Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence (IJCAI 2015)

An Extension-Based Approach to Belief Revision in Abstract Argumentation*

Martin Diller, Adrian Haret, Thomas Linsbichler, Stefan Rümmele, and Stefan Woltran
{diller,haret,linsbich,ruemmele,woltran}@dbai.tuwien.ac.at
Institute of Information Systems
Vienna University of Technology, Austria

Abstract

Argumentation is an inherently dynamic process. Consequently, recent years have witnessed tremendous research efforts towards an understanding of how the seminal AGM theory of belief change can be applied to argumentation, in particular for Dung's abstract argumentation frameworks (AFs). However, none of the attempts has yet succeeded in handling the natural situation where the revision of an AF is guaranteed to be representable by an AF as well. In this work, we present a generic solution to this problem which applies to many prominent I-maximal argumentation semantics. In order to prove a full representation theorem, we make use of recent advances in both areas of argumentation and belief change. In particular, we utilize the concepts of realizability in argumentation and the notion of compliance as used in Horn revision.

1 Introduction

Argumentation has become a major research area in Artificial Intelligence (AI) over the last two decades [Bench-Capon and Dunne, 2007; Rahwan and Simari, 2009]. This is not only because of the intrinsic interest of this topic and recent applications (e.g. in legal reasoning [Bench-Capon *et al.*, 2009] and E-Governance [Cartwright and Atkinson, 2009]) but also because there are fundamental connections between argumentation and other areas of AI, mainly non-monotonic reasoning.

The work by Dung [1995] on abstract argumentation, in particular, is usually seen as a significant landmark in the consolidation of the field of argumentation in AI. The central concern of abstract argumentation is the evaluation of a set of arguments and their relations in order to be able to extract subsets of the arguments, so called "extensions", that can all be accepted together from some point of view. Dung's argumentation frameworks (AFs), which are still the most widely used and investigated among the several argumentation formalisms, are directed graphs where nodes represent arguments and links correspond to one argument attacking another. The criteria or methods used to settle the acceptance of arguments, on the other hand, are called "semantics" (see [Baroni *et al.*, 2011] for a recent overview).

Given that argumentation can be viewed as a *process* as well as a product, recent years have seen an increasing number of studies on different problems in the *dynamics* of argumentation frameworks [Baumann, 2012; Bisquert *et al.*, 2011; 2013; Boella *et al.*, 2009; Booth *et al.*, 2013; Cayrol *et al.*, 2010; Doutre *et al.*, 2014; Kontarinis *et al.*, 2013; Krümpelmann *et al.*, 2012; Nouioua and Würbel, 2014; Sakama, 2014]. The problem we tackle here is how to revise an AF when some new information is provided. Along the lines of the AGM theory [Alchourrón *et al.*, 1985; Gärdenfors, 1988], by revision we mean an operation that incorporates the new information while bringing minimal change to the extensions of the original AF.

To the best of our knowledge, this has first been considered for AFs explicitly in [Coste-Marquis *et al.*, 2014a], where the problem of revision of AFs is defined as follows: given a semantics, an AF and a revision formula expressing how the status of some arguments has to be changed, find *a set of AFs* which satisfy the revision formula and whose extensions are as close as possible to the extensions of the input AF. Following the AGM approach, rationality postulates for a revision operator on AFs can be formulated and Coste-Marquis *et al.* [2014a] also provide a *representation theorem*. Such a result establishes a close link between obeying the postulates and exploiting a particular type of ranking on extensions of AFs in order to compute the output of revision. This approach is thus similar to the one by Katsuno and Mendelzon [1991].

In this work we study revision operators that produce a *single* AF as output. First, this is in accordance with the standard way of defining revision in the AGM theory where the result of revising an input theory by a revision formula is another theory. Second, revision yielding a single AF also makes concepts of iterated revision [Darwiche and Pearl, 1997; Spohn, 1988] amenable to argumentation. More specifically, we study two types of revision operators. The first type considers the new information represented as a propositional formula expressing the desired change in the extensions of the original AF. The second type is revision by an AF. Hence, the new information is restricted in the sense that it can only stem from another AF's outcome. While the first type follows the framework of [Coste-Marquis *et al.*, 2014a], the latter assumes that the knowledge to be incorporated (e.g. another agent's beliefs) is in the form of an AF. It is more in line with work on Horn

*Supported by the Austrian Science Fund (FWF) under grants I1102, P25518, P25521, W1255, and Y698.

revision [Delgrande and Peppas, 2015] where all involved formulas stem from the same fragment.

Our main contributions are as follows:

• We derive full *representation theorems* for both mentioned types of revision; our results are, moreover, *generic* in the sense that they hold for a wide range of semantics including preferred, semi-stable, stage, and stable semantics.

• For the revision-by-formula approach, we give novel notions of *compliance* [Delgrande and Peppas, 2015] in order to restrict the rankings. This is required to guarantee that the outcome of the corresponding operators can be expressed by an AF under a given semantics. To this end, exact knowledge about the expressiveness of argumentation semantics is needed. For most of the standard semantics, [Dunne *et al.*, 2014] provides the necessary results.

• In the revision-by-AF approach, we show that the concept of compliance can be dropped, thus standard revision operators satisfying all postulates like Dalal's [1988] operator can be directly applied to revision of AFs. However, an additional postulate (again borrowed from [Delgrande and Peppas, 2015]) is needed for the representation theorem.

2 Preliminaries

We first recall basic notions of Dungs's abstract frameworks (the reader is referred to [Dung, 1995; Baroni *et al.*, 2011] for further background), then present recent results from [Dunne *et al.*, 2014] which we require for our results and finally define rankings as used in belief change in terms of extensions.

We assume an arbitrary but finite domain $\mathfrak{A}$ of arguments. An argumentation framework (AF) is a pair $F = (A, R)$ where $A \subseteq \mathfrak{A}$ is non-empty, and $R \subseteq A \times A$ is the attack relation. The collection of all AFs is given as $AF_{\mathfrak{A}}$.

Given $F = (A, R)$, an argument $a \in A$ is *defended* (in F) by a set $S \subseteq A$ if for each $b \in A$ such that $(b, a) \in R$, there is a $c \in S$ with $(c, b) \in R$. A set T of arguments is defended (in F) by S if each $a \in T$ is defended by S (in F). A set $S \subseteq A$ is *conflict-free* (in F), if there are no arguments $a, b \in S$, such that $(a, b) \in R$. We denote the set of all conflict-free sets in F as $cf(F)$. A set $S \in cf(F)$ is called *admissible* (in F) if S defends itself. We denote the set of admissible sets in F as $adm(F)$. For $S \subseteq A$, the range of S (wrt. F), denoted S^+_F, is the set $S \cup \{a \mid \exists s \in S : (s, a) \in R\}$.

A *semantics* maps each $F \in AF_{\mathfrak{A}}$ to a set of extensions $\mathbb{S} \subseteq 2^{\mathfrak{A}}$. For the stable, preferred, stage, and semi-stable semantics respectively, the extensions are defined as follows:

- $S \in stb(F)$, if $S \in cf(F)$ and $S^+_F = A$;
- $S \in prf(F)$, if $S \in adm(F)$ and $\nexists T \in adm(F)$ s.t. $T \supset S$;
- $S \in stg(F)$, if $S \in cf(F)$ and $\nexists T \in cf(F)$ with $T^+_F \supset S^+_F$;
- $S \in sem(F)$, if $S \in adm(F)$ and $\nexists T \in adm(F)$ s.t. $T^+_F \supset S^+_F$.

The signature Σ_σ of a semantics σ is defined as $\Sigma_\sigma = \{\sigma(F) \mid F \in AF_{\mathfrak{A}}\}$, containing exactly those sets of extension which can be *realized* under σ. Exact characterizations of those sets for the aforementioned semantics have been given in [Dunne *et al.*, 2014]. If S_1 and S_2 are two extensions such that $S_1 \neq S_2$, we say that S_1 and S_2 are *⊆-comparable* if $S_1 \subseteq S_2$ or $S_2 \subseteq S_1$. We say that S_1 and S_2 are *⊆-incomparable* if they are not ⊆-comparable. Some of our results will apply to semantics for which the following properties hold in terms of realizability:

Definition 1. A semantics σ is called *proper I-maximal* if for each $\mathbb{S} \in \Sigma_\sigma$ it holds that (i) for any $S_1, S_2 \in \mathbb{S}$, $S_1 \subseteq S_2$ implies $S_1 = S_2$; (ii) $\mathbb{S}' \in \Sigma_\sigma$ for any $\mathbb{S}' \subseteq \mathbb{S}$ with $\mathbb{S}' \neq \emptyset$; and (iii) for any ⊆-incomparable $S_1, S_2 \in 2^{\mathfrak{A}}$ it holds that $\{S_1, S_2\} \in \Sigma_\sigma$.

In words, an I-maximal [Baroni and Giacomin, 2007] semantics σ is proper if, for any AF F, we can realize any non-empty subset of $\sigma(F)$ under σ and on the other hand, for any pair of ⊆-incomparable sets of arguments, we can find an AF having exactly these as extensions under σ.

The next observation follows from [Dunne *et al.*, 2014].

Proposition 1. *Preferred, stable, semi-stable and stage semantics are proper I-maximal.*

Definition 2. Given a semantics σ, we define the function $f_\sigma : 2^{2^{\mathfrak{A}}} \mapsto AF_{\mathfrak{A}}$ mapping sets of extensions to AFs such that $f_\sigma(\mathbb{S}) = F$ with $\sigma(F) = \mathbb{S}$ if $\mathbb{S} \in \Sigma_\sigma$ and $f_\sigma(\mathbb{S}) = (\emptyset, \emptyset)$ otherwise.

Note that $\mathbb{S} \in \Sigma_\sigma$ guarantees that we can find an AF which, when evaluated under σ results in $\mathbb{S}$. We leave the exact specifications of such AFs open; canonical constructions for the semantics we consider can be found in [Dunne *et al.*, 2014]. Such constructions may result in AFs with additional arguments to those contained in a $\mathbb{S} \in \Sigma_\sigma$, although recent work on realizability in compact AFs [Baumann *et al.*, 2014] could pave the way for constructions of AFs without new arguments. In general, f_σ is not unique. Nevertheless, throughout the paper we assume f_σ to be fixed for every σ.

By $\mathcal{P}_{\mathfrak{A}}$ we denote the set of propositional formulas over $\mathfrak{A}$, where the arguments in $\mathfrak{A}$ represent propositional variables. A set of arguments $E \subseteq \mathfrak{A}$ can be seen as an interpretation, where $a \in E$ means that a is assigned *true* and $a \notin E$ means that a is assigned *false*. If a formula $\varphi \in \mathcal{P}_{\mathfrak{A}}$ evaluates to *true* under an interpretation E, E is a model of φ. $[\varphi]$ denotes the set of models of φ. Moreover, $\varphi_1 \equiv \varphi_2$ if $[\varphi_1] = [\varphi_2]$.

A *pre-order* $\preceq$ *on* $2^{\mathfrak{A}}$ is a reflexive, transitive binary relation on $2^{\mathfrak{A}}$. If $E_1 \preceq E_2$ or $E_2 \preceq E_1$ for any $E_1, E_2 \in 2^{\mathfrak{A}}$, the pre-order $\preceq$ is *total*. Moreover, for $E_1, E_2 \in 2^{\mathfrak{A}}$, $E_1 \prec E_2$ denotes the strict part of $\preceq$, i.e. $E_1 \preceq E_2$ and $E_2 \not\preceq E_1$. We write $E_1 \approx E_2$ the case $E_1 \preceq E_2$ and $E_2 \preceq E_1$. An *I-total pre-order on* $2^{\mathfrak{A}}$ is a pre-order on $2^{\mathfrak{A}}$ such that $E_1 \preceq E_2$ or $E_2 \preceq E_1$ for any pair E_1, E_2 of ⊆-incomparable extensions. Finally, for a set of sets of arguments $\mathbb{S} \subseteq 2^{\mathfrak{A}}$ and a pre-order $\preceq$, $\min(\mathbb{S}, \preceq) = \{E_1 \in \mathbb{S} \mid \nexists E_2 \in \mathbb{S} : E_2 \prec E_1\}$.

3 Representation Theorems

A key insight in belief change is the realization that any belief revision operator can be characterized using rankings on the possible worlds described by the language. Intuitively, the rankings can be thought of as plausibility relations on possible worlds. Revision by a formula φ then amounts to choosing the most plausible worlds among the models of φ. The fact that this strategy is sound with respect to the postulates is guaranteed by a representation result [Katsuno and Mendelzon, 1991].

In our approach, subsets of $2^{\mathfrak{A}}$ play the role both of extensions for AFs and models of propositional formulas, and will be the possible worlds that a revision operator chooses from. Consequently, we use rankings on $2^{\mathfrak{A}}$ to characterize the class of revision operators for AFs. We employ two main types of rankings, introduced below.

Definition 3. Given a semantics σ, an *(I-)faithful assignment* maps every $F \in AF_{\mathfrak{A}}$ to an (I-)total pre-order $\preceq_F$ on $2^{\mathfrak{A}}$ such that, for any ($\subseteq$-incomparable) $E_1, E_2 \in 2^{\mathfrak{A}}$ and $F, F_1, F_2 \in AF_{\mathfrak{A}}$, it holds that:

(i) if $E_1, E_2 \in \sigma(F)$, then $E_1 \approx E_2$,

(ii) if $E_1 \in \sigma(F)$ and $E_2 \notin \sigma(F)$, then $E_1 \prec_F E_2$,

(iii) if $\sigma(F_1) = \sigma(F_2)$, then $\preceq_{F_1} = \preceq_{F_2}$.

The pre-order $\preceq_F$ assigned to F by an (I-)faithful assignment is referred to as the *(I-)faithful ranking associated with* F.

Faithful assignments approximate the notion with the same name introduced by Katsuno and Mendelzon [1991]. I-faithful assignments differ in that they require the rankings to be I-total, thus allowing (but not requiring) them to be partial with respect to $\subseteq$-comparable pairs of extensions. Section 3.2 explains why we need to introduce this second type.

3.1 Revision by Propositional Formulas

We first consider revision of an AF by a propositional formula. Hence we are interested in operators of the form $\star_\sigma : AF_{\mathfrak{A}} \times \mathcal{P}_{\mathfrak{A}} \mapsto AF_{\mathfrak{A}}$ mapping an AF F and a consistent[1] propositional formula φ to a revised AF $F \star_\sigma \varphi$. Intuitively, φ describes information which should be incorporated in F. To this end, the operator revises F such that the σ-extensions of F change according to the models of φ. We define the revision postulates adjusted to the AF scenario in a similar manner to [Coste-Marquis *et al.*, 2014a].

(P$\star$1) $\sigma(F \star_\sigma \varphi) \subseteq [\varphi]$.

(P$\star$2) If $\sigma(F) \cap [\varphi] \neq \emptyset$ then $\sigma(F \star_\sigma \varphi) = \sigma(F) \cap [\varphi]$.

(P$\star$3) If $[\varphi] \neq \emptyset$ then $\sigma(F \star_\sigma \varphi) \neq \emptyset$.

(P$\star$4) If $\varphi \equiv \psi$ then $\sigma(F \star_\sigma \varphi) = \sigma(F \star_\sigma \psi)$.

(P$\star$5) $\sigma(F \star_\sigma \varphi) \cap [\psi] \subseteq \sigma(F \star_\sigma (\varphi \wedge \psi))$.

(P$\star$6) If $\sigma(F \star_\sigma \varphi) \cap [\psi] \neq \emptyset$ then $\sigma(F \star_\sigma (\varphi \wedge \psi)) \subseteq \sigma(F \star_\sigma \varphi) \cap [\psi]$.

P$\star$1 says that when we revise by φ, the extensions of the revision output should be among the interpretations that satisfy φ. P$\star$2 specifies that if φ is consistent with F (in the sense that they share models), revision amounts to nothing more than taking the common models. P$\star$3 says that if φ is a consistent formula, then revision by φ should also be consistent. P$\star$4 guarantees irrelevance of syntax. P$\star$5 and P$\star$6 ensure that revision is performed with minimal change to the AF F. For further discussion on the motivation of the postulates, see [Alchourrón *et al.*, 1985; Gärdenfors, 1988; Katsuno and Mendelzon, 1991].

Due to reasons pointed out in the introduction we require the result of the revision to be a single AF. For this reason we face a similar challenge to that encountered in Horn revision [Delgrande and Peppas, 2015]. Specifically, it may be the case that a set of extensions $\mathbb{S}$ that is the desired outcome of the revision under a semantics σ cannot be realized under σ. To overcome this problem we use Σ_σ to define the following restriction on pre-orders, which we will need to obtain our representation theorem:

Definition 4. A pre-order $\preceq$ is σ-compliant if for every consistent formula $\varphi \in \mathcal{P}_{\mathfrak{A}}$ it holds that $\min([\varphi], \preceq) \in \Sigma_\sigma$.

Example 1. Consider the pre-order $\preceq$ defined as $\{a,b,c\} \prec \{a,b\} \approx \{a,c\} \approx \{b,c\} \prec \{a\} \approx \{b\} \approx \{c\} \prec \emptyset$. Now let $\varphi = \neg(a \wedge b \wedge c)$ and observe that $\min([\varphi], \preceq) = \{\{a,b\},\{a,c\},\{b,c\}\}$. From the results in [Dunne *et al.*, 2014] we know that for $\sigma \in \{\mathit{stb}, \mathit{prf}, \mathit{stg}, \mathit{sem}\}$, $\{\{a,b\},\{a,c\},\{b,c\}\} \notin \Sigma_\sigma$, hence $\preceq$ is not σ-compliant.

On the other hand let $\preceq'$ be the pre-order defined as $\{a,b,c\} \prec' \{a\} \approx' \{b\} \approx' \{c\} \prec' \{a,b\} \prec' \{a,c\} \prec' \{b,c\} \prec' \emptyset$ followed by an arbitrary $\prec'$-chain of the remaining elements of $2^{\mathfrak{A}}$. One can check that $\preceq'$ is σ-compliant. For instance, $\min([\varphi], \preceq') = \{\{a\},\{b\},\{c\}\} \in \Sigma_\sigma$.

Note that for semantics considered in [Dunne *et al.*, 2014], their results imply that checking whether a given pre-order is σ-compliant can be done in polynomial time. Using the notion of σ-compliance enables us to extend the AGM approach to capture the revision of argumentation frameworks under proper I-maximal semantics by propositional formulas.

Theorem 1. *Let σ be a proper I-maximal semantics, $F \in AF_{\mathfrak{A}}$ and $\preceq_F$ a σ-compliant and faithful ranking associated with F. Define an operator $\star_\sigma : AF_{\mathfrak{A}} \times \mathcal{P}_{\mathfrak{A}} \mapsto AF_{\mathfrak{A}}$ by*

$$F \star_\sigma \varphi = f_\sigma(\min([\varphi], \preceq_F)).$$

Then $\star_\sigma$ satisfies postulates P$\star$1 – P$\star$6.

Proof. First of all, by the definition of f_σ and due to the fact that $\preceq_F$ is σ-compliant, we have that $\sigma(f_\sigma(\min([\varphi], \preceq_F))) = \min([\varphi], \preceq_F)$, hence $\sigma(F \star_\sigma \varphi) = \min([\varphi], \preceq_F)$. Therefore postulates P$\star$1 and P$\star$4 follow immediately.

If $\sigma(F) \cap [\varphi] \neq \emptyset$, it follows from $\preceq_F$ being faithful that $\min([\varphi], \preceq_F) = \sigma(F) \cap [\varphi]$, satisfying P$\star$2.

P$\star$3 follows since $\preceq_F$ is transitive and $\mathfrak{A}$ is finite and therefore if $[\varphi] \neq \emptyset$ then $[\varphi]$ has minimal elements, hence $\min([\varphi], \preceq_F) \neq \emptyset$.

P$\star$5 and P$\star$6 are trivially satisfied if $\sigma(F \star_\sigma \varphi) \cap [\psi] = \emptyset$. Assume $\sigma(F \star_\sigma \varphi) \cap [\psi] \neq \emptyset$ and, towards a contradiction, that there is some $E \in \min([\varphi], \preceq_F) \cap [\psi]$ with $E \notin \min([\varphi \wedge \psi], \preceq_F)$. Since $E \in [\varphi \wedge \psi]$ there must be some $E' \in [\varphi \wedge \psi]$ with $E' \prec_F E$, a contradiction to $E \in \min([\varphi], \preceq_F)$. Therefore $\sigma(F \star_\sigma \varphi) \cap [\psi] \subseteq \sigma(F \star_\sigma (\varphi \wedge \psi))$. To show that also $\sigma(F \star_\sigma (\varphi \wedge \psi)) \subseteq \sigma(F \star_\sigma \varphi) \cap [\psi]$, assume $E \in \min([\varphi \wedge \psi], \preceq_F)$ and $E \notin \min([\varphi], \preceq_F) \cap [\psi]$. Since $E \in [\psi]$, $E \notin \min([\varphi], \preceq_F)$. Let $E' \in \min([\varphi], \preceq_F) \cap [\psi]$ (assumed to be nonempty). Then $E' \in [\varphi \wedge \psi]$ holds. As $E \in \min([\varphi \wedge \psi], \preceq_F)$ and $\preceq_F$ is total, $E \preceq_F E'$. Hence $E \in \min([\varphi], \preceq_F)$ follows from $E' \in \min([\varphi], \preceq_F)$, a contradiction. □

[1] The restriction to consistent formulas is due to the fact that argumentation semantics usually are not capable of expressing the empty set of extensions. For semantics which can realize the empty set, such as the stable semantics, our results in this section apply even without this restriction.

Theorem 2. *Let $\star_\sigma : AF_{\mathfrak{A}} \times \mathcal{P}_{\mathfrak{A}} \mapsto AF_{\mathfrak{A}}$ be an operator satisfying postulates P$\star$1 – P$\star$6 for a proper I-maximal semantics σ. Then, there exists a faithful assignment mapping every $F \in AF_{\mathfrak{A}}$ to a faithful ranking $\preceq_F$ on $2^{\mathfrak{A}}$ such that $\preceq_F$ is σ-compliant and $\sigma(F \star_\sigma \varphi) = \min([\varphi], \preceq_F)$ for every $\varphi \in \mathcal{P}_{\mathfrak{A}}$.*

Proof. For a set of interpretations $\mathbb{S}$, we denote by $\phi(\mathbb{S})$ a formula with $[\phi(\mathbb{S})]=\mathbb{S}$. If the elements of $\mathbb{S} = \{E_1, \dots, E_n\}$ are given explicitly we also write $\phi(E_1, \dots, E_n)$ for $\phi(\mathbb{S})$.

We define the binary relation $\preceq_F$ on $2^{\mathfrak{A}}$ as follows:

$$E \preceq_F E' \text{ iff } E \in \sigma(F \star_\sigma \phi(E, E')).$$

We begin by showing that $\preceq_F$ is a total pre-order. It follows from P$\star$1 and P$\star$3 that $\sigma(F \star_\sigma \phi(E, E'))$ is a non-empty subset of $\{E, E'\}$. Therefore $\preceq_F$ is total. Moreover, if $E = E'$ then $\sigma(F \star_\sigma \phi(E)) = \{E\}$. Hence $E \preceq_F E$ holds for each $E \in 2^{\mathfrak{A}}$, i.e., $\preceq_F$ is reflexive.

In order to show transitivity of $\preceq_F$, let $E_1, E_2, E_3 \in 2^{\mathfrak{A}}$ and assume $E_1 \preceq_F E_2$ and $E_2 \preceq_F E_3$. By P$\star$1 and P$\star$3, $\sigma(F \star_\sigma \phi(E_1, E_2, E_3))$ is a non-empty subset of $\{E_1, E_2, E_3\}$. First assume $\sigma(F \star_\sigma \phi(E_1, E_2, E_3)) \cap \{E_1, E_2\} = \emptyset$. Then $\sigma(F \star_\sigma \phi(E_1, E_2, E_3)) = \{E_3\}$. Knowing that $\phi(E_2, E_3) \equiv \phi(E_2, E_3) \wedge \phi(E_1, E_2, E_3)$, we obtain from P$\star$5 and P$\star$6 that $\sigma(F \star_\sigma \phi(E_1, E_2, E_3)) \cap \{E_2, E_3\} = \sigma(F \star_\sigma \phi(E_2, E_3))$. But this implies $\sigma(F \star_\sigma \phi(E_2, E_3)) = \{E_3\}$, a contradiction to $E_2 \preceq_F E_3$.

On the other hand assume $\sigma(F \star_\sigma \phi(E_1, E_2, E_3)) \cap \{E_1, E_2\} \neq \emptyset$. Since $E_1 \preceq_F E_2$ we know that $E_1 \in \sigma(F \star_\sigma \phi(E_1, E_2))$ holds. Considering the fact that $\phi(E_1, E_2) \equiv \phi(E_1, E_2) \wedge \phi(E_1, E_2, E_3)$, we obtain from P$\star$5 and P$\star$6 that $\sigma(F \star_\sigma \phi(E_1, E_2, E_3)) \cap \{E_1, E_2\} = \sigma(F \star_\sigma \phi(E_1, E_2))$. Thus, $E_1 \in \sigma(F \star_\sigma \phi(E_1, E_2, E_3)) \cap \{E_1, E_2\}$. Also $E_1 \in \sigma(F \star_\sigma \phi(E_1, E_2, E_3)) \cap \{E_1, E_3\}$ holds. By $\phi(E_1, E_3) \equiv \phi(E_1, E_3) \wedge \phi(E_1, E_2, E_3)$, we obtain from P$\star$5 and P$\star$6 that $\sigma(F \star_\sigma \phi(E_1, E_2, E_3)) \cap \{E_1, E_3\} = \sigma(F \star_\sigma \phi(E_1, E_3))$. Therefore $E_1 \in \sigma(F \star_\sigma \phi(E_1, E_3))$, meaning that $E_1 \preceq_F E_3$.

Having shown that $\preceq_F$ is total, reflexive and transitive, it follows that $\preceq_F$ is a total pre-order. The following lemmata show that $\star_\sigma$ can indeed be simulated by $\preceq_F$.

Lemma 1. *Let $E_1, E_2 \in 2^{\mathfrak{A}}$ such that $E_1 \preceq_F E_2$. Then for all formulas $\varphi \in \mathcal{P}_{\mathfrak{A}}$, if $E_1 \in [\varphi]$ and $E_2 \in \sigma(F \star_\sigma \varphi)$ then $E_1 \in \sigma(F \star_\sigma \varphi)$.*

Proof. Let φ be a formula such that $E_1 \in [\varphi]$ and $E_2 \in \sigma(F \star_\sigma \varphi)$. Then from P$\star$5 and P$\star$6 it follows that $\sigma(F \star_\sigma (\varphi \wedge \phi(E_1, E_2))) = \sigma(F \star_\sigma \varphi) \cap [\phi(E_1, E_2)]$. Moreover, from $E_2 \in \sigma(F \star_\sigma \varphi)$ and P$\star$1 we derive that $E_2 \in [\varphi]$, hence $[\phi(E_1, E_2)] \subseteq [\varphi]$. By P$\star$4 we now get $\sigma(F \star_\sigma (\varphi \wedge \phi(E_1, E_2))) = \sigma(F \star_\sigma \phi(E_1, E_2))$. Therefore, by P$\star$5 and P$\star$6, $\sigma(F \star_\sigma \phi(E_1, E_2)) = \sigma(F \star_\sigma \varphi) \cap [\phi(E_1, E_2)]$. This, together with the assumption that $E_1 \preceq_F E_2$ (and therefore $E_1 \in \sigma(F \star_\sigma \phi(E_1, E_2))$), entails $E_1 \in \sigma(F \star_\sigma \varphi)$. ☐

Lemma 2. *For $\varphi \in \mathcal{P}_{\mathfrak{A}}$, $\min([\varphi], \preceq_F) = \sigma(F \star_\sigma \varphi)$ holds.*

Proof. $\subseteq$: Let $\varphi \in \mathcal{P}_{\mathfrak{A}}$. Towards a contradiction assume that there is some $E_1 \in \min([\varphi], \preceq_F)$ such that $E_1 \notin \sigma(F \star_\sigma \varphi)$. Since $[\varphi] \neq \emptyset$ it follows by P$\star$3 that $\sigma(F \star_\sigma \varphi) \neq \emptyset$. Let $E_2 \in \sigma(F \star_\sigma \varphi)$. Lemma 1 entails that $E_1 \not\preceq_F E_2$. This also means, recalling that $E_1 \in \min([\varphi], \preceq_F)$, that $E_2 \not\preceq_F E_1$. But this means by the definition of $\preceq_F$ that $\sigma(F \star_\sigma \phi(E_1, E_2)) = \emptyset$, a contradiction to $\star_\sigma$ satisfying P$\star$3.

$\supseteq$: Let $\varphi \in \mathcal{P}_{\mathfrak{A}}$ and $E_1 \in \sigma(F \star_\sigma \varphi)$. We show that for all $E_2 \in [\varphi]$ it holds that $E_1 \preceq_F E_2$. To this end let $E_2 \in [\varphi]$. From $E_1 \in \sigma(F \star_\sigma \varphi)$ we know that $\sigma(F \star_\sigma \varphi) \cap [\phi(E_1, E_2)] \neq \emptyset$. By P$\star$5 and P$\star$6 we get $\sigma(F \star_\sigma \varphi) \cap [\phi(E_1, E_2)] = \sigma(F \star_\sigma (\varphi \wedge \phi(E_1, E_2)))$. Since $E_1, E_2 \in [\varphi]$ it follows by P$\star$4 that $\sigma(F \star_\sigma (\varphi \wedge \phi(E_1, E_2))) = \sigma(F \star_\sigma \phi(E_1, E_2))$. Now as $E_1 \in \sigma(F \star_\sigma \varphi)$ by assumption, also $E_1 \in \sigma(F \star_\sigma \phi(E_1, E_2))$, meaning that $E_1 \preceq_F E_2$. Since E_2 was chosen arbitrarily, $E_1 \in \min([\varphi], \preceq_F)$ follows. ☐

From Lemma 2 it follows that $\preceq_F$ is σ-compliant. It remains to show that $\preceq_F$ is faithful wrt. F. If $\sigma(F) = \emptyset$ this is trivially the case. Therefore assume $\sigma(F) \neq \emptyset$. By P$\star$2 we get $\sigma(F \star_\sigma \top) = \sigma(F)$ (note that $[\top] = 2^{\mathfrak{A}}$). Hence $\sigma(F) = \min([\top], \preceq_F) = \min(2^{\mathfrak{A}}, \preceq_F)$, meaning that for $E_1, E_2 \in 2^{\mathfrak{A}}$, $E_1 \approx_F E_2$ if $E_1, E_2 \in \sigma(F)$ and $E_1 \prec_F E_2$ if $E_1 \in \sigma(F)$ and $E_2 \notin \sigma(F)$. Therefore conditions (i) and (ii) from Definition 3 are fulfilled. Condition (iii) holds since $\preceq_F$ is defined with respect to the extensions of F, hence $\preceq_F$ is faithful wrt. F. ☐

The attentive reader might have noticed that we did not make explicit use of the restriction to proper I-maximal semantics in the proofs above. It is rather used implicitely since, in general, rankings which are both faithful and σ-compliant only exist if σ fulfills property (ii) of proper I-maximality.

We can use the representation results obtained from Theorems 1 and 2 to define concrete operators via faithful and compliant rankings. For instance, the ranking $\preceq_F$ where the σ-extensions of F are the minimal elements and the remaining candidates in $2^{\mathfrak{A}}$ are ordered as a $\prec$-chain leads to a simple but natural operator for any semantics σ. The concrete contents of Σ_σ will be the crucial aspect to consider when defining more refined operators under a certain semantics σ.

3.2 Revision by Argumentation Frameworks

In this section we investigate operators $*_\sigma : AF_{\mathfrak{A}} \times AF_{\mathfrak{A}} \mapsto AF_{\mathfrak{A}}$. Such operators map an AF F and an AF G to an AF $F *_\sigma G$. The underlying concept of a model is given by the argumentation semantics σ. As before, we show a correspondence between a set of postulates and a class of rankings on $2^{\mathfrak{A}}$. The revision postulates, in the manner of [Katsuno and Mendelzon, 1991], are formulated as follows:

(A$*$1) $\sigma(F *_\sigma G) \subseteq \sigma(G)$.

(A$*$2) If $\sigma(F) \cap \sigma(G) \neq \emptyset$, then $\sigma(F *_\sigma G) = \sigma(F) \cap \sigma(G)$.

(A$*$3) If $\sigma(G) \neq \emptyset$, then $\sigma(F *_\sigma G) \neq \emptyset$.

(A$*$4) If $\sigma(G) = \sigma(H)$, then $\sigma(F *_\sigma G) = \sigma(F *_\sigma H)$.

(A$*$5) $\sigma(F *_\sigma G) \cap \sigma(H) \subseteq \sigma(F *_\sigma f_\sigma(\sigma(G) \cap \sigma(H)))$.

(A$*$6) If $\sigma(F *_\sigma G) \cap \sigma(H) \neq \emptyset$, then $\sigma(F *_\sigma f_\sigma(\sigma(G) \cap \sigma(H))) \subseteq \sigma(F *_\sigma G) \cap \sigma(H)$.

(Acyc) If for $0 \leq i \leq n$, $\sigma(F *_\sigma G_{i+1}) \cap \sigma(G_i) \neq \emptyset$ and $\sigma(F *_\sigma G_0) \cap \sigma(G_n) \neq \emptyset$ then $\sigma(F *_\sigma G_n) \cap \sigma(G_0) \neq \emptyset$.

Postulate *Acyc* is borrowed from [Delgrande and Peppas, 2015]. Its addition to the set of postulates is motivated by the following problem. Suppose that for an AF F we have a ranking $\prec_F$ on $2^{\mathfrak{A}}$ which behaves as in the following figure for the extensions $\{a\}, \{b, c\}, \{a, c\}$ and $\{b\}$, and as a faithful ranking otherwise.

$$\begin{tikzcd} & \{a,c\} \arrow[dr] & \\ \{b,c\} \arrow[ur] & & \{b\} \arrow[dl] \\ & \{a\} \arrow[ul] & \end{tikzcd}$$

The arrow means that the relation is strict: thus $\{a\} \preceq_F \{b, c\}$ and $\{b, c\} \npreceq_F \{a\}$. The relation $\prec_F$, then, contains a non-transitive cycle and is not a pre-order. However, quick inspection of the figure shows that for any non-empty subset $\mathbb{S}$ of extensions that can be represented by an AF under a proper I-maximal semantics σ, $\min(\mathbb{S}, \prec_F)$ is still well defined and non-empty. Thus it is possible to define an operator $*_\sigma$ in the usual way: $F *_\sigma G = f_\sigma(\min(\sigma(G), \prec_F))$, and – crucially – the operator $*_\sigma$ satisfies postulates A$*$1 – A$*$6.

Additionally, there is no transitive ranking $\preceq^\star_F$ which is transitive and yields the same revision operator. To see this, notice that if $\preceq^\star_F$ existed, it would have to satisfy $\min(\{\{a\}, \{b, c\}\}, \preceq^\star_F) = \{\{a\}\}$, because we know that $\sigma(F *_\sigma f_\sigma(\{\{a\}, \{b, c\}\})) = \{\{a\}\}$. Thus it would hold that $\{a\} \prec^*_F \{b, c\}$. Similarly, we get that $\{b, c\} \prec^*_F \{a, c\} \prec^*_F \{b\} \prec^*_F \{a\}$, and the cycle is reiterated.

Nonetheless, non-transitive cycles are something we want to avoid: since a natural reading of the rankings on $2^{\mathfrak{A}}$ is as plausibility relations, we would like these rankings to be transitive, and it is thus undesirable to have revision operators that characterize non-transitive rankings. In order to prevent this situation we make use of *Acyc*.

The second detail that needs to be mentioned is our use of I-faithful assignments, motivated by the way in which proper I-maximal semantics work. Given a revision operator $*_\sigma$ and $F \in AF_{\mathfrak{A}}$, the natural way to rank two extensions E_1 and E_2 is by appeal to $F *_\sigma f_\sigma(\{E_1, E_2\})$: if $E_1 \in \sigma(F *_\sigma f_\sigma(\{E_1, E_2\}))$, then E_1 is considered 'more plausible' than E_2 and it should hold that $E_1 \preceq_F E_2$. However, by proper I-maximality of σ, $f_\sigma(\{E_1, E_2\})$ exists only if E_1 and E_2 are $\subseteq$-incomparable. Thus if E_1 and E_2 are $\subseteq$-comparable, $*_\sigma$ might not have any means to adjudicate between E_1 and E_2, hence it is natural to allow them to be incomparable with respect to $\preceq_F$.

Given these preliminaries, we can now state our main representation results.

Theorem 3. *Let σ be a proper I-maximal semantics, $F \in AF_{\mathfrak{A}}$ and $\preceq_F$ an I-faithful ranking associated with F. Define an operator $*_\sigma \colon AF_{\mathfrak{A}} \times AF_{\mathfrak{A}} \mapsto AF_{\mathfrak{A}}$ by:*

$$F *_\sigma G = f_\sigma(\min(\sigma(G), \preceq_F)).$$

*Then $*_\sigma$ satisfies postulates A$*$1 – A$*$6 and Acyc.*

Proof. Since σ is proper I-maximal, any non-empty subset of $\sigma(G)$ (in particular, $\min(\sigma(G), \preceq_F)$) is realizable under σ. Thus $*_\sigma$ is well-defined and we do not need to add any extra condition on $\preceq_F$, such as σ-compliance. Keeping this in mind, the proof that A$*$1 – A$*$6 hold is entirely similar to Theorem 1. In the following we show that *Acyc* also holds.

Let $G_0, G_1, \ldots, G_n$ be a sequence of AFs such that for all $i \in \{1, \ldots, n\}$, $(F *_\sigma G_i) \cap \sigma(G_{i-1}) \neq \emptyset$ and $(F *_\sigma G_0) \cap \sigma(G_n) \neq \emptyset$ holds. From $(F *_\sigma G_1) \cap \sigma(G_0) \neq \emptyset$ we derive by proper I-maximality of σ that $\min(\sigma(G_1), \preceq_F) \cap \sigma(G_0) \neq \emptyset$. Hence there is an extension $E'_0 \in \sigma(G_0)$ such that $E'_0 \preceq_F E_1$ for all $E_1 \in \sigma(G_1)$. Likewise we get from $(F *_\sigma G_2) \cap \sigma(G_1) \neq \emptyset$ that there is an extension $E'_1 \in \sigma(G_1)$ such that $E'_1 \preceq_F E_2$ for all $E_2 \in \sigma(G_2)$, $\ldots$, and from $(F * G_n) \cap \sigma(G_{n-1}) \neq \emptyset$ that there is an extension $E'_{n-1} \in \sigma(G_{n-1})$ such that $E'_{n-1} \preceq_F E_n$ for all $E_n \in \sigma(G_n)$. From transitivity of $\preceq_F$ we get $E'_0 \preceq_F E_n$ for all $E_n \in \sigma(G_n)$. Finally, from $(F *_\sigma G_0) \cap \sigma(G_n) \neq \emptyset$ it follows that there is some $E''_0 \in \sigma(G_0)$ with $E''_0 \preceq_F E_0$ for all $E_0 \in \sigma(G_0)$ and $E''_0 \in \sigma(G_n)$. Now from $E''_0 \preceq_F E'_0 \preceq_F E_n$ (for all $E_n \in \sigma(G_n)$) it follows that $E''_0 \in \min(\sigma(G_n), \preceq_F)$. Hence $\sigma(F *_\sigma G_n) \cup \sigma(G_0) \neq \emptyset$. ☐

Theorem 4. *Let $*_\sigma \colon AF_{\mathfrak{A}} \times AF_{\mathfrak{A}} \mapsto AF_{\mathfrak{A}}$ be an operator satisfying postulates A$*$1 – A$*$6 and Acyc for a proper I-maximal semantics σ. Then, there exists an I-faithful assignment mapping every $F \in AF_{\mathfrak{A}}$ to an I-faithful ranking $\preceq_F$ on $2^{\mathfrak{A}}$ such that $F *_\sigma G = f_\sigma(\min(\sigma(G), \preceq_F))$ for any $G \in AF_{\mathfrak{A}}$.*

Proof. Assume there is $*_\sigma \colon AF_{\mathfrak{A}} \times AF_{\mathfrak{A}} \mapsto AF_{\mathfrak{A}}$ satisfying postulates A$*$1-A$*$6 and *Acyc*, and take an $F \in AF_{\mathfrak{A}}$. We construct $\preceq_F$ in two steps. First we define a relation $\preceq'_F$ on $2^{\mathfrak{A}}$ by saying that for any two $\subseteq$-incomparable $E, E' \in 2^{\mathfrak{A}}$:

$$E \preceq'_F E' \text{ iff } E \in \sigma(F *_\sigma f_\sigma(\{E, E'\})).$$

The relation $\preceq'_F$ is reflexive, as A$*$1 and A$*$3 imply that $E \in \sigma(F *_\sigma f_\sigma(\{E\}))$, but not necessarily total. In the next step we take $\preceq_F$ to be the transitive closure of $\preceq'_F$. In other words:

$$E \preceq_F E' \text{ iff there exist } E_1, \ldots, E_n \text{ such that: } E_1 = E, E_n = E' \text{ and } E_1 \preceq'_F \cdots \preceq'_F E_n.$$

The remainder of the proof shows that $\preceq_F$ is the desired I-faithful ranking. First, notice that if $E_1 \preceq'_F E_2$ then $E_1 \preceq_F E_2$. Hence $\preceq_F$ is reflexive and, by construction, it is transitive, which makes it a pre-order on $2^{\mathfrak{A}}$. Additionally, for any two $\subseteq$-incomparable extensions E_1, E_2, proper I-maximality of σ guarantees that $f_\sigma(\{E_1, E_2\})$ exists. By A$*$1 and A$*$3, $\sigma(F *_\sigma f_\sigma(\{E_1, E_2\}))$ is a non-empty subset of $\{E_1, E_2\}$, thus $E_1 \preceq'_F E_2$ or $E_2 \preceq'_F E_1$ and $\preceq'_F$ is I-total. Next we argue that $\preceq$ is an I-faithful ranking.

Due to proper I-maximality of σ, a set $\{E_1, E_2\}$ is realizable whenever E_1 and E_2 are $\subseteq$-incomparable. Thus, we usually write simply $\{E_1, E_2\}$ instead of $\sigma(f_\sigma(\{E_1, E_2\}))$.

Lemma 3. *If $E_1, E_2 \in \sigma(F)$, then $E_1 \approx_F E_2$.*

Proof. From A$*$2 and proper I-maximality of σ, we get $\sigma(F * f_\sigma(\{E_1, E_2\})) = \sigma(F) \cap \{E_1, E_2\} = \{E_1, E_2\}$. Thus $E_1 \preceq'_F E_2$ and $E_2 \preceq'_F E_1$, which implies $E_1 \approx_F E_2$. ☐

Lemma 3 shows that $\preceq_F$ satisfies property (i) of I-faithful assignments. For property (ii) we make use of the following lemmas. It is in this context that *Acyc* proves crucial.

Lemma 4. *If $E_1, \ldots, E_n$ are pairwise distinct extensions with $E_1 \preceq'_F E_2 \preceq'_F \cdots \preceq'_F E_n \preceq'_F E_1$, then $E_1 \preceq'_F E_n$.*

Proof. If $n = 2$ the conclusion follows immediately. In the following we assume that $n > 2$. From the hypothesis we have that $E_i \in \sigma(F *_\sigma f_\sigma(\{E_i, E_{i+1}\}))$, for $i \in \{1, n-1\}$, and $E_n \in \sigma(F *_\sigma f_\sigma(\{E_n, E_1\}))$. It follows that $E_1 \in \sigma(F *_\sigma f_\sigma(\{E_1, E_2\})) \cap \{E_n, E_1\}$, $E_i \in \sigma(F *_\sigma f_\sigma(\{E_i, E_{i+1}\})) \cap \{E_{i-1}, E_i\}$, for $i \in \{2, \dots, n-1\}$, and $E_n \in \sigma(F *_\sigma f_\sigma(\{E_n, E_1\})) \cap \{E_{n-1}, E_n\}$. Applying *Acyc*, we get that $\sigma(F *_\sigma f_\sigma(\{E_n, E_1\})) \cap \{E_1, E_2\} \neq \emptyset$. From A$*$5 and A$*$6 it follows that $\sigma(F *_\sigma f_\sigma(\{E_n, E_1\})) \cap \{E_1, E_2\} = \sigma(F *_\sigma f_\sigma(\{E_n, E_1\} \cap \{E_1, E_2\}))$. Since $\{E_n, E_1\} \cap \{E_1, E_2\} = \{E_1\}$ we get by A$*$4 that $\sigma(F *_\sigma f_\sigma(\{E_n, E_1\} \cap \{E_1, E_2\})) = \sigma(F *_\sigma f_\sigma(E_1))$. Finally, using A$*$1 and A$*$3 we conclude that $\sigma(F *_\sigma f_\sigma(E_1)) = \{E_1\}$, and thus $E_1 \in \sigma(F *_\sigma f_\sigma(\{E_n, E_1\}))$, implying $E_1 \preceq'_F E_n$. $\square$

Lemma 5. *For any extensions E and E', if $E \prec'_F E'$ then $E \prec_F E'$.*

Proof. Suppose, on the contrary, that $E' \preceq_F E$. Then there exist $E_1, \dots, E_n$ such that $E_1 = E'$, $E_n = E$ and $E_1 \preceq'_F \cdots \preceq'_F E_n$. Since we also have $E \prec'_F E'$, we can apply Lemma 4 to get $E_1 \preceq'_F E_n$, a contradiction. $\square$

Lemma 6. *If E_1 and E_2 are $\subseteq$-incomparable extensions and $E_1 \in \sigma(F)$, $E_2 \notin \sigma(F)$, then $E_1 \prec_F E_2$.*

Proof. By proper I-maximality of σ and A$*$2 we get $\sigma(F *_\sigma f_\sigma(\{E_1, E_2\})) = \sigma(F) \cap \{E_1, E_2\} = \{E_1\}$. This implies that $E_1 \prec'_F E_2$ and by Lemma 5 $E_1 \prec_F E_2$. $\square$

Lemma 6 gives us property (ii). Property (iii) is true because $\preceq_F$ extends $\preceq'_F$, which is defined solely with respect to the extensions of F.

Lastly, we show that the extensions of $F *_\sigma G$, for any $G \in AF_{\mathfrak{A}}$, are the minimal elements of $\sigma(G)$ under $\preceq_F$.

Lemma 7. *For any two extensions E_1, E_2 and any $G \in AF_{\mathfrak{A}}$, if $E_1 \in \sigma(G)$, $E_2 \in \sigma(F *_\sigma G)$ and $E_1 \preceq'_F E_2$, then $E_1 \in \sigma(F *_\sigma G)$.*

Proof. Since $E_2 \in \sigma(F *_\sigma G)$, by A$*$1 we get that $E_2 \in \sigma(G)$. Thus E_1 and E_2 are both σ-extensions of G, and by proper I-maximality of σ, $f_\sigma(\{E_1, E_2\})$ exists. Given this, the rest of the proof is similar to the one for Lemma 1. $\square$

Lemma 8. *For any $G \in AF_{\mathfrak{A}}$, $\min(\sigma(G), \preceq'_F) = \sigma(F *_\sigma G)$.*

Proof. Keeping in mind that for any two σ-extensions E_1, E_2 of G, by proper I-maximality of σ), $E_1 \preceq'_F E_2$ or $E_2 \preceq'_F E_1$, the proof is similar to the one for Lemma 2. $\square$

Lemma 9. *For any argumentation framework $G \in AF_{\mathfrak{A}}$, it holds that $\min(\sigma(G), \preceq_F) = \min(\sigma(G), \preceq'_F)$.*

Proof. $\subseteq$: Let $E_1 \in \min(\sigma(G), \preceq_F)$ and suppose there exists $E_2 \in \sigma(G)$ with $E_2 \prec'_F E_1$. By Lemma 5, this implies that $E_2 \prec_F E_1$, a contradiction to $E_1 \in \min(\sigma(G), \preceq_F)$. It follows that $E_1 \preceq'_F E_2$, thus $E_1 \in \min(\sigma(G), \preceq'_F)$.

$\supseteq$: Take $E_1 \in \min(\sigma(G), \preceq'_F)$ and any $E_2 \in \sigma(G)$. If $E_2 = E_1$, it follows that $E_1 \preceq'_F E_2$. If $E_2 \neq E_1$, then by proper I-maximality of σ, E_1 and E_2 are $\subseteq$-incomparable and thus $E_1 \preceq'_F E_2$ or $E_2 \preceq'_F E_1$. We cannot have that $E_2 \prec'_F E_1$, since this would contradict the hypothesis that $E_1 \in \min(\sigma(G), \preceq'_F)$, therefore $E_1 \preceq'_F E_2$. In both cases it follows that $E_1 \preceq_F E_2$, hence $E_1 \in \min(\sigma(G), \preceq_F)$. $\square$

Lemmata 8 and 9 imply that for any $G \in AF_{\mathfrak{A}}$, $\sigma(F * G) = \min(\sigma(F), \preceq_F)$. This concludes the proof. $\square$

Finally, notice that any faithful assignment on $2^{\mathfrak{A}}$ can be used, *via* Theorem 3, to construct a revision operator $\sigma_* \colon AF_{\mathfrak{A}} \times AF_{\mathfrak{A}} \to AF_{\mathfrak{A}}$. The reason why revision by argumentation frameworks is easier than revision by propositional formulas is the fact that any subset of $\sigma(F)$ is realizable under σ, for any proper I-maximal semantics σ and $F \in AF_{\mathfrak{A}}$. Also, any faithful assignment is an I-faithful assignment in our sense, which implies, by Theorem 3, that $*_\sigma$ satisfies A$*$1 – A$*$6 and *Acyc*. Thus, any model-based revision operator from the standard literature on belief change (e.g., Dalal's operator [Dalal, 1988]) can be used as a revision operator of AFs by AFs.

4 Discussion

Related Work As we already indicated in the introduction to this work, there has been a substantial amount of research in the dynamics of argumentation frameworks although the problems investigated and approaches that have been developed to tackle these differ considerably. We briefly discuss the most closely related work.

Booth *et al.* [2013] develop a general AGM-like framework for modelling the dynamics of AFs based on a labelling approach. Particularly relevant for their work is the notion of "fall back beliefs" which represent the most rational outcome of the input AF given a constraint. Their revision operator is defined in terms of conflict free labellings and their notion of "fall back beliefs" is developed only for complete semantics.

We have already mentioned the work presented in [Coste-Marquis *et al.*, 2014a] which is the starting point of our investigation. On the semantic side, their revision functions are defined following a two step process: first a counter-part to the notion of faithful assignment on the models of the revision operators is defined; secondly, a set of AFs that generate such extensions is constructed using different criteria, e.g. minimizing the changes in the attack relation of the input AF vs. minimizing the number of AFs generated. The main difference between the work presented in [Coste-Marquis *et al.*, 2014a] to our work is that we consider the issue of revision of AFs as minimal change in the arguments statuses under the constraint that a single AF has to be produced. The price we pay is that in our solution the revised AF may have new arguments while in [Coste-Marquis *et al.*, 2014a] only the attack relation can be modified. Recent work on realizability in compact AFs [Baumann *et al.*, 2014] could pave the way for revision in the setting where, the result is a single AF, but no additional arguments are allowed to come into play.

Finally, in [Coste-Marquis *et al.*, 2014b] the revision problem for AFs is translated into propositional logic thus enabling the use of classical AGM revision operators to carry out revision in the AF setting. On the other hand, the revision formulas are defined in terms of the skeptical acceptance of arguments. Nevertheless, the output of revision is still a set of AFs rather than a single AF.

Conclusion and Outlook We have presented a generic solution to the problem of revision for argumentation frameworks which applies to many prominent I-maximal argumentation semantics. The key to obtain our AGM-style representation theorems was the combination of recent advances from argumentation theory and belief change. We identify several directions for future work: (1) extend our results to semantics which are not proper I-maximal; (2) identify operators based on σ-compliant rankings for specific semantics σ; (3) analyze whether our insights can be extended to a broader theory of belief change within fragments; (4) apply our findings to other belief change operations. In particular, iterated belief revision seems to have natural applications in the argumentation domain and we believe that the understanding of revision yielding a single AF is fundamental for this purpose; (5) take the syntactic form of the AF into account. One possibility would be a two-step approach, where our abstract revision is the first step. Based on this result, a second step would revise the syntactic structure of the AF.

References

[Alchourrón *et al.*, 1985] Carlos E. Alchourrón, Peter Gärdenfors, and David Makinson. On the logic of theory change: partial meet contraction and revision functions. *J. Symb. Log.*, 50(2):510–530, 1985.

[Baroni and Giacomin, 2007] Pietro Baroni and Massimiliano Giacomin. On principle-based evaluation of extension-based argumentation semantics. *Artif. Intell.*, 171(10-15):675–700, 2007.

[Baroni *et al.*, 2011] Pietro Baroni, Martin Caminada, and Massimiliano Giacomin. An introduction to argumentation semantics. *Knowledge Eng. Review*, 26(4):365–410, 2011.

[Baumann *et al.*, 2014] Ringo Baumann, Wolfgang Dvořák, Thomas Linsbichler, Hannes Strass, and Stefan Woltran. Compact argumentation frameworks. In *Proc. ECAI*, pages 69–74, 2014.

[Baumann, 2012] Ringo 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] Trevor J. M. Bench-Capon and Paul E. Dunne. Argumentation in artificial intelligence. *Artif. Intell.*, 171(10-15):619–641, 2007.

[Bench-Capon *et al.*, 2009] Trevor J. M. Bench-Capon, Henry Prakken, and Giovanni Sartor. Argumentation in legal reasoning. In *Argumentation in Artificial Intelligence*, pages 363–382. Springer, 2009.

[Bisquert *et al.*, 2011] Pierre Bisquert, Claudette Cayrol, Florence Dupin de Saint-Cyr, and Marie-Christine Lagasquie-Schiex. Change in argumentation systems: exploring the interest of removing an argument. In *Proc. SUM*, pages 275–288, 2011.

[Bisquert *et al.*, 2013] Pierre Bisquert, Claudette Cayrol, Florence Dupin de Saint-Cyr, and Marie-Christine Lagasquie-Schiex. Enforcement in argumentation is a kind of update. In *Proc. SUM*, pages 30–43, 2013.

[Boella *et al.*, 2009] Guido Boella, Souhila Kaci, and Leendert van der Torre. Dynamics in argumentation with single extensions: attack refinement and the grounded extension (extended version). In *Proc. ArgMAS*, pages 150–159, 2009.

[Booth *et al.*, 2013] Richard Booth, Souhila Kaci, Tjitze Rienstra, and Leendert van der Torre. A logical theory about dynamics in abstract argumentation. In *Proc. SUM*, pages 148–161, 2013.

[Cartwright and Atkinson, 2009] Dan Cartwright and Katie Atkinson. Using computational argumentation to support e-participation. *IEEE Intell. Syst.*, 24(5):42–52, 2009.

[Cayrol *et al.*, 2010] Claudette Cayrol, Florence Dupin de Saint-Cyr, and Marie-Christine Lagasquie-Schiex. Change in abstract argumentation frameworks: adding an argument. *J. Artif. Intell. Res.*, 38:49–84, 2010.

[Coste-Marquis *et al.*, 2014a] Sylvie Coste-Marquis, Sébastien Konieczny, Jean-Guy Mailly, and Pierre Marquis. On the revision of argumentation systems: minimal change of arguments statuses. In *Proc. KR*, pages 72–81, 2014.

[Coste-Marquis *et al.*, 2014b] Sylvie Coste-Marquis, Sébastien Konieczny, Jean-Guy Mailly, and Pierre Marquis. A translation-based approach for revision of argumentation frameworks. In *Proc. JELIA*, pages 397–411, 2014.

[Dalal, 1988] Mukesh Dalal. Investigations into a theory of knowledge base revision. In *Proc. AAAI*, pages 475–479, 1988.

[Darwiche and Pearl, 1997] Adnan Darwiche and Judea Pearl. On the logic of iterated belief revision. *Artif. Intell.*, 89(1-2):1–29, 1997.

[Delgrande and Peppas, 2015] James P. Delgrande and Pavlos Peppas. Belief revision in Horn theories. *Artif. Intell.*, 218:1–22, 2015.

[Doutre *et al.*, 2014] Sylvie Doutre, Andreas Herzig, and Laurent Perrussel. A dynamic logic framework for abstract argumentation. In *Proc. KR*, pages 62–71, 2014.

[Dung, 1995] Phan Minh Dung. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. *Artif. Intell.*, 77(2):321–357, 1995.

[Dunne *et al.*, 2014] Paul E. Dunne, Wolfgang Dvorák, Thomas Linsbichler, and Stefan Woltran. Characteristics of multiple viewpoints in abstract argumentation. In *Proc. KR*, pages 52–61, 2014.

[Gärdenfors, 1988] Peter Gärdenfors. *Knowledge in flux. Modeling the dynamics of epistemic states*. Cambridge: MIT Press, 1988.

[Katsuno and Mendelzon, 1991] Hirofumi Katsuno and Alberto O. Mendelzon. Propositional knowledge base revision and minimal change. *Artif. Intell.*, 52(3):263–294, 1991.

[Kontarinis *et al.*, 2013] Dionysios Kontarinis, Elise Bonzon, Nicolas Maudet, Alan Perotti, Leon van der Torre, and Serena Villata. Rewriting rules for the computation of goal-oriented changes in an argumentation system. In *Proc. CLIMA*, pages 51–68, 2013.

[Krümpelmann *et al.*, 2012] Patrick Krümpelmann, Matthias Thimm, Marcelo A. Falappa, Alejandro J. García, Gabriele Kern-Isberner, and Guillermo R. Simari. Selective revision by deductive argumentation. In *Proc. TAFA*, pages 147–162, 2012.

[Nouioua and Würbel, 2014] Farid Nouioua and Éric Würbel. Removed set-based revision of abstract argumentation frameworks. In *Proc. ICTAI*, pages 784–791, 2014.

[Rahwan and Simari, 2009] Iyad Rahwan and Guillermo R. Simari, editors. *Argumentation in Artificial Intelligence*. Springer, 2009.

[Sakama, 2014] Chiaki Sakama. Counterfactual reasoning in argumentation frameworks. In *Proc. COMMA*, pages 385–396, 2014.

[Spohn, 1988] Wolfgang Spohn. Ordinal conditional functions: A dynamics theory of epistemic states. In *Causation in decision, belief change and statistics, vol. 2*, pages 105–134. Kluwer Academic Publishers, 1988.