

Two Sides of the Same Coin: Belief Revision and Enforcing Arguments*

Adrian Haret, Johannes P. Wallner, and Stefan Woltran
Institute of Logic and Computation, TU Wien, Austria

Abstract

We study a type of change on knowledge bases inspired by the dynamics of formal argumentation systems, where the goal is to enforce acceptance of certain arguments. We put forward that enforcing acceptance of arguments can be viewed as a member of the wider family of belief change operations, and that an axiomatic treatment of it is therefore desirable. In our case, laying down axioms enables a precise account of the close connection between enforcing arguments and belief revision. Our analysis of enforcing arguments proceeds by (i) axiomatizing it as an operation in propositional logic and providing a representation result in terms of rankings on sets of interpretations, (ii) showing that it stands in close relationship to belief revision, and (iii) using it as a gateway towards a principled treatment of enforcement in abstract argumentation.

1 Introduction

The axiomatic approach to belief change [Alchourrón *et al.*, 1985; Gärdenfors, 1988; Katsuno and Mendelzon, 1992] is a powerful tool in the study of dynamically evolving knowledge. Its unified methodology has proven useful for the formalization and comparison of various change operations (e.g., revision or update), and has opened up an avenue for the foundational study of other change operations of relevance to the AI community. The present work looks at enforcing acceptance of arguments, an operation from the rapidly growing field of formal argumentation [Bench-Capon and Dunne, 2007], whose focus on formalized debates lends itself naturally to an analysis in terms of beliefs and their dynamics. Indeed, interest in connecting argumentation and belief change has steadily risen over the last few years (see [Falappa *et al.*, 2011; Paglieri and Castelfranchi, 2004] for an overview, and the dedicated workshop [Fermé *et al.*, 2013], which was held again in end of 2017).

Enforcing arguments is a core operator in formal argumentation, whose aim is to force a reasoner to accept (a set

of) arguments. Typically, this aim must be achieved through minimal modifications to the reasoner’s current state of arguments [Baumann and Brewka, 2010; Baumann, 2012b; de Saint-Cyr *et al.*, 2016], with the intuition that enforcing should be done in an “economic” manner. While several aspects of enforcing arguments have been investigated, a general axiomatic study of the semantic changes required by this operation (as carried out for operators in belief change) has so far received little attention. The multitude of formal models in argumentation, coupled with a plethora of semantics, means that such a general investigation is helpful for avoiding ad-hoc approaches focused on single formalisms.

In this paper we provide an axiomatization of enforcement with the novel feature of accommodation, or preservation, of current points of view. This comes in contrast to existing enforcement, which may disrupt beliefs arbitrarily. In line with our vision of a principled approach, we begin by providing an axiomatization of our enforcement operation in propositional logic, which (1) places enforcement on the map of generic belief change operations and highlights connections to other operations in this field, and (2) lays the groundwork for deriving foundational representation results. We argue that this provides a general framework for enforcement, and take first steps in applying it to the context of abstract argumentation.

We find that viewing enforcement as a dual to revision provides a natural way for axiomatizing these requirements: while revising a knowledge base κ with μ returns a non-empty subset of the models of μ , a dual axiomatization of revision represents a family of operators that return a non-tautological superset of the models of μ . In this way models of μ are enforced, while models of κ ’s are, if possible, included, all the while avoiding an uninformative tautology.

Our main contributions are summarized below:

- We provide an alternative representation result for revision that uses rankings on *sets of* interpretations, rather than the more usual rankings on interpretations, which are better suited for relating revision and enforcement.
- We axiomatize enforcement as a natural dual to revision, prove the connection between revision and enforcement and show a representation result for enforcement based on ranking sets of interpretations.
- Based on the formalization in propositional logic, we axiomatize our enforcement operation in the context of argumentation frameworks (AFs) [Dung, 1995], a major

*This work has been supported by the Austrian Science Fund (FWF): P30168-N31, P25521, I2854, W1255-N23 and a Marietta Blau grant (OeAD-GmbH).

formalism in argumentation and their admissible semantics. We show that certain “idealized” axioms cannot be applied to AFs, yet our axiomatization in logic, aided by a characterization of the admissible semantics, leads to a principled approach to modify the axioms to suit AFs.

- We show how distance-based revision operators (e.g., [Dalal, 1988]) can be adapted to propositional enforcement and to AFs under admissible semantics.
- We argue that, when permitting only certain structural changes to an AF, axioms have to be adapted to such constraints, yet this endeavor requires semantical characterizations of the restrictions.

2 Belief Revision

We assume a finite set \mathcal{P} of propositional atoms and \mathcal{L} is the set of formulas generated from \mathcal{P} with the usual connectives. A knowledge base κ is a finite set of formulas, typically interpreted conjunctively (i.e., as $\bigwedge_{\varphi \in \kappa} \varphi$). The set of propositional knowledge bases is $2^{\mathcal{L}}$. The universe \mathcal{U} is the set of possible interpretations for formulas in \mathcal{L} . The models of a propositional formula μ are the interpretations which satisfy it, and we write $[\mu]$ (respectively, $[\kappa]$) for the set of models of μ (respectively, for $\bigcap_{\varphi \in \kappa} [\varphi]$). We typically write models as words where letters are the atoms assigned to true, e.g., $\{\{a, b\}, \{b, c\}\}$ is written as $\{ab, bc\}$. If $\mu_1, \mu_2 \in \mathcal{L}$, we say that $\mu_1 \models \mu_2$, if $[\mu_1] \subseteq [\mu_2]$, and that $\mu_1 \equiv \mu_2$, if $[\mu_1] = [\mu_2]$. A formula μ (a knowledge base κ) is consistent if $[\mu] \neq \emptyset$ ($[\kappa] \neq \emptyset$). The set of consistent knowledge bases is $2_c^{\mathcal{L}}$. The set of refutable (i.e., non-tautological) knowledge bases is $2_r^{\mathcal{L}}$.

A propositional revision operator \circ maps a $\kappa \in 2_c^{\mathcal{L}}$ and a $\mu \in \mathcal{L}$ to $\kappa \circ \mu \in \mathcal{L}$. The intention is that $\kappa \circ \mu$ is the result of modifying existing beliefs κ such that new, trusted information μ is accepted. The following list of axioms is typically acknowledged as a basic core of rational properties for revision [Katsuno and Mendelzon, 1992]:

- (R₁) $\kappa \circ \mu \models \mu$.
- (R₂) If $\kappa \wedge \mu$ is consistent, then $\kappa \circ \mu \equiv \kappa \wedge \mu$.
- (R₃) If μ is consistent, then $\kappa \circ \mu$ is consistent.
- (R₄) If $\kappa_1 \equiv \kappa_2$ and $\mu_1 \equiv \mu_2$, then $\kappa_1 \circ \mu_1 \equiv \kappa_2 \circ \mu_2$.
- (R₅) $(\kappa \circ \mu) \wedge \varphi \models \kappa \circ (\mu \wedge \varphi)$.
- (R₆) If $(\kappa \circ \mu) \wedge \varphi$ is cons., then $\kappa \circ (\mu \wedge \varphi) \models (\kappa \circ \mu) \wedge \varphi$.

Reflection on axioms R_{1–6} reveals that, if faced with inconsistency, an operator \circ satisfying them has to give up information from κ and that, in doing so, \circ behaves as if it has preferences over units of information. This is formalized using the following notions. A *preorder* \leq on a set W is a reflexive, transitive binary relation on W . The set of transitive relations on W is denoted as $\mathcal{T}(W)$. We write $<$ for the strict part of \leq , i.e., $x < x'$ if $x \leq x'$ and $x' \not\leq x$. We write $x \approx x'$ if $x \leq x'$ and $x' \leq x$. The \leq -minimal elements of W with respect to \leq are defined as $\min_{\leq} W = \{x \in W \mid \nexists x' \in W \text{ such that } x' < x\}$. An *assignment from W_1 to W_2* is a function $\alpha: W_1 \rightarrow \mathcal{T}(W_2)$.

A useful result in belief revision connects axioms R_{1–6} to assignments from $2_c^{\mathcal{L}}$ to \mathcal{U} , for which we may write \leq_{κ} instead of $\alpha(\kappa)$ if there is no danger of ambiguity. We say that

an assignment from $2_c^{\mathcal{L}}$ to \mathcal{U} represents a revision operator \circ (and that \circ is represented by an assignment from $2_c^{\mathcal{L}}$ to \mathcal{U}) if, for any $\kappa \in 2_c^{\mathcal{L}}$ and $\mu \in \mathcal{L}$, it holds that $[\kappa \circ \mu] = \min_{\leq_{\kappa}} [\mu]$. Given an assignment from $2_c^{\mathcal{L}}$ to \mathcal{U} , $\kappa, \kappa_1, \kappa_2 \in 2_c^{\mathcal{L}}$ and $w_1, w_2 \in \mathcal{U}$, the following properties are of particular interest: (f₁) if $w_1, w_2 \in [\kappa]$, then $w_1 \approx_{\kappa} w_2$; (f₂) if $w_1 \in [\kappa]$ and $w_2 \notin [\kappa]$, then $w_1 <_{\kappa} w_2$; (f₃) if $\kappa_1 \equiv \kappa_2$, then $\leq_{\kappa_1} = \leq_{\kappa_2}$. A *faithful assignment* is an assignment from $2_c^{\mathcal{L}}$ to \mathcal{U} such that \leq_{κ} is a total preorder and properties f_{1–3} are satisfied. Faithful assignments provide the semantic counterpart to axioms R_{1–6}, as shown by the following representation result.

Theorem 1 ([Katsuno and Mendelzon, 1992]). A revision operator satisfies axioms R_{1–6} iff there exists a faithful assignment which represents it.

Revision by ranking sets of interpretations. By Theorem 1, revision according to axioms R_{1–6} is equivalent to ranking interpretations in a way that depends on κ , and choosing the most plausible allowed interpretations. Alternatively, we show, revision can be seen as choosing one element from a set of formulas (in semantic terms, from a set of *sets of interpretations*). If \mathcal{M} is a set, we write $\text{sub}(\mathcal{M}) = 2^{\mathcal{M}} \setminus \{\emptyset\}$ for the set of non-empty subsets of \mathcal{M} . We now switch to working with assignments from $2_c^{\mathcal{L}}$ to $\text{sub}(\mathcal{U})$. Given such an assignment, $\kappa, \kappa_1, \kappa_2 \in 2_c^{\mathcal{L}}$, and $W_1, W_2 \in \text{sub}(\mathcal{U})$, the following properties are of particular interest:

- (r₁) If $W_1 \neq W_2$, there is $W_3 \in \text{sub}(\mathcal{U})$ such that $W_3 \subseteq W_1 \cup W_2$ and $\min_{\leq_{\kappa}} \{W_1, W_2, W_3\} = \{W_3\}$.
- (r₂) If $W_1, W_2 \in \text{sub}([\kappa])$ and $W_1 \subseteq W_2$, then $W_2 \leq_{\kappa} W_1$.
- (r₃) If $W_1 \in \text{sub}([\kappa])$ and $W_2 \notin \text{sub}([\kappa])$, then $W_2 \not\leq_{\kappa} W_1$.
- (r₄) If $\kappa_1 \equiv \kappa_2$, then $\leq_{\kappa_1} = \leq_{\kappa_2}$.
- (r₅) If $W_1 \leq_{\kappa} W_3 \leq_{\kappa} W_2$ and $W_3 \subseteq W_2$, then $W_1 \cap W_2 \subseteq W_3$.
- (r₆) If $W_1 \leq_{\kappa} W_3 \leq W_2$, $W_3 \subseteq W_2$ and $W_1 \cap W_2 \neq \emptyset$, then $W_3 \subseteq W_1 \cap W_2$.

An *r-assignment* is an assignment from $2_c^{\mathcal{L}}$ to $\text{sub}(\mathcal{U})$ such that \leq_{κ} is a partial transitive relation and r_{1–6} are satisfied. The intention is to use properties r_{1–6} to represent revision operators satisfying axioms R_{1–6}: intuitively, sets of interpretations W_1 and W_2 stand for formulas, and $W_1 <_{\kappa} W_2$ encodes the fact that W_1 gets chosen over W_2 if a choice is to be made among them. Property r₁ says that, for any sets of interpretations W_1 and W_2 , there is $W_3 \subseteq W_1 \cup W_2$ strictly preferred among the three, i.e., W_3 would get chosen if the choice were between W_1, W_2 and W_3 . This property is equivalent to $\min_{\leq} \text{sub}([\mu])$ being a singleton, for any $\mu \in \mathcal{L}$: assuming two distinct \leq_{κ} -minimal subsets of $[\mu]$ generates a contradiction with r₁. Since we only consider assignments satisfying r₁, we henceforth identify $\min_{\leq} \text{sub}([\mu])$ with its unique element. Thus we can define a revision operator whose output is exactly this minimal element. For further intuitions, (r₂) states that supersets are preferred (if models of κ), (r₃) states that models of κ are preferred, (r₄) denotes syntax independence, and (r₅) and (r₆) can be seen as semantic counterparts of axioms R_{5–6}. We say that an assignment from $2_c^{\mathcal{L}}$ to $\text{sub}(\mathcal{U})$ satisfying r₁ represents a revision operator \circ if, for any $\kappa \in 2_c^{\mathcal{L}}$ and $\mu \in \mathcal{L}$, it holds that $[\kappa \circ \mu] = \min_{\leq_{\kappa}} \text{sub}([\mu])$.




Figure 1: Rankings associated to κ by a faithful assignment and an r-assignment; an arrow from x to y indicates that $x \leq_{\kappa} y$.

Example 1. Consider the alphabet $\mathcal{P} = \{a, b\}$, a knowledge base $\kappa = \{\neg a\}$ and an r-assignment that associates to κ the ranking \leq_{κ} in Figure 1-(b). Revision of κ by $\mu = a$ yields $\min_{\leq_{\kappa}} \text{sub}(\{\mu\}) = \min_{\leq_{\kappa}} \{\{a\}, \{ab\}, \{a, ab\}\} = \{a, ab\}$.

Finally, we convince ourselves that things are properly arranged with a representation result involving r-assignments.

Theorem 2. A revision operator satisfies R_{1-6} iff there exists an r-assignment which represents it.

Faithful assignments and r-assignments represent the same axioms, so one would expect there to be a correspondence between them. This correspondence can be made precise, as follows. If \leq^g is a relation on \mathcal{M} , the *lifting of \leq^g to $\text{sub}(\mathcal{M})$* is a relation \leq^l on $\text{sub}(\mathcal{M})$ defined, for any $W_1, W_2 \in \text{sub}(\mathcal{M})$, by taking $W_1 \leq^l W_2$ iff $(l_1) \min_{\leq^g} W_1 = W_1$, and either $(l_{21}) W_2 \subseteq W_1$, or $(l_{22}) w_1 <^g w_2$, for any $w_1 \in W_1$ and $w_2 \in W_2 \setminus W_1$. Conversely, if \leq^l is a relation on $\text{sub}(\mathcal{M})$, the *grounding of \leq^l to \mathcal{M}* is a relation \leq^g on \mathcal{M} defined, for any $w_1, w_2 \in \mathcal{M}$, by taking $w_1 \leq^g w_2$ iff $w_1 \in \min_{\leq^l} \text{sub}(\{w_1, w_2\})$. Notably, lifting a faithful assignment results in an r-assignment, while grounding an r-assignment results in a faithful assignment.

Example 2. For $\kappa = \{\neg a\}$, \leq_{κ}^r in Figure 1-(b) is obtained by lifting \leq_{κ}^f in Figure 1-(a): notice, e.g., that $\{\emptyset, b\} \leq_{\kappa}^r \{\emptyset, b\}$, since $\emptyset \approx_{\kappa}^f b$ (giving us l_1) and $\{\emptyset, b\} \subseteq \{\emptyset, b\}$ (giving us l_{21}); we also have $\{\emptyset\} <_{\kappa}^r \{\emptyset, a\}$, because $\emptyset \approx_{\kappa}^f \emptyset$ (giving us l_1) and $\emptyset <_{\kappa}^f a$ (giving us l_{22}). Conversely, \leq_{κ}^f is the grounding of \leq_{κ}^r , we get, e.g., that $\emptyset <_{\kappa}^f a$, because $\min_{\leq_{\kappa}^r} \text{sub}(\{\emptyset, a\}) = \min_{\leq_{\kappa}^r} \{\{\emptyset\}, \{a\}, \{\emptyset, a\}\} = \emptyset$.

3 Propositional Enforcement

Axiomatization. Enforcement stems from abstract argumentation, where new information can be brought to a debate, requiring reconsideration of the acceptance status of certain arguments [de Saint-Cyr *et al.*, 2016]. Concretely, an argumentation framework undergoes change such that its extensions (in this section, its *models*) include some distinguished elements. More broadly, daily life provides us with interactions where agents balance their positions in order to take into account differing points of view. Accommodation is achieved not necessarily through unconditional acceptance of one of the viewpoints (as in revision), but by finding a middle ground between the different sources. Such interactions often have a strategic flavor, and success consists in finding a position that

allows the possibility of having one's preferred options, while avoiding an uninformative solution such as a tautology.

Example 3. The prime-minister is weighing two responses to an economic downturn: a combination of increased taxes and spending cuts, i.e., austerity measures (a), and borrowing money (b). The prime-minister is privately in favor of implementing austerity ($\mu = a$), but polling reveals that the public is in favor of borrowing ($\kappa_1 = \{b\}$). The prime-minister wants to put forward a plan that leaves the option of austerity on the table, while at the same time appeasing the public. The most inclusive solution in this situation is $\kappa_1 \vee \mu \equiv a \vee b$.

Suppose that further polling reveals the public to be opposed to austerity ($\kappa_2 = \{\neg a\}$). In this case we have $\kappa_2 \vee \mu \equiv \neg a \vee a$, and the prime-minister would like to avoid putting forward a vacuous proposal. Nonetheless, the same considerations from before apply and we reason that the prime-minister's options range as follows: she could put forward a , if she is confident she can override public opinion; $b \rightarrow a$, for a compromise solution; or $a \vee b$, for a non-committal position.

We formalize these intuitions through an enforcement operator \triangleright , mapping a refutable (i.e., non-tautological) knowledge base $\kappa \in 2^{\mathcal{L}}$ and a formula $\mu \in \mathcal{L}$ to a knowledge base $\kappa \triangleright \mu \in 2^{\mathcal{L}}$. We restrict to refutable knowledge bases, since tautological formulas are a special case (no useful information and all interpretations are models). Following common belief change practice, we start by proposing a set of axioms:

- (E₁^P) $\mu \models \kappa \triangleright \mu$.
- (E₂^P) If $\kappa \vee \mu$ is refutable, then $\kappa \triangleright \mu \equiv \kappa \vee \mu$.
- (E₃^P) If μ is refutable, then $\kappa \triangleright \mu$ is refutable.
- (E₄^P) If $\kappa_1 \equiv \kappa_2$ and $\mu_1 \equiv \mu_2$, then $\kappa_1 \triangleright \mu_1 \equiv \kappa_2 \triangleright \mu_2$.
- (E₅^P) $\kappa \triangleright (\mu \vee \varphi) \models (\kappa \triangleright \mu) \vee \varphi$.
- (E₆^P) If $(\kappa \triangleright \mu) \vee \varphi$ is refutable, then $(\kappa \triangleright \mu) \vee \varphi \models \kappa \triangleright (\mu \vee \varphi)$.

Axioms E₁₋₆^P can be seen as dual to axioms R₁₋₆ for revision, and their logic is as follows. Axiom E₁^P formalizes the idea that the enforcement result takes μ into account, by requiring that μ implies it. Axiom E₂^P says that unless $\kappa \vee \mu$ is a tautology, in which case it would carry no useful information, we are to simply take the disjunction of κ and μ . Thus, ideally, enforcement shall accommodate (preserve) κ 's models. Axioms E₃^P and E₄^P say that enforcement should yield an informative result when μ is informative, and that this result should not depend on the syntax of κ and μ . Finally, axioms E₅^P and E₆^P ensure that the enforcement operator has a stable behavior when the enforcement formula is allowed to vary. For an intuition, consider Example 3: if the policy proposal ($\kappa_1 \triangleright \mu$) were to be augmented with a third option (c), then the axioms ensure equivalence to $\kappa_1 \triangleright (\mu \vee c)$ (the result is the same as if c was part of original proposal and public opinion needs to be accommodated either way).

Relating revision and enforcement. It is intended that the enforcement axioms E₁₋₆^P bear some resemblance to the revision axioms R₁₋₆, which raises the question of what connection there is between the two operations. The following theorem makes this connection precise.




Figure 2: Order associated with κ by an e-assignment.

Theorem 3. A propositional enforcement operator \triangleright satisfies E_{1-6}^P iff there is a propositional revision operator \circ satisfying R_{1-6} s.t., for any $\kappa \in 2_r^{\mathcal{L}}$ and $\mu \in \mathcal{L}$: $\kappa \triangleright \mu = \neg(\neg\kappa \circ \neg\mu)$.

Theorem 3 reveals that enforcement is a dual to revision. Intuitively, in order to accommodate μ , an enforcement operator *gives up* possible worlds that do not satisfy μ , in a manner that shows some preference for preserving as much as possible of the information in κ . Thus enforcement operators act as choice functions, choosing the ‘best’ elements from feasible options: the menu of feasible options here is the set of refutable supersets of $[\mu]$, defined as $\text{sup}([\mu]) = \{W \neq \mathcal{U} \mid [\mu] \subseteq W\}$. Deciding which of these supersets is best requires an underlying ranking over sets of interpretations, and to this purpose we employ assignments from $2_r^{\mathcal{L}}$ to $2_r^{\mathcal{U}}$, where $2_r^{\mathcal{U}} = 2^{\mathcal{U}} \setminus \{\mathcal{U}\}$. Given such an assignment, $\kappa, \kappa_1, \kappa_2 \in 2_r^{\mathcal{L}}$, and $W_1, W_2 \in 2_r^{\mathcal{U}}$, consider the following properties:

- (e₁) If $W_1 \neq W_2$, there is $W_3 \in \text{sup}(\mathcal{U})$ such that $W_1 \cap W_2 \subseteq W_3$ and $\min_{\leq_{\kappa}} \{W_1, W_2, W_3\} = \{W_3\}$.
- (e₂) If $W_1, W_2 \in \text{sup}([\kappa])$ and $W_1 \subseteq W_2$, then $W_1 \leq_{\kappa} W_2$.
- (e₃) If $W_1 \in \text{sup}([\kappa])$, $W_2 \notin \text{sup}([\kappa])$, then $W_2 \not\leq_{\kappa} W_1$.
- (e₄) If $\kappa_1 \equiv \kappa_2$, then $\leq_{\kappa_1} = \leq_{\kappa_2}$.
- (e₅) If $W_1 \leq_{\kappa} W_3 \leq_{\kappa} W_2$ and $W_2 \subseteq W_3$, then $W_3 \subseteq W_1 \cup W_2$.
- (e₆) If $W_1 \leq_{\kappa} W_3 \leq_{\kappa} W_2$, $W_2 \subseteq W_3$ and $W_1 \cup W_2 \neq \mathcal{U}$, then $W_1 \cup W_2 \subseteq W_3$.

An *e-assignment* α^e is an assignment from $2_r^{\mathcal{L}}$ to $2_r^{\mathcal{U}}$ such that \leq_{κ} is a partial transitive relation and e₁₋₆ are satisfied. An e-assignment α^e *represents an enforcement operator* \triangleright (alternatively, \triangleright is *represented by* α^e) if, for any $\kappa \in 2_r^{\mathcal{L}}$ and $\mu \in \mathcal{L}$, it holds that $[\kappa \triangleright \mu] = \min_{\leq_{\kappa}} \text{sup}([\mu])$.

Example 4. For the alphabet $\mathcal{P} = \{a, b\}$, take a knowledge base $\kappa = \{a\}$ and the ranking \leq_{κ} in Figure 2. Enforcement by $\mu = \neg a$ yields $[\kappa \triangleright \mu] = \min_{\leq_{\kappa}} \text{sup}([\mu]) = \min_{\leq_{\kappa}} \{\{\emptyset, b\}, \{\emptyset, a, b\}, \{\emptyset, b, ab\}\} = \{\emptyset, b\}$.

To establish a representation result, we use the connection to revision presented in Theorem 3. We show a similar correspondence between r-assignments and e-assignments.

Proposition 4. For any r-assignment α^r there is an e-assignment α^e such that, for any $\kappa \in 2_r^{\mathcal{L}}$ and $W_1, W_2 \in 2_r^{\mathcal{U}}$, it holds that $W_1 \leq_{\kappa}^e W_2$ iff $(\mathcal{U} \setminus W_1) \leq_{\neg\kappa}^r (\mathcal{U} \setminus W_2)$. Conversely, for any e-assignment α^e there is an r-assignment α^r such that $W_1 \leq_{\kappa}^r W_2$ iff $(\mathcal{U} \setminus W_1) \leq_{\neg\kappa}^e (\mathcal{U} \setminus W_2)$.

This correspondence facilitates the following result.

Theorem 5. A propositional enforcement operator satisfies axioms E_{1-6}^P iff there is an e-assignment representing it.

It is here that our representation of revision operators using rankings on sets of formulas pays off richly, as it highlights in a very clear manner the correspondence between revision and enforcement at the *semantic level*. A representation for enforcement using rankings on interpretations is unwieldy and lacks the intuitive appeal of an e-assignment.

Concrete enforcement operators. The *drastic enforcement operator* \triangleright^D , is defined, for any $\kappa \in 2_r^{\mathcal{L}}$ and $\mu \in \mathcal{L}$, by: $\kappa \triangleright^D \mu = \kappa \vee \mu$ if $\kappa \vee \mu$ is refutable, and μ otherwise.

A more fine-grained operator uses the Hamming distance between interpretations (i.e., the number of atoms on which they differ) to choose a set of interpretations situated between $[\mu]$ and $[\kappa \vee \mu]$. We call this *Dalal’s enforcement operator* \triangleright^H , as it uses the identity in Theorem 3 to adapt Dalal’s revision operator [Dalal, 1988] into an enforcement operator. Define, first, the minimum cardinality of the symmetric difference of models of two formulas as $|\Delta|^{min}(\varphi, \psi) = \min\{|w \Delta w'| \mid w \in [\varphi], w' \in [\psi]\}$, for $\varphi, \psi \in \mathcal{L}$. The operator \triangleright^H is now defined by taking $[\kappa \triangleright^H \mu] = \{w \in \mathcal{U} \mid \nexists w' \in [\neg\kappa] \text{ such that } |w \Delta w'| = |\Delta|^{min}(\neg\kappa, \neg\mu)\}$. Intuitively, $[\kappa \triangleright^H \mu]$ is obtained by removing the interpretations from $[\neg\mu]$ at minimum distance to $[\neg\kappa]$.

Example 5. For an alphabet $\mathcal{P} = \{a, b\}$, take $\kappa = \{\neg a \vee \neg b\}$, and suppose we want to enforce $\mu = a \wedge b$. For \triangleright^D we get $\kappa \triangleright^D \mu = \mu$, since $\kappa \vee \mu$ is a tautology. For \triangleright^H , we have $[\neg\kappa] = \{ab\}$ and $[\neg\mu] = \{\emptyset, a, b\}$. The models from $\neg\mu$ that have minimum distance to $\neg\kappa$ are a and b . Therefore, $[\kappa \triangleright^H \mu] = \mathcal{U} \setminus \{a, b\} = \{\emptyset, ab\} = [a \leftrightarrow b]$.

Drastic and Hamming enforcement operators satisfy E_{1-6}^P (via Theorem 3, any revision operator satisfying R_{1-6} can be converted into an enforcement operator satisfying E_{1-6}^P).

4 Enforcement in Argumentation

Existing forms of enforcement focus on the formal model of argumentation frameworks (AFs) [Dung, 1995], representable as directed graphs $F = (A, R)$ with A a set of arguments and R a directed conflict relation (*attacks*) between arguments. The study of AFs is mainly concerned with defining subsets of arguments (*extensions*) that can be accepted together when taking into account the graph structure. As a result, the argumentation literature offers a wide range of criteria (*AF semantics*) for establishing jointly acceptable arguments. Enforcing a set S with respect to F typically amounts to changing F such that F ’s set of extensions includes S [Baumann and Brewka, 2010; Baumann, 2012b].

In this section we focus on the *admissible semantics*, whose definition requires the following notions. For an AF $F = (A, R)$, a set $E \subseteq A$ is *conflict-free in F* if there is no directed edge between any arguments in E ; E is an *admissible extension* if E is conflict-free, and for each argument c attacking $a \in E$, there is an argument $b \in E$ that attacks c (b counters c , and can be thought of as defending a). If F is an AF, then $\text{adm}(F)$ is the set of its admissible extensions.

Example 6. Let $F = (\{a, b, c\}, \{(b, c)\})$ be an AF. We have $\text{adm}(F) = \{\emptyset, a, b, ab\}$. Enforcing $\{ac\}$ requires finding an AF F' such that $ac \in \text{adm}(F')$.

Notice the similarity between the extensions of an AF and the models of a propositional formula: though syntactically different, formulas and AFs share a common semantic representation, and it is this insight that raises the prospect of modelling enforcement as a belief change operation.

We wish to apply the axiomatic approach developed for propositional enforcement (see previous section) to AFs. In doing so, we will exploit the analogy between extensions and propositional models. However, we immediately run into a problem. In contrast to propositional logic, there are sets S of sets of arguments for which there is *no* AF F such that $adm(F) = S$. A set S of sets of arguments is *adm-realizable* if there exists an AF F such that $adm(F) = S$ [Dunne *et al.*, 2015]. Fortunately, we know exactly what properties S has to satisfy in order for such an F to exist. The *admissible closure* of a set S of extensions [Dunne *et al.*, 2015; Niskanen *et al.*, 2016] characterizes all *adm-realizable* sets, and is defined by taking $Cl(S) = \{E \mid E = \bigcup_{E' \in S'} E', S' \subseteq S, \forall a, b \in E \text{ it holds that } \exists E'' \in S \text{ s.t. } \{a, b\} \subseteq E''\}$. Intuitively, in the admissible closure of S , we add E , if E is the union of some sets in S , and each pair of arguments a, b from E is conflict-free, i.e., both occur together in some set in S . It holds that $Cl(S)$ is the smallest superset of S that is *adm-realizable*, and that S is *adm-realizable* iff $Cl(S) = S$. The *signature of the admissible semantics* $\Sigma_{adm} = \{adm(F) \mid F \text{ is an AF}\}$ thus can be characterized by $\Sigma_{adm} = \{S \mid Cl(S) = S\}$.

Example 7. If $M = \{\emptyset, a, ab, ac, bc\}$, then $Cl(M) = M \cup \{abc\}$ and thus M is not *adm-realizable*.

An *adm-enforcement* operator \triangleright_{adm} is a function which, given an AF F and $\mu \in \mathcal{L}$, maps $adm(F)$ and $[\mu]$ to a set of extensions $adm(F \triangleright_{adm} \mu)$. The propositional formula μ encodes the extensions to be enforced, under the assumption that sets of arguments are identified with interpretations and admissible extensions with models. The notation makes it clear that the intention is to realize $adm(F \triangleright_{adm} \mu)$ as an AF under the admissible semantics, and it is worth noting how the issue of realizability has a role to play: a set of extensions is *adm-realizable* as an AF only if it is in the signature Σ_{adm} of the admissible semantics; given our identification of extensions with interpretations, one can think of all *adm-realizable* sets as a subset of the semantics of propositional logic. In this way, enforcement on *adm-realizable* sets can be viewed as operating in a “fragment” of propositional logic.

Going from propositional enforcement to *adm-enforcement* of AFs, we say a set S is *refutable* if it is a proper subset of $2^{\mathcal{A}}$, with \mathcal{A} the set of all arguments occurring in F and μ . We propose the following axioms:

- (E₁^A) $[\mu] \subseteq adm(F \triangleright_{adm} \mu)$.
- (E₂^A) $adm(F \triangleright_{adm} \mu) \subseteq Cl(adm(F) \cup [\mu])$.
- (E₃^A) If $Cl(adm(F) \cup [\mu])$ is refutable, then $adm(F) \cup [\mu] \subseteq adm(F \triangleright_{adm} \mu)$.
- (E₄^A) If $adm(F_1) = adm(F_2)$ and $[\mu_1] = [\mu_2]$, then $adm(F_1 \triangleright_{adm} \mu_1) = adm(F_2 \triangleright_{adm} \mu_2)$.
- (E₅^A) $adm(F \triangleright_{adm} (\mu \vee \varphi)) \subseteq Cl(adm(F \triangleright_{adm} \mu) \cup [\varphi])$.
- (E₆^A) If $adm(F \triangleright_{adm} \mu) \cup [\varphi]$ is refutable, then $adm(F \triangleright_{adm} \mu) \cup [\varphi] \subseteq adm(F \triangleright_{adm} (\mu \vee \varphi))$.¹

¹Refutable refers here to arguments occurring in F , μ and φ .

Comparison with axioms E_{1-6}^P reveals differences, imposed by the particularities of AF semantics. One change concerns axiom E_2^A , which reflects a no-effort case: a straight adaptation of the propositional case suggests $adm(F) \cup [\mu]$ as default output, but this is not guaranteed to be *adm-realizable*.

Example 8. Let $\mathcal{A} = \{a, b, c\}$, $adm(F) = \{\emptyset, a, ab\}$ and $[\mu] = \{ac, bc\}$, we have $adm(F) \cup [\mu]$ is not *adm-realizable*.

Since $adm(F \triangleright_{adm} \mu)$ includes $[\mu]$ (by E_1^A), and is *adm-realizable*, we have $Cl([\mu]) \subseteq adm(F \triangleright_{adm} \mu)$. What we need from E_2^A is an *adm-realizable* set that acts as an approximation of $adm(F) \cup [\mu]$, and we suggest $Cl(adm(F) \cup [\mu])$ as an upper bound. To justify the choice of axioms E_3^A and E_5^A , notice that straightforward adaptations of E_3^P and E_5^P , given by E_7^A and E_8^A below, clash with the other postulates:

- (E₇^A) If $Cl([\mu])$ is refutable, $adm(F \triangleright_{adm} \mu)$ is refutable.
- (E₈^A) $adm(F \triangleright_{adm} \mu) \subseteq adm(F \triangleright_{adm} \mu) \cup [\varphi]$.

Proposition 6. There is no *adm-enforcement* operator satisfying axioms $E_{2-4,6-7}^A$, or $E_{1-2,4,8}^A$.

Thus, we have to weaken $E_{7,8}^A$. Axiom E_3^A sets a lower bound for $adm(F \triangleright_{adm} \mu)$ if $Cl(adm(F \triangleright_{adm} \mu))$ is refutable, and axiom E_5^A replaces $adm((F \triangleright_{adm} \mu) \cup [\varphi])$ with its closure.

Next, we ask whether there exist *adm-enforcement* operators satisfying axioms E_{1-6}^A . One natural idea is to use existing propositional enforcement operators, but this founders on the discovery that they typically do not guarantee an *adm-realizable* result, as witnessed by Examples 8 and 9.

Example 9. If $\mathcal{A} = \{a, b, c\}$, $adm(F) = \{\emptyset, b, c, ac, bc, abc\}$, and $[\mu] = \{a, b, ab, ac\}$, then \triangleright^H yields $\{a, b, c, ab, ac, bc\}$, which is not *adm-closed*, and thus not *adm-realizable*.

To ensure that enforcement yields an *adm-realizable* result, we propose to use propositional operators and coerce their results into Σ_{adm} of the *adm* semantics via the *adm-closure*.

Definition 1. Let \mathcal{F} be the set of all AFs, α an assignment from \mathcal{F} to $2^{\mathcal{U}}$ mapping every $F \in \mathcal{F}$ to a transitive order \leq_F such that \leq_F satisfies e_{1-6} . Then *the adm-refinement w.r.t.* α is an *adm-enforcement* operator $\triangleright_{adm}^{ref}$ and is defined as $adm(F \triangleright_{adm}^{ref} \mu) = Cl(\min_{\leq_F} \sup(\mu))$ for any $F \in \mathcal{F}, \mu \in \mathcal{L}$.

By Theorem 5, properties e_{1-6} are used to represent axioms E_{1-6}^P . Definition 1 can thus be seen as taking a propositional enforcement operator and refining its output such that it falls inside the signature Σ_{adm} . This is possible because of our semantic approach where we treat $adm(F)$ as the models of F , i.e., as a set of sets of atoms in the base alphabet. Interestingly, it turns out that refinements satisfy axioms E_{1-6}^A .

Proposition 7. If α is an assignment satisfying e_{1-6} , the *adm-refinement* with respect to α satisfies axioms E_{1-6}^A .

An equivalent result does not hold if we now switch to the duals of axioms E_{1-6}^A . This shows that the restriction to admissible sets makes revision and enforcement diverge: in this context they are no longer mirror images of each other.

Structural constraints and existing enforcement. We argue that imposing structural constraints on how an AF may be modified, axiomatization of enforcement requires semantical

characterizations of such constraints. To illustrate one such restriction, consider normal expansions [Baumann, 2012a]: given an AF $F = (A, R)$, another AF $F' = (A', R')$ is a normal expansion of F if $A \subseteq A'$, $R \subseteq R'$ and for each new attack $(a, b) \in R' \setminus R$ we have that a or b is an added argument (in $A' \setminus A$). (Attacks between existing arguments must stay untouched.) The signature of normal expansions is, for an AF $F = (A, R)$ with expanded arguments A' , $\Sigma_F^{A'} = \{adm(F') \mid F' = (A', R') \text{ a normal expansion of } F\}$.

Example 10. Assume an AF $F = (\{b, c\}, \{(b, c)\})$ and the goal to enforce $\{ac\}$ to be admissible under normal expansions that may add a ($A' = \{a, b, c\}$). There are only two sets of sets of arguments that include $\{ac\}$, are *adm*-realizable, and realizable via normal expansions, namely $W_1 = \{\emptyset, a, ac\}$ (adding (a, b)), and $W_2 = \{\emptyset, a, b, ac\}$ (adding (a, b) and (b, a)): $\Sigma_F^{A'} \cap \text{sup}(\{ac\}) = \{W_1, W_2\}$. Note that $adm(F) = \{\emptyset, b\}$ and $Cl_{adm}(adm(F) \cup \{ac\}) = adm(F) \cup \{ac\}$.

As illustrated by the example, there are cases where operator $\triangleright_{adm}^{\text{ref}}$ yields an *adm*-realizable set, yet not realizable via normal expansions. We now show axiom satisfaction when we define an operator that may choose arbitrarily from $\Sigma_F^{A'}$ as the outcome of enforcement, i.e., for an AF $F = (A, R)$, we define $adm(F \triangleright^s \mu) \in \Sigma_F^{A'}$, for pre-specified expanded arguments $A' \supseteq A$, and only require that $[\mu] \subseteq adm(F \triangleright^s \mu)$ (enforcement is successful). The alphabet (arguments) and \mathcal{A} (for defining refutability) for such operators is then equal to A' . For instance, in Example 10, such an operator may either choose W_1 or W_2 . We assume here that there always is a superset of $[\mu]$ in $\Sigma_F^{A'}$. Let \mathcal{S} be all such operators \triangleright^s . Independently of the choice, two axioms cannot be satisfied.

Proposition 8. If \triangleright^s is an operator such that $\triangleright^s \in \mathcal{S}$, then \triangleright^s satisfies axiom E_1^A , but neither E_2^A nor E_4^A .

In strict extension enforcement for normal expansions [Baumann and Brewka, 2010] for an AF $F = (A, R)$ and $S \subseteq 2^A$, the output is $F' = (A', R')$ s.t. F' is a normal expansion of F , $S \subseteq adm(F')$ and $|R \Delta R'|$ is minimal, breaking ties arbitrarily. We assume that F' is uniquely determined.

Example 11. Continuing Ex. 10, enforcing $\{ac\}$ under strict extension enforcement adds (a, b) , corresponding to W_1 .

This shows that axioms for enforcement under *adm* cannot be preserved under normal expansions. Generally, strict extension enforcement satisfies axiom E_1^A , but not E_{2-6}^A . Other variants of enforcement [Coste-Marquis *et al.*, 2015], under minimal changes to attacks, can be shown to violate at least one axiom. Rather than a shortcoming of the axiomatic approach, this shows that semantical characterizations of structural constraints are needed: e.g., under normal expansions admissible extensions can be preserved (see Ex. 10). However, axiom E_2^A is tailored to unrestricted enforcement under *adm*. What is needed to adapt a property such as $adm(F) \cup [\mu] \subseteq adm(F \triangleright^s \mu) \subseteq U$ to other types of enforcement is a new bound U . While semantic effects (e.g., adding an argument [Cayrol *et al.*, 2010]), have been studied, characterizations for structural constraints are missing. Writing good axioms for enforcement, we find, goes hand in hand with understanding the semantic effects of changing an AF's

structure, and the challenge for axiomatization is to take these effects into account while preserving the intuitions.

5 Related Work

Theorem 3 makes it clear that propositional enforcement is closely related to revision, deserving to be treated as a member of the wider family of belief change operators. The idea that enforcement accommodates the two sources recalls the operation of non-prioritized revision [Hansson, 1999], though our axiomatization sets it apart from these proposals. Our choice of representing the operators using rankings on *sets* of interpretations is reminiscent of other similar approaches [Grove, 1988; Gärdenfors and Makinson, 1988; Lindström and Rabinowicz, 1989; Rott, 1992; Meyer *et al.*, 2000] and bears close similarity to the monoselective approach [Hansson, 2015]. The axioms for propositional enforcement are similar to axioms proposed for revising desires [Dubois *et al.*, 2017], though used to different effect.

Adapting propositional enforcement to an AF setting recalls work on revision in fragments of propositional logic, e.g., Horn logic [Delgrande and Peppas, 2015]; in our case, the fragment is characterized in terms of Σ_{adm} . Our approach is most similar to [Creignou *et al.*, 2014] through the usage of a closure function that takes us into the fragment we are interested in, while the most striking contrast is our finding that a refined operator satisfies a set of intuitive axioms.

In [de Saint-Cyr *et al.*, 2016], enforcement is modeled as change on FO formulas, whose semantics are AFs. In [Doutre *et al.*, 2014], enforcing arguments is encoded in a logic for dynamics. Axioms in [de Saint-Cyr *et al.*, 2016] relate AFs encoded in the formulas, while Doutre *et al.* [2014] show that axioms can be defined via the logic. Booth *et al.* [2013] consider an agent that has beliefs and an AF. If inconsistent together, they study AF expansions to restore consistency. In [Moguillansky and Simari, 2016; Moguillansky *et al.*, 2013], acceptance revision operators are defined that incorporate a new logic (rule)-based argument, and they give axioms ensuring the argument's acceptance. Further, studies on enforcement in AFs concern characterizations [Baumann and Brewka, 2010; Baumann, 2012b; Coste-Marquis *et al.*, 2015] and computational results [Coste-Marquis *et al.*, 2015; Wallner *et al.*, 2017]. In contrast, we axiomatize the semantic change induced by our novel enforcement operator.

6 Conclusions

In this paper we have axiomatized enforcement in the spirit of belief change operators, and showed its duality to revision. We have made the case that the connection to revision is best seen through the lens of a new type of representation result, using rankings on sets of interpretations. Moreover, we have modified the axioms to fit AFs, showing that operators satisfying the modified axioms can be obtained in a seamless way if no structural restrictions are present. For future work, we think that our results should be extended by considering further semantics and characterizations of structural changes.

References

- [Alchourrón *et al.*, 1985] C. E. Alchourrón, P. Gärdenfors, and D. Makinson. On the logic of theory change: Partial meet contraction and revision functions. *J. Symb. Log.*, 50(2):510–530, 1985.
- [Baumann and Brewka, 2010] R. Baumann and G. Brewka. Expanding argumentation frameworks: Enforcing and monotonicity results. In *Proc. COMMA*, pages 75–86, 2010.
- [Baumann, 2012a] R. Baumann. Normal and strong expansion equivalence for argumentation frameworks. *Artif. Intell.*, 193:18–44, 2012.
- [Baumann, 2012b] R. Baumann. What does it take to enforce an argument? Minimal change in abstract argumentation. In *Proc. ECAI*, pages 127–132, 2012.
- [Bench-Capon and Dunne, 2007] T. J. M. Bench-Capon and P. E. Dunne. Argumentation in artificial intelligence. *Artif. Intell.*, 171(10-15):619–641, 2007.
- [Booth *et al.*, 2013] R. Booth, S. Kaci, T. Rienstra, and L. W. N. van der Torre. A logical theory about dynamics in abstract argumentation. In *Proc. SUM*, pages 148–161, 2013.
- [Cayrol *et al.*, 2010] C. Cayrol, F. Dupin de Saint-Cyr, and M.-C. Lagasquie-Schiex. Change in abstract argumentation frameworks: Adding an argument. *J. Artif. Intell. Res.*, 38:49–84, 2010.
- [Coste-Marquis *et al.*, 2015] S. Coste-Marquis, S. Konieczny, J.-G. Mailly, and P. Marquis. Extension enforcement in abstract argumentation as an optimization problem. In *Proc. IJCAI*, pages 2876–2882, 2015.
- [Creignou *et al.*, 2014] N. Creignou, O. Papini, R. Pichler, and S. Woltran. Belief revision within fragments of propositional logic. *J. Comput. Syst. Sci.*, 80(2):427–449, 2014.
- [Dalal, 1988] M. Dalal. Investigations into a theory of knowledge base revision. In *Proc. AAAI*, pages 475–479, 1988.
- [de Saint-Cyr *et al.*, 2016] F. Dupin de Saint-Cyr, P. Bisquert, C. Cayrol, and M.-C. Lagasquie-Schiex. Argumentation update in YALLA (Yet Another Logic Language for Argumentation). *Int. J. Approx. Reasoning*, 75:57–92, 2016.
- [Delgrande and Peppas, 2015] J. P. Delgrande and P. Peppas. Belief revision in Horn theories. *Artif. Intell.*, 218:1–22, 2015.
- [Doutre *et al.*, 2014] S. Doutre, A. Herzig, and L. Perrussel. A dynamic logic framework for abstract argumentation. In *Proc. KR*, pages 62–71, 2014.
- [Dubois *et al.*, 2017] D. Dubois, E. Lorini, and H. Prade. The strength of desires: A logical approach. *Minds and Machines*, 27(1):199–231, 2017.
- [Dung, 1995] P. M. Dung. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. *Artif. Intell.*, 77(2):321–358, 1995.
- [Dunne *et al.*, 2015] P. E. Dunne, W. Dvořák, T. Linsbichler, and S. Woltran. Characteristics of multiple viewpoints in abstract argumentation. *Artif. Intell.*, 228:153–178, 2015.
- [Falappa *et al.*, 2011] M. A. Falappa, A. J. García, G. Kern-Isberner, and G. R. Simari. On the evolving relation between belief revision and argumentation. *Knowledge Eng. Review*, 26(1):35–43, 2011.
- [Fermé *et al.*, 2013] E. L. Fermé, D. M. Gabbay, and G. R. Simari, editors. *Trends in Belief Revision and Argumentation Dynamics*. 2013.
- [Gärdenfors and Makinson, 1988] P. Gärdenfors and D. Makinson. Revisions of Knowledge Systems Using Epistemic Entrenchment. In *Proc. TARK*, pages 83–95, 1988.
- [Gärdenfors, 1988] P. Gärdenfors. *Knowledge in Flux: Modelling the Dynamics of Epistemic States*. 1988.
- [Grove, 1988] A. Grove. Two modellings for theory change. *J. Philosophical Logic*, 17(2):157–170, 1988.
- [Hansson, 1999] S. O. Hansson. A survey of non-prioritized belief revision. *Erkenntnis*, 50(2):413–427, 1999.
- [Hansson, 2015] S. O. Hansson. A monoselective presentation of AGM revision. *Studia Logica*, 103(5):1019–1033, 2015.
- [Katsuno and Mendelzon, 1992] H. Katsuno and A. O. Mendelzon. Propositional Knowledge Base Revision and Minimal Change. *Artif. Intell.*, 52(3):263–294, 1992.
- [Lindström and Rabinowicz, 1989] S. Lindström and W. Rabinowicz. Epistemic entrenchment with incomparabilities and relational belief revision. In *The Logic of Theory Change*, pages 93–126, 1989.
- [Meyer *et al.*, 2000] T. A. Meyer, W. A. Labuschagne, and J. Heidema. Refined Epistemic Entrenchment. *Journal of Logic, Language and Information*, 9(2):237–259, 2000.
- [Moguillansky and Simari, 2016] M. Moguillansky and G. R. Simari. A specialized set theoretic semantics for acceptability dynamics of arguments. In *Proc. COMMA*, pages 391–402, 2016.
- [Moguillansky *et al.*, 2013] M. O. Moguillansky, N. D. Rotstein, M. A. Falappa, A. J. García, and G. R. Simari. Dynamics of knowledge in *DeLP* through argument theory change. *TPLP*, 13(6):893–957, 2013.
- [Niskanen *et al.*, 2016] A. Niskanen, J. P. Wallner, and M. Järvisalo. Synthesizing argumentation frameworks from examples. In *Proc. ECAI*, pages 551–559, 2016.
- [Paglieri and Castelfranchi, 2004] F. Paglieri and C. Castelfranchi. Revising beliefs through arguments: Bridging the gap between argumentation and belief revision in MAS. In *Proc. ArgMAS*, pages 78–94, 2004.
- [Rott, 1992] H. Rott. Preferential belief change using generalized epistemic entrenchment. *Journal of Logic, Language and Information*, 1(1):45–78, 1992.
- [Wallner *et al.*, 2017] J. P. Wallner, A. Niskanen, and M. Järvisalo. Complexity results and algorithms for extension enforcement in abstract argumentation. *J. Artif. Intell. Res.*, 60:1–40, 2017.