dbai

TECHNICAL REPORT

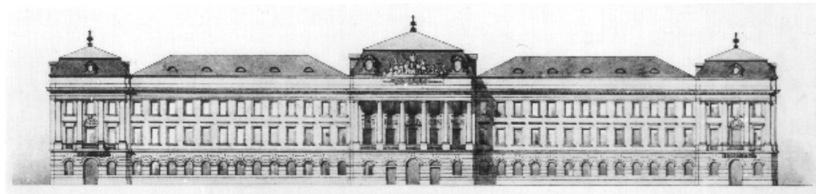

INSTITUT FÜR INFORMATIONSSYSTEME

ABTEILUNG DATENBANKEN UND ARTIFICIAL INTELLIGENCE

An Extension-Based Approach to Belief Revision in Abstract Argumentation

DBAI-TR-2017-106

Martin Diller **Adrian Haret** **Thomas Linsbichler**
Stefan Rümmele **Stefan Woltran**

Institut für Informationssysteme
Abteilung Datenbanken und
Artificial Intelligence
Technische Universität Wien
Favoritenstr. 9
A-1040 Vienna, Austria
Tel: +43-1-58801-18403
Fax: +43-1-58801-18493
sekret@dbai.tuwien.ac.at
www.dbai.tuwien.ac.at

DBAI TECHNICAL REPORT
2017

TECHNISCHE UNIVERSITÄT WIEN
Vienna University of Technology

DBAI TECHNICAL REPORT

DBAI TECHNICAL REPORT DBAI-TR-2017-106, 2017

An Extension-Based Approach to Belief Revision in Abstract Argumentation

Martin Diller [1] **Adrian Haret** [2] **Thomas Linsbichler** [3]

Stefan Rümmele [4] **Stefan Woltran** [5]

Abstract. Argumentation is an inherently dynamic process, and 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 to Dung's abstract argumentation frameworks (AFs). However, none of the attempts have yet succeeded in solving the natural situation where the revision of an AF is guaranteed to be representable by a single AF. Here we present a generic solution to this problem, which applies to many prominent argumentation semantics. To prove a full representation theorem, we make use of recent advances in both areas of argumentation and belief change. In particular, we use the concept of realizability in argumentation and the concept of compliance as introduced in Horn revision. We also present a family of concrete belief change operators tailored specifically for AFs and analyse their computational complexity.

[1]TU Wien, Austria. diller@dbai.tuwien.ac.at

[2]TU Wien, Austria. haret@dbai.tuwien.ac.at

[3]TU Wien, Austria. linsbich@dbai.tuwien.ac.at

[4]UNSW Sydney and University of Sydney, Australia. s.rummele@unsw.edu.au

[5]TU Wien, Austria. woltran@dbai.tuwien.ac.at

Acknowledgements: This research has been supported by the Austrian Science Fund (FWF) under grants I1102, I2854, P25518, P25521, P30168, W1255, and Y698.

1 Introduction

Argumentation has emerged, over the last two decades, as a major research area in Artificial Intelligence (AI) [8, 50]. This is due not just to the intrinsic interest of the topic and to its recent applications (e.g., in legal reasoning [9] and e-governance [16]), but also because of fundamental connections between argumentation and other areas of AI, mainly non-monotonic reasoning.

The significant landmark in the consolidation of argumentation as a distinct field of AI has been the introduction of abstract argumentation frameworks (AFs) [31], which are directed graphs whose nodes represent arguments and where links correspond to attacks between arguments. To this day AFs remain the most widely used and investigated among the several argumentation formalisms. The study of AFs is mainly concerned with finding subsets of arguments (called *extensions*) that can all be accepted together when taking into consideration the structure encoded in the graph. As a result, the argumentation literature offers a wide range of criteria (called *semantics of AFs*) for establishing which arguments are jointly acceptable [3].

Our work fits into the growing number of studies on the *dynamics* of argumentation frameworks [5, 10, 11, 12, 14, 18, 30, 39, 40, 49, 52]. This line of research is motivated by a realization that, as part of interactive reasoning processes, argumentation frameworks have to undergo change when new information becomes available. Particularly important in this respect is change with respect to the acceptability of certain arguments: one can imagine that increased knowledge of facts would settle certain issues, with the effect that arguments pertaining to them would have to either become part of, or be excluded from any extension of our AF. Thus, we would expect such increased knowledge to be reflected in a new AF which managed to preserve as much semantic information from the original one, while making sure that its extensions satisfy the added constraints. The main issue, in this setting, is finding appropriate ways of formalizing such a notion of minimal change at the semantic level, with the understanding that the graph structure of the revised AF is then constructed around the semantic information. Settling on a specific graph structure for the revised AF is an interesting problem in its own right, though it is a separate issue from the one concerning us here, and left for future work.

We look at the problem in the spirit of the AGM paradigm in propositional belief revision [1, 36], and our analysis is in particular modelled after the model-based approach of Katsuno and Mendelzon [38]. The issue was first tackled in this way by Coste-Marquis et al. [19],[1] where AF revision was defined as follows: given a semantics, an AF and a revision formula encoding desired changes in the status of some arguments, find *a set of AFs* satisfying the revision formula, whose extensions are as close as possible to the extensions of the input AF. Remarkably, a *representation theorem* illuminates the problem by co-opting the semantics of AFs: performing AF revision in accordance with some rationality postulates was shown to be equivalent to choosing among possible extensions of AFs, according to a particular type of rankings on extensions.

A notable difference between this work and our approach is that we study AF revision operators producing a *single* AF as output. The motivation for this is twofold. First, such a restriction is more in line with standard AGM revision, where revising an input propositional theory by a revision formula produces a single propositional theory. Second, revision yielding a single AF

[1] Other recent work in this direction includes [6, 11, 25, 44, 45]; we discuss these papers in Section 6.

makes concepts of iterated revision [24, 54] amenable to argumentation: indeed, for an iterated application of belief change operations, it is desirable that the input and output are of the same type. The same point holds for persuasion, where some current state of discourse needs to be updated such that an agent is convinced to accept a certain argument: it has been emphasized that modelling persuasion can benefit from applying change operations in argumentation [37]. Thus, understanding belief change of abstract argumentation formalisms can pave the way towards a general theory of formal persuasion.

Restricting the output of AF revision operators to single AFs poses significant challenges, as standard operators from the propositional belief change literature are not easily applicable in the new context and familiar representation results break down. A conspicuous problem is ensuring that what we get from an AF revision operator, which is typically a set of extensions, *can* be expressed as a single AF under some given semantics. The problem is exacerbated by the variety of semantics on offer and their expressive particularities.

We study two types of revision operators. The first type considers the new information represented as a propositional formula. This formula encodes, by its models, a set of extensions representing the change (in terms of extensions) we want to induce in the original AF. The second type is revision by an AF, where 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 [19], the latter assumes that the knowledge to be incorporated (for example another agent's beliefs) is in the form of an AF. This is more in line with work on Horn revision [26], where all involved formulas belong to a single fragment of propositional logic. The two types of revision present interesting differences. Revision by a propositional formula is characterizable using standard revision postulates, as long as rankings on extensions satisfy a *compliance* restriction. Revision by an AF, on the other hand, turns out not to require compliance, but is only characterizable using an extra postulate called Acyc and what we call *proper I-maximal semantics*.

Our main contributions can be summarized as follows.

- We obtain full *representation theorems* for the two types of revision mentioned. Notably, our results are *generic* in the sense that they hold for a wide range of argumentation semantics, including preferred, semi-stable, stage, and stable semantics.

- For the revision-by-formula approach, we give novel notions of *compliance* [26] to restrict the rankings (Section 3.1). This is required to guarantee that the outcome of the corresponding operators can be realized as 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, the necessary results are known [33]. It turns out that standard revision operators such as Dalal's operator [23] do not satisfy the required compliance. We thus introduce a new class of AF revision operators, following the intuition of minimal-distance based revision in a similar way to Dalal's operator (Section 3.2).

- In the revision-by-AF approach, we show that the concept of compliance can be dropped and standard revision operators satisfying all postulates like Dalal's operator can be directly applied to revision of AFs. However, an additional postulate (borrowed from Horn revision [26]) is needed for the representation theorem (Section 4). This amended set of postulates, together with an explicit commitment to what we call *proper I-maximal semantics*, turns out to characterize a class of I-faithful rankings on extensions.

- Finally, we analyse the computational complexity of some specific revision operators when using stable and preferred semantics. For the revision-by-AF approach our result of Θ_2^{P}-completeness for stable semantics matches the known complexity for Dalal's revision in (fragments of) classical logic [34, 42, 22], while it turns out that the intrinsically higher complexity of preferred semantics [32] is also reflected in the revision task for which we show Θ_3^{P}-completeness (Section 5). For the refinement of Dalal's operator in the revision-by-formula approach our results indicate a slight increase in complexity to Δ_2^{P} for stable and to Δ_3^{P} for preferred semantics.

The paper is structured as follows. In Section 2 we provide background notions and results for argumentation and belief revision. In Section 3 we study revision of AFs by propositional formulas, introduce the concept of σ-compliance and faithful assignments, prove a representation result and introduce novel revision operators that satisfy all postulates in this setting. In Section 4 we switch to revision of AFs by other AFs, introduce I-faithful assignments and prove a representation result. Section 5 provides a complexity analysis of operators introduced in previous sections. Section 6 discusses related work. Section 7 contains the conclusion and outlines directions for future work.

This article is an extended version of [28]. Additional material includes the specific revision operators for revision by propositional formulas in Section 3 and the complexity analysis in Section 5.

2 Preliminaries

We first recall basic notions of Dung's abstract argumentation frameworks (Dung [31], Baroni et al. [3] provide further background), present recent results on signatures of semantics [33], and then recall the basic concepts of belief revision.

Argumentation 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 a set of arguments 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 $T \subseteq A$ is defended (in F) by S if each $a \in T$ is defended (in F) by S. 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 (with respect to 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 [56], and semi-stable [15] semantics, the extensions are defined as follows:

$S \in \mathit{stb}(F)$, if $S \in \mathit{cf}(F)$ and $S^+_F = A$ (stable)
$S \in \mathit{prf}(F)$, if $S \in \mathit{adm}(F)$ and $\nexists T \in \mathit{adm}(F)$ s.t. $T \supset S$ (preferred)
$S \in \mathit{stg}(F)$, if $S \in \mathit{cf}(F)$ and $\nexists T \in \mathit{cf}(F)$ with $T^+_F \supset S^+_F$ (stage)
$S \in \mathit{sem}(F)$, if $S \in \mathit{adm}(F)$ and $\nexists T \in \mathit{adm}(F)$ s.t. $T^+_F \supset S^+_F$ (semi-stable)

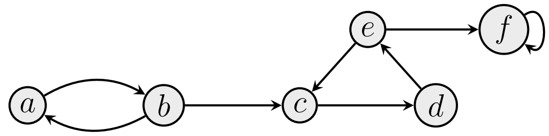

Figure 1: AF discussed in Example 1.

Example 1. To illustrate the semantics, consider the following AF F:

$$F = (\{a, b, c, d, e, f\}, \{(a, b), (b, a), (b, c), (c, d), (d, e), (e, c), (e, f), (f, f)\}).$$

In Figure 1 we depict this AF as a directed graph, with the arguments as nodes and the attacks as directed edges. It can be checked, by direct inspection, that there is no conflict-free set of arguments in F attacking all other arguments, hence $\mathit{stb}(F) = \emptyset$. The admissible sets of F are given by $\mathit{adm}(F) = \{\emptyset, \{a\}, \{b\}, \{b, d\}\}$, and hence $\mathit{prf}(F) = \{\{a\}, \{b, d\}\}$. By $\{a\}^+_F = \{a, b\} \subset \{a, b, c, d, e\} = \{b, d\}^+_F$ we get that $\{b, d\}$ is the only semi-stable extension of F, i.e. $\mathit{sem}(F) = \{\{b, d\}\}$. Finally, it holds that $\mathit{stg}(F) = \{\{a, e\}, \{b, e\}, \{b, d\}\}$.

A set of extensions $\mathbb{S}$ can be *realized under a semantics* σ if there exists an AF $F \in \mathit{AF}_{\mathfrak{A}}$ such that $\sigma(F) = \mathbb{S}$. The signature Σ_σ of semantics σ is defined as $\Sigma_\sigma = \{\sigma(F) \mid F \in \mathit{AF}_{\mathfrak{A}}\}$, containing exactly those sets of extension which can be realized under σ. Exact characterizations of the signatures of the introduced semantics are known [33]. 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 $\subseteq$*-comparable* if $S_1 \subset S_2$ or $S_2 \subset S_1$. We say that S_1 and S_2 are $\subseteq$*-incomparable* if they are not $\subseteq$-comparable.[2] A set of extensions $\mathbb{S} \subseteq 2^{\mathfrak{A}}$ is *incomparable* if all its elements are pairwise $\subseteq$-incomparable. A set of extensions $\mathbb{S} \subseteq 2^{\mathfrak{A}}$ is *tight* if for all extensions $S \in \mathbb{S}$ and arguments $a \in \bigcup_{S \in \mathbb{S}} S$, it holds that: if $S \cup \{a\} \notin \mathbb{S}$, then there exists an $s \in S$ such that $\{a, s\} \not\subseteq S'$ for any $S' \in \mathbb{S}$.

The signatures of the semantics of interest have precise characterizations using the notions just introduced. For the stable and stage semantics the characterizations are as follows:

$$\begin{aligned}\Sigma_{\mathit{stb}} &= \{\mathbb{S} \subseteq 2^{\mathfrak{A}} \mid \mathbb{S} \text{ is incomparable and tight}\},\\ \Sigma_{\mathit{stg}} &= \{\mathbb{S} \subseteq 2^{\mathfrak{A}} \mid \mathbb{S} \neq \emptyset \text{ and } \mathbb{S} \text{ is incomparable and tight}\}.\end{aligned}$$

Regarding the other semantics, it suffices for our purposes to state that $\Sigma_{\mathit{stg}} \subset \Sigma_{\mathit{sem}} = \Sigma_{\mathit{prf}}$ [33]. We will make use of these results in Sections 3 and 4. Also, some of our results will apply to semantics for which the following properties hold.

Definition 1. A semantics σ is called *proper I-maximal* if for each $\mathbb{S} \in \Sigma_\sigma$ it holds that

[2]Note that a set S of arguments is $\subseteq$-incomparable to itself.

1. $\mathbb{S}$ is incomparable,

2. $\mathbb{S}' \in \Sigma_\sigma$ for any $\mathbb{S}' \subseteq \mathbb{S}$ with $\mathbb{S}' \neq \emptyset$, and

3. for any $S_1, S_2 \in 2^{\mathfrak{A}}$ it holds that $\{S_1, S_2\} \in \Sigma_\sigma$.

In words, an I-maximal [2] semantics σ is proper if, on the one hand, it holds that for any AF F we can realize under σ any non-empty subset of $\sigma(F)$, and, on the other hand, any pair of $\subseteq$-incomparable sets of arguments, is realizable under σ. The next observation follows from the characterizations of the signatures [33], and shows that the semantics we are interested in are all proper I-maximal.

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

Proof. We need to show that properties (1) to (3) from Definition 1 hold. (1) is well-known. Properties (2) and (3) were already shown in the following results of Dunne et al. [33]: (2) follows directly from Lemma 2.2 for stable and stage semantics and from Lemma 4.2 for preferred and semi-stable semantics; Proposition 10 contains (3). □

Definition 2. Given a semantics σ, we define the function $f_\sigma : 2^{2^{\mathfrak{A}}} \to AF_{\mathfrak{A}}$ 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.

By definition, $\mathbb{S} \in \Sigma_\sigma$ guarantees that we can find an AF which, when evaluated under σ, results in having $\mathbb{S}$ as set of σ-extensions. We leave the exact specifications of such AFs open but canonical constructions for the semantics we consider have been published [33]. Such constructions may result in AFs with additional arguments to those contained in $\mathbb{S}$, though recent work on realizability in compact AFs [7] could pave the way for constructions of AFs without new arguments. The function f_σ is not necessarily unique. Nevertheless, throughout the paper we assume f_σ to be fixed for every σ.

For an AF $F = (B, S)$ we use A_F to refer to B and R_F refers to S. Finally, given AFs F and G and arguments $X \subseteq A_F$, we define $F - X = (A_F \setminus X, \{(a, b) \in R_F \mid a, b \in (A_F \setminus X)\})$ and $F \cup G = (A_F \cup A_G, R_F \cup R_G)$.

Belief revision By $\mathsf{P}_{\mathfrak{A}}$ we denote the set of propositional formulas over $\mathfrak{A}$, where the arguments in $\mathfrak{A}$ are taken to be 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 $\varphi \in \mathsf{P}_{\mathfrak{A}}$ evaluates to *true* under an interpretation E, E is a model of φ. The set of models of φ is denoted by $[\varphi]$. We write $\varphi_1 \equiv \varphi_2$ if $[\varphi_1] = [\varphi_2]$. A formula φ is *consistent* if $[\varphi] \neq \emptyset$. We will identify a finite set K of propositional formulas with $\bigwedge K$, such that $[K] = [\bigwedge K]$ and K is consistent if $\bigwedge K$ is consistent.

A propositional revision operator $\circ$ maps a set K of propositional formulas and a propositional formula φ to a propositional formula $K \circ \varphi$. The set K, called a knowledge base, is the theory to be revised, φ is the revision formula representing new information which needs to be incorporated into K, and $K \circ \varphi$ is the revision outcome. The revision outcome is constrained by rationality postulates, a core set of which [38] we present below:

	$\{a,b\}$	$\{a,b,c\}$	$d(E,K)$
$\{a,c\}$	2	1	1
$\{c\}$	3	2	2

Table 1: $d(E,K)$, for $E \in [\varphi]$.

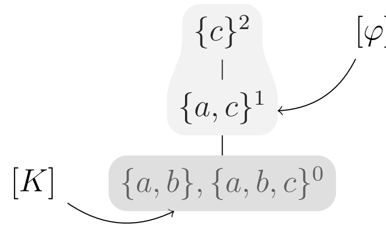

Figure 2: Preorder on interpretations.

(KM∘1) $K \circ \varphi \models \varphi$.

(KM∘2) If $K \wedge \varphi$ is consistent, then $K \circ \varphi \equiv K \wedge \varphi$.

(KM∘3) If φ is consistent, then $K \circ \varphi$ is consistent.

(KM∘4) If $K_1 \equiv K_2$ and $\varphi_1 \equiv \varphi_2$, then $K_1 \circ \varphi_1 \equiv K_2 \circ \varphi_2$.

(KM∘5) $(K \circ \varphi_1) \wedge \varphi_2 \models K \circ (\varphi_1 \wedge \varphi_2)$.

(KM∘6) If $(K \circ \varphi_1) \wedge \varphi_2$ is consistent, then $K \circ (\varphi_1 \wedge \varphi_2) \models (K \circ \varphi_1) \wedge \varphi_2$.

A key insight in belief change is that any propositional revision operator satisfying postulates KM∘1–KM∘6 can be characterized using rankings on the possible worlds described by the language. Intuitively, such rankings can be thought of as plausibility relations, whereby possible states of affairs are ordered according to how 'close' they are to K. Revising a knowledge base K by a formula φ then amounts to selecting the models of φ closest to K.

The natural way of parsing the idea of closeness is to use some distance between interpretations. A common choice is *Hamming distance* d_H, defined as the number of atoms on which two interpretations differ. For example, $d_H(\{a,b,c\},\{b,c,d\}) = |\{a,d\}| = 2$. Known in propositional revision as *Dalal's operator* [23], this approach consists in first defining the distance between an interpretation E and a knowledge base K as $d(E,K) = \min\{d_H(E,E') \mid E' \in [K]\}$. Then, to revise K by φ, one selects the models of φ with minimal distance to K.

Example 2 (Dalal operator for propositional revision)**.** Consider the knowledge base $K = \{a \wedge b\}$, which we want to revise by $\varphi = \neg b \wedge c$. The models of φ are $[\varphi] = \{\{a,c\},\{c\}\}$, and Dalal's approach gives us that $d(\{a,c\},K) = 1$, while $d(\{c\},K) = 2$. The distances from each model of φ to each model of K are shown in Table 1, with models of φ as row names and models of K as column names. Models in $[\varphi]$ can now be ordered according to their distance to K, visualized in Figure 2. The revision operator selects the models of φ with minimal distance to K as the models of the revision outcome. Intuitively, these are the models of φ 'closest' to K, to be ultimately converted back to a propositional formula. In our case we get the single interpretation $\{a,c\}$, which corresponds to $K \circ \varphi \equiv a \wedge \neg b \wedge c$.

To apply this approach to AF revision we will use a unified semantic representation of AFs and logical formulas. Thus, in our approach, sets of arguments from $\mathfrak{A}$ play the role both of extensions of AFs and of models of propositional formulas, and will be the possible worlds a revision operator chooses from. In the following we define the kinds of rankings on $2^{\mathfrak{A}}$ which will be used to characterize the class of AF revision operators.

A *preorder* $\preceq$ *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 preorder $\preceq$ is *total*. Moreover, for $E_1, E_2 \in 2^{\mathfrak{A}}$, $E_1 \prec E_2$ denotes the strict part of $\preceq$, that is $E_1 \preceq E_2$ and $E_2 \not\preceq E_1$. We write $E_1 \approx E_2$ to abbreviate the case when $E_1 \preceq E_2$ and $E_2 \preceq E_1$. An *I-total preorder on* $2^{\mathfrak{A}}$ is a preorder 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 $\subseteq$-incomparable extensions. Finally, for a set of extensions $\mathbb{S} \subseteq 2^{\mathfrak{A}}$ and a preorder $\preceq$, $\min(\mathbb{S}, \preceq) = \{E_1 \in \mathbb{S} \mid \nexists E_2 \in \mathbb{S} : E_2 \prec E_1\}$.

A general way of mapping every knowledge base K to a preorder $\preceq_K$ on interpretations is called an *assignment*, and in propositional revision one typically works with assignments that are *faithful*. Example 2 illustrates one way of generating faithful preorders on interpretations for propositional revision. We will formally introduce faithful assignments in Section 3.1. Here we mention that assignments provide the opportunity of a model-based characterization of revision operators. We say that an assignment *represents an operator* $\circ$ (or, alternatively, that $\circ$ *is represented by an assignment*) if, for any knowledge base K and formula φ, it holds that $[K \circ \varphi] = \min([\varphi], \preceq_K)$. In the case of propositional revision, working with faithful assignments satisfies postulates KM$\circ$1–KM$\circ$6. The following representation result expresses this fact.

Theorem 1 ([38])**.** *If $\circ$ is a propositional revision operator, then $\circ$ satisfies postulates KM$\circ$1–KM$\circ$6 if and only if there exists a faithful assignment which represents it.*

The notion of revision operators being represented by assignments stays the same when K or φ are replaced by AFs. In the remaining sections we obtain similar representation results for AF revision, use these results to construct concrete AF revision operators and, finally, analyse the computational complexity of our proposed operators.

3 Revision by Propositional Formulas

We first consider revision of an AF by a propositional formula, performed through operators of the type $\star_\sigma \colon AF_{\mathfrak{A}} \times \mathsf{P}_{\mathfrak{A}} \to AF_{\mathfrak{A}}$. Given a semantics σ, such operators map an AF F and a consistent propositional formula φ to a revised AF, denoted $F \star_\sigma \varphi$.[3] As mentioned in the introduction, we insist that the revision outcome should be a single AF rather than a set of AFs: this requirement is reflected in our definition of AF revision operators.

Intuitively, the revision formula φ encodes information which should be incorporated in F. More concretely, a revision operator $\star_\sigma$ is expected to change F such that the σ-extensions of $F \star_\sigma \varphi$ come to incorporate the models of φ. At the same time, F should not suffer more change than is strictly necessary. This requirement of minimal change, along with other natural requirements expected from an AF revision operator, are captured by the logical postulates presented in Section 3.1.

[3] We restrict the second argument to consistent formulas because 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.

Figure 3: F is to be revised with $\varphi = c \wedge d$, which may result in G.

E	$\{a, d\}$	$\{b, d\}$	$d_{prf}(E, F)$
$\{c, d\}$	2	2	2
$\{a, c, d\}$	1	3	1
$\{b, c, d\}$	3	1	1
$\{a, b, c, d\}$	2	2	2

Table 2: $d(E, F)$, for $E \in [\varphi]$.

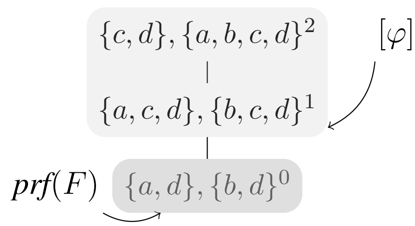

Figure 4: Preorder on extensions.

Example 3. Take $\mathfrak{A} = \{a, b, c, d\}$ and the AF F in Figure 3, whose preferred extensions are $prf(F) = \{\{a, d\}, \{b, d\}\}$. Revising F by a formula $\varphi = c \wedge d$ would mean finding an AF F', guaranteed to contain the arguments c and d in each of its preferred extensions. In other words, we need to find an AF F' such that $prf(F') \subseteq [\varphi]$, where $[\varphi] = \{\{c, d\}, \{a, c, d\}, \{b, c, d\}, \{a, b, c, d\}\}$. The AF G in Figure 3 fits this requirement, as $prf(G) = \{\{a, c, d\}, \{b, c, d\}\}$. However, this is not the only property that the revised AF is expected to satisfy: we also want to ensure that there is no other such AF which is 'closer' to F, according to some measure of closeness on extensions considered suitable.

It is straightforward to adapt Dalal's operator for the context of AF revision. Thus, we define the distance between a set of arguments E and an AF F with respect to a semantics σ as $d_\sigma(E, F) = \min\{d_H(E, E') \mid E' \in \sigma(F)\}$. The AF revision operator then selects the models of φ with minimal distance to F. See Example 4 below for a concrete application.

Example 4. Using Dalal's (adapted) approach for the AF F in Example 3 under preferred semantics, we obtain the distances in Table 2, which generate the preorder $\preceq_F^D$ partly depicted in Figure 4. The interpretations with minimal distance to F are $\{a, c, d\}$ and $\{b, c, d\}$, which gives us that $\sigma(F \star_{prf}^D \varphi) = \{\{a, c, d\}, \{b, c, d\}\}$. Hence the AF G in Example 3 is a suitable revision outcome according to Dalal's (adapted) operator.

Unfortunately, as we show in Section 3.1, Dalal's approach does not work in general for defining AF revision operators, as the semantic output is not guaranteed to be realizable under a semantics σ. We show this on a concrete example and outline our approach for overcoming these difficulties in the rest of this section.

3.1 Postulates and Representation Result

We adapt the revision postulates [38] to the context of AF revision, in a manner similar to work by Coste-Marquis et al. [19].

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

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

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

(P⋆4) If $\sigma(F_1) = \sigma(F_2)$ and $\varphi \equiv \psi$ then $\sigma(F_1 \star_\sigma \varphi) = \sigma(F_2 \star_\sigma \psi)$.

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

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

A few words are in order regarding the meaning of the postulates. P⋆1 says that when we revise an AF F by a formula φ, the extensions of the revision output should be among the models of φ. P⋆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⋆3 says that if φ is a consistent formula, then revision by φ should also be consistent. P⋆4, known as irrelevance of syntax, guarantees that the output of revision does not depend on how the revision formula is specified. P⋆5 and P⋆6 ensure that revision is performed with minimal change to the AF F. Motivation for the postulates has been extensively discussed in the belief revision literature ([1, 36, 38]).

We next define faithful assignments for AFs, adapting the concept with the same name from propositional revision [38], which will be used to characterize AF revision operators.

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

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

(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 preorder $\preceq_F$ assigned to F by a faithful assignment is called the *faithful ranking associated with* F.

At this point, one natural strategy would be to import existing propositional revision operators and use them for AF revision. Such a move was gestured towards in Example 4. However, it turns out that doing so is not possible, because the outcome of the revision under a semantics σ could be a set of extensions $\mathbb{S}$ which cannot be realized under σ (see Example 5 below). This is significant, since as argued in Section 1, we require the output of an AF revision operator to be a single AF. In this, we face a similar challenge to that encountered in Horn revision [26]. To overcome this problem we use Σ_σ to define the following restriction on preorders, which we will need to obtain our representation theorem.

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

The following example shows why the adapted Dalal operator does not work for AF revision, the reason being that it generates preorders which are not σ-compliant.

Example 5. Let $\mathfrak{A} = \{a, b, c\}$, a semantics σ and an AF F such that $\sigma(F) = \{\{a, b, c\}\}$. Dalal's approach, using Hamming distance,[4] generates the following preorder $\preceq$:

$$\{a, b, c\} \prec \{a, b\} \approx \{a, c\} \approx \{b, c\} \prec \{a\} \approx \{b\} \approx \{c\} \prec \emptyset.$$

Now take $\varphi = \neg(a \wedge b \wedge c)$. We obtain that $\min([\varphi], \preceq) = \{\{a, b\}, \{a, c\}, \{b, c\}\}$. Observe, now, that $\{a, b\} \cup \{c\} \notin \{\{a, b\}, \{a, c\}, \{b, c\}\}$, but the arguments c and b appear together in some extension in $\{\{a, b\}, \{a, c\}, \{b, c\}\}$, and the same holds true of c and a. This means that $\{\{a, b\}, \{a, c\}, \{b, c\}\}$ is not tight, and given the characterization of the signatures of semantics introduced in Section 2, it follows that $\{\{a, b\}, \{a, c\}, \{b, c\}\} \notin \Sigma_\sigma$, for $\sigma \in \{\mathit{stb}, \mathit{stg}\}$.[5] Hence, $\preceq$ is not σ-compliant.

On the other hand, let $\preceq'$ be the preorder defined as $\{a, b, c\} \prec' \{a\} \approx' \{b\} \approx' \{c\} \prec' \{a, b\} \prec' \{a, c\} \prec' \{b, c\} \prec' \emptyset$. It is straightforward to verify that, for $\sigma \in \{\mathit{stb}, \mathit{prf}, \mathit{stg}, \mathit{sem}\}$, $\preceq'$ is σ-compliant. As an example, for φ from above we get that $\min([\varphi], \preceq') = \{\{a\}, \{b\}, \{c\}\} \in \Sigma_\sigma$.

For the semantics of interest (stable, preferred, stage and semi-stable) and a set of extensions $\mathbb{S}$, we can check in polynomial time whether $\mathbb{S} \in \Sigma_\sigma$ [33]. Hence, we can decide in polynomial time whether a given preorder is σ-compliant.

The concept of σ-compliance makes an AGM-style representation result possible for AF revision by propositional formulas under arbitrary semantics. The following two results make this characterization precise.

Theorem 2. *If, for some semantics σ, there exists a faithful assignment mapping any $F \in AF_{\mathfrak{A}}$ to a σ-compliant and faithful ranking $\preceq_F$, let $\star_\sigma : AF_{\mathfrak{A}} \times \mathsf{P}_{\mathfrak{A}} \to AF_{\mathfrak{A}}$ be a revision operator defined as follows:*

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

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

Proof. Consider an arbitrary $F \in AF_{\mathfrak{A}}$. Since $\preceq_F$ is σ-compliant, we have $\min([\varphi], \preceq_F) \in \Sigma_\sigma$ for every $\varphi \in \mathsf{P}_{\mathfrak{A}}$. Therefore, by definition of f_σ (see Definition 2), it holds that $\sigma(f_\sigma(\min([\varphi], \preceq_F))) = \min([\varphi], \preceq_F)$. Hence, $\sigma(F \star_\sigma \varphi) = \min([\varphi], \preceq_F)$ and postulate P$\star$1 follows immediately. We will use this equality for arbitrary formulas as a shortcut in the rest of the proof.

If $\sigma(F) \cap [\varphi] \neq \emptyset$, it follows from $\preceq_F$ being faithful that $\min([\varphi], \preceq_F) = \sigma(F) \cap [\varphi]$, and thus P$\star$2 is satisfied. Postulate P$\star$3 holds because $\preceq_F$ is transitive and $\mathfrak{A}$ is finite and therefore if $[\varphi] \neq \emptyset$ then $[\varphi]$ has minimal elements, hence $\min([\varphi], \preceq_F) \neq \emptyset$.

[4] See Example 4.

[5] The characterizations of signatures [33] show that also $\{\{a, b\}, \{a, c\}, \{b, c\}\} \notin \Sigma_\tau$ for $\tau \in \{\mathit{prf}, \mathit{sem}\}$.

The preorder $\preceq_F$ being a faithful ranking means it has been obtained from a faithful assignment. Therefore, for any AF G with $\sigma(F) = \sigma(G)$ it must hold that $\preceq_F = \preceq_G$ (cf. (iii) from Definition 3). Therefore, for formulas $\varphi \equiv \psi$, $\min([\varphi], \preceq_F) = \min([\psi], \preceq_G)$. It follows that $\sigma(F \star_\sigma \varphi) = \sigma(G \star_\sigma \psi)$, showing that $\star_\sigma$ satisfies P$\star$4.

Postulates 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 $\sigma(F \star_\sigma (\varphi \wedge \psi)) \subseteq \sigma(F \star_\sigma \varphi) \cap [\psi]$ also holds, assume $E \in \min([\varphi \wedge \psi], \preceq_F)$ and $E \notin \min([\varphi], \preceq_F) \cap [\psi]$. Since $E \in [\psi]$, it follows that $E \notin \min([\varphi], \preceq_F)$. Let $E' \in \min([\varphi], \preceq_F) \cap [\psi]$ (assumed to be non-empty). 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 from $E' \in \min([\varphi], \preceq_F)$ it follows that $E \in \min([\varphi], \preceq_F)$, a contradiction. □

Theorem 2 shows that σ-compliant preorders can be used to obtain AF revision operators which meet our requirement of having their output expressible as a single AF. We specify some concrete ways of constructing σ-compliant preorders in Section 3.2. The next result shows that all operators satisfying postulates P$\star$1–P$\star$6 are represented by some σ-compliant assignment.

Theorem 3. *If $\star_\sigma\colon AF_{\mathfrak{A}} \times \mathsf{P}_{\mathfrak{A}} \to AF_{\mathfrak{A}}$ is an operator satisfying postulates P$\star$1–P$\star$6, for some semantics σ, then there exists a faithful assignment mapping every $F \in AF_{\mathfrak{A}}$ to a σ-compliant faithful ranking $\preceq_F$ on $2^{\mathfrak{A}}$ such that $\sigma(F \star_\sigma \varphi) = \min([\varphi], \preceq_F)$, for every $\varphi \in \mathsf{P}_{\mathfrak{A}}$.*

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

Let $F \in AF_{\mathfrak{A}}$ be an arbitrary AF. We define the binary relation $\preceq_F$ on $2^{\mathfrak{A}}$ as follows:

$$E \preceq_F E' \quad \text{if and only if} \quad E \in \sigma(F \star_\sigma \phi(E, E')).$$

We begin by showing that $\preceq_F$ is a total preorder. 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, also by P$\star$1 and P$\star$3, $\sigma(F \star_\sigma \phi(E)) = \{E\}$. Hence $E \preceq_F E$ holds for each $E \in 2^{\mathfrak{A}}$. In other words, $\preceq_F$ is reflexive.

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 $\{E_1, E_2, E_3\}$. We reason by case analysis. *Case 1.* Assume, first, that $\sigma(F \star_\sigma \phi(E_1, E_2, E_3)) \cap \{E_1, E_2\} = \emptyset$. 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 get from P$\star$4 that $\sigma(F \star_\sigma \phi(E_2, E_3)) = \sigma(F \star_\sigma \phi(E_2, E_3) \wedge \phi(E_1, E_2, E_3))$. By P$\star$5 and P$\star$6 we obtain $\sigma(F \star_\sigma \phi(E_1, E_2, E_3)) \cap \{E_2, E_3\} = \sigma(F \star_\sigma (\phi(E_2, E_3) \wedge \phi(E_1, E_2, E_3)))$. Combining the last two equalities, we get $\sigma(F \star_\sigma \phi(E_1, E_2, E_3)) \cap \{E_2, E_3\} = \sigma(F \star_\sigma \phi(E_2, E_3))$. But this implies that $\sigma(F \star_\sigma \phi(E_2, E_3)) = \{E_3\}$, contradicting the fact that $E_2 \preceq_F E_3$. *Case 2.* Assume, next, that $\sigma(F \star_\sigma \phi(E_1, E_2, E_3)) \cap \{E_1, E_2\} \neq \emptyset$. Since $E_1 \preceq_F E_2$ we know that $E_1 \in \sigma(F \star_\sigma \phi(E_1, E_2))$ holds. Considering 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$4, P$\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, $E_1 \in \sigma(F \star_\sigma$

$\phi(E_1, E_2, E_3)) \cap \{E_1, E_2\}$, and also $E_1 \in \sigma(F \star_\sigma \phi(E_1, E_2, E_3)) \cap \{E_1, E_3\}$ holds. Considering the fact that $\phi(E_1, E_3) \equiv \phi(E_1, E_3) \wedge \phi(E_1, E_2, E_3)$, we obtain from P$\star$4, P$\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 preorder. The following lemmata show that $\star_\sigma$ is indeed represented by the assignment mapping F to $\preceq_F$.

Lemma 1. *If $E_1, E_2 \in 2^{\mathfrak{A}}$ such that $E_1 \preceq_F E_2$, then for all formulas $\varphi \in \mathsf{P}_{\mathfrak{A}}$, it holds that 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 we can simplify the equation we derived from P$\star$5 and P$\star$6 to $\sigma(F \star_\sigma \phi(E_1, E_2)) = \sigma(F \star_\sigma \varphi) \cap [\phi(E_1, E_2)]$. 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. *If $\varphi \in \mathsf{P}_{\mathfrak{A}}$, then it holds that $\min([\varphi], \preceq_F) = \sigma(F \star_\sigma \varphi)$.*

Proof. We show the double inclusion. For the $\subseteq$-direction, take $\varphi \in \mathsf{P}_{\mathfrak{A}}$ and an extension $E_1 \in \min([\varphi], \preceq_F)$. Since $E_1 \in [\varphi]$ and $[\varphi] \neq \emptyset$, we get by P$\star$3 that $\sigma(F \star_\sigma \varphi) \neq \emptyset$. Take, therefore, an extension $E_2 \in \sigma(F \star_\sigma \varphi)$. By P$\star$1 we have that $E_2 \in [\varphi]$ and hence $E_1 \preceq_F E_2$. By Lemma 1, it follows that $E_1 \in \sigma(F \star_\sigma \varphi)$.

For the $\supseteq$-direction, take $\varphi \in \mathsf{P}_{\mathfrak{A}}$ and $E_1 \in \sigma(F \star_\sigma \varphi)$. By P$\star$1, we know that $E_1 \in [\varphi]$. We show that for all $E_2 \in [\varphi]$ it holds that $E_1 \preceq_F E_2$. To this end, take an arbitrary $E_2 \in [\varphi]$: 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, $E_1 \in \sigma(F \star_\sigma \phi(E_1, E_2))$ must also hold, meaning that $E_1 \preceq_F E_2$. Since E_2 was chosen arbitrarily, it follows that $E_1 \in \min([\varphi], \preceq_F)$. □

It is uncontroversial that $\sigma(F \star_\sigma \varphi) \in \Sigma_\sigma$, so by Lemma 2 it follows that $\preceq_F$ is σ-compliant. What is left to show is that the definition of $\preceq_F$ gives rise to a faithful assignment for AFs. We begin by showing that properties (i) and (ii) of Definition 3 hold. If $\sigma(F) = \emptyset$ this is trivially the case, hence let us assume that $\sigma(F) \neq \emptyset$. By P$\star$2 we get $\sigma(F \star_\sigma \top) = \sigma(F)$ (since $[\top] = 2^{\mathfrak{A}}$ and therefore $[\top] \cap \sigma(F) = \sigma(F)$). Hence $\sigma(F) = \min([\top], \preceq_F) = \min(2^{\mathfrak{A}}, \preceq_F)$, meaning that for $E_1, E_2 \in 2^{\mathfrak{A}}$, $E_1 \not\prec_F E_2$ 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. Finally, condition (iii) holds since, for any AFs $F, G \in AF_{\mathfrak{A}}$ with $\sigma(F) = \sigma(G)$ and any sets of arguments $E, E' \subseteq \mathfrak{A}$, P$\star$4 ensures that $\sigma(F \star_\sigma \phi(E, E')) = \sigma(G \star_\sigma \phi(E, E'))$, hence $\preceq_F = \preceq_G$. It follows that $\preceq_F$ gives rise to a faithful assignment. □

Theorems 2 and 3 are very general in capturing any possible semantics of AFs. However, rather implicitly, the results impose an important property on a semantics σ: namely, that for each

AF F, every non-empty subset $\mathbb{S}$ of $\sigma(F)$ is again realizable under σ.[6] In Theorem 2 this is by the preorder being both faithful and σ-compliant, while in Theorem 3 it is ensured by the operator satisfying P$\star$2. The following result shows that no rational operators exist for semantics not having this property.

Proposition 2. *Let τ be a semantics such that Property 2 from Definition 1 does not hold. Then there is no operator $\star_\tau : AF_{\mathfrak{A}} \times \mathsf{P}_{\mathfrak{A}} \rightarrow AF_{\mathfrak{A}}$ satisfying P$\star$2.*

Proof. Semantics τ not fulfilling Property 2 from Definition 1 means that are is some $\mathbb{S} \in \Sigma_\tau$ and some $\mathbb{S}' \subseteq \mathbb{S}$ such that $\mathbb{S}' \neq \emptyset$ and $\mathbb{S}' \notin \Sigma_\tau$. Now let F be an AF such that $\tau(F) = \mathbb{S}$ (it exists by $\mathbb{S} \in \Sigma_\tau$) and consider the formula $\phi(\mathbb{S}')$ having $[\phi(\mathbb{S}')] = \mathbb{S}'$. By $\mathbb{S}' \subseteq \mathbb{S}$ and $\mathbb{S}' \neq \emptyset$ we have $\tau(F) \cap [\phi(\mathbb{S}')] \neq \emptyset$. Therefore any operator $\star_\tau$ would be required to give $\tau(F \star_\tau \varphi) = \tau(F) \cap [\phi(\mathbb{S}')] = \mathbb{S}'$, which is not possible since $\mathbb{S}' \notin \Sigma_\tau$. □

3.2 Concrete Operators

With Theorems 2 and 3, finding concrete AF revision operators comes down to defining appropriate rankings on extensions, where by appropriate we mean faithful and σ-compliant. When dealing with proper I-maximal semantics, an easy and immediate way to construct such rankings is to use linear orders on extensions.

Proposition 3. *Consider a faithful assignment from AFs to faithful rankings which, for any semantics σ, $F \in AF_{\mathfrak{A}}$, and $E_1, E_2 \in 2^{\mathfrak{A}}$, satisfies the following additional property:*

(iv) *if $E_1, E_2 \notin \sigma(F)$, then either $E_1 \prec_F E_2$ or $E_2 \prec_F E_1$.*

If σ is proper I-maximal, any revision operator $\star_\sigma$ represented by this assignment satisfies postulates P$\star$1–P$\star$6.

Proof. Considering Theorem 2, all that is left to show is that $\preceq_F$ is σ-compliant. If $\sigma(F) \cap [\varphi] \neq \emptyset$, it follows that $\min([\varphi], \preceq_F) = \sigma(F) \cap [\varphi] \subseteq \sigma(F)$. By Property 2 of proper I-maximal semantics, $\sigma(F) \cap [\varphi]$ is realizable under σ. If $\sigma(F) \cap [\varphi] = \emptyset$, notice first that condition (iv) is equivalent to saying that for the interpretations outside $\sigma(F)$, $\preceq_F$ behaves like a linear order. This means that $\min([\varphi], \preceq_F)$ is a singleton, and thus realizable under σ (cf. Property 3 of proper I-maximal semantics). □

Example 6. Take the AF F from Example 3 with *prf*$(F) = \{\{a, d\}, \{b, d\}\}$, which we revise by $\varphi = c \wedge d$. Suppose that we have an assignment which maps F to a preorder $\preceq_F$ where $\{c, d\} \prec_F \{a, c, d\} \prec_F \{b, c, d\} \prec_F \{a, b, c, d\}$, being in line with Property (iv) from Proposition 3. We get that $\min([\varphi], \preceq_F) = \{\{c, d\}\}$, hence for an operator $\star_{prf}$ represented by this assignment, $F \star_{prf} \varphi$ corresponds to an AF which has only one preferred extension, namely $\{c, d\}$. Thus, $F \star_{prf} \varphi$ could be the AF G' in Figure 5.

[6] As it turns out, this coincides with Property 2 of proper I-maximal semantics, the class of semantics we will focus on in Section 4.

Figure 5: F is to be revised with $\varphi = c \wedge d$, resulting in G'.

As mentioned in Section 2, for any σ-realizable set of extensions $\mathbb{S}$ there are infinitely many AFs F such that $\sigma(F) = \mathbb{S}$. Thus, in Example 6 we could have chosen any AF whose set of preferred extensions is the singleton $\{\{c, d\}\}$, and it is a legitimate question which of the possible AFs to choose as the revision outcome. We do not touch on this issue other than by saying that *some* AF shall be chosen. Nevertheless, it would be natural to enforce some kind of minimality requirement at this step as well, for instance by expecting $F \star_\sigma \varphi$ to modify the attack relation as little as possible. Such a requirement is to be explored in future work, where it would be integrated with recent results in this direction [20, 48].

Though easy to define, AF revision operators based on linear preorders can be excessively discriminating in their choice of extensions. Using a more familiar option such as Dalal's operator also does not work, since the rankings obtained with Hamming distance are usually not σ-compliant for the semantics σ under consideration and argumentation semantics in general (see Example 5).

To obtain alternative revision operators, we first introduce some new notions. In the following, we assume that arguments in $\mathfrak{A}$ are strictly ranked according to something like *an alphabetical order*, such that a is preferred to b, a_1 is preferred to a_2, and so on. The exact choice of this ranking does not matter so much, just that it orders the arguments linearly. For an extension E, $\overrightarrow{E}$ is a vector obtained by ordering the arguments in E in descending order according to the ranking just introduced. Thus, if $E = \{c, a, d, b\}$, then $\overrightarrow{E} = (a, b, c, d)$. We are then able to compare such vectors according to the lexicographic order $\preceq_{\text{lex}}$ in the obvious way. Thus, we have that $(a, c) \prec_{\text{lex}} (b, c)$ and $(a, b) \prec_{\text{lex}} (a, c)$. If the length of $\overrightarrow{E}$ is k, then *the prefix of* $\overrightarrow{E}$, denoted $\overrightarrow{E}^{\#}$, is the vector containing the first $k-1$ elements of $\overrightarrow{E}$. For example, if $\overrightarrow{E} = (a, b, c, d)$, then $\overrightarrow{E}^{\#} = (a, b, c)$. By convention, if $|\overrightarrow{E}| = 1$, then $\overrightarrow{E}^{\#} = \overrightarrow{\emptyset}$.

Next, we show that any set of extensions can be partitioned in such a way that elements of the partition are σ-realizable, at least for any semantics σ such that $\Sigma_{stg} \subseteq \Sigma_\sigma$. This partition then provides the means to define a broad range of faithful rankings.

Definition 5. If $\mathbb{S}$ is a finite set of extensions, *the indexed preorder* $\preceq^{\mathbb{S}}$ *on* $\mathbb{S}$ is defined, for any $E_1, E_2 \in \mathbb{S}$, as follows:

$$E_1 \preceq^{\mathbb{S}} E_2 \text{ if and only if } \quad |E_1| < |E_2| \text{ or, } \\ |E_1| = |E_2| \text{ and } \overrightarrow{E_1}^{\#} \preceq_{\text{lex}} \overrightarrow{E_2}^{\#}.$$

It is straightforward to see that $\preceq^{\mathbb{S}}$ is reflexive, transitive and total, and is thus a total preorder on $\mathbb{S}$. In the following we will refer to it as *the indexed preorder on* $\mathbb{S}$. Moreover, $\preceq^{\mathbb{S}}$ partitions $\mathbb{S}$ into sets of extensions, which can be visualized as distinct levels of $\mathbb{S}$ (see Example 7). We call

$\preceq^{\mathbb{S}_1}$	$\preceq^{\mathbb{S}_2}$
$\{a, b, c\}$	$\{b, c, d, e, f\}$
$\|$	$\|$
$\{b, c\}$	$\{a, c, d, e, f\}, \{a, c, d, e, g\}$
$\|$	$\|$
$\{a, b\}, \{a, c\}$	$\{a, b, c, d, e\}, \{a, b, c, d, f\}, \{a, b, c, d, g\}$
$\|$	$\|$
$\{a\}, \{b\}, \{c\}$	$\{d, e, f\}$
$\|$	$\|$
$\emptyset$	$\{c, d, e\}, \{c, d, g\}$

Figure 6: The indexed levels of $\mathbb{S}_1$ and $\mathbb{S}_2$.

the sets in this partition of $\mathbb{S}$ *the indexed levels of* $\mathbb{S}$. The indexed levels are such that for any two extensions $E_1, E_2 \in \mathbb{S}$, if E_1 and E_2 are in the same indexed level, then $E_1 \approx^{\mathbb{S}} E_2$, and if E_1 and E_2 belong to different indexed levels, then either $E_1 \prec^{\mathbb{S}} E_2$, or $E_2 \prec^{\mathbb{S}} E_1$.

Example 7. Figure 6 depicts two sets of extensions $\mathbb{S}_1$ and $\mathbb{S}_2$ arranged according to their indexed levels. The convention is that the more preferred a level is, the lower it is displayed in the preorder. Extensions with greater cardinality are strictly less preferred than extensions with smaller cardinality, and among extensions of equal cardinality tie-breaking occurs according to the lexicographic order applied to the prefixes. Thus, indexed levels consist of extensions of equal cardinality with the same prefix.

Intuitively, the indexed preorder gives precedence to extensions with fewer elements, and then to elements placed earlier in the alphabetical order. This approach fits nicely with the idea that arguments have a certain priority, such that if forced to choose between extensions, then we will choose extensions with higher priority arguments. And, while the specific choice of a priority relation may seem arbitrary here without more concrete information about the argumentation context, it seems clear that we need to select among extensions if revision is to occur. For instance, consider revising an AF F by a formula φ, where the models of φ are $[\varphi] = \{\{a, b\}, \{b, c\}, \{a, c\}\}$ and $\sigma(F) \cap [\varphi] = \emptyset$. We cannot accept $[\varphi]$ as the outcome of the revision operator, because $[\varphi]$ is not σ-realizable under any of the semantics among stable, preferred, semi-stable and stage. Hence, one way or another, some kind of choice has to be made among the models of φ, and it seems natural to assume that a revision operator would choose according to some implicit preference over arguments. In our case, the indexed preorder gives us that $\{a, b\} \approx^{[\varphi]} \{a, c\} \prec^{[\varphi]} \{b, c\}$, and hence $\{a, b\}$ and $\{a, c\}$ are chosen, while $\{b, c\}$ is left out.

Sets of extensions constituting a indexed level of any indexed preorder turn out to have beneficial properties.

Proposition 4. *If $\mathbb{S}$ is a set of extensions and $\mathbb{S}_i$ is one of its indexed levels, then any set of extensions $\mathbb{S}' \subseteq \mathbb{S}_i$ is tight and incomparable.*

Proof. Suppose $\mathbb{S}' = \{E_1, \dots, E_n\}$. Since $E_1, \dots, E_n$ are on the same level, they have the same cardinality and their prefixes coincide. Let us then write $E_i = \{a_1, \dots, a_k, b_i\}$, for $i \in \{1, \dots, n\}$. One can immediately see that any two distinct extensions E_i and E_j in $\mathbb{S}'$ are $\subseteq$-incomparable, since they differ in arguments b_i and b_j. Moreover, take an extension E_i from $\mathbb{S}'$ and an argument $a \in \bigcup_{S \in \mathbb{S}} S$ such that $E_i \cup \{a\} \notin \mathbb{S}$. The only way this can happen is if $a = b_j$, for $j \neq i$. But then b_j and b_i never appear together in any of the extensions of $\mathbb{S}'$, which shows that $\mathbb{S}'$ is tight. □

Since, for a non-empty set of extensions $\mathbb{S}$, tight and incomparable are sufficient conditions for $\mathbb{S} \in \Sigma_{stg}$, the insight gained from Proposition 4 allows us to define σ-compliant preorders based on the indexed preorder for any semantics σ which has $\Sigma_\sigma \supseteq \Sigma_{stg}$. In the remainder of this section, we assume σ to be an arbitrary such semantics. In particular, as pointed out in Section 2, this includes also stable, preferred and semi-stable semantics.

Definition 6. For an AF F and a proper I-maximal semantics σ such that $\Sigma_\sigma \supseteq \Sigma_{stg}$, the *canonical preorder* $\preceq_F^{\text{can}}$ on $2^{\mathfrak{A}}$ is defined, for any $E_1, E_2 \in 2^{\mathfrak{A}}$, as follows:

$$E_1 \preceq_F^{\text{can}} E_2 \text{ if and only if } \begin{cases} E_1 \in \sigma(F) \text{ or,} \\ E_1, E_2 \notin \sigma(F) \text{ and } E_1 \preceq^{2^{\mathfrak{A}} \setminus \sigma(F)} E_2. \end{cases}$$

Definition 7. For an AF F, a proper I-maximal semantics σ such that $\Sigma_\sigma \supseteq \Sigma_{stg}$, and a given faithful ranking $\preceq_F$, the *indexed refinement of* $\preceq_F$ is a preorder defined, for any $E_1, E_2 \in 2^{\mathfrak{A}}$, as follows:

$$E_1 \preceq_F^{\text{ir}} E_2 \text{ if and only if } \begin{cases} E_1 \in \sigma(F) \text{ or,} \\ E_1, E_2 \notin \sigma(F), E_1 \approx_F E_2 \text{ and } E_1 \preceq^{2^{\mathfrak{A}} \setminus \sigma(F)} E_2 \text{ or,} \\ E_1, E_2 \notin \sigma(F) \text{ and } E_1 \prec_F E_2. \end{cases}$$

In the canonical preorder $\preceq_F$ we have the extensions of F as the minimal elements, while the remaining extensions in $2^{\mathfrak{A}}$ are ordered according to the indexed preorder. The indexed refinement is obtained by taking an existing faithful ranking $\preceq_F$ (which, recall, may not be σ-compliant) and rearranging its levels according to the indexed preorder, leaving the inter-level ranking unchanged. The up-shot is that the new levels will be σ-compliant (see Example 8).

Example 8. Let F be an AF such that $stb(F) = \{\{b, c\}\}$. Figure 7 depicts the canonical preorder $\preceq_F^{\text{can}}$ and Figure 8 shows the ranking $\preceq_F^{D,\text{ir}}$, obtained by refining the ranking $\preceq_F^D$. The latter, in turn, is generated with Hamming distance and is not σ-compliant, for any semantics $\sigma \in \{stb, prf, sem, stg\}$. Notice, on the other hand, that both $\preceq_F^{\text{can}}$ and $\preceq_F^{D,\text{ir}}$ are σ-compliant. Also notice how the levels of $\preceq_F^D$ get split according to the indexed preorder to obtain $\preceq_F^{D,\text{ir}}$.

Using the canonical and the refined preorders, we can define AF revision operators in the familiar way, by taking $F *_\sigma \varphi = f_\sigma(\min([\varphi], \preceq_F))$. We will call the operator defined using the canonical preorder *the canonical operator*, and denote it by $\star_\sigma^{\text{can}}$. If $\star_\sigma^x$ is an existing AF revision operator, we will call the operator defined using the indexed preorder the *indexed-refined revision*

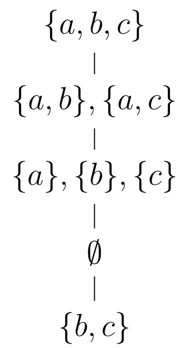

Figure 7: $\preceq_F^{\text{can}}$.

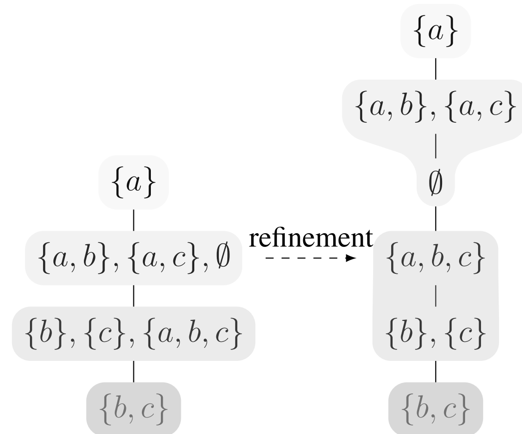

Figure 8: $\preceq_F^D$ and $\preceq_F^{D,\,\text{ir}}$.

operator, and denote it by $\star_\sigma^{x,\,\text{ir}}$. Notice that Definition 7 can be used to refine any existing, standard revision operator, by defining a new assignment on top of the standard one. In particular, we get operators such as $\star_\sigma^{D,\,\text{ir}}$, obtained by refining Dalal's operator. We show next that they also satisfy postulates P⋆1–P⋆6.

Theorem 4. *For a proper I-maximal semantics σ such that $\Sigma_{stg} \subseteq \Sigma_\sigma$, the revision operator $\star_\sigma^{can}$ and the family of revision operators $\star_\sigma^{x,\text{ir}}$ are well defined and they satisfy postulates P⋆1–P⋆6.*

Proof. By Proposition 4, the canonical and refined assignments are σ-compliant on $2^{\mathfrak{A}} \backslash \sigma(F)$. By proper I-maximality of σ, any proper subset of $\sigma(F)$ is also σ-realizable. Therefore the operators are well-defined. They are also faithful, hence by Theorem 2 the operators satisfy postulates P⋆1–P⋆6. □

4 Revision by Argumentation Frameworks

Next, we consider revision of an AF by another AF, performed through operators of the type $*_\sigma \colon AF_{\mathfrak{A}} \times AF_{\mathfrak{A}} \to AF_{\mathfrak{A}}$. Such operators map an AF F and an AF G to an AF $F *_\sigma G$, the intuitive idea being that we want to change F minimally, in order to incorporate the models of G. The underlying concept of a model is given, as before, by the argumentation semantics σ. We consider here the class of proper I-maximal semantics including stable, preferred, stage and semi-stable semantics. 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 [38], 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 $\sigma(F_1) = \sigma(F_2)$ and $\sigma(G) = \sigma(H)$, then $\sigma(F_1 *_\sigma G) = \sigma(F_2 *_\sigma H)$.

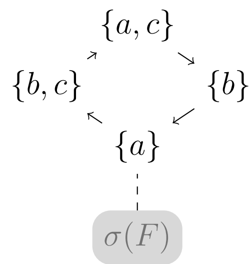

Figure 9: Cycles in rankings on extensions.

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

The postulate Acyc is adapted from [26] and is motivated by the realization that, without it, postulates A∗1–A∗6 can characterize revision operators generated with unsuitable rankings (see Example 9).

Example 9. Suppose that for an AF F we have a ranking $\preceq_F$ on $2^{\mathfrak{A}}$ which behaves as in Figure 9 for the extensions $\{a\}, \{b, c\}, \{a, c\}$ and $\{b\}$, and as a faithful ranking otherwise. An arrow means that the relation is strict: for example, $\{a\} \preceq_F \{b, c\}$ and $\{b, c\} \not\preceq_F \{a\}$, abbreviated as $\{a\} \prec_F \{b, c\}$. The relation $\preceq_F$, then, contains a non-transitive cycle and is not a preorder. However, quick inspection of the figure reveals that for any non-empty σ-realizable set $\mathbb{S}$, $\min(\mathbb{S}, \preceq_F)$ is still well defined and non-empty (recall that we are assuming σ to be proper I-maximal; therefore elements of $\mathbb{S}$ are pairwise $\subseteq$-incomparable). For instance, if $\mathbb{S} = \{\{a\}, \{b, c\}\}$, then $\min(\mathbb{S}, \preceq_F) = \{\{a\}\}$. Thus we can define an operator $*_\sigma$ in the familiar way, by taking $F *_\sigma G = f_\sigma(\min(\sigma(G), \preceq_F))$, and it is then straightforward to verify that this operator $*_\sigma$ is well-defined and satisfies postulates A∗1–A∗6.

Additionally, there is no ranking $\preceq'_F$ which is transitive and yields the same revision operator. To see this, notice that if such a ranking (call it $\preceq'_F$) existed, it would have to satisfy $\min(\{\{a\}, \{b, c\}\}, \preceq'_F) = \{\{a\}\}$, 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, we want to avoid non-transitive cycles: since a natural reading of the rankings on $2^{\mathfrak{A}}$ is that they are plausibility relations, one would expect them to be transitive, and it is thus undesirable to have revision operators that characterize non-transitive rankings. To prevent this situation we make use of the additional postulate Acyc.

On the ranking side we define a less demanding version of faithful assignments, which is adjusted to the nature of (proper) I-maximal semantics.

Definition 8. Given a semantics σ, an *I-faithful assignment* maps every $F \in AF_{\mathfrak{A}}$ to an I-total preorder $\preceq_F$ 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 $E_1 \approx_F E_2$,

(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 preorder $\preceq_F$ assigned to F by an I-faithful assignment is called the *I-faithful ranking associated with* F.

I-faithful assignments differ from faithful assignments in that they require the rankings to be only I-total, thus allowing (but not requiring) them to be partial with respect to $\subseteq$-comparable pairs of extensions. Our use of I-faithful assignments is motivated by how 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 'at least as plausible' as 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 decide between E_1 and E_2, hence it is natural to allow them to be incomparable with respect to $\preceq_F$. With these preliminaries, we can now state our main representation results.

Theorem 5. *If, for some proper I-maximal semantics σ, there exists an I-faithful assignment mapping any $F \in AF_{\mathfrak{A}}$ to an I-faithful ranking $\preceq_F$, let $*_\sigma \colon AF_{\mathfrak{A}} \times AF_{\mathfrak{A}} \to AF_{\mathfrak{A}}$ be a revision operator defined as follows:*

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

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

Proof. Since σ is proper I-maximal, any non-empty subset of $\sigma(G)$ (and 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. Specifically, for any AF G, $\sigma(F *_\sigma G) = \min(\sigma(G), \preceq_F)$, which we use without further comment in the remainder of the proof.

It is straightforward to see that A$*$1 is satisfied. Postulate A$*$2 holds, since the elements of $\sigma(F)$ are the minimal elements of $\preceq_F$, as $\preceq_F$ is I-faithful. Postulate A$*$3 follows from transitivity of $\preceq_F$ and finiteness of G. Postulate A$*$4 follows from Property (iii) of I-faithful assignments. Postulates A$*$5 and A$*$6 can be shown analogously to P$\star$5 and P$\star$6 in Theorem 2.

It remains to be shown that Acyc also holds. Let $G_0, G_1, \ldots, G_n$ be a sequence of AFs such that for all $0 \leq i < n$, $\sigma(F *_\sigma G_{i+1}) \cap \sigma(G_i) \neq \emptyset$ and $\sigma(F *_\sigma G_0) \cap \sigma(G_n) \neq \emptyset$ holds. From $\sigma(F *_\sigma G_1) \cap \sigma(G_0) \neq \emptyset$ 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 $\sigma(F *_\sigma G_2) \cap \sigma(G_1) \neq \emptyset$ that there is an extension $E_1' \in \sigma(G_1)$ such that $E_1' \preceq_F E_2$ for all

$E_2 \in \sigma(G_2), \dots$, 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 $\sigma(F *_\sigma G_0) \cap \sigma(G_n) \neq \emptyset$ it follows that there is some $E'_n \in \sigma(G_n)$ with $E'_n \preceq_F E_0$ for all $E_0 \in \sigma(G_0)$ (in particular for E'_0). Now from $E'_n \preceq_F E'_0 \preceq_F E_n$ (for all $E_n \in \sigma(G_n)$) it follows that $E'_n \in \min(\sigma(G_n), \preceq_F)$. Hence $\sigma(F *_\sigma G_n) \cap \sigma(G_0) \neq \emptyset$. □

Theorem 6. *If $*_\sigma : AF_{\mathfrak{A}} \times AF_{\mathfrak{A}} \to AF_{\mathfrak{A}}$ is 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 $\sigma(F *_\sigma G) = \min(\sigma(G), \preceq_F)$, for any $G \in AF_{\mathfrak{A}}$.*

Proof. Assume there is an operator $*_\sigma : AF_{\mathfrak{A}} \times AF_{\mathfrak{A}} \to AF_{\mathfrak{A}}$ satisfying postulates A∗1–A∗6 and Acyc, and take an arbitrary $F \in AF_{\mathfrak{A}}$. 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 $E, E' \in 2^{\mathfrak{A}}$:

$E \preceq'_F E'$ if and only if $E \in \sigma(F *_\sigma f_\sigma(\{E, E'\}))$.

The relation $\preceq'_F$ is reflexive, as A∗1 and A∗3 imply that $E \in \sigma(F *_\sigma f_\sigma(\{E\}))$. In the next step we take $\preceq_F$ to be the transitive closure of $\preceq'_F$. In other words:

$E \preceq_F E'$ if and only if there exist $E_1, \dots, E_n$ 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 preorder 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 $\preceq_F$ 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 $\{E_1, E_2\}$ 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 $\sigma(F *_\sigma f_\sigma(\{E_1, E_2\})) = \sigma(F) \cap \{E_1, E_2\} = \{E_1, E_2\}$. 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 $E_1, \dots, E_n$ 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. From the definition of $\preceq_F$ it is clear that $E \preceq_F E'$. It remains to be shown that $E' \not\preceq_F E$. Suppose, towards a contradiction, 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 $E \prec'_F E'$ by assumption, 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). For Property (iii) assume an AF $F' \in AF_{\mathfrak{A}}$ with $\sigma(F) = \sigma(F')$. A$*$4 ensures that $\preceq'_F = \preceq'_{F'}$ and therefore it also holds that $\preceq_F = \preceq_{F'}$.

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

Lemma 7. *For any two extensions $E_1, E_2 \subseteq \mathfrak{A}$ 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. From the assumption that $E_2 \in \sigma(F *_\sigma G)$, we have $\sigma(F *_\sigma G) \cap \{E_1, E_2\} \neq \emptyset$. By A$*$5 and A$*$6 we get $\sigma(F *_\sigma G) \cap \{E_1, E_2\} = \sigma(F *_\sigma f_\sigma(\sigma(G) \cap \{E_1, E_2\}))$. Moreover, by A$*$1, we get that $E_2 \in \sigma(G)$. We also know that $E_1 \in \sigma(G)$, so $\{E_1, E_2\} \subseteq \sigma(G)$. Thus $\sigma(G) \cap \{E_1, E_2\} = \{E_1, E_2\}$ and from this and A$*$4 it follows that $\sigma(F *_\sigma f_\sigma(\sigma(G) \cap \{E_1, E_2\})) = \sigma(F *_\sigma f_\sigma(\{E_1, E_2\}))$. Putting these results together with the fact that $E_1 \in \sigma(F *_\sigma f_\sigma(\{E_1, E_2\}))$ (since $E_1 \preceq'_F E_2$), we get that $E_1 \in \sigma(F *_\sigma G)$. $\square$

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 resembles the one for Lemma 2. $\square$

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

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 $E_1 \in \min(\sigma(G), \preceq_F)$. □

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

Regarding concrete operators, notice that any faithful assignment for AFs can be used, *via* Theorem 5, to represent a revision operator $*_\sigma : AF_{\mathfrak{A}} \times AF_{\mathfrak{A}} \to AF_{\mathfrak{A}}$. Remarkably, then, for revision by AFs we do not need a restriction on rankings such as σ-compliance to ensure that operators are well defined. The reason revision by AFs 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 5, that $*_\sigma$ satisfies A$*$1–A$*$6 and Acyc. Thus, any model-based revision operator from the standard literature on belief change (for example Dalal's operator [23]) can be used as a revision operator of AFs by AFs.

Example 10. Consider an AF F as in Example 5, with $stb(F) = \{\{a, b, c\}\}$, for instance $F = (\{a, b, c\}, \emptyset)$. The corresponding ranking obtained with Hamming distance, $\{a, b, c\} \prec^D_F \{a, b\} \approx^D_F \{a, c\} \approx^D_F \{b, c\} \prec^D_F \{a\} \approx^D_F \{b\} \approx^D_F \{c\} \prec^D_F \emptyset$, was problematic when revising by a propositional formula, because the desired outcome of a revision operator could turn out to be $\{\{a, b\}, \{b, c\}, \{a, c\}\}$, which usually is not σ-realizable (see Example 5). We cannot, however, run into this problem when revising by an AF G, since the outcome of revision will, by definition, be a proper subset of $\sigma(G)$, namely $\min(\sigma(G), \preceq^D_F)$. Due to the proper I-maximality of σ, any proper subset of $\sigma(G)$ is also σ-realizable. It follows that Dalal's operator and, by the same token, any other standard revision operator, can be applied in this setting.

5 Complexity

Next, we study the complexity of Dalal's operator and its refinement in the argumentation setting. We will consider the following decision problem for semantics $\sigma \in \{stb, prf\}$:

GIVEN:	the original AF F, the revising AF G (or formula φ), and a set of arguments E,
DECIDE:	whether E is a σ-extension of the revision of F by G (or φ).

In particular, the problem is closely related to model checking in propositional logic revision, the complexity of which was studied by Liberatore and Schaerf [42]. We will first show the exact complexity of Dalal revision by AFs and then give complexity bounds for the refinement of Dalal's operator for revision by formulas.

We assume familiarity with standard complexity concepts, such as P, NP and completeness. Given a complexity class $\mathcal{C}$, a $\mathcal{C}$ oracle decides a given sub-problem from $\mathcal{C}$ in one computation step. The class Σ^{P}_k (and Δ^{P}_k) contains the problems that can be decided in polynomial time by a non-deterministic (deterministic) Turing machine with unrestricted access to a $\Sigma^{\mathsf{P}}_{k-1}$ oracle. In particular, $\Sigma_0^{\mathsf{P}} = \mathsf{P}$, $\Sigma_1^{\mathsf{P}} = \mathsf{NP}$, and $\Delta_2^{\mathsf{P}} = \mathsf{P}^{\mathsf{NP}}$. The classes Δ_k^{P} have been refined by the classes Θ_k^{P} (also denoted $\Delta_k^{\mathsf{P}}[\mathcal{O}(\log m)]$), in which the number of oracle calls in bounded by $\mathcal{O}(\log m)$, where m is the input size.

The complexity classes introduced above have complete problems involving quantified Boolean formulas (QBFs). By a k-existential QBF we denote a QBF of the form $Q_1X_1 \dots Q_kX_k\varphi(X_1, \dots, X_k)$ with $Q_1 = \exists$, $Q_2, \dots, Q_k \in \{\exists, \forall\}$, $Q_i \neq Q_{i+1}$ for $1 \leq i < k$, and (i) if $Q_k = \forall$ then φ is in DNF containing no monoms which are trivial for $X_1 \cup \dots \cup X_{k-1}$, (ii) if $Q_k = \exists$ then φ is in CNF containing no clauses which are trivial for $X_1 \cup \dots \cup X_{k-1}$. A monom m (or clause c) is trivial for X if all atoms occurring in m (or c) are contained in X. In particular, a 1-existential QBF is of the form $\exists X\varphi(X)$ with φ being in CNF without empty clauses. It is true if and only if $\varphi(X)$ is satisfiable. For a set of arguments $X = \{x_1, \dots, x_n\}$ we denote by $\overline{X}$ the set of arguments $\{\overline{x_1}, \dots, \overline{x_n}\}$.

The classes $\Theta_{k+1}^{\mathsf{P}}$ (for $k \geq 1$) have the following complete problems [35, 57, 55], which we will make use of in the subsequent hardness proofs:

GIVEN:	k-existential QBFs $\Phi_1, \dots, \Phi_m$ such that Φ_i being false implies Φ_{i+1} being false for $1 \leq i < m$,
DECIDE:	whether $\max\{1 \leq i \leq m \mid \Phi_i \text{ is true}\}$ is odd.

Given an AF F and a set of arguments E, deciding whether $E \in \mathit{stb}(F)$ is in P and deciding whether $E \in \mathit{prf}(F)$ is coNP-complete [29].

We begin with the complexity of Dalal's operator for revision by AFs under stable semantics. We will make use of the following construction, which is adapted from reductions used in proofs by Dimopoulos and Torres [29] and Dunne and Bench-Capon [32]. In what follows, we consider a more general setting by giving up the restriction that $\mathfrak{A}$ is finite.

Definition 9. Given a propositional formula $\varphi(X) = \bigwedge_{c\in C} c$ with each $c \in C$ a disjunction of literals from X, we define $F_\varphi = (A_\varphi, R_\varphi)$ as:

$$\begin{aligned}
A_\varphi =& X \cup \overline{X} \cup C \cup \{\varphi, \overline{\varphi}\},\\
R_\varphi =& \{(x, \overline{x}), (\overline{x}, x) \mid x \in X\} \cup \{(c, c') \mid c, c' \in C, c \neq c'\} \cup\\
& \{(x, c) \mid x \text{ occurs in } c\} \cup \{(\overline{x}, c) \mid \neg x \text{ occurs in } c\} \cup\\
& \{(c, \varphi) \mid c \in C\} \cup \{(\varphi, \overline{\varphi})\}.
\end{aligned}$$

Figure 10 depicts F_φ for an exemplary CNF formula $\varphi(X)$.

Lemma 10. *Given a propositional formula $\varphi(X) = \bigwedge_{c\in C} c$ with each $c \in C$ a disjunction of literals from X, it holds that:*

1. *φ is satisfiable if and only if there exists $E \in \mathit{stb}(F_\varphi)$ such that $\overline{\varphi} \notin E$;*
2. *for each $E, E' \in \mathit{stb}(F_\varphi)$ such that $\overline{\varphi} \notin E$ and $\overline{\varphi} \in E'$ it holds that $|E| + 1 = |E'|$;*
3. *for each $E \in \mathit{stb}(F_\varphi)$ such that $\overline{\varphi} \notin E$ and each $E' \in \mathit{stb}(F_\varphi - (C \cup \{\overline{\varphi}\}))$ it holds that $|E| = |E'|$.*

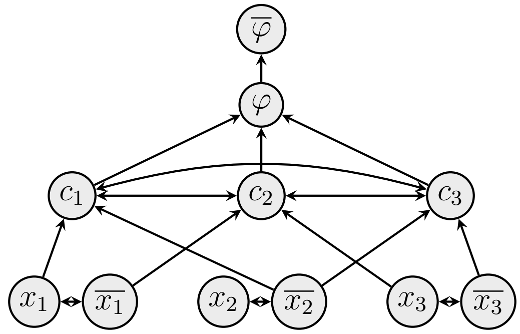

Figure 10: AF F_φ for $\varphi(X) = (x_1 \vee \neg x_2) \wedge (\neg x_1 \vee x_3) \wedge (\neg x_2 \vee \neg x_3)$.

Proof. We begin with the observation that every stable extension of F_φ or $F_\varphi - (C \cup \{\overline{\varphi}\})$ contains $S \cup (\overline{X \setminus S})$ for some $S \subseteq X$, since each argument $x \in X$ is in symmetric conflict with $\overline{x}$ and neither receives any further attacks.

1. ($\Rightarrow$): Assume φ is satisfiable, hence there exists $S \subseteq X$ such that for each $c \in C$, $S \models c$. Therefore, by construction of F_φ, $S \cup (\overline{X \setminus S})$ attacks all $c \in C$. Thus $S \cup (\overline{X \setminus S}) \cup \{\varphi\} \in \mathit{stb}(F_\varphi)$. ($\Leftarrow$): Let $E \in \mathit{stb}(F_\varphi)$ with $\overline{\varphi} \notin E$. Moreover let $S \subseteq X$ for which $S \cup (\overline{X \setminus S}) \subseteq E$ (recall from before that such an S must exist). Since φ is the only attacker of $\overline{\varphi}$ it follows that $\varphi \in E$ and further $c \notin E$ for all $c \in C$. Therefore $S \cup (\overline{X \setminus S})$ must attack each $c \in C$, meaning by construction of F_φ that $S \models c$ for each $c \in C$, hence $S \models \varphi$; that is, φ is satisfiable.

2. From the ($\Leftarrow$)-direction of (1) we get that each $E \in \mathit{stb}(F_\varphi)$ with $\overline{\varphi} \notin E$ has $|E| = |X| + 1$. For an arbitrary $E' \in \mathit{stb}(F_\varphi)$ with $\overline{\varphi} \in E'$ it must hold that $\varphi \notin E'$, hence for at least one $c \in C$ we must have $c \in E'$. Since, as we know, $S \cup (\overline{X \setminus S}) \subseteq E'$ for some $S \subseteq X$, and by C forming a clique, $c \in E$ for at most one $c \in C$, it follows that $|E'| = |X| + 2$, that is $|E| + 1 = |E'|$.

3. Obviously, $|E'| = |X| + 1$ for each $E' \in \mathit{stb}(F_\varphi - (C \cup \{\overline{\varphi}\}))$. Hence, from the observation in (2), the result follows.

This concludes the proof. □

Given these observations we can show the exact complexity of Dalal's operator for revision under stable semantics.

Theorem 7. *Given AFs $F, G \in AF_{\mathfrak{A}}$ and $E \subseteq \mathfrak{A}$, deciding whether $E \in \mathit{stb}(F *^D_{stb} G)$ is Θ_2^{P}-complete.*

Proof. For membership in Θ_2^{P} we sketch an algorithm that decides $E \in \mathit{stb}(F *^D_{stb} G)$ in polynomial time with $\mathcal{O}(\log m)$ calls to an NP oracle, where $m = |A_F| + |A_G|$. First we check whether $E \in \mathit{stb}(G)$ (in P); if no, then we return with a negative answer. Then the minimal distance $z = \min\{d_{stb}(T, F) \mid T \in \mathit{stb}(G)\}$ is determined. It holds that $z \leq m$, since $S \subseteq A_F$ (resp.

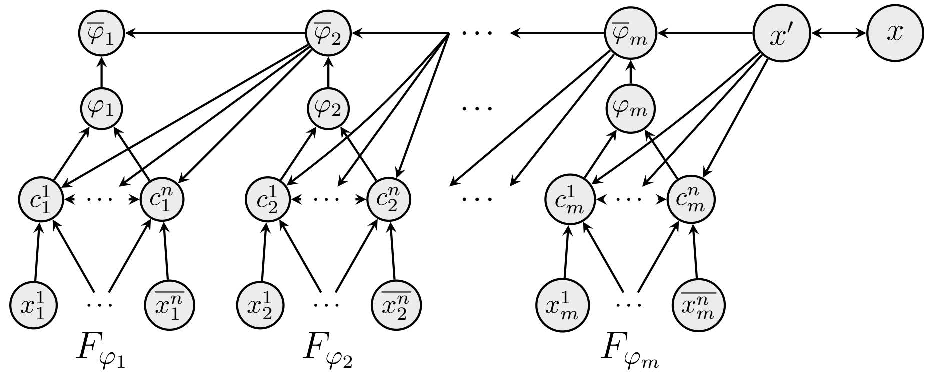

Figure 11: Illustration of the AF obtained from the reduction in the proof of Theorem 7.

$T \subseteq A_G$) for each $S \in stb(F)$ (resp. $T \in stb(G)$). Now z can be computed by binary search with $\mathcal{O}(\log m)$ calls to the following NP procedure: guess $S \subseteq A_F$, $T \subseteq A_G$ and check whether $S \in stb(F)$, $T \in stb(G)$ and $d_H(S,T) < z$ (checking this is in P). Having obtained z, we finally call another NP oracle to check whether there is an $S \in stb(F)$ such that $d_H(S,E) = z$. If such an S does exist, then $E \in stb(F *_{stb}^{D} G)$, otherwise not.

To show Θ_2^{P} hardness we give a polynomial-time reduction from the following problem (recall that a 1-existential QBF being false is equivalent to a propositional formula being unsatisfiable):

GIVEN: propositional formulas $\varphi_1(X_1), \dots, \varphi_m(X_m)$ such that
φ_i unsatisfiable implies φ_{i+1} unsatisfiable, for $1 \leq i < m$,
DECIDE: whether $k = \max\{1 \leq i \leq m \mid \varphi_i \text{ is satisfiable}\}$ is odd.

Without loss of generality we can assume that: (i) $X_i \cap X_j = \emptyset$ for all $1 \leq i, j \leq m, i \neq j$, (ii) $n = |X_i| = |X_j|$ for all $1 \leq i, j \leq m$, (iii) each φ_i is in CNF with C_i denoting the set of clauses of φ_i, and (iv) m is odd. Now given an instance of this problem, define $F = \bigcup_{1 \leq i \leq m} F_{\varphi_i} \cup F_i$ where F_{φ_i} is given by Definition 9 and:

$$F_i = \big(\{\overline{\varphi}_i, \overline{\varphi}_{i+1}\} \cup C_i, \{(\overline{\varphi}_{i+1}, \overline{\varphi}_i)\} \cup \{(\overline{\varphi}_{i+1}, c) \mid c \in C_i\}\big) \qquad 1 \leq i < m$$
$$F_m = (\{\overline{\varphi}_m, x, x'\} \cup C_m, \{(x, x'), (x', x), (x', \overline{\varphi}_m)\} \cup \{(x', c) \mid c \in C_m\}).$$

Intuitively, F contains the frameworks F_{φ_i} constructed according to Definition 9 together with "connecting frameworks" F_i which make $\overline{\varphi}_{i+1}$ attack $\overline{\varphi}_i$ and all clause-arguments C_i. F_m can be seen as the "starting framework". A schematic illustration of F can be seen in Figure 11. Moreover, we define $G = (\{x, x'\}, \{(x, x'), (x', x)\})$ and $E = \{x\}$.

Due to the splitting Property [4], the stable extensions of F are composed of the union of stable extensions of its components, where the computation of $stb(F_{\varphi_i})$ has to take into account $stb(F_{\varphi_{i+1}})$. That is, $stb(F) = \{\{\alpha\} \cup \bigcup_{1 \leq i \leq m} E_i \mid \alpha \in \{x, x'\}, E_i \in stb(F'_{\varphi_i})\}$ where

- $F'_{\varphi_m} = F_{\varphi_m}$ if $\alpha = x$ and $F'_{\varphi_m} = F_{\varphi_m} - (C_m \cup \{\overline{\varphi}_m\})$ if $\alpha = x'$, and
- $F'_{\varphi_i} = F_{\varphi_i}$ if $\overline{\varphi}_{i+1} \notin E_{i+1}$ and $F'_{\varphi_i} = F_{\varphi_i} - (C_i \cup \{\overline{\varphi}_i\})$ if $\overline{\varphi}_{i+1} \in E_{i+1}$ for $1 \leq i < m$.

Recall that k is the highest index such that φ_k is satisfiable. Consider an i with $k < i \leq m$. If $F'_{\varphi_i} = F_{\varphi_i}$ then we know, by Lemma 10.1 and φ_i being unsatisfiable, that $\overline{\varphi}_i \in E_i$, hence

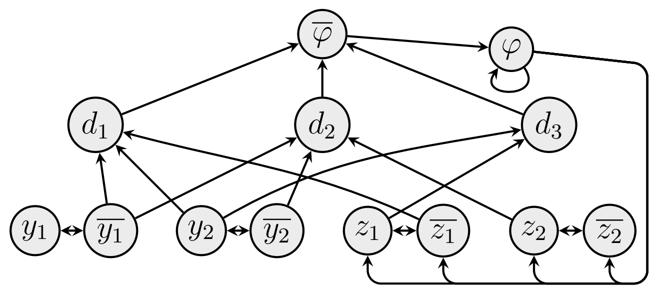

Figure 12: AF F_Φ for the QBF $\Phi = \exists y_1, y_2 \forall z_1, z_2 : (y_1 \wedge \neg y_2 \wedge z_1) \vee (y_1 \wedge y_2 \wedge \neg z_2) \vee (\neg y_2 \wedge \neg z_1)$.

$F'_{\varphi_{i-1}} = F_{\varphi_{i-1}} - (C_{i-1} \cup \{\overline{\varphi}_{i-1}\})$. On the other hand if $F'_{\varphi_i} = F_{\varphi_i} - (C_i \cup \{\overline{\varphi}_i\})$ then obviously $\overline{\varphi}_i \notin E_i$, hence $F'_{\varphi_{i-1}} = F_{\varphi_{i-1}}$. Now consider an i with $1 \leq i \leq k$. Again from Lemma 10.1, we get that there is some $E \in$ *stb*(F_{φ_i}) with $\overline{\varphi}_i \notin E$. Therefore, by Lemma 10.2 and 10.3, for $\alpha \in \{x, x'\}$ the extension $S^*_\alpha = \{\alpha\} \cup \bigcup_{1 \leq i \leq m} E_i$ with $\overline{\varphi}_i \notin E_i$ for $1 \leq i \leq k$ is the one with the minimal distance to $\{\alpha\}$ among all elements of *stb*(F) (recall the assumption that $|X_i| = |X_j|$ for all $1 \leq i, j \leq m$). Now if k is odd, we get, by the assumption that m is odd, that $m - k$ is even. Hence $d_H(S^*_x, \{x\}) = d_H(S^*_{x'}, \{x'\})$ and furthermore *stb*$(F *^D_{stb} G) = \{\{x\}, \{x'\}\}$, that is $E \in$ *stb*$(F *^D_{stb} G)$. If, on the other hand, k is even, then $m - k$ is odd and, by Lemma 10.2 and 10.3, $d_H(S^*_x, \{x\}) = d_H(S^*_{x'}, \{x'\}) + 1$, hence $E \notin$ *stb*$(F *^D_{stb} G) = \{\{x'\}\}$.

□

Now we turn to preferred semantics, where we will make use of the following construction.

Definition 10. Given a 2-existential QBF $\Phi = \exists Y \forall Z \varphi(Y, Z)$ where φ is a DNF $\bigvee_{d \in D} d$ with each d a conjunction of literals from $X = Y \cup Z$, we define $F_\Phi = (A_\Phi, R_\Phi)$ as:

$$\begin{aligned}
A_\Phi =& X \cup \overline{X} \cup D \cup \{\varphi, \overline{\varphi}\},\\
R_\Phi =& \{(x, \overline{x}), (\overline{x}, x) \mid x \in X\} \cup\\
& \{(\overline{x}, d) \mid x \text{ occurs in } d\} \cup \{(x, d) \mid \neg x \text{ occurs in } d\} \cup\\
& \{(d, \overline{\varphi}) \mid d \in D\} \cup \{(\overline{\varphi}, \varphi), (\varphi, \varphi)\} \cup \{(\varphi, z) \mid z \in Z\}.
\end{aligned}$$

The construction is illustrated on an exemplary 2-existential QBF Φ in Figure 12. We show two technical lemmata before giving the actual complexity result.

Lemma 11. *Let $\Phi = \exists Y \forall Z \varphi(Y, Z)$ where φ is a DNF $\bigvee_{d \in D} d$. For each $d \in D$, $S \subseteq Y$ and $T \subseteq Z$ it holds that:*

- *$S \cup T \models d$ if and only if d is defended by $S \cup (\overline{Y \setminus S}) \cup T \cup (\overline{Z \setminus T})$;*
- *$S \cup T \not\models d$ if and only if d is attacked by $S \cup (\overline{Y \setminus S}) \cup T \cup (\overline{Z \setminus T})$.*

Proof. If $S \cup T \models d$, then the set of arguments attacking d is, according to Definition 10, contained in $\overline{S} \cup (Y \setminus S) \cup \overline{T} \cup (Z \setminus T)$. Therefore, it is not attacked and even defended by $S \cup (\overline{Y \setminus S}) \cup T \cup (\overline{Z \setminus T})$.

If $S \cup T \not\models d$, then there is some argument attacking d which is not contained in $\overline{S} \cup (Y \setminus S) \cup \overline{T} \cup (Z \setminus T)$. Therefore, it is attacked and, consequently, not defended by $S \cup (\overline{Y \setminus S}) \cup T \cup (\overline{Z \setminus T})$. □

Lemma 12. *Consider the* 2*-existential QBF* $\Phi = \exists Y \forall Z \varphi(Y, Z)$ *where* φ *is a DNF* $\bigvee_{d \in D} d$*. It holds that:*

1. Φ *is true if and only if there exists* $E \in prf(F_\Phi)$ *such that* $\overline{\varphi} \notin E$*;*
2. *for each* $E \in prf(F_\Phi)$ *it holds that (a)* $|E| = |Y| + |Z| + 1$ *if* $\overline{\varphi} \in E$ *and (b)* $|E| = |Y|$ *if* $\overline{\varphi} \notin E$*;*
3. *for each* $E \in prf(F_\Phi - \{\overline{\varphi}\})$ *is holds that* $|E| = |Y|$*.*

Proof. 1. $\Rightarrow$: Assume Φ is true. That is, there is some $S \subseteq Y$ such that for all $T \subseteq Z$ it holds that $\varphi(S, T)$ is true. We show that $E = S \cup (\overline{Y \setminus S}) \in prf(F_\Phi)$. First, E is easily checked to be admissible. Towards a contradiction, assume there is some $E' \in adm(F_\Phi)$ with $E' \supset E$. Further assume there is some $d \in D$ included in $E' \setminus E$. Due to the non-triviality of d there is at least one $z \in Z \cup \overline{Z}$ attacking d and, consequently, it must hold that $\overline{z} \in E'$. Then, due to φ attacking all $Z \cup \overline{Z}$, $\overline{\varphi} \in E'$, we get a contradiction to conflict-freeness of E' since also $d \in D$. Knowing that $d \notin E'$ for all $d \in D$, assume that $\overline{\varphi} \in E'$. To this end $\overline{\varphi}$ has to be defended by E' from each $d \in D$. This means that there must be some $T \subseteq Z$ such that $T \cup (\overline{Z \setminus T}) \subseteq E'$ and each $d \in D$ is attacked by $S \cup (\overline{Y \setminus S}) \cup T \cup (\overline{Z \setminus T})$. But then, by Lemma 11, $S \cup T \not\models d$ for each $d \in D$, a contradiction to $\varphi(S, T)$ being true.

   $\Leftarrow$: We show the contrapositive, that if Φ is false then all $E \in prf(F_\Phi)$ have $\overline{\varphi} \in E$. Observe that for any $S \subseteq Y$, $S \cup (\overline{Y \setminus S})$ is admissible in F_Φ, hence $S \cup (\overline{Y \setminus S})$ is contained in some preferred extension. Moreover, each preferred extension must contain $S \cup (\overline{Y \setminus S})$ for some $S \subseteq Y$. Consider an arbitrary $S \subseteq Y$. As, by assumption, Φ is false, there must be some $T \subseteq Z$ such that $\varphi(S, T)$ is false. Hence for every $d \in D$ it must hold that $S \cup T \not\models d$ and consequently, by Lemma 11, d is attacked by $X_S = S \cup (\overline{Y \setminus S}) \cup T \cup (\overline{Z \setminus T})$. Hence $X_S \cup \{\overline{\varphi}\}$ is admissible and, by attacking all other arguments, also preferred in F_Φ. Now assume there is an $E' \in prf(F_\Phi)$ with $S \subseteq E'$ and $\overline{\varphi} \notin E'$. By the latter no argument among $Z \cup \overline{Z}$ can be in E' as it cannot be defended from φ. Hence, to be $\subseteq$-incomparable to all the preferred extensions which do include $\overline{\varphi}$, E' must include some $d \in D$. But also this in not possible as by assumption there must be some $T \subseteq Z$ making $S \cup T \not\models d$, meaning, by Lemma 11, that d is attacked by $S \cup (\overline{Y \setminus S}) \cup T \cup (\overline{Z \setminus T})$. If it is attacked by $S \cup (\overline{Y \setminus S})$ then E' is not conflict-free; if it is attacked by $T \cup (\overline{Z \setminus T})$ then E' is not admissible. We conclude that all $E \in prf(F_\Phi)$ have $\overline{\varphi} \in E$.

2. Consider some $E \in prf(F_\Phi)$. (a) If $\overline{\varphi} \in E$ then $d \notin E$ for all $d \in D$, hence a maximal conflict-free selection of arguments among $Y \cup \overline{Y} \cup Z \cup \overline{Z}$ must be included in E, therefore $S \cup (\overline{Y \setminus S}) \cup T \cup (\overline{Z \setminus T}) \subseteq E$ for some $S \subseteq Y$ and $T \subseteq Z$. Hence $|E| = |Y| + |Z| + 1$. (b) If $\overline{\varphi} \notin E$ then no argument among $Z \cup \overline{Z}$ can be defended. Moreover, as φ does not contain monoms which are trivial for Y, it follows by Lemma 11 that no $d \in D$ can be defended. On the other hand, E must include a maximal conflict-free selection of arguments among $Y \cup \overline{Y}$, hence $|E| = |Y|$.

3. Let $F'_\Phi = F_\Phi - \{\overline{\varphi}\}$ and observe that the self-attacking argument φ is unattacked in F'_Φ. Hence none of the arguments $Z \cup \overline{Z}$ can be defended. Moreover, as φ does not contain

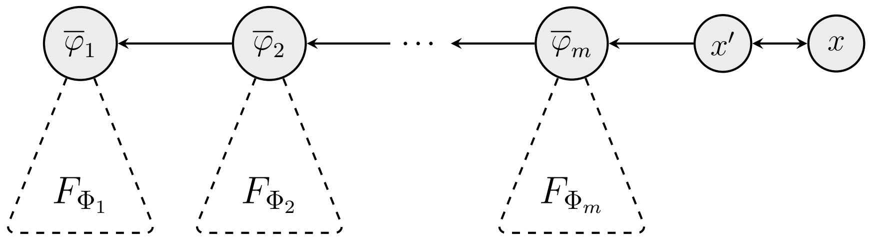

Figure 13: Illustration of the AF obtained from the reduction in the proof of Theorem 8.

monoms which are trivial for Y, each argument $d \in D$ is attacked by $Z \cup \overline{Z}$ and can therefore also not be defended. It follows that the preferred extensions of F'_Φ are given by $S \cup (\overline{Y \setminus S})$ for each $S \subseteq Y$, each containing $|Y|$ arguments.

□

Theorem 8. *Given AFs $F, G \in AF_\mathfrak{A}$ and $E \subseteq \mathfrak{A}$, deciding whether $E \in prf(F *^D_{prf} G)$ is Θ^P_3-complete.*

Proof. To show membership in Θ^P_3 we sketch an algorithm that decides $E \in prf(F *^D_{prf} G)$ in polynomial time with $\mathcal{O}(\log m)$ calls to a Σ^P_2 oracle, where $m = |A_F| + |A_G|$. First, we check whether $E \in prf(G)$ (in coNP); if no we return with a negative answer. Second, the minimal distance $z = \min\{d_{prf}(T, F) \mid T \in prf(G)\}$ is determined. Since $S \subseteq A_F$ (resp. $T \subseteq A_G$) for each $S \in prf(F)$ (resp. $T \in prf(G)$), it holds that $d \leq m$. Therefore it can be computed by binary search with $\mathcal{O}(\log m)$ oracle calls to the following Σ^P_2 procedure: Guess $S \subseteq A_F$, $T \subseteq A_G$ and check (in coNP) whether $S \in prf(F)$, $T \in prf(G)$ and $d_H(S, T) < z$. Having obtained z, we finally call the oracle once again to check whether there is an $S \in prf(F)$ with $d_H(S, E) = z$. If such an S does exist then $E \in prf(F *^D_{prf} G)$, otherwise not.

To show Θ^P_3-hardness we give a polynomial-time reduction from the following problem: Given 2-existential QBFs $\Phi_1, \ldots, \Phi_m$ such that Φ_i being false implies Φ_{i+1} being false for $1 \leq i < m$, decide whether $k = \max\{1 \leq i \leq m \mid \Phi_i \text{ is true}\}$ is odd. We use the following notation to identify the elements of QBFs: $\Phi_i = \exists Y_i \forall Z_i \varphi_i$. W.l.o.g. we can assume that (i) the variables of the QBFs are pairwise distinct, (ii) $|Y_i| = |Y_j|$ and $|Z_i| = |Z_j|$ for all $1 \leq i, j \leq m$, and (iii) m is odd. Due to (ii) we will use $|Y|$ to denote $|Y_i|$ and $|Z|$ to denote $|Z_i|$ for any i. Now for each $\Phi_i = \exists Y_i \forall Z_i \varphi_i$, let F_{Φ_i} be as given in Definition 10. We define $F = \bigcup_{1 \leq i \leq m} F_{\Phi_i} \cup F_i$ where:

$$F_i = \big(\{\overline{\varphi}_i, \overline{\varphi}_{i+1}\}, \{(\overline{\varphi}_{i+1}, \overline{\varphi}_i)\}\big) \qquad 1 \leq i < m$$
$$F_m = (\{\overline{\varphi}_m, x, x'\}, \{(x, x'), (x', x), (x', \overline{\varphi}_m)\})\,.$$

Figure 13 depicts a schematic example of F. The subframeworks F_i can be regarded as "connecting frameworks", adding just an attack from $\overline{\varphi}_{i+1}$ to $\overline{\varphi}_i$. F_m is the "starting framework". Moreover, we define $G = (\{x, x'\}, \{(x, x'), (x', x)\})$ and $E = \{x\}$. We show that $E \in prf(F *^D_{prf} G)$ if and only if k is odd.

Due to the splitting property of preferred semantics [4], the preferred extensions of F are composed as $prf(F) = \{\{\alpha\} \cup \bigcup_{1 \leq i \leq m} E_i \mid \alpha \in \{x, x'\}, E_i \in prf(F'_{\Phi_i})\}$, where:

- $F'_{\Phi_m} = F_{\Phi_m}$ if $\alpha = x$ and $F'_{\Phi_m} = (F_{\Phi_m} - \{\overline{\varphi}_m\})$ if $\alpha = x'$, and
- $F'_{\Phi_i} = F_{\Phi_i}$ if $\overline{\varphi}_{i+1} \notin E_{i+1}$ and $F'_{\Phi_i} = F_{\Phi-1} - \{\overline{\varphi}_i\})$ if $\overline{\varphi}_{i+1} \in E_{i+1}$ for $1 \leq i < m$.

Recall that k is the highest index such that Φ_k is true. Due to Lemma 12 it holds that:

- $1 \leq i \leq k$: we have either $|E_i| = |Y|$ or $|E_i| = |Y| + |Z| + 1$;
- $k < i \leq m$: if $\alpha = x$ we have $|E_i| = |Y| + |Z| + 1$ for $i \in \{m, m-2, \dots\}$ and $|E_i| = |Y|$ for $i \in \{m-1, m-3, \dots\}$; otherwise we have $|E_i| = |Y|$ for $i \in \{m, m-2, \dots\}$ and $|E_i| = |Y| + |Z| + 1$ for $i \in \{m-1, m-3, \dots\}$.

Moreover, we get from Lemma 12 that each F_{Φ_i} with $1 \leq i \leq k$ has an extension $E^*_i \in \mathit{prf}(F_{\Phi_i})$ with $\overline{\varphi}_i \notin E^*_i$, hence $|E^*_i| = |Y|$. Let $S^*_\alpha \in \mathit{prf}(F)$ be now such that $E_i = E^*_i$ for all $1 \leq i \leq k$. By the observations above and assumption (ii), S^*_α has minimal distance to $\{\alpha\}$ among all preferred extensions containing α, for $\alpha \in \{x, x'\}$.

If k is odd, we get, by the assumption that m is odd, that $m - k$ is even, hence $d_H(S^*_\alpha, \{\alpha\}) = m|Y| + \frac{m-k}{2}(|Z| + 1) + 1$ for both $\alpha \in \{x, x'\}$. Therefore $\mathit{prf}(F *^D_{\mathit{prf}} G) = \{\{x\}, \{x'\}\}$, i.e. $E \in \mathit{prf}(F *^D_{\mathit{prf}} G)$.

If k is even, then $m - k$ is odd. We get $d_H(S^*_{x'}, \{x'\}) = m|Y| + \lfloor\frac{m-k}{2}\rfloor(|Z| + 1) + 1 < m|Y| + \lceil\frac{m-k}{2}\rceil(|Z| + 1) + 1 = d_H(S^*_x, \{x\})$, hence $E \notin \mathit{prf}(F *^D_{\mathit{prf}} G) = \{\{x'\}\}$.

□

As elaborately discussed in Section 3.2, Dalal's operator cannot be directly applied to revision of AFs by propositional formulas, as the rankings obtained from Hamming distance do not meet the requirements for inducing rational operators. Therefore we consider here the refinement of Dalal's operator $\star^{D,\mathrm{ir}}_\sigma$, as introduced in Definition 7. We begin by showing that hardness carries over from the operator $\star^D_\sigma$ for revision by AFs.

Theorem 9. *Given an AF $F \in AF_{\mathfrak{A}}$, $\varphi \in \mathsf{P}_{\mathfrak{A}}$ and $E \subseteq \mathfrak{A}$, then:*

- *deciding whether $E \in stb(F \star^{D,\mathrm{ir}}_{stb} \varphi)$ is Θ^{P}_2-hard;*
- *deciding whether $E \in prf(F \star^{D,\mathrm{ir}}_{prf} \varphi)$ is Θ^{P}_3-hard.*

Proof. Let $\sigma \in \{\mathit{stb}, \mathit{prf}\}$. Further, let $G = (\{x, x'\}, \{(x, x'), (x', x)\})$ and $\phi(\sigma(G))$ be the formula having exactly $\sigma(G)$ as its models. We will give a polynomial time reduction from the problem, given $F \in AF_{\mathfrak{A}}$ and $E \subseteq \mathfrak{A}$, whether $E \in \sigma(F *^D_\sigma G)$. Inspecting the hardness proofs of Theorems 7 and 8, we see that this problem is Θ^{P}_2-hard for $\sigma = \mathit{stb}$ and Θ^{P}_3-hard for $\sigma = \mathit{prf}$ even for this fixed G.[7] Hence the reduction will give the desired result.

[7] For the sake of interest, we give the reduction of an arbitrary, but fixed, AF G.

Consider some $F \in AF_{\mathfrak{A}}$ and $E \subseteq \mathfrak{A}$. W.l.o.g. assume that $n = |E|$ is even and that the elements of E are the alphabetically minimal arguments. We define:

$$F' = F \cup (\{y_1, \dots, y_{\frac{n}{2}}\}, \emptyset), \text{ and}$$
$$\varphi = \phi(\sigma(G)) \wedge \bigwedge_{1 \le i \le n} \left(\neg \left(\bigwedge_{a \in E} a \wedge \bigwedge_{a' \in (A_G \setminus E)} \neg a' \right) \leftrightarrow y_i \right),$$

with $\{y_1, \dots, y_n\}$ being newly introduced arguments. We show that $E \in \sigma(F *^D_\sigma G)$ if and only if $E \in \sigma(F' \star^{D,\text{ir}}_\sigma \varphi)$. First recall that $[\phi(\sigma(G))] = \sigma(G)$. Now let $S \in \sigma(G)$. The second part of φ then ensures that if $S = E$ then $S \in [\varphi]$ and $S \cup Y \notin [\varphi]$ for any $Y \subseteq \{y_1, \dots, y_n\}$ ($Y \neq \emptyset$), and if $S \neq E$ then $S \cup \{y_1, \dots, y_n\} \in [\varphi]$ and $S \cup Y \notin [\varphi]$ for any $Y \subset \{y_1, \dots, y_n\}$. Therefore we derive the following:

- denoting $S' = S \cup \{y_1, \dots, y_n\}$ for every $S \in (\sigma(G) \setminus \{E\})$ and denoting $E' = E$, it holds that $[\varphi] = \{S' \mid S \in \sigma(G)\}$;
- denoting $T' = T \cup \{y_1, \dots, y_{\frac{n}{2}}\}$ for every $T \in \sigma(F)$, it holds that $\sigma(F') = \{T' \mid T \in \sigma(F)\}$.

Therefore, it holds for every $S \in [\varphi]$ that $d_\sigma(S', F') = d_\sigma(S, F) + \frac{n}{2}$ (note the initial assumption that n is even), that is $S_1 \preceq^D_F S_2$ if and only if $S'_1 \preceq^D_{F'} S'_2$.

Now first assume $E \notin \sigma(F *^D_\sigma G)$. That means there is some $S \in \sigma(G)$ such that $S \prec^D_F E$. But then, by our last observation, also $S' \prec^D_{F'} E'$ and, since the refinement only affects extensions on the same level w.r.t. $\preceq^D_{F'}$, also $S' \prec^{D,\text{ir}}_{F'} E'$. Therefore $E \notin \sigma(F \star^{D,\text{ir}}_\sigma \varphi)$.

On the other hand assume $E \in \sigma(F *^D_\sigma G)$. That means that for all $S \in \sigma(G)$ it holds that $E \preceq^D_F S$. For those $S \in \sigma(G)$ with $E \prec^D_F S$ we get $E' \prec^{D,\text{ir}}_{F'} S'$ as before. Consider an $S \in \sigma(G)$ with $E \approx^D_F S$. From $n = |E|$ we get that $|E| \le |S'|$. This together with the assumption that E contains the alphabetically smallest arguments, we get that $E \preceq^{2^{\mathfrak{A}} \setminus \sigma(F)} S'$ (cf. Definition 5). Therefore, by Definition 7, $E \preceq^{D,\text{ir}}_{F'} S'$. Since this holds for every $S \in \sigma(G)$ we conclude that $E \in \sigma(F' \star^{D,\text{ir}}_\sigma \varphi)$. □

As an upper bound for the complexity, we show membership in Δ^{P}_2 for revision with respect to stable semantics and membership in Δ^{P}_3 for preferred semantics.

Theorem 10. *Given an AF $F \in AF_{\mathfrak{A}}$, $\varphi \in \mathsf{P}_{\mathfrak{A}}$ and $E \subseteq \mathfrak{A}$, then:*

- *deciding whether $E \in stb(F \star^{D,\text{ir}}_{stb} \varphi)$ is in Δ^{P}_2;*
- *deciding whether $E \in prf(F \star^{D,\text{ir}}_{prf} \varphi)$ is in Δ^{P}_3.*

Proof. We show the result for *stb* and then argue how to adapt the proof to obtain the result for *prf*. To this end we sketch an algorithm that decides $E \in stb(F \star^{D,\text{ir}}_{stb} \varphi)$ in polynomial time with access to an NP oracle. Let $m = |A_F| + |var(\varphi)|$.[8] First we check whether $E \in [\varphi]$ (in P); if no we return with a negative answer. Then the minimal distance of a model of φ to F, that is

[8] We denote by $var(\varphi)$ the set of variables occurring in φ here.

$z = \min\{d_{stb}(T, F) \mid T \in [\varphi]\}$ is determined. As argued in the membership-part of the proof of Theorem 7, this requires at most $\mathcal{O}(\log m)$ calls to an NP procedure. Knowing the minimal distance z, we have to determine the minimum indexed level of extensions with distance z to F, where a model of φ is contained. There is one level for each e with $0 \leq e \leq m$ (the size of an extension) and each prefix $p = (a_1, \ldots, a_n)$ with $n < m$ and $a_i \in A_F \cup \mathit{var}(\varphi)$. Hence, each level can be identified by a pair (e, p) and the number of levels is at most exponential in m. We can now determine the minimum (e, p)-level containing a model of φ by binary search with $\mathcal{O}(\log 2^m) = \mathcal{O}(m)$ calls to the following NP procedure: guess $S \subseteq A_F$, $T \subseteq \mathit{var}(\varphi)$, and check whether $S \in \mathit{stb}(F)$, $T \in [\varphi]$, $d_H(S, T) = z$, $|T| \leq e$ and $\overrightarrow{T}^{\#} \preceq_{\text{lex}} p$. These checks can be computed in polynomial time. Having obtained z, e, and p, we finally check $|E| = e$, $\overrightarrow{E}^{\#} = p$, and, by another NP oracle call, whether there is an $S \in \mathit{stb}(F)$ such that $d_H(S, E) = z$; if these checks turn out positive, $E \in \mathit{stb}(F \star^{D,\text{ir}}_{\mathit{stb}} \varphi)$, otherwise not.

The proof for Δ^{P}_3-membership of deciding whether $E \in \mathit{prf}(F \star^{D,\text{ir}}_{\mathit{prf}} \varphi)$ uses the same polynomial time procedure, now with access to a Σ^{P}_2 oracle. That is, every oracle call involving a check of containment in the stable extensions of an AF now has to check containment in the preferred extensions of the AF, which is not in P but in NP. Therefore whenever the procedure for *stb* calls an NP oracle, the procedure for *prf* has to make use of a Σ^{P}_2 oracle. □

We have to leave the exact complexity for the refined version of Dalal's operator for revision by formulas open, but Theorem 10 suggests that the indexed refinement of the ranking obtained from Hamming distance prevents us from determining the level of interest (which is the minimal one where models of the revision formula occur) with logarithmically many oracle calls. Therefore we tend to assume that the indexed refinement indeed leads to a computationally slightly more complex operator.

6 Related Work

As indicated in the introduction, there has been a substantial amount of research in the dynamics of argumentation frameworks, even though the problems investigated and approaches that have been developed to address them differ considerably. For instance, a number of studies look at simple modifications of AFs (e.g., adding/removing an argument/attack) and how they affect evaluation via different semantics [10, 12, 13, 17, 18, 41].

In the following we describe those studies more closely related to the revision of AFs as considered in this work, more or less in the order of publication. Most of these studies deal with revision of AFs in scenarios that are either more restrictive than our own, or otherwise approach the problem from a slightly different perspective. Also, it is worth noting that no general results on the complexity of revision of AFs have as yet been presented.

The focus of Baumann [5] is on whether one can modify an AF such that a certain subset of arguments is contained in some extension (w.r.t. a semantics of interest) and, if so, what the number of minimal modifications is. On the other hand, Kontarinis et al. [39] propose a strategy in terms of rewriting rules to compute the minimal number of modifications on the attack relation of an AF to enforce a desired acceptance status of an argument. Booth et al. [14] give an AGM-like

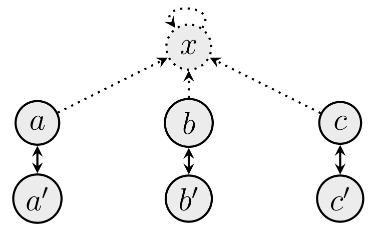

Figure 14: AFs discussed in Example 11.

characterization of revision of AFs when certain logical "constraints" expressing beliefs regarding the labellings of the AFs are "strengthened" to incorporate newly held beliefs. But the focus is on determining certain "fall back beliefs" when the newly held beliefs are inconsistent with those held previously. How to compute the fall back beliefs is developed in detail for the complete semantics.

Our starting point was the work on AF revision by Coste-Marquis et al. [19], where revision functions are defined following a two step process: first a counter-part to the concept 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, for example minimizing the changes in the attack relation of the input AF vs. minimizing the number of AFs generated. The main difference between the work by Coste-Marquis et al. [19] and our approach is that we consider the issue of revision of AFs as minimal change in the extensions of the original AF under the constraint that a single AF has to be produced. As already mentioned previously and showcased in Example 5, this constraint requires us to take into account the expressive peculiarities of the different semantics. Also, to realize the desired outcome by a single AF, the introduction of additional arguments is inevitable in certain cases.

Example 11. Consider the AF F depicted in Figure 14 (without the dotted part) and observe that $\sigma(F) = \{\{a,b,c\}, \{a,b,c'\}, \{a',b,c\}, \{a,b',c\}, \{a,b',c'\}, \{a',b,c'\}, \{a',b',c\}, \{a',b',c'\}\}$ for $\sigma \in \{\mathit{stb}, \mathit{prf}, \mathit{sem}, \mathit{stg}\}$. Now let $\star : AF_{\mathfrak{A}} \times \mathsf{P}_{\mathfrak{A}} \rightarrow AF_{\mathfrak{A}}$ be an arbitrary revision operator satisfying the rationality postulates. Then the revision of F by the formula $\neg(a' \wedge b' \wedge c')$ must, by postulate P$\star$2, result in an AF F' having $\sigma(F') = (\sigma(F) \setminus \{\{a',b',c'\}\})$. If we want F' to contain only arguments $\{a,b,c,a',b',c'\}$, it can be verified that all attacks which occur in F must also be present in F' and no other attack among the original arguments can be added. Hence we necessarily end up having $\sigma(F') = \sigma(F)$ when disallowing additional arguments. With the use of the new argument x, we can, however, realize $(\sigma(F) \setminus \{\{a',b',c'\}\})$ by the AF in Figure 14 including the dotted part (under $\sigma \in \{\mathit{stb}, \mathit{sem}, \mathit{stg}\}$; a different AF does the job for *prf*).

As the previous example shows, the choice of Coste-Marquis et al. [19] to let the revision result in a set of AFs is indeed substantiated if a fixed set of arguments is assumed. But if the result is to be instantiated as a single AF, as in our approach, then we have a good argument to allow the advancement of new arguments as part of the dynamic process. Recent work by Baumann et al. [7] looks at realizability in *compact* AFs, which could pave the way for revision where the result is a single AF and no additional arguments are allowed to come into play.

An issue related to revision of AFs is enforcement of arguments through minimal modifications to the attack relation. This is taken up in Doutre et al. [30], where enforcement is encoded in the framework of Dynamic Logic of Propositional Assignments (DL-PA). In the same direction, work by Nouioua and Würbel [49] provides an adaptation of the Removed-Set-Revision approach in propositional logic for the situation when adding attack relations and arguments to an AF results in the AF having no stable extension. Coste-Marquis et al. [20] translate the revision problem for AFs into propositional logic, thus enabling the use of classical AGM revision operators. However, revision formulas are defined in terms of the sceptical acceptance of arguments and the output of revision is still a set of AFs rather than a single AF. Coste-Marquis et al. [21] define operators to enforce that a set of arguments is a subset of an extension of an AF, Implementing this as pseudo-Boolean optimization problem leads to promising results.

Reasoning about the dynamics of AFs under different semantics is formalized in Baumann and Brewka [6] by means of a monotonic logic (Dung logic), based on the notion of *k-models*. This logic allows formulation of AGM-like postulates but, as with our results in Section 3 on revision by propositional formulas, realizability issues prevent standard distance-based revision operators from being applicable in this context. As a response, an alternative syntactic-based revision operator for the stable semantics is developed, and this operator returns a unique AF as output. For the other semantics, several other ideas for revision operators, selection functions from a set of possible AFs, are sketched.

Moguillansky [44] develops a theory of remainder sets for abstract argumentation, where revision is defined via expansion and contraction. A representation result for the basic postulates (success, consistency, inclusion, vacuity and core-retainment) is obtained, but this is nonetheless a more syntax-based approach to belief change in argumentation. Also, postulates in this approach are formulated with respect to the acceptance of an argument, rather than, as we interpret them, with respect to sets of extensions. An approach similar to ours, focused on postulates and representation results, and which also highlights the subtleties of instantiating the output as a single AF, looks at merging AFs in the presence of integrity constraints [27]. Merging differs from revision in that it attempts to integrate different sources of information, none of which is taken to have any priority.

Finally, we refer to recent work likewise inspired by the AGM theory of belief change, but which goes well beyond our work. In [25] (see also preceding work [11]) a very general theory to model dynamics of AFs is proposed. This theory makes it possible to express how an agent who has beliefs in the form of her own argumentation system can interact on a target argumentation system that may represent the state of knowledge at a given stage of a debate. Here AFs (and the dynamics of AFs) are encoded within the general, tailor made first order language YALLA. Further afield, both Moguillansky and Simari [45] and Snaith and Reed [53] present models of dynamics in *structured* (as opposed to abstract) argumentation. The former offers a model building on results by Moguillansky [44] (see also previous work from this group [46, 51, 47]), while the latter is a model for ASPIC+, one of the main existing formalisms for structured argumentation.

7 Conclusion and Outlook

In this work we have presented a generic solution to the problem of revision of AFs, which applies to many prominent I-maximal argumentation semantics. Compared to previous attempts in the literature, we aimed for revision operators which guarantee that the result is representable by a single AF. The key to obtaining our AGM-style representation theorems was the combination of recent advances from argumentation theory [33] and belief change [26]. We have considered two different approaches to revision of AFs. For revision by propositional formulas we have given a representation result which applies to arbitrary argumentation semantics in conjunction with compliant rankings on extensions. This compliance requirement has led us to develop general refinements of rankings, which in turn permitted us to obtain novel concrete operators for a wide range of semantics. For revision by AFs, on the other hand, the representation result has been restricted to proper I-maximal semantics, a class including standard semantics such as stable, preferred, semi-stable and stage. This result is nonetheless significant, as it allows any revision operator from the propositional setting to be applied in the AF context. Finally, we analysed the computational complexity of (a refinement of) Dalal's operator, where hardness goes up to Θ_3^{P} for revision under preferred semantics.

We identify several directions for future work. Fist, we want to extend our results in the revision-by-AF approach to semantics which are not proper I-maximal. Another interesting issue is the combination of semantics in the definition of revision operators for AFs, as done in the hybrid approach to revision of abstract dialectical frameworks [43]. Moreover, meaningful revision operators will have to 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. On a more general level, we want to analyse whether our insights can be extended to a broader theory of belief change within fragments. Finally we plan to 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.

References

[1] 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.

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

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

[4] Ringo Baumann. Splitting an argumentation framework. In James P. Delgrande and Wolfgang Faber, editors, *Proceedings of the 11th International Conference on Logic Programming*

and Nonmonotonic Reasoning (LPNMR 2011), volume 6645 of *Lecture Notes in Computer Science*, pages 40–53. Springer, 2011.

[5] Ringo Baumann. What does it take to enforce an argument? Minimal change in abstract argumentation. In Luc De Raedt, Christian Bessière, Didier Dubois, Patrick Doherty, Paolo Frasconi, Fredrik Heintz, and Peter J. F. Lucas, editors, *Proceedings of the 20th European Conference on Artificial Intelligence (ECAI 2012)*, volume 242 of *Frontiers in Artificial Intelligence and Applications*, pages 127–132. IOS Press, 2012.

[6] Ringo Baumann and Gerhard Brewka. AGM meets abstract argumentation: Expansion and revision for dung frameworks. In Qiang Yang and Michael Wooldridge, editors, *Proceedings of the 24th International Joint Conference on Artificial Intelligence (IJCAI 2015)*, pages 2734–2740. AAAI Press, 2015.

[7] Ringo Baumann, Wolfgang Dvořák, Thomas Linsbichler, Christof Spanring, Hannes Strass, and Stefan Woltran. On rejected arguments and implicit conflicts: The hidden power of argumentation semantics. *Artif. Intell.*, 241:244–284, 2016.

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

[9] Trevor J. M. Bench-Capon, Henry Prakken, and Giovanni Sartor. Argumentation in legal reasoning. In Guillermo Simari and Iyad Rahwan, editors, *Argumentation in Artificial Intelligence*, pages 363–382. Springer, 2009. ISBN 978-0-387-98196-3.

[10] 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 Salem Benferhat and John Grant, editors, *Proceedings of the 5th International Conference on Scalable Uncertainty Management (SUM 2011)*, volume 6929 of *Lecture Notes in Computer Science*, pages 275–288. Springer, 2011.

[11] Pierre Bisquert, Claudette Cayrol, Florence Dupin de Saint-Cyr, and Marie-Christine Lagasquie-Schiex. Enforcement in argumentation is a kind of update. In Weiru Liu, V. S. Subrahmanian, and Jef Wijsen, editors, *Proceedings of the 7th International Conference on Scalable Uncertainty Management (SUM 2013)*, volume 8078 of *Lecture Notes in Computer Science*, pages 30–43. Springer, 2013.

[12] Guido Boella, Souhila Kaci, and Leendert van der Torre. Dynamics in argumentation with single extensions: attack refinement and the grounded extension (extended version). In Peter McBurney, Iyad Rahwan, Simon Parsons, and Nicolas Maudet, editors, *Proceedings of the 6th International Workshop on Argumentation in Multi-Agent Systems (ArgMAS 2009), Revised Selected and Invited Papers*, volume 6057 of *Lecture Notes in Computer Science*, pages 150–159. Springer, 2009.

[13] Guido Boella, Souhila Kaci, and Leendert van der Torre. Dynamics in argumentation with single extensions: abstraction principles and the grounded extension. In Claudio Sossai and Gaetano Chemello, editors, *Procedings of the 10th European Conference on Symbolic and Quantitative Approaches to Reasoning with Uncertainty, (ECSQARU 2009)*, volume 5590 of *Lecture Notes in Computer Science*, pages 107–118. Springer, 2009.

[14] Richard Booth, Souhila Kaci, Tjitze Rienstra, and Leendert van der Torre. A logical theory about dynamics in abstract argumentation. In Weiru Liu, V. S. Subrahmanian, and Jef Wijsen, editors, *Proceedings of the 7th International Conference on Scalable Uncertainty Management (SUM 2013)*, volume 8078 of *Lecture Notes in Computer Science*, pages 148–161. Springer, 2013.

[15] Martin Caminada, Walter Alexandre Carnielli, and Paul E. Dunne. Semi-stable semantics. *J. Log. Comput.*, 22(5):1207–1254, 2012.

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

[17] Claudette Cayrol and Marie-Christine Lagasquie-Schiex. Change in abstract bipolar argumentation systems. In Christoph Beierle and Alex Dekhtyar, editors, *Proceedings of the 9th International Conference on Scalable Uncertainty Management (SUM 2015)*, volume 9310 of *Lecture Notes in Computer Science*, pages 314–329. Springer, 2015.

[18] 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.

[19] Sylvie Coste-Marquis, Sébastien Konieczny, Jean-Guy Mailly, and Pierre Marquis. On the revision of argumentation systems: minimal change of arguments statuses. In Chitta Baral, Giuseppe De Giacomo, and Thomas Eiter, editors, *Proceedings of the 14th International Conference on Principles of Knowledge Representation and Reasoning (KR 2014)*, pages 72–81. AAAI Press, 2014.

[20] Sylvie Coste-Marquis, Sébastien Konieczny, Jean-Guy Mailly, and Pierre Marquis. A translation-based approach for revision of argumentation frameworks. In Eduardo Fermé and João Leite, editors, *Proceedings of the 14th European Conference on Logics in Artificial Intelligence (JELIA 2014)*, volume 8761 of *Lecture Notes in Computer Science*, pages 397–411. Springer, 2014.

[21] Sylvie Coste-Marquis, Sébastien Konieczny, Jean-Guy Mailly, and Pierre Marquis. Extension enforcement in abstract argumentation as an optimization problem. In Qiang Yang and Michael Wooldridge, editors, *Proceedings of the 24th International Joint Conference on Artificial Intelligence (IJCAI 2015)*, pages 2876–2882. AAAI Press, 2015.

[22] Nadia Creignou, Reinhard Pichler, and Stefan Woltran. Do hard SAT-related reasoning tasks become easier in the Krom fragment? In Francesca Rossi, editor, *Proceedings of the 23rd International Joint Conference on Artificial Intelligence (IJCAI 2013)*, pages 824–831. IJCAI/AAAI, 2013.

[23] Mukesh Dalal. Investigations into a theory of knowledge base revision. In Howard E. Shrobe, Tom M. Mitchell, and Reid G. Smith, editors, *Proceedings of the 7th National Conference on Artificial Intelligence (AAAI 1988)*, pages 475–479. AAAI Press / The MIT Press, 1988.

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

[25] Florence Dupin de Saint-Cyr, Pierre Bisquert, Claudette Cayrol, and Marie-Christine Lagasquie-Schiex. Argumentation update in YALLA (Yet Another Logic Language for Argumentation). *Int. J. Approx. Reasoning*, 75:57–92, 2016.

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

[27] Jérôme Delobelle, Adrian Haret, Sébastien Konieczny, Jean-Guy Mailly, Julien Rossit, and Stefan Woltran. Merging of abstract argumentation frameworks. In Chitta Baral, James P. Delgrande, and Frank Wolter, editors, *Proceedings of the 15th International Conference on Principles of Knowledge Representation and Reasoning (KR 2016)*, pages 33–42. AAAI Press, 2016.

[28] Martin Diller, Adrian Haret, Thomas Linsbichler, Stefan Rümmele, and Stefan Woltran. An extension-based approach to belief revision in abstract argumentation. In Qiang Yang and Michael Wooldridge, editors, *Proceedings of the 24th International Joint Conference on Artificial Intelligence, (IJCAI 2015)*, pages 2926–2932. AAAI Press, 2015.

[29] Yannis Dimopoulos and Alberto Torres. Graph theoretical structures in logic programs and default theories. *Theor. Comput. Sci.*, 170(1-2):209–244, 1996.

[30] Sylvie Doutre, Andreas Herzig, and Laurent Perrussel. A dynamic logic framework for abstract argumentation. In Chitta Baral, Giuseppe De Giacomo, and Thomas Eiter, editors, *Proceedings of the 14th International Conference on Principles of Knowledge Representation and Reasoning (KR 2014)*, pages 62–71. AAAI Press, 2014.

[31] 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.

[32] Paul E. Dunne and Trevor J. M. Bench-Capon. Coherence in finite argument systems. *Artif. Intell.*, 141(1/2):187–203, 2002.

[33] Paul E. Dunne, Wolfgang Dvořák, Thomas Linsbichler, and Stefan Woltran. Characteristics of multiple viewpoints in abstract argumentation. *Artif. Intell.*, 228:153–178, 2015.

[34] Thomas Eiter and Georg Gottlob. On the complexity of propositional knowledge base revision, updates, and counterfactuals. *Artif. Intell.*, 57(2-3):227–270, 1992.

[35] Thomas Eiter and Georg Gottlob. The complexity of logic-based abduction. *J. ACM*, 42(1): 3–42, 1995.

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

[37] Anthony Hunter. Computational persuasion with applications in behaviour change. In Pietro Baroni, Thomas F. Gordon, Tatjana Scheffler, and Manfred Stede, editors, *Proceedings of the 6th International Conference on Computational Models of Argument (COMMA 2016)*, volume 287 of *Frontiers in Artificial Intelligence and Applications*, pages 5–18. IOS Press, 2016.

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

[39] 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 João Leite, Tran Cao Son, Paolo Torroni, Leon van der Torre, and Stefan Woltran, editors, *Proceedings of the 14th International Workshop on Computational Logic in Multi-Agent Systems (CLIMA XIV)*, volume 8143 of *Lecture Notes in Computer Science*, pages 51–68. Springer, 2013.

[40] Patrick Krümpelmann, Matthias Thimm, Marcelo A. Falappa, Alejandro J. García, Gabriele Kern-Isberner, and Guillermo R. Simari. Selective revision by deductive argumentation. In Sanjay Modgil, Nir Oren, and Francesca Toni, editors, *Proceedings of the 1st International Workshop on Theory and Applications of Formal Argumentation - (TAFA 2011), Revised Selected Papers*, volume 7132 of *Lecture Notes in Computer Science*, pages 147–162. Springer, 2012.

[41] Bei Shui Liao, Li Jin, and Robert C. Koons. Dynamics of argumentation systems: A division-based method. *Artif. Intell.*, 175(11):1790–1814, 2011.

[42] Paolo Liberatore and Marco Schaerf. Belief revision and update: Complexity of model checking. *J. Comput. Syst. Sci.*, 62(1):43–72, 2001.

[43] Thomas Linsbichler and Stefan Woltran. Revision of abstract dialectical frameworks: Preliminary report. In Sarah Gaggl, Juan Carlos Nieves, and Hannes Strass, editors, *Proceedings of the 1st international Workshop on Argumentation in Logic Programming and Non-Monotonic Reasoning (Arg-LPNMR 2016)*, pages 15–22, 2016.

[44] Martín O. Moguillansky. A study of argument acceptability dynamics through core and remainder sets. In Marc Gyssens and Guillermo Ricardo Simari, editors, *Proceedings of the*

9th International Symposium on Foundations of Information and Knowledge Systems (FoIKS 2016), volume 9616 of *Lecture Notes in Computer Science*, pages 3–23. Springer, 2016.

[45] Martín O. Moguillansky and Guillermo Ricardo Simari. A specialized set theoretic semantics for acceptability dynamics of arguments. In Pietro Baroni, Thomas F. Gordon, Tatjana Scheffler, and Manfred Stede, editors, *Proceedings of the 6th International Conference on Computational Models of Argument (COMMA 2016)*, volume 287 of *Frontiers in Artificial Intelligence and Applications*, pages 391–402. IOS Press, 2016.

[46] Martín O. Moguillansky, Nicolás D. Rotstein, Marcelo A. Falappa, Alejandro Javier García, and Guillermo Ricardo Simari. Argument theory change applied to defeasible logic programming. In Dieter Fox and Carla P. Gomes, editors, *Proceedings of the 23rd Conference on Artificial Intelligence (AAAI 2008)*, pages 132–137. AAAI Press, 2008.

[47] Martín O. Moguillansky, Renata Wassermann, and Marcelo A. Falappa. Inconsistent-tolerant base revision through argument theory change. *Logic Journal of the IGPL*, 20(1):154–186, 2012.

[48] Andreas Niskanen, Johannes Peter Wallner, and Matti Järvisalo. Synthesizing argumentation frameworks from examples. In Gal A. Kaminka, Maria Fox, Paolo Bouquet, Eyke Hüllermeier, Virginia Dignum, Frank Dignum, and Frank van Harmelen, editors, *Proceedings of the 22nd European Conference on Artificial Intelligence (ECAI 2016)*, volume 285 of *Frontiers in Artificial Intelligence and Applications*, pages 551–559. IOS Press, 2016.

[49] Farid Nouioua and Éric Würbel. Removed set-based revision of abstract argumentation frameworks. In *Proceedings of the 26th International Conference on Tools with Artificial Intelligence (ICTAI 2014)*, pages 784–791. IEEE Computer Society, 2014.

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

[51] Nicolás D. Rotstein, Martín O. Moguillansky, Marcelo A. Falappa, Alejandro Javier García, and Guillermo Ricardo Simari. Argument theory change: Revision upon warrant. In Philippe Besnard, Sylvie Doutre, and Anthony Hunter, editors, *Proceedings of the 2nd International Conference on Computational Models of Argument (COMMA 2008)*, volume 172 of *Frontiers in Artificial Intelligence and Applications*, pages 336–347. IOS Press, 2008.

[52] Chiaki Sakama. Counterfactual reasoning in argumentation frameworks. In Simon Parsons, Nir Oren, Chris Reed, and Federico Cerutti, editors, *Proceedings of the 5th International Conference on Computational Models of Argument (COMMA 2014)*, volume 266 of *Frontiers in Artificial Intelligence and Applications*, pages 385–396. IOS Press, 2014.

[53] Mark Snaith and Chris Reed. Argument revision. *J. Log. Comput.*, 2016. In Press. Available online at http://dx.doi.org/10.1093/logcom/exw028.

[54] 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.

[55] Larry J. Stockmeyer and Albert R. Meyer. Word problems requiring exponential time: Preliminary report. In Alfred V. Aho, Allan Borodin, Robert L. Constable, Robert W. Floyd, Michael A. Harrison, Richard M. Karp, and H. Raymond Strong, editors, *Proceedings of the 5th Annual ACM Symposium on Theory of Computing*, pages 1–9. ACM, 1973.

[56] Bart Verheij. Two approaches to dialectical argumentation: admissible sets and argumentation stages. In John-Jules C. Meyer and Linda C. van der Gaag, editors, *Proceedings of the 8th Dutch Conference on Artificial Intelligence (NAIC'96)*, pages 357–368, 1996.

[57] Klaus W. Wagner. Bounded query classes. *SIAM J. Comput.*, 19(5):833–846, 1990.