \mathsf{N}\psi \qquad \Gamma \Rightarrow \phi \qquad \psi \Rightarrow \neg\psi' \qquad \Gamma',\, \phi' \supset \mathsf{N}'\psi' \Rightarrow \psi''}{\Gamma',\, \phi' \supset \mathsf{N}'\psi' \not\Rightarrow \psi''}$$

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

$$\text{NCTD}\quad \frac{\begin{matrix}\Gamma,\, \phi \supset \mathsf{N}\psi \\ \Rightarrow \mathsf{N}\psi\end{matrix} \qquad \Gamma \Rightarrow \phi \quad \Gamma' \Rightarrow \phi' \quad \phi \Rightarrow \phi' \quad \psi \Rightarrow \neg\psi' \qquad \begin{matrix}\Gamma',\, \phi' \supset \mathsf{O}\psi' \\ \Rightarrow \mathsf{O}\psi'\end{matrix}}{\Gamma',\, \phi' \supset \mathsf{O}\psi' \not\Rightarrow \mathsf{O}\psi'}$$

$$\text{NN}'\text{SPECU}\quad \frac{\begin{matrix}\Gamma,\, \phi \supset \mathsf{N}\psi \\ \Rightarrow \neg(\phi' \supset \mathsf{N}'\psi')\end{matrix} \qquad \Gamma \Rightarrow \phi \quad \phi \Rightarrow \phi' \quad \psi \Rightarrow \neg\psi' \qquad \begin{matrix}\Gamma',\, \phi' \supset \mathsf{N}'\psi' \\ \Rightarrow \psi''\end{matrix}}{\Gamma',\, \phi' \supset \mathsf{N}'\psi' \not\Rightarrow \psi''}$$

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

Fig. 2. Some more sequent elimination rules for normative reasoning (where $\mathsf{NN}' \in \{\mathsf{OO}, \mathsf{OP}, \mathsf{PO}\}$ and $\mathsf{N} \in \{\mathsf{O}, \mathsf{P}\}$)

Attacks between arguments are defined with reference to some $A \in \mathsf{Arg}(\Sigma)$ as follows:

- $\hat{A}$ denotes the top sequent in the proof A;[5]
- We say that a sequent $\Gamma \Rightarrow \psi$ is a *subsequent of* A if it is contained in A, and $\mathsf{Prem}(\hat{A}) \vdash_{\mathsf{SDL}} \bigwedge \Gamma$ (or, equivalently, if $\mathsf{Prem}(\hat{A}) \Rightarrow \bigwedge \Gamma \in \mathsf{Arg}(\Sigma)$).[6]

According to the next definition, an argument is attacked in some of its subsequents (including its top sequent).

Definition 3. *Let* $R = \frac{\Gamma_1 \Rightarrow \phi_1 \ \ldots \ \Gamma_n \Rightarrow \phi_n}{\Gamma_n \not\Rightarrow \phi_n}$ *be a sequent elimination rule in Figure 2, 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* $\mathfrak{s} = \theta(\Gamma_1) \Rightarrow \theta(\phi_1)$ *and* $\mathfrak{s}' = \theta(\Gamma_n) \Rightarrow \theta(\phi_n)$*. We say that* $\mathfrak{s}$ $\mathcal{R}$-attacks $\mathfrak{s}'$ *if* $\mathfrak{s}$ R*-attacks* $\mathfrak{s}'$ *for some* $R \in \mathcal{R}$*.*
- *An argument* $A \in \mathsf{Arg}(\Sigma)$ R-attacks *an argument* $B \in \mathsf{Arg}(\Sigma)$ *if* $\hat{A}$ R*-attacks some subsequent of* B*. Similarly,* A $\mathcal{R}$-attacks B *if* A R*-attacks* B *for some* $R \in \mathcal{R}$*.*

[5] The top sequent is the top of the proof tree A if we conceive of proofs as trees, or the last line in the proof A if proofs are considered to be lists of lines.

[6] Intuitively speaking, the second condition warrants that the subsequents of a proof A of $\mathfrak{s} = \Gamma \Rightarrow \psi$ are only those sequents whose premises are charged in the proof of $\mathfrak{s}$. Take for instance the proof of $\Rightarrow \phi \supset \phi$ from $\phi \Rightarrow \phi$ by [$\Rightarrow\supset$]. This prevents for instance attacks on A by $\neg\phi \Rightarrow \neg\phi$.

Definition 4. *A* normative argumentation framework *induced by a set of elimination rules $\mathcal{R}$ is the logical argumentation framework $\mathcal{AF}_{\mathcal{R}}(\Sigma) = \langle \mathsf{Arg}(\Sigma), Attack \rangle$ in which $(A, B) \in Attack$ iff A $\mathcal{R}$-attacks B.*

Normative Entailments Induced by Argumentation Frameworks

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 [9], 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 $\mathcal{AF}$, and a maximal complete subset of Args is called a* preferred extension *of $\mathcal{AF}$.*

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

- $\Sigma \mid\!\sim_{\mathsf{gr}} \psi$ if there is $A \in \mathsf{Arg}(\Sigma)$ in the grounded extension of $\mathcal{AF}_{\mathcal{R}}(\Sigma)$ such that $\hat{A} = \Gamma \Rightarrow \psi$.[7]
- $\Sigma \mid\!\sim^{\cap}_{\mathsf{pr}} \psi$ [$\Sigma \mid\!\sim^{\cup}_{\mathsf{pr}} \psi$] if in every [some] preferred extension of $\mathcal{AF}_{\mathcal{R}}(\Sigma)$ there is $A \in \mathsf{Arg}(\Sigma)$ with $\hat{A} = \Gamma \Rightarrow \psi$.[8,9]

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

4 Some Examples

In this section we will demonstrate our argumentative model for normative reasoning by means of various examples.

Example 4. Let us recall Example 1, where $\Sigma_1 = \{m, a, m \supset \mathsf{O}\neg f, (m \wedge a) \supset \mathsf{O}f\}$. Some arguments in $\mathsf{Arg}(\Sigma_1)$ are listed in Figure 3 (right). We do not spell out the very simple proofs given by each argument but only list the top sequents and subsequent relationships. For instance, arguments A, B, C, D and E are one-liner proofs, argument F is obtained from B and C by weakening, etc. Figure 3 (left) shows an attack diagram where the only attack rule is $\mathsf{OOSPECU}$.

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

[8] A more cautious approach is to define: $\Sigma \mid\!\sim^{\cap}_{\mathsf{pr}} \psi$ [$\Sigma \mid\!\sim^{\cup}_{\mathsf{pr}} \psi$] if there is an $A \in \mathsf{Arg}(\Sigma)$ with $\hat{A} = \Gamma \Rightarrow \psi$ that is in every [some] preferred extension of $\mathcal{AF}_{\mathcal{R}}(\Sigma)$.

[9] Similar entailment relations may of-course be defined for other semantics of abstract argumentation such as [semi-]stable semantics, ideal semantics, etc.

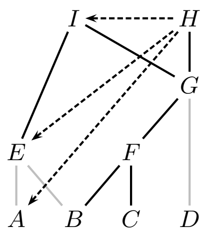

$\hat{A} = m \supset \mathsf{O}\neg f \Rightarrow m \supset \mathsf{O}\neg f$
$\hat{B} = m \Rightarrow m$
$\hat{C} = a \Rightarrow a$
$\hat{D} = (m \wedge a) \supset \mathsf{O}f \Rightarrow (m \wedge a) \supset \mathsf{O}f$
$\hat{E} = m, m \supset \mathsf{O}\neg f \Rightarrow \mathsf{O}\neg f$
$\hat{F}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

Fig. 3. (Part of) the normative argumentation framework of Example 4: dashed arrows are OOSPECU-attacks, solid black lines indicate subsequents (the top sequents of lower arguments are subsequents of higher ones) and the gray line merely helps the reader to see which sequents share premises.

We observe that H OOSPECU-attacks A and E, and since $\hat{E}$ is a subsequent of I, the latter is also attacked by H. It follows that, as expected, we have the following deductions:

- $\Sigma_1 \not\mid\!\sim \mathsf{O}\neg f$. 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.[10]
- $\Sigma_1 \mid\!\sim \mathsf{O}f$. Indeed, G is not OOSPECU-attackable by an argument in $\mathsf{Arg}(\Sigma)$, thus it is part of every grounded and preferred extension of the underlying normative argumentation framework, and so its descendent follows from Σ_1.[11]

Example 5. Caminada [20] 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 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: $\Sigma_{\mathsf{pro}} = \{s, p, s \supset m, p \supset \neg\mathsf{P}r, m \supset \mathsf{P}r, m \supset \mathsf{O}f\}$. We can proof the two sequents $\mathfrak{s}_1 = p, p \supset \mathsf{O}\neg r \Rightarrow \mathsf{O}\neg r$ and $\mathfrak{s}_2 = s, s \supset m, m \supset \mathsf{P}r \Rightarrow \mathsf{P}r$ which NN'CONF-attack each other (where $\mathsf{NN}' \in \{\mathsf{OP}, \mathsf{PO}\}$).

Caminada uses this example to illustrate what is sometimes considered a shortcoming of deductive approaches to defeasible reasoning [20, 21]. Given two conflicting inference steps described schematically by $\phi \rightsquigarrow \neg\theta$ and $\psi_1 \rightsquigarrow \psi_2 \rightsquigarrow \theta$.

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

[11] It is important to note that G *is* OOSPECU-attackable by SDL-derivable arguments, but none of them is in $\mathsf{Arg}(\Sigma_1)$. For instance, since 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 attackable by an argument with, say, the SDL-derivable top 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 $m \wedge a \supset \mathsf{O}\neg f \not\in \Sigma_1$, this argument is not in $\mathsf{Arg}(\Sigma_1)$. 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 OOSPECU-attack $\hat{G}$ and $\hat{H}$ though it is attacked by $\hat{H}$.

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 $\mathfrak{s}_4 = p, p \supset \mathsf{O}\neg r, m \supset \mathsf{P}r \Rightarrow \neg m$ is provable and conflicts with $\mathfrak{s}_3 = s, s \supset m \Rightarrow m$. Caminada argues that this violates the principle 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' $\mathfrak{s}_4$ by attack rules such as FCONF as follows: 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 $\mathfrak{s}_7 = s, p, s \supset m \Rightarrow \neg((m \supset \mathsf{P}r) \wedge (p \supset \mathsf{O}\neg r))$. Now, $\mathfrak{s}_7$ attacks all sequents that have $p \supset \mathsf{O}\neg r$ and $m \supset \mathsf{P}r$ as premises, such as $\mathfrak{s}_4$.

In Figure 4 we depict an excerpt of an attack diagram for Σ_{pro} with the attack rules FCONF and NN'CONF. We get, e.g., $\Sigma_{\mathsf{pro}} \mid\!\sim_{\mathsf{gr}} \mathsf{O}f$ and $\Sigma_{\mathsf{pro}} \not\mid\!\sim_{\mathsf{gr}} \mathsf{P}r$. If we use only FCONF, we also get $\Sigma_{\mathsf{pro}} \mid\!\sim_{\mathsf{gr}} \mathsf{P}r$ and $\Sigma_{\mathsf{pro}} \mid\!\sim_{\mathsf{gr}} \mathsf{O}\neg r$ but we are still not able to accept arguments with both conflicting conditional obligations as premise (e.g., A_4 and A_6) since such arguments are FCONF-attacked by A_7.[12]

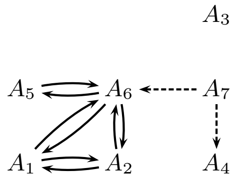

$\mathfrak{s}_1 = p, p \supset \mathsf{O}\neg r \Rightarrow \mathsf{O}\neg r$
$\mathfrak{s}_2 = s, s \supset m, m \supset \mathsf{P}r \Rightarrow \mathsf{P}r$
$\mathfrak{s}_3 = s, s \supset m \Rightarrow m$
$\mathfrak{s}_4 = p, p \supset \mathsf{O}\neg r, m \supset \mathsf{P}r \Rightarrow \neg m$
$\mathfrak{s}_5 = s, s \supset m, m \supset \mathsf{O}f \Rightarrow \mathsf{O}f$
$\mathfrak{s}_6 = p, p \supset \mathsf{O}\neg r, m \supset \mathsf{P}r, s, s \supset m \Rightarrow \mathsf{O}\bot$
$\mathfrak{s}_7 = s, s \supset m, p \Rightarrow \neg((m \supset \mathsf{P}r) \wedge (p \supset \mathsf{O}\neg r))$

Fig. 4. An attack diagram for Example 5 where $\hat{A}_i = \mathfrak{s}_i$. FCONF attacks are dashed, NN'CONF-attacks are solid ($\mathsf{N}, \mathsf{N}' \in \{\mathsf{O}, \mathsf{P}\}$).

Example 6. In the next example we take a look at contrary-to-duty (in short, CTD) obligations. A paradigmatic example is Forrester's Gentle Murderer scenario [18]: generally, one ought not to kill ($\top \supset \mathsf{O}\neg f$). However, upon killing, this should be done gently ($k \supset \mathsf{O}(k \wedge g)$). Let $\Sigma_2 = \{k, \top \supset \mathsf{O}\neg k, k \supset \mathsf{O}(k \wedge g)\}$.

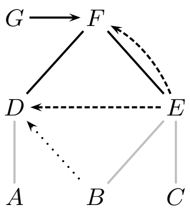

$\hat{A} = \top \supset \mathsf{O}\neg k \Rightarrow \top \supset \mathsf{O}\neg k$
$\hat{B} = k \Rightarrow k$
$\hat{C} = k \supset \mathsf{O}(k \wedge g) \Rightarrow k \supset \mathsf{O}(k \wedge g)$
$\hat{D} = \top \supset \mathsf{O}\neg k \Rightarrow \mathsf{O}\neg k$
$\hat{E} = k, k \supset \mathsf{O}(k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

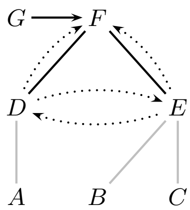

Fig. 5. Two modelings of Forrester's Gentle Murderer

[12] This is in line with Goble's [6, p. 27/28] analysis of an enriched version of Horty's Smith argument [7]: given $\{\mathsf{O}\neg f, \mathsf{O}(f \vee s), \mathsf{O}\neg s\}$ (f is fighting in the army, s is performing civil service) he advocates to let both $\mathsf{O}s$ and $\mathsf{O}\neg s$ be derivable without aggregating them to $\mathsf{O}(s \wedge \neg s)$.

Van der Torre and Tan [22] distinguish CTD-obligations from cases of specificity. In the former the general obligations are not canceled or overridden but have still normative force (despite the fact that they are violated), while in cases of specificity the more general conditional obligations are canceled and thus deprived of normative force. There are various ways in which in our framework this distinction can be taken into account. One way of doing so is as follows. Instead of using strong rules such as $\mathsf{OOSPECU}$ in Example 4 that 'destroy' overridden conditional obligations in the sense that they do not appear in the consequence set, we can make use of rules such as OCTD (Figure 2) that preserve 'overshadowed' conditional CTD obligations despite the fact that detachment is blocked, or incorporate OIC that blocks detachment from violated obligations. This is illustrated in Figure 5 (left) with the attack rules OCTD (dashed arrow), OIC (dotted arrow) and CON (solid arrow). Alternatively, we could model overshadowing by means of OOCONF instead of OCTD. This is illustrated in Figure 5 (right) with attack rules OOCONF (dotted arrows) and CON (solid arrow). Where $\Xi = \{A, B, C, G\}$, we have two preferred extensions: $\Xi \cup \{D\}$ and $\Xi \cup \{E\}$. Hence, $\Sigma_2 \mid\!\sim^{\cup}_{\mathsf{pr}} \mathsf{O}\neg k$ and $\Sigma_2 \mid\!\sim^{\cup}_{\mathsf{pr}} \mathsf{O}(k \wedge g)$. In the skeptical approach we get $\Sigma_2 \mid\!\sim^{\cap}_{\mathsf{pr}} \mathsf{O}(\neg k \vee (k \wedge g))$ and $\Sigma_2 \mid\!\sim^{\cap}_{\mathsf{pr}} \mathsf{O}\neg k \vee \mathsf{O}(k \wedge g)$. Yet another option is to use a very liberal approach with CON only. This will block arguments with inconsistent premises such as F but otherwise allows e.g., to derive both $\mathsf{O}\neg k$ and $\mathsf{O}(g \wedge k)$ even via the grounded approach: $\Sigma_2 \mid\!\sim_{\mathsf{gr}} \mathsf{O}\neg k$ and $\Sigma_2 \mid\!\sim_{\mathsf{gr}} \mathsf{O}(k \wedge g)$.

Example 7. Consider the next paradigmatic CTD-case (Chisholm paradox, [23]):
- 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 modeling 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 modeling 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.[13]

Example 8. Let us consider a variant of Example 4. Suppose that beside the obligation not to eat with your fingers we have the permission to do so in case asparagus is served, but it is considered impolite to eat asparagus with your fingers if there is guest who considers this rude. The enriched set of premises

[13] An alternative modeling of 2 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.

may look as follows: $\Sigma_3 = \{a, m, c, m \supset \mathsf{O}\neg f, (m \wedge a) \supset \mathsf{P}f, (m \wedge a \wedge c) \supset \mathsf{O}\neg f\}$. The situation is depicted in Figure 6, where the attack rules OPSPECU (dotted arrows) and POSPECU (dashed arrows).

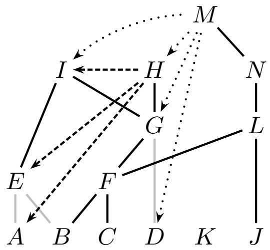

$\hat{D} = (m \wedge a) \supset \mathsf{P}f \Rightarrow (m \wedge a) \supset \mathsf{P}f$
$\hat{G} = m, a, (m \wedge a) \supset \mathsf{P}f \Rightarrow \mathsf{P}f$
$\hat{H} = m, a, (m \wedge a) \supset \mathsf{P}f \Rightarrow \neg(m \supset \mathsf{O}\neg f)$
$\hat{I} = m, a, m \supset \mathsf{O}\neg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

Fig. 6. A normative argumentation framework for Example 8 (arguments A, B, C, E, F are as in Figure 3)

Thus, $\Sigma_3 \mid\!\sim \mathsf{O}\neg f$ (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.

Example 9. Next we take a look at a simple conflict that is neither a specificity nor a CTD-case. Let $\Sigma_4 = \{a, b, a \supset \mathsf{O}(c \wedge d), b \supset \mathsf{O}(\neg c \wedge d)\}$. Figure 7 shows the situation for the attack rule OOCONFU (dotted arrows).

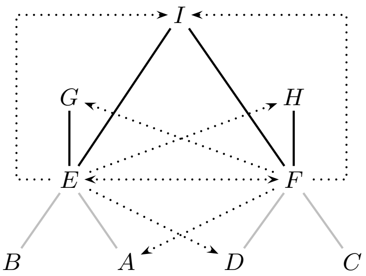

$\hat{A} = a \supset \mathsf{O}(c \wedge d) \Rightarrow a \supset \mathsf{O}(c \wedge d)$
$\hat{B} = a \Rightarrow a$
$\hat{C} = b \Rightarrow b$
$\hat{D} = b \supset \mathsf{O}(\neg c \wedge d) \Rightarrow b \supset \mathsf{O}(\neg c \wedge d)$
$\hat{E} = a, a \supset \mathsf{O}(c \wedge d) \Rightarrow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

Fig. 7. A simple conflict

We have the following preferred extensions: $\{A, B, E, G\}$ and $\{C, D, F, H\}$. Note that we have the 'floating conclusion'[14] $\Sigma_4 \mid\!\sim^{\cap}_{\mathsf{pr}} \mathsf{O}d$ since one of G and H is in every preferred extension.

Example 10. The next example illustrates a conflict between three obligations. Let $\Sigma_5 = \{c, c \supset \mathsf{O}(a \vee b), c \supset \mathsf{O}(\neg a \vee b), c \supset \mathsf{O}\neg a\}$. It is interesting to note that modeling this scenario with OOCONFU is problematic. In this case no conflicts are triggered since the triple-conflict is not reducible to a binary conflict that fits the attack rule OOCONFU. This may be avoided by using OCONFU' instead of OOCONFU, as we get for instance $\Sigma_5 \mid\!\sim^{\cap}_{\mathsf{pr}} \mathsf{O}a \vee \mathsf{O}(\neg a \wedge b) \vee \mathsf{O}(\neg a \wedge \neg b)$. This example shows that elimination rules should be carefully chosen.[15]

[14] In nonmonotonic reasoning *floating conclusions* are conclusions that are obtained from each of a set of otherwise conflicting arguments.

[15] In Section 5 we will prove that OCONFU' is rather well-behaved and can be used to give an argumentative account of a specific Input/Output logic.

5 Basic Properties and Relation to Input/Output Logic

We start with two basic observations which can easily be verified by the reader:

1. For any set of attack rules previously defined: whenever Σ is SDL-consistent (i.e., $\Sigma \nvdash_{\mathsf{SDL}} \bot$) then $\Sigma \vdash_{\mathsf{SDL}} \psi$ iff $\Sigma \mid\!\sim \psi$. It is easy to verify that in this case all arguments in $\mathsf{Arg}(\Sigma)$ are selected since no argumentative attacks occur.
2. Where CON is part of the attack rules, (i) $\Sigma \mid\!\sim \phi$ implies that ϕ is SDL-consistent (i.e., $\phi \nvdash_{\mathsf{SDL}} \bot$) and, consequently, (ii) $\mid\!\sim$ is strongly paraconsistent (i.e., for all Σ, $\Sigma \not\mid\!\sim \bot$).

The framework of Input/Output logics [24] represents one of the standard approaches in conditional deontic logic. Many logics devised in this framework come with a simple and intuitive syntactic as well as a semantic characterization. The framework has been extended in order to deal with conflicts among conditionals such as in Contrary-to-Duty obligations [25]. There are in-depth studies concerning the modeling of permissions [19]. Moreover, there are links to other frameworks such as default logic (see [25]), logic programming [26] and adaptive logics [27]. In view of this it is interesting to notice that I/O logics can also be related to our framework, as will be established below (see Theorem 1 and Corollary 1).

In the following we focus on premise sets Σ 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, let Σ_F be a set of non-modal propositional formulas, Σ_O a set of pairs of non-modal formulas (ψ, ϕ) ('I/O-pairs') and $\Sigma_O^* = \{\psi \supset \mathsf{O}\phi \mid (\psi, \phi) \in \Sigma_O\}$. Let also CPL be classical propositional logic, and denote by $\mathrm{Cn}_{\mathsf{CPL}}(\Gamma)$ the transitive closure of Γ with respect to $\vdash_{\mathsf{CPL}}$. The following definitions describe the 'out' and the 'out_2'-function in [24]:

Definition 6. $\mathsf{out}(\Sigma_F, \Sigma_O) = \{\psi \mid (\phi, \psi) \in \Sigma_O,\ \Sigma_F \vdash_{\mathsf{CPL}} \phi\}$.

Definition 7. *$\phi \in \mathsf{out}_2(\Sigma_F, \Sigma_O)$ iff $\phi \in \mathrm{Cn}_{\mathsf{CPL}}(\mathsf{out}(\Xi, \Sigma_O))$ for all CPL-maximal consistent extensions Ξ of Σ_F. In the degenerated case in which Σ_F is CPL-inconsistent, we define $out_2(\Sigma_F, \Sigma_O)$ to be $\mathrm{Cn}_{\mathsf{CPL}}(\{\psi \mid (\psi', \psi) \in \Sigma_O\})$.*

The following is a corollary of Observation 4 in [24]:[16]

Lemma 1. *$\Sigma_F \cup \Sigma_O^* \vdash_{\mathsf{SDL}} \mathsf{O}\phi$ iff $\phi \in \mathsf{out}_2(\Sigma_F, \Sigma_O)$.*

In order to deal with situations in which $\mathsf{out}(\Sigma_F, \Sigma_O)$ is inconsistent, Makinson and Van Der Torre [25] 'contextualize' their output-functions to maximal sets of conditionals that are consistent with Σ_F, so-called maxfamilies:[17]

Definition 8. *We consider the following sets:*

[16] In [24] the authors show the correspondence for all normal modal logics L, for which $\vdash_{\mathsf{K}} \subseteq \vdash_{\mathsf{L}} \subseteq \vdash_{\mathsf{K45}}$.

[17] The approach in [25] is more general since it takes into account sets of additional constraints beside our requirement of consistency.

– $\Gamma_O \in \mathsf{maxfamily}(\Sigma_F, \Sigma_O)$ *iff* $\mathsf{out}_2(\Sigma_F, \Gamma_O)$ *is* CPL*-consistent and for all* $(\psi, \phi) \in \Sigma_O \setminus \Gamma_O$, $\mathsf{out}_2(\Sigma_F, \Gamma_O \cup \{(\psi, \phi)\})$ *is not* CPL*-consistent.*
– $\psi \in \mathsf{out}_2^{\cup}(\Sigma_F, \Sigma_O)$ *iff* $\psi \in \bigcup_{\Gamma_O \in \mathsf{maxfamily}(\Sigma_F, \Sigma_O)} \mathsf{out}_2(\Sigma_F, \Gamma_O)$.
– $\psi \in \mathsf{out}_2^{\cap}(\Sigma_F, \Sigma_O)$ *iff* $\psi \in \bigcap_{\Gamma_O \in \mathsf{maxfamily}(\Sigma_F, \Sigma_O)} \mathsf{out}_2(\Sigma_F, \Gamma_O)$.

We now show that in our argumentative approach the Input/Output logics in Definition 8 can be characterized by means of the attack rule OCONFU'.

Theorem 1. *If Σ_F is* CPL*-consistent, then the set of all the preferred extensions of $\mathcal{AF}_{\mathsf{OCONFU'}}(\Sigma_F \cup \Sigma_O^*)$ is $\{\mathsf{Arg}(\Sigma_F \cup \Gamma_O^*) \mid \Gamma_O \in \mathsf{maxfamily}(\Sigma_F, \Sigma_O)\}$.*

Proof (Sketch). Let $\Gamma_O \in \mathsf{maxfamily}(\Sigma_F, \Sigma_O)$. By Lemma 1, $\Sigma_F \cup \Gamma_O^*$ is SDL-consistent and hence $\mathsf{Arg}(\Sigma_F \cup \Gamma_O^*)$ is conflict-free. Thus, each argument A attacking any argument in $\mathsf{Arg}(\Sigma_F \cup \Gamma_O^*)$ is such that $A \notin \mathsf{Arg}(\Sigma_F \cup \Gamma_O^*)$. Let $A \in \mathsf{Arg}(\Sigma_F \cup \Sigma_O^*) \setminus \mathsf{Arg}(\Sigma_F \cup \Gamma_O^*)$. This means that there is a $\psi \supset \mathsf{O}\phi \in \mathsf{Prem}(\hat{A}) \cap (\Sigma_O^* \setminus \Gamma_O^*)$. Since $\mathsf{out}_2(\Sigma_F, \Gamma_O \cup \{(\psi, \phi)\})$ is CPL-inconsistent we have by Lemma 1 that $\Sigma_F \cup \Gamma_O^* \cup \{\psi \supset \mathsf{O}\phi\}$ is SDL-inconsistent. Thus, there is a finite $\Theta \subseteq \Sigma_F \cup \Gamma_O^*$ such that $\Theta, \psi \supset \mathsf{O}\phi \Rightarrow \perp$ is $\mathcal{C}_{\mathsf{SDL}}$-provable. By $[\Rightarrow\supset]$, we derive $\mathfrak{s} = \Theta \Rightarrow \neg(\psi \supset \mathsf{O}\phi)$. Let C be the corresponding proof with $\hat{C} = \mathfrak{s}$. Then $C \in \mathsf{Arg}(\Sigma_F \cup \Gamma_O^*)$ and C OCONFU'-attacks A. We have shown that $\mathsf{Arg}(\Sigma_F \cup \Gamma_O^*)$ is defended and that it is maximally so.

Now assume there is an admissible extension Ξ of $\mathcal{AF}_{\mathsf{OCONFU'}}(\Sigma_F \cup \Sigma_O^*)$ such that there is no $\Gamma_O \in \mathsf{maxfamily}(\Sigma_F, \Sigma_O)$ for which $\Xi \subseteq \mathsf{Arg}(\Sigma_F \cup \Gamma_O^*)$. Hence, there is no $\Gamma_O \in \mathsf{maxfamily}(\Sigma_F, \Sigma_O)$ for which $\Gamma_\Xi = \bigcup_{A \in \Xi} \{(\psi, \phi) \mid \psi \supset \mathsf{O}\phi \in \mathsf{Prem}(\hat{A})\} \subseteq \Gamma_O$. This means $\mathsf{out}_2(\Sigma_F, \Gamma_\Xi)$ is CPL-inconsistent. By Lemma 1, $\Sigma_F \cup \Gamma_\Xi^*$ is SDL-inconsistent. Hence, there are finite $\Theta_F \subseteq \Sigma_F$ and $\Theta_O^* \subseteq \Gamma_\Xi^*$ such that $\Theta_F, \Theta_O^* \Rightarrow \perp$ is $\mathcal{C}_{\mathsf{SDL}}$-derivable. Since Σ_F is CPL-consistent, $\Theta_O^* \neq \emptyset$. With Weakening and $[\Rightarrow\supset]$ we have an argument C, with $\hat{C} = \Theta_F, \Theta_O^* \setminus \{\psi \supset \mathsf{O}\phi\} \Rightarrow \neg(\psi \supset \mathsf{O}\phi)$ where $\psi \supset \mathsf{O}\phi \in \Theta_O^*$. By the definition of Γ_Ξ, there is an $A \in \Xi$ for which $\psi \supset \mathsf{O}\phi \in \mathsf{Prem}(\hat{A})$. By the subformula property of SDL (see [17]) we can suppose that:

(†) for all $\gamma \supset \mathsf{O}\gamma'$ that occur in subsequents of C, $(\gamma, \gamma') \in \Theta_O$.

Then C OCONFU'-attacks A. Also, by (†), the only way to attack C leads to an attack on Ξ as well. Thus, Ξ cannot be defended from C. □

Corollary 1. *Where the only attack rule is* OCONFU'*, for every $\lambda \in \{\cup, \cap\}$ it holds that $\psi \in \mathsf{out}_2^{\lambda}(\Sigma_F, \Sigma_O)$ iff $\Sigma_F \cup \Sigma_O^* \mathrel{|\!\!\sim}^{\lambda}_{\mathsf{pr}} \mathsf{O}\psi$.*

Example 11. Suppose that $\Sigma_O = \{(p_1, q_1 \wedge q_2), (p_2, \neg q_1 \wedge q_2)\}$ and $\Sigma_F = \{p_1, p_2\}$. We have $\mathsf{maxfamily}(\Sigma_F, \Sigma_O) = \{\{(p_1, q_1 \wedge q_2)\}, \{(p_2, \neg q_1 \wedge q_2)\}\}$. Since $q_2 \in \mathsf{out}_2(\Sigma_F, \{(p_1, q_1 \wedge q_2)\}) \cap \mathsf{out}_2(\Sigma_F, \{(p_2, \neg q_1 \wedge q_2)\})$, also $q_2 \in \mathsf{out}_2^{\cap}(\Sigma_F, \Sigma_O)$.

In the normative argumentation framework $\mathcal{AF}_{\mathsf{OCONFU'}}(\Sigma_F, \Sigma_O^*)$ we have two preferred extensions: one with e.g. arguments with top sequents $p_1, p_1 \supset \mathsf{O}(q_1 \wedge q_2) \Rightarrow \neg(p_2 \supset \mathsf{O}(\neg q_1 \wedge q_2))$, $p_1, p_1 \supset \mathsf{O}(q_1 \wedge q_2) \Rightarrow \mathsf{O}(q_1 \wedge q_2)$, and $p_1, p_1 \supset \mathsf{O}(q_1 \wedge q_2) \Rightarrow \mathsf{O}q_2$, and another one with e.g. arguments with top sequents $p_2, p_2 \supset \mathsf{O}(\neg q_1 \wedge q_2) \Rightarrow \neg(p_1 \supset \mathsf{O}(q_1 \wedge q_2))$, $p_2, p_2 \supset \mathsf{O}(\neg q_1 \wedge q_2) \Rightarrow \mathsf{O}(\neg q_1 \wedge q_2)$, and $p_2, p_2 \supset \mathsf{O}(\neg q_1 \wedge q_2) \Rightarrow \mathsf{O}q_2$. Thus, $\Sigma_F \cup \Sigma_O^* \mathrel{|\!\!\sim}^{\cap}_{\mathsf{pr}} \mathsf{O}q_2$.

Further investigations of entailment relations resulting from the application of attack rules other than OCONFU' will be considered in a future work.

6 Discussion and Outlook

The idea to use argumentation and abstract argumentation in particular to model normative reasoning is not new. Two examples are [28, 29]. The approach in [28] is based on bipolar abstract argumentation frameworks: beside an attack arrow a support arrow is used to express conditional obligations. Also in [29] Dung's framework is enhanced by a support relation this time signifying evidential support. Prolog-like predicates are used to encode argument schemes of normative reasoning and an algorithm is provided to translate them into an argumentation framework. One of the main differences in our approach based on logical argumentation 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 proofs in SDL. A by-product of this is that our approach is more tightly linked to deontic logic.

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 cases for instance call for weakened principles of strengthening the antecedent which are still strong enough to support many intuitively valid inferences. E.g., the principle of Rational Monotonicity has been challenged in [30] and replaced by a weakened version which itself has been criticized in [31]. In contrast, our base logic uses the standard implication of CPL to model conditional obligations and allows for full strengthening of the antecedent. Unwanted 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. There are other non-monotonic accounts of normative reasoning such as [7] based on default logic, Input/Output logic [25], or adaptive logics [3, 5, 6, 32]. Due to space restrictions we postpone a more elaborate comparison with these frameworks to future work.

In future work we also plan to investigate ways to combine and prioritize among attack rules, to distinguish preferences/priorities among obligations and permissions, and to relate our work to different accounts of permission [19, 33], as demonstrated next.

Example 12. Let us add the facts r_1, r_2 and the conditional obligations $r_1 \supset \mathsf{O}\neg s$ and $(r_1 \wedge r_2) \supset \mathsf{O}s$ to Σ_F and Σ_O (respectively) in Example 11. This results in the premise set $\Sigma = \{p_1, p_2, r_1, r_2, p_1 \supset \mathsf{O}(q_1 \wedge q_2), p_2 \supset \mathsf{O}(\neg q_1 \wedge q_2), r_1 \supset \mathsf{O}\neg s, (r_1 \wedge r_2) \supset \mathsf{O}s\}$. Let us also add the attack rule $\mathsf{OOSPECU}$ to the previously used OCONFU'. One would expect that we get the consequence $\mathsf{O}s$ since arguments for the conflicting $\mathsf{O}\neg s$ such as proofs with the top sequent $r_1, r_1 \supset \mathsf{O}\neg s \Rightarrow \mathsf{O}\neg s$ are $\mathsf{OOSPECU}$-attacked by arguments with the top sequent $r_1, r_2, (r_1 \wedge r_2) \supset \mathsf{O}s \Rightarrow \neg(r_1 \supset \mathsf{O}\neg s)$. However, the latter arguments are

OCONFU'-attacked by arguments with head $r_1, r_2, r_1 \supset \mathsf{O}\neg s \Rightarrow \neg((r_1 \wedge r_2) \supset \mathsf{O}s)$. More generally, adding OOSPECU to OCONFU' doesn't alter the semantic selections. Hence, in order to model configurations in which arbitrary deontic conflicts occur together with specificity cases, we may need to prioritize OOSPECU-attacks over OCONFU'. The details of this are left for future research.

There are various other resolution principles for conflicts besides specificity and the latter does not apply in all cases or may be in conflict with other principles. For instance, "lex posterior derogat legi priori" may apply expressing that more recent laws override older ones. In order to model this we need to express temporal information and hence enhance our language.

Finally, we plan to investigate whether other nonmonotonic approaches and non truth-functional logics can be expressed in our framework.[18] Also, we shall examine base logics that are obtained from SDL by removing some of the inference rules in $\mathcal{C}_{\mathsf{SDL}}$, and so such logics may not have deterministic matrices. There is also forthcoming work on dynamic proofs for sequent-based argumentation [14], which may be useful to automatize normative reasoning as modeled in this paper.

References

1. Arieli, O.: A sequent-based representation of logical argumentation. In: Proc. CLIMA'13. LNCS 8143, Springer (2013) 69–85
2. Aqvist, L.: Deontic logic. In Gabbay, D.M., Guenthner, F., eds.: Handbook of Philosophical Logic, Vol.8 (2nd edition). Kluwer (2002) 147–264
3. Beirlaen, M., Straßer, C., Meheus, J.: An inconsistency-adaptive deontic logic for normative conflicts. Journal of Philosophical Logic **2**(42) (2013) 285–315
4. Carmo, J., Jones, A.J.I.: Deontic logic and contrary-to-duties. In Gabbay, D.M., Guenthner, F., eds.: Handbook of Philosophical Logic, Vol.8 (2nd edition). Kluwer (2002) 265–343
5. Straßer, C.: A deontic logic framework allowing for factual detachment. Journal of Applied Logic **9**(1) (2010) 61–80
6. Goble, L.: Deontic logic (adapted) for normative conflicts. Logic Journal of the IGPL (2013)
7. Horty, J.F.: Reasoning with moral conflicts. Noûs **37**(4) (2003) 557–605
8. Straßer, C., Beirlaen, M.: An Andersonian deontic logic with contextualized sanctions. In: Deontic Logic in Computer Science. Springer (2012) 151–169
9. Dung, P.M.: On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. Artifical Intelligence **77** (1995) 321–358
10. Dung, P.M., Kowalski, R.A., Toni, F.: Assumption-based argumentation. Argumentation in Artificial Intelligence (2009) 199–218
11. Prakken, H.: An abstract framework for argumentation with structured arguments. Argument and Computation **1**(2) (2011) 93–124
12. Batens, D.: A universal logic approach to adaptive logics. Logica Universalis **1** (2007) 221–242

[18] For instance, we will check whether out_4 from [25] may be expressible by using $\Gamma_O^\star = \Gamma_O^* \cup \{\mathsf{O}\phi \supset \mathsf{O}\psi \mid \phi \supset \psi \in \Gamma_O\}$ instead of Γ_O^* as in Corollary 1.

13. Straßer, C.: Adaptive Logic and Defeasible Reasoning. Applications in Argumentation, Normative Reasoning and Default Reasoning. Trends in Logic, Volume 38. Springer (2014)
14. Arieli, O., Straßer, C.: Dynamic derivations for sequent-based logical argumentation (Submitted).
15. Gentzen, G.: 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).
16. Horty, J.F.: Moral dilemmas and nonmonotonic logic. Journal of Philosophical Logic **23** (1994) 35–65
17. Valentini, S.: The sequent calculus for the modal logic D. Unione Matematica Italiana. Bollettino. A. Serie **7**(7) (1993) 3
18. Forrester, J.: Gentle murder, or the adverbial samaritan. Journal of Philosophy **81** (1984) 193–197
19. Stolpe, A.: A theory of permission based on the notion of derogation. Journal of Applied Logic **8**(1) (2010) 97–113
20. Caminada, M.: On the issue of contraposition of defeasible rules. Frontiers in Artificial Intelligence and Applications **172** (2008) 109–115
21. Prakken, H.: Some reflections on two current trends in formal argumentation. Logic Programs, Norms and Action (2012) 249–272
22. Van Der Torre, L., Tan, Y.H.: Cancelling and overshadowing: two types of defeasibility in defeasible deontic logic. In: Proc. IJCAI'95. (1995) 1525–1532
23. Chisholm, R.M.: Contrary-to-duty imperatives and deontic logic. Analysis **24** (1963) 33–36
24. Makinson, D., Van Der Torre, L.: Input/Output logics. Journal of Philosophical Logic **29** (2000) 383–408
25. Makinson, D., Van Der Torre, L.: Constraints for Input/Output logics. Journal of Philosophical Logic **30**(2) (2001) 155–185
26. Gonçalves, R., Alferes, J.J.: An embedding of input-output logic in deontic logic programs. In: Deontic Logic in Computer Science. Springer (2012) 61–75
27. Straßer, C., Beirlaen, M., Putte, F.V.D.: Dynamic proof theories for input/output logic. Under review. (2014)
28. Gabbay, D.: Bipolar argumentation frames and contrary to duty obligations, preliminary report. In: Proc. CLIMA'12. Springer (2012) 1–24
29. Oren, N., Luck, M., Miles, S., Norman, T.: An argumentation inspired heuristic for resolving normative conflict. In: Proc. AAMAS-08. (2008) 41–56
30. Goble, L.: A proposal for dealing with deontic dilemmas. In: Proc. DEON'04. LNCS 3065, Springer (2004) 74–113
31. Straßer, C.: An adaptive logic framework for conditional obligations and deontic dilemmas. Logic and Logical Philosophy **19**(1–2) (2010) 95–128
32. Meheus, J., Beirlaen, M., Van De Putte, F.: Avoiding deontic explosion by contextually restricting aggregation. In: Proc. DEON'10. LNAI 6181, Springer (2010) 148–165
33. Makinson, D., Van Der Torre, L.: Permission from an input/output perspective. Journal of Philosophical Logic **32**(4) (2003) 391–416