An Extension-Based Approach to Belief Revision in Abstract Argumentation

Martin Diller[a], Adrian Haret[a], Thomas Linsbichler[a], Stefan Rümmele[b], Stefan Woltran[a]

[a] *Institute of Information Systems, TU Wien, Vienna, Austria*
{diller,haret,linsbich,woltran}@dbai.tuwien.ac.at
[b] *UNSW Sydney and University of Sydney, Australia*
s.rummele@unsw.edu.au

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 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 analyze their computational complexity.

1. Introduction

Argumentation has emerged, over the last two decades, as a major research area in Artificial Intelligence (AI) [10, 58]. This is due not just to the intrinsic interest of the topic and to its recent applications (see [50] and [2] for surveys emphasizing applications of argumentation in areas such as legal reasoning, medicine, and e-governance) 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) [33], 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 evaluating the acceptance of arguments when taking into consideration the structure encoded in the graph. A common approach to this is finding subsets of arguments (called *extensions*) that can all be accepted together. As a result, the argumentation literature offers a wide range of criteria (called *semantics of AFs*) for establishing which arguments are jointly acceptable [5].

Our work fits into the growing number of studies on the *dynamics* of argumentation frameworks [7, 11, 12, 13, 15, 19, 32, 44, 45, 57, 60]. This line of research is motivated by the 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: it is to be expected that increased knowledge of facts settles certain issues, with the effect that arguments pertaining to them have to either become part of, or be excluded from any extension of our AF. Thus, such increased knowledge must be reflected in a new AF which manages 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 to find appropriate ways of formalizing the 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 through the lens of the semantic approach to propositional belief revision [42], where a knowledge base has a finite representation in a formal language, and this representation is used to encode a finite set of models. In our setup, AFs play the role of knowledge bases and their extensions under a certain semantics are the models. Thus, given a semantics σ, an AF F and a revision formula φ encoding desired changes in the status of some arguments, the task of a revision operator $\circ_\sigma$ is defined as follows: find an AF $F \circ_\sigma \varphi$ which manages to both satisfy φ and preserve as much useful information from F as possible. Example 1 illustrates the main steps in this process.

Example 1. Consider a propositional knowledge base $K = \{a \leftrightarrow \neg b, \neg c, d\}$, to undergo revision by $\varphi = c \wedge d$. A propositional revision operator $\circ$ would be expected to return a knowledge base $K \circ \varphi$ which implies φ: one can envision many ways to do this, but an approach based on minimizing information loss such as Dalal's operator (see [25] or Section 2) would pick the models of φ considered most plausible from the point of view of the models of K, and return the knowledge base $\{a \leftrightarrow \neg b, c, d\}$.

Consider, now, the AF F depicted in Figure 1, where some semantics σ has singled out the extensions $\{a, d\}$ and $\{b, d\}$ as jointly acceptable: we think of these sets as the models of F. Suppose, next, that in light of new information (in the form of a propositional formula or another AF), we learn that c and d must be accepted. If F is to undergo revision by a formula $\varphi = c \wedge d$, this is interpreted as asking for an AF $F \circ_\sigma \varphi$ whose extensions satisfy certain constraints, e.g., they are models of φ. A strategy of minimizing information loss such as the one mentioned above would return $\{\{a, c, d\}, \{b, c, d\}\}$ as a suitable set of candidates. In the final step, a function f_σ constructs an AF $F \circ_\sigma \varphi$ with precisely this set of extensions.

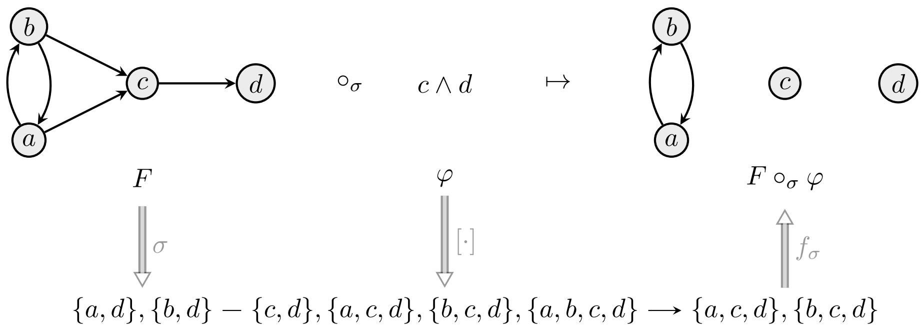

Figure 1: F undergoes revision by φ

In the paper we fill out this picture by formulating rationality constraints to guide the revision process, distinguishing between different forms which the new information can take (a formula or another AF), and making sure the resulting set of extensions can be represented by a single AF. The latter step turns out to be sensitive to the semantics used and poses non-trivial challenges. Remarkably, a *representation theorem* illuminates the problem, by showing that performing AF revision in accordance with some rationality postulates is equivalent to choosing among possible extensions of AFs, according to a particular type of rankings on extensions.

For the rationality constraints, we adapt a well-known core set of postulates from the literature on propositional revision [42]. In keeping with the different ways in which new information can be expressed, we study two types of revision operators. The first 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) to be induced 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 is similar to existing work [21], the latter assumes that new information stems from another agent's beliefs, and is in the form of an AF. This is more in line with work on Horn revision [28], where all involved formulas belong to some fragment of propositional logic. The two types of revision present interesting differences, particularly when considering the realizability of the result as an AF. 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*. Finally, we analyze the computational complexity of the main revision tasks.

The issue was first tackled in this manner by Coste-Marquis et al. [21],[1] with

[1]Other recent work in this direction includes [8, 12, 27, 51, 52] and is discussed Section 6.

the notable difference that there the revision result is defined to be *a set of AFs*. For instance, in Example 1 this would amount to finding an AF for each of the extensions in the set $\{\{a, c, d\}, \{b, c, d\}\}$ and offering this collection of AFs as output. This way of presenting the result circumvents an expressibility problem: as soon revision is thought of at the semantic level, one stumbles upon the issue that standard operators produce results which cannot be represented by a single AF. Thus, returning multiple AFs instead of just one is a way of avoiding the problem.

By contrast, we study AF revision operators producing a *single* AF as output, the motivation for which is twofold. First, such a restriction is more in line with standard propositional revision, where revising an input propositional theory by a formula produces a single propositional theory. Second, this is a natural condition to impose for the kind of applications where revision is most relevant. In the scenarios we have in mind new information is streamed in as it becomes available, which leads to continuous refinement of the original AF. Such repeated application of the revision function is possible only if one has an AF to work with at all times, and if this AF reflects the changes prompted by the sequence of inputs. Thus, characterizing revision operators which produce a single AF is the first essential step in understanding a dynamic process reminiscent of iterated propositional revision [26, 63].

At the same time, restricting the output of AF revision operators to single AFs poses significant challenges and is not just a special case of the approach where the result is a set of AFs. 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 the chosen semantics. It turns out that standard operators from the propositional belief change literature are not easily applicable in the new context and familiar representation results break down. The problem is exacerbated by the variety of semantics on offer and their expressive particularities.

Our main contributions can be summarized as follows.

- We obtain full *representation theorems* for the two types of revision mentioned. Notably, our results hold for a wide range of argumentation semantics including preferred, semi-stable, stage, and stable semantics. For other prominent semantics, such as the complete semantics, we show the impossibility of finding operators adhering to the full set of revision postulates.

- For the revision-by-formula approach, we give novel notions of *compliance* [28] to restrict the rankings (Section 3.2). 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 [36] is needed. For most of the standard semantics, the necessary results have been shown in [36]. It turns out that standard revision operators such as Dalal's operator [25] 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 3.3).

- 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 [28]) 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 certain type of rankings on extensions.

- Finally, we analyze the computational complexity of some specific revision operators when using stable and preferred semantics.[2] 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 [37, 47, 24], while it turns out that the intrinsically higher complexity of preferred semantics [35] 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 motivation, 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 [30]. 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 [5, 33], present recent results on signatures of semantics [36], and then recall the basic concepts of belief revision [1, 39, 42].

Argumentation. Real life instances of argumentation involve working with defeasible arguments in an environment of uncertainty and changing information,

[2] We consider these two semantics as representative and refer to [48] for complexity results for the semi-stable and stage semantics, which turn out to be similar to preferred semantics.

a: Employers move abroad.
b: Employers invest in technology.
c: More people join the workforce.
d: Unemployment rises.

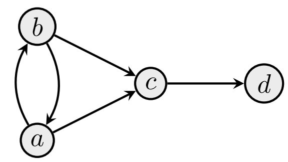

F

Figure 2: F models the interaction between arguments over the minimum wage.

and this calls for formal models that are sensitive to the dialectical nature of such reasoning [10]. The following example illustrates the use of abstract argumentation models towards this purpose.

Example 2. An agent tries to predict whether raising the minimum wage will benefit society. As described in [65], such complex issues are best navigated by taking into account multiple points of view (in a on-the-one-hand/on-the-other-hand style of thinking), and being responsive to new information. In the minimum wage case, an agent might reason as follows:

> One the one hand, it is possible that raising the minimum wage leads to employers hiring fewer workers, which would raise overall unemployment; on the other hand, effects on unemployment might not be that high, and it will definitely increase the labor supply; but then again, it might determine employers to move their business abroad, or to invest more in automation.

Such a chain of reasoning is difficult to capture through a series of logical entailments. Nonetheless, one can discern the main points on which the agent could be persuaded to change its mind. To model this argument in an abstract argumentation framework, we zoom out from the specific content of the issues raised and focus on what the agent thinks are the lines of tension between them. Concretely, we use four abstract arguments whose meanings and attacks are depicted in Figure 2. Drawing an attack from c to d indicates the perception of a conflict between the two arguments: the conviction that they cannot be jointly accepted. The question of which arguments can be jointly accepted given this argumentation framework will be taken up when we discuss the semantics of abstract argumentation.

Formally, 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 $AF_{\mathfrak{A}}$. For an AF $F = (B, S)$ we use A_F to refer to B and R_F to refer 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 $F \cup G = (A_F \cup A_G, R_F \cup R_G)$.

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 $\mathbb{S} \subseteq 2^{\mathfrak{A}}$. The intended meaning of a semantics is to provide criteria for selecting sets of jointly acceptable arguments. Prominent proposals are the stable, preferred, stage [66], semi-stable [17], and complete semantics, defined as in Table 1. We will not formally introduce unique-status semantics like the grounded, ideal [34], and eager [16] semantics. The fact that they have exactly one extension for each AF is what is important for the remainder of the paper.

$S \in stb(F)$, if $S \in cf(F)$ and $S_F^+ = A$	(stable)
$S \in prf(F)$, if $S \in adm(F)$ and $\nexists T \in adm(F)$ s.t. $T \supset S$	(preferred)
$S \in stg(F)$, if $S \in cf(F)$ and $\nexists T \in cf(F)$ with $T_F^+ \supset S_F^+$	(stage)
$S \in sem(F)$, if $S \in adm(F)$ and $\nexists T \in adm(F)$ s.t. $T_F^+ \supset S_F^+$	(semi-stable)
$S \in com(F)$, if $S \in adm(F)$ and $a \in S$ for all $a \in A$ defended by S	(complete)

Table 1: Definitions of the main semantics.

Example 3. For the AF F from Example 2 (see Figure 2), we can imagine the agent wanting to extract a maximal set of arguments that can be jointly accepted. This warrants the usage of the preferred semantics, which gives us $prf(F) = \{\{a, d\}, \{b, d\}\}$. When evaluated on F, all other semantics give the same results, except the complete semantics, which has the empty set as an additional extension (i.e., $com(F) = \{\{a, d\}, \{b, d\}, \emptyset\}$).

To do justice to the full range of semantics, we need a larger AF. To that end, consider the AF G:

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

depicted in Figure 3. It can be checked, by direct inspection, that there is no conflict-free set of arguments in G attacking all other arguments, hence $stb(G) = \emptyset$. The admissible sets of G are given by $adm(G) = \{\emptyset, \{a\}, \{b\}, \{b, d\}\}$, and hence $prf(G) = \{\{a\}, \{b, d\}\}$. Since $\{b\}$ defends argument d, we have $com(G) = adm(G) \setminus \{\{b\}\}$. By $\{a\}_G^+ = \{a, b\} \subset \{a, b, c, d, e\} = \{b, d\}_G^+$ we get that $\{b, d\}$ is the only semi-stable extension of G, i.e. $sem(G) = \{\{b, d\}\}$. Finally, it holds that $stg(G) = \{\{a, e\}, \{b, e\}, \{b, d\}\}$.

A set of extensions $\mathbb{S}$ can be *realized under a semantics* σ if there exists an AF $F \in AF_{\mathfrak{A}}$ such that $\sigma(F) = \mathbb{S}$. The signature Σ_σ of semantics σ is defined as $\Sigma_\sigma = \{\sigma(F) \mid F \in AF_{\mathfrak{A}}\}$, containing exactly those sets of extension which can be realized under σ. Exact characterizations of the signatures of the introduced semantics are known [36]. 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

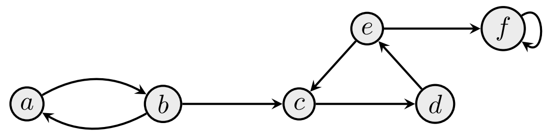

Figure 3: AF G discussed in Example 3.

and S_2 are $\subseteq$-*incomparable* if they are not $\subseteq$-comparable.[3] 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 following example illustrates these concepts.

Example 4. Consider the set $\mathbb{S} = \{\{a, b\}, \{a, c, e\}, \{b, d, e\}\}$. First, it is easy to check that elements of $\mathbb{S}$ are pairwise $\subseteq$-incomparable and, consequently, $\mathbb{S}$ is incomparable. On the other hand, $\mathbb{S}$ can be shown to be not tight: consider $\{a, b\} \in \mathbb{S}$ and $e \in \bigcup_{S \in \mathbb{S}} S$. It turns out that $\{a, b\} \cup \{e\} \notin \mathbb{S}$, but both $\{a, e\} \subseteq S'$ and $\{b, e\} \subseteq S''$ for some $S', S'' \in \mathbb{S}$. The modified set $\mathbb{S}' = \mathbb{S} \setminus \{\{a, b\}\}$, on the other hand, is indeed tight.

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, as shown in [36, Theorem 1]:

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

Regarding the other semantics, it suffices for our purposes to state that $\Sigma_{stg} \subset \Sigma_{sem} = \Sigma_{prf}$ [36, Theorem 2]. 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:

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 $\subseteq$-incomparable $S_1, S_2 \in 2^{\mathfrak{A}}$ it holds that $\{S_1, S_2\} \in \Sigma_\sigma$.

In words, an I-maximal [4] 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 (including singletons), is realizable under σ. The next observation follows from the characterizations of the signatures [36], and shows that the semantics we are interested in are all proper I-maximal.

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

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. Property (1) is known (e.g., from [4], where it is named I-maximality). Properties (2) and (3) appear in [36]: (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, a *realizing function* $f_\sigma : 2^{2^{\mathfrak{A}}} \to AF_{\mathfrak{A}}$ maps 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. Canonical realizing functions for the semantics we consider have been published [36]. Such functions may yield AFs with additional arguments to those contained in $\mathbb{S}$, though recent work on realizability in compact AFs [9] could pave the way for constructions of AFs without new arguments. The realizing function f_σ is not necessarily unique, but as a simplifying assumption we will assume throughout the paper that f_σ is fixed for every σ.

Belief revision. Revision occurs when, upon having access to new information, an agent changes some of its beliefs.

Example 5. Suppose the agent from Example 2 believes, at some point, that businesses are moving abroad (a) and that they are investing massively in new technology (b), its epistemic state described by the knowledge base $K = \{a \wedge b\}$. Now, if the agent hears from a trusted source that rumours about the exodus of companies are vastly exaggerated, it is forced to revise its beliefs by $\varphi = \neg a$. We expect that the agent gives up a and adds $\neg a$ to K. We also expect that the agent modifies its beliefs in the most economic way possible, such that it does not give up more than it strictly needs to.

Formally, we denote by $\mathsf{P}_{\mathfrak{A}}$ 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 finite 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 K needs to adapt to, and $K \circ \varphi$ is the revision outcome. The revision outcome is constrained by rationality postulates, a core set of which [42] are the following:

(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 $K \circ (\varphi_1 \wedge \varphi_2) \models (K \circ \varphi_1) \wedge \varphi_2$.

Postulate KM∘1 formalizes the idea of success, by requiring the revision result to imply the new information φ. Postulate KM∘2 says that if the new information is consistent with existing beliefs, revision simply amounts to adding the new information to the existing beliefs. Postulate KM∘3 says that revision produces a consistent result, if the revision formula is consistent. Postulate KM∘4 formalizes the idea of irrelevance of syntax, by requiring the result to be independent of how information is formulated. Postulates KM∘5 and KM∘6 introduce further coherence constraints on the selection of information from varying revision formulas (i.e., the formulas φ_1 and φ_2). The postulates have been extensively discussed in the belief revision literature (see [1, 39, 42]). More intuition about postulates KM∘5 and KM∘6 is given in Example 6.

A key insight of 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 plausible they seem from the point of view encoded by K. Revising a knowledge base K by a formula φ then amounts to selecting the models of φ most plausible according to K.

Example 6. Take propositional formulas φ_1 and φ_2 and a knowledge base K such that $[\varphi_1] = \{w_1, w_2, w_3\}$, $[\varphi_2] = \{w_1, w_2\}$ and $[K \circ \varphi_1] = \{w_1\}$. According to the plausibility-ranking interpretation of revision, this means that from the point of view of K the interpretation w_1 is considered strictly more plausible than interpretations w_2 and w_3. Consider, next, revision by $\varphi_1 \wedge \varphi_2$ which, according to postulate KM∘4, is equivalent to revision by φ_2. If it would be the case that $w_2 \in [K \circ \varphi_2]$, this would amount to now saying that w_2 is at least as plausible as w_1, which violates the intuition of a stable underlying plausibility relation. Postulates KM∘5 and KM∘6 implement this stability, by making sure that the same elements are selected across varying menus of available options.

A natural way of parsing the idea of plausibility 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* [25], 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. Dalal's operator is illustrated in Example 7.

	$\overbrace{\{a,b\}}^{[K]}$	$\overbrace{\{a,b,c\}}$	$d(E,K)$
$[\varphi]$ $\{a,c\}$	2	1	1
$[\varphi]$ $\{c\}$	3	2	2

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

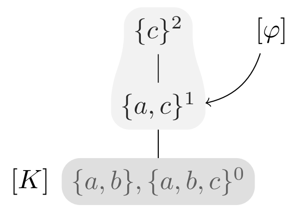

Figure 4: Preorder on interpretations.

Example 7. 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 2. Models in $[\varphi]$ can now be ordered according to their distance to K, visualized in Figure 4. 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*. In propositional revision one typically works with assignments that are *faithful*, meaning that models of K are ranked as the most plausible elements, while non-models are less plausible. Example 7 illustrates one way of generating faithful preorders on interpretations for propositional revision. We will formally introduce faithful assignments in Section 3.2. 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 propositional revision working with faithful assignments turns out to be closely linked to postulates KM$\circ$1–KM$\circ$6, as settled by the following classical representation result.

Theorem 1 ([42])**.** *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.*

In the remaining sections we obtain similar representation results for AF revision, use these results to construct concrete AF revision operators and, finally, analyze the computational complexity of our proposed operators.

3. Revision by Propositional Formulas

3.1. Motivation and Notation

Revision of an AF is required when new information becomes available, challenging an agent's construal of how different aspects of the problem affect each other. This information may come in the form of hard, non-negotiable facts, which is the type of scenario we look at in this section. The following example provides some intuition.

Example 8. Consider the agent from Example 2 again, with doubts about the social benefits of raising the minimum wage, and an inclination to extract as much information from its epistemic state as the epistemic state allows. The AF modelling the agent's reasoning is depicted in Figure 2, and its preferred extensions are $prf(F) = \{\{a, d\}, \{b, d\}\}$. Suppose, now, that an empirical study comes out detailing the impacts of minimum wage on employers and workers, and the agent is persuaded that raising the minimum wage raises the labour supply while, at the same time, lowering the demand for labour. In our framework, the agent accepts that both c and d must hold. What the agent has to change its mind about, in this case, is the fact that c and d cannot appear together in an extension, which we model by saying that the agent has to revise its AF by $\varphi = c \wedge d$.

We are in the situation first described in Example 1, particularized to the preferred semantics. Such a revision, we claim, requires 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 the models of φ are $[\varphi] = \{\{c, d\}, \{a, c, d\}, \{b, c, d\}, \{a, b, c, d\}\}$. The AF G in Figure 5 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 more plausible than G according to some measure of plausibility considered suitable.

Formally, we model this type of scenario with 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 $F \star_\sigma \varphi$.[4] Intuitively, the revision formula φ encodes information which F is required to imply. More concretely,

[4] We restrict the second argument of $\star_\sigma$ to consistent formulas because argumentation se-

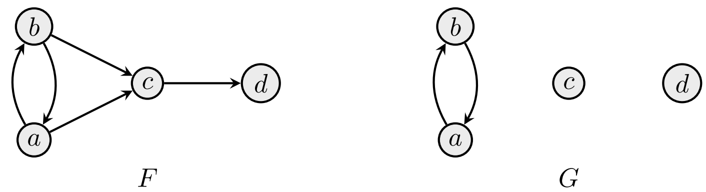

Figure 5: F is to be revised with $\varphi = c \wedge d$, which may result in G.

a revision operator $\star_\sigma$ is expected to change F such that the σ-extensions of $F \star_\sigma \varphi$ are a subset of the models of φ. At the same time, F should not suffer more change than is strictly necessary. This requirement of minimal change is captured, along with other natural requirements expected from an AF revision operator, by the logical postulates presented in Section 3.2. Note that the amount of change to be minimized occurs on the semantic level, i.e., in terms of the extensions of the original AF. Minimal change on the syntactic level is handled in other work [7, 22].

As mentioned in Section 1, we believe it is more natural to represent the resulting set of extensions by a single AF rather than a set of AFs. The latter option, pursued in other existing approaches [21], is a departure from standard accounts of revision, where the result is typically of the same type as the input: think of propositional revision, where the resulting set of interpretations is packaged back into a propositional formula. The need to match input and output arises when revision is applied sequentially, as one expects will be the case in a full-scale real world scenario: if an expression such as '$(F \star_\sigma \varphi_1) \star_\sigma \varphi_2$' is to make sense, then the result of the first revision must be expressed by a single AF and it is preferable if a revision operator can deliver that. However, imposing such a condition on a revision operator adds a layer of difficulty, given by the fact that not any set of extensions is realizable under a given semantics. This means that, although our approach is a special case of [21] in the sense that outcomes are restricted to those which can be expressed by a single AF, we require an alternative treatment in order to obtain representation results.

The difficulties imposed by this condition on the shape of the outcome become apparent once one tries to use standard revision operators for AF revision. On the one hand, it is straightforward to adapt a standard operator, like Dalal, to this purpose. 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 9 below for a concrete application.

Example 9. Using Dalal's adapted approach for the AF F in Example 8 under preferred semantics, we obtain the distances in Table 3, which generate

mantics usually are not capable of expressing the empty set of extensions: among the semantics considered in this paper, only the stable semantics can realize the empty set. However, the results in this section apply even without this restriction for stable semantics.

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

Table 3: $d(E,F)$, for $E \in [\varphi]$.

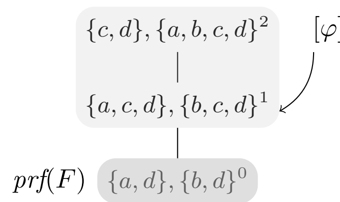

Figure 6: Preorder on extensions.

the preorder $\preceq_F^D$ partly depicted in Figure 6. 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 5 is a suitable revision outcome according to Dalal's adapted operator.

As straightforward as this seems, it turns out that it does not work in general, as the semantic output is not guaranteed to be realizable under the semantics σ we consider. We show this on a concrete example and outline our approach for overcoming these difficulties in Section 3.2.

3.2. Postulates and Representation Result

We adapt the revision postulates [42] to the context of AF revision, in a manner similar to work by Coste-Marquis et al. [21].

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

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

Notice that the postulates are expressed in semantic terms, and they are parameterized by the argumentation semantics σ used to evaluate AFs. The semantics of propositional formulas remains unchanged. Having this semantic perspective of revision, i.e., as a choice function selecting from the models of the incoming information, enables a common perspective on revision for various types of formalisms, whether it is propositional theories, or AFs. In particular, the coherence constraints encoded by postulates P⋆5 and P⋆6 (see Example 6) make as much sense in this setting as in the propositional setting. For the general meaning and motivation of the postulates, see the discussion of the posulates for propositional revision in Section 2.

We next define faithful assignments for AFs, adapting the concept with the same name from propositional revision [42], 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.

We say that an assignment *represents an operator* $\star_\sigma$ (or, alternatively, that $\star_\sigma$ *is represented by an assignment*) if, for any AF F and formula φ, it holds that $\sigma(F \star_\sigma \varphi) = \min([\varphi], \preceq_K)$.

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 9. 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 10 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 [28]. To overcome the problem we use the signature Σ_σ of a semantics 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 10. 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,[5] 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,

[5] See Example 9.

it follows that $\{\{a,b\},\{a,c\},\{b,c\}\} \notin \Sigma_\sigma$, for $\sigma \in \{\mathit{stb}, \mathit{stg}\}$.[6] 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 $\mathbb{S} \in \Sigma_\sigma$ [36]. Hence, we can decide in polynomial time whether a given preorder is σ-compliant.

The concept of σ-compliance makes a 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 $\star_\sigma\colon AF_\mathfrak{A} \times \mathsf{P}_\mathfrak{A} \to AF_\mathfrak{A}$ 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 $\min([\varphi], \preceq_F) \neq \emptyset$.

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

[6] The characterizations of signatures [36] show that also $\{\{a,b\},\{a,c\},\{b,c\}\} \notin \Sigma_\tau$ for $\tau \in \{\mathit{prf}, \mathit{sem}\}$.

total, $E \preceq_F E'$. Hence from $E' \in \min([\varphi], \preceq_F)$ it follows that $E \in \min([\varphi], \preceq_F)$, a contradiction. $\square$

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 3.3. The next result shows that all operators satisfying postulates P$\star$1–P$\star$6 are represented by some σ-compliant assignment. Note that we will, as parts of the proof, show intermediate results as lemmas – therein the assumptions made in the proof until then also apply.

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 $E_1 \in \sigma(F \star_\sigma \varphi)$. $\square$

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 $E_1 \in \min([\varphi], \preceq_F)$. $\square$

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. $\square$

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.[7] In Theorem 2 this is by the preorder being both

[7] 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.

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 $\star_\tau : AF_{\mathfrak{A}} \times \mathsf{P}_{\mathfrak{A}} \to AF_{\mathfrak{A}}$ satisfying P$\star$2.*

Proof. Semantics τ not fulfilling Property 2 from Definition 1 means that there are 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 $\tau(F \star_\tau \phi(\mathbb{S}')) = \tau(F) \cap [\phi(\mathbb{S}')] = \mathbb{S}'$, which is not possible since $\mathbb{S}' \notin \Sigma_\tau$. $\square$

In particular, this applies to the complete semantics which, as an immediate consequence of [36, Theorem 4], does not satisfy this property of closure under subset.

Corollary 1. *There is no operator $\star_{com} : AF_{\mathfrak{A}} \times \mathsf{P}_{\mathfrak{A}} \to AF_{\mathfrak{A}}$ for complete semantics satisfying P$\star$2.*

Finally, note that Theorems 2 and 3 also apply to unique-status semantics such as grounded, ideal and eager semantics. Operators for these semantics will, when revising an AF F by a formula φ, always select a single model of φ and let this be the single extension of the revised AF.

3.3. 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 due to Property 3 of proper I-maximal semantics (note that a extensions is $\subseteq$-incomparable to itself). $\square$

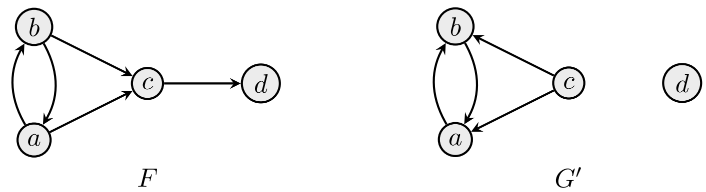

Figure 7: F is to be revised with $\varphi = c \wedge d$, resulting in G'.

Example 11. Take the AF F from Example 8 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 7.

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 11 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. Ideas to a similar effect considered by Coste-Marquis et al. [21] are relevant here, and can be used to complement our approach.

Indeed, as has already been indicated, revision as defined in [21] is also a two-step process. In the first step a set of extensions is produced that satisfies the revision formula and is as close as possible to the extensions of the input AF. In a second step, a set of AFs needs to be selected such that the union of their extensions coincides with the extensions selected in the first step. Coste-Marquis et al. elaborate on concrete “generation operators” that, for a fixed set of arguments (those in the extensions of interest), can be used to select the relevant set of AFs in the second step. Concretely, in their work they consider generation of the set of AFs by minimizing the change on the attack relation, minimizing the number of selected AFs, as well as a combination of these.

While in our work we do not face the problem of minimizing the number of selected AFs, Coste-Marquis et al.'s proposals for minimizing change on the attack relation are clearly relevant. Concretely, the authors consider defining the notion of minimal change on the attack relation through a notion of pseudo-distance, such as the Hamming distance given by $d_{g_H}(F, F') = |(R \setminus R') \cup (R' \setminus R)|$ (for AFs $F = (A, R), F' = (A', R')$), that induces a preorder between AFs that are to be selected (in [21] the Hamming distance is generalised to set of AFs). More elaborate pseudo-distances such as those considered in [20] can also be used. Finally, if a unique AF cannot be selected based on minizing change on the attack relation only, a tie-break rule may be in order. We leave a more detailed exploration of the issue of selecting a concrete AF for future work, where it would be integrated also with more recent results in this direction [22, 56].

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 10).

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:

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

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 12). We call 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 12. Figure 8 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

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

Figure 8: The indexed levels of $\mathbb{S}_1$ and $\mathbb{S}_2$.

approach fits nicely with the wider literature focused on ways to lift rankings on objects to rankings on sets of objects [3]. In our case, the construction of the ranking on extensions assumes that arguments have a certain priority, and that if forced to choose between extensions, we choose extensions with higher priority arguments. And it seems clear that, if revision is to occur, we will need to select among extensions. Example 13 illustrates this.

Example 13. Consider AF F to be revised 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. At the same time, postulates P$\star$1 and P$\star$3 force us to select some non-empty subset of $[\varphi]$.

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. For φ from Example 13, 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. $\square$

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* $\preceq_F^{\mathrm{can}}$ on $2^{\mathfrak{A}}$ is defined, for any $E_1, E_2 \in 2^{\mathfrak{A}}$, as follows:

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

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:

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

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 effect of this is that each new level is σ-compliant (see Example 14).

Example 14. Let F be an AF such that $stb(F) = \{\{b, c\}\}$. Figure 9 depicts the canonical preorder $\preceq_F^{\mathrm{can}}$ and Figure 10 shows the ranking $\preceq_F^{D,\mathrm{ir}}$, 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 $\preceq_F^{\mathrm{can}}$ and $\preceq_F^{D,\mathrm{ir}}$ are σ-compliant. Also notice how the levels of $\preceq_F^D$ get split according to the indexed preorder to obtain $\preceq_F^{D,\mathrm{ir}}$.

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 $\star_\sigma^{\mathrm{can}}$. If $\star_\sigma^x$ is an existing AF revision operator, we will call the operator defined using the indexed preorder the *indexed-refined revision operator*, and denote it by $\star_\sigma^{x,\mathrm{ir}}$. 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 $\star_\sigma^{D,\mathrm{ir}}$, obtained by refining Dalal's operator. We show next that they also satisfy postulates P$\star$1–P$\star$6.

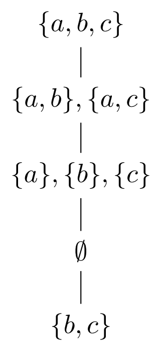

Figure 9: $\preceq_F^{\text{can}}$.

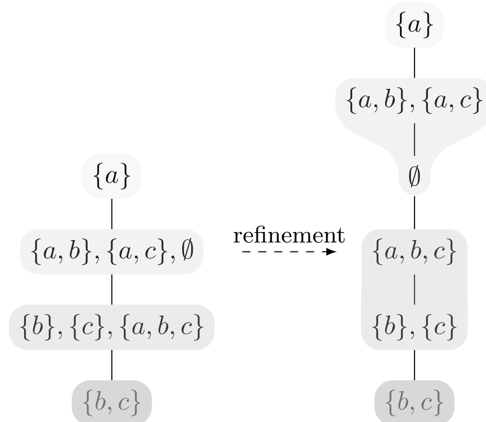

Figure 10: $\preceq_F^{D}$ and $\preceq_F^{D,\,\text{ir}}$.

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$\star$1–P$\star$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$\star$1–P$\star$6. □

The choice of giving priority to extensions with fewer elements makes sense if we think that smaller extensions are more discriminating than larger ones: intuitively, an extension that contained every argument in $\mathfrak{A}$ would mean that all arguments are jointly accepted, which would translate back into an AF that carried very little information about the argumentative structure of the issue at hand. Notwithstanding, one should not place too much weight on this detail: one could, for instance, invert the inequality in the definition of indexed preorders, such that larger extensions are given priority, and the result would still be a σ-compliant preorder. The force of our argument lies more in showing that such rankings exist. Finding rankings that are more realistic, or useful, though an interesting topic, is beyond the scope of this paper.

4. Revision by Argumentation Frameworks

4.1. Motivation and Notation

New information comes not just in the form of bare facts, but often as the outcome of sustained argumentative efforts from another agent. This mandates an approach to revision where AFs are revised by other AFs.

Example 15. Consider the agent in Example 2 and his view of the minimum wage situation encoded by the AF F in Figure 2. Suppose, now, that this agent

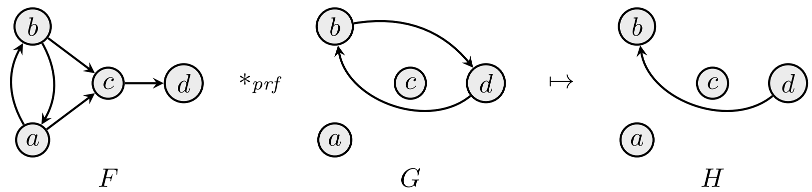

Figure 11: F is revised by G with respect to preferred semantics, resulting in H

reads the works of an eminent economist whose view on the matter is that, except for the major impact of technology, none of the other factors have any effect on each other. Our agent believes that the economist is a trusted source and is persuaded to accept the conclusions of their argument. In other words, the agent accepts the gist of the economist's argument and is now faced with the task of revising by the AF G in Figure 11. The agent, we can imagine, reasons as follows:

> The economist makes a convincing case that things are as in G. From this it follows that $\{\{a, b, c\}, \{a, b, d\}\}$ can be jointly accepted. Between these two possibilities, $\{a, c, d\}$ comes closest to what I thought was the case, and is what I would be inclined to accept now.

As a result, the agent brings its epistemic state in line with the AF H in Figure 11, whose single preferred extension is $\{a, c, d\}$.

Formally, we model revision of an AF by another AF by 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 conform with 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.

4.2. Postulates, Representation Result and Concrete Operators

As before, we show a correspondence between a set of postulates and an assignment mapping AFs to rankings on $2^{\mathfrak{A}}$. The revision postulates, in the manner of [42], 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 $\sigma(F_1 *_\sigma G) = \sigma(F_2 *_\sigma H)$.

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

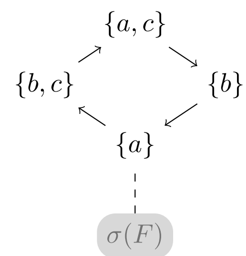

Figure 12: Cycles in rankings on extensions.

(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 $\sigma(F *_\sigma G_n) \cap \sigma(G_0) \neq \emptyset$.

Similarly to previous sections, we say that an assignment *represents an operator* $*_\sigma$ (or, alternatively, that $*_\sigma$ *is represented by an assignment*) if, for any AFs F and G, it holds that $\sigma(F \star_\sigma G) = \min(\sigma(G), \preceq_K)$.

Though we are no longer in the realm of propositional logic, the common semantic approach we adopt affords a unified view of all types of revision studied in the paper. In particular, since the semantics of AFs can be seen as a restricted fragment of the semantics of propositional theories, the interpretation of operator revising an AF F by an AF G carries over from the propositional case: we view such an operator as implicitly choosing among the extensions G the ones that make most sense according to an underlying plausibility ranking. For motivation about postulates A∗1–A∗6, see the discussion of the postulates for propositional revision in Section 2. The one addition to the core set postulates, i.e., the postulate Acyc, is adapted from [28] and is motivated by the realization that, without it, postulates A∗1–A∗6 are represented by assignments containing non-transitive rankings. Example 16 makes this point clear.

Example 16. Suppose that for an AF F we have a ranking $\preceq_F$ on $2^{\mathfrak{A}}$ which behaves as in Figure 12 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.[8]

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 (cf. Definition 3) 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 given the initial standpoint encoded in F, 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:*

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

[8]Notice that we do not run into this problem when revising by a formula, since in such a setting we can always revise by a formula φ whose models are $\{a\}$, $\{b\}$, $\{b,c\}$, $\{a,c\}$ (which is not available in this setting), and a well-defined operator satisfying the postulates is forced to give us a non-empty answer, thereby eliminating the non-transitive cycle. Thus, Acyc is needed only for this variant of revision.

*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.

For postulate A$*$3 first note that, since σ is proper I-maximal, any $S, S' \in \sigma(G)$ are $\subseteq$-incomparable. Hence, as $\preceq_F$ is I-faithful, $S \preceq_F S'$ or $S' \preceq_F S$ (or both) for any $S, S' \in \sigma(G)$. This together with the fact that $\sigma(G)$ is finite, means that $\min(\sigma(G), \preceq_F) \neq \emptyset$.

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 $\min(\sigma(G_1), \preceq_F) \cap \sigma(G_0) \neq \emptyset$. Hence there is an extension $E_0' \in \sigma(G_0) \cap \sigma(G_1)$ 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 $E_1' \in \sigma(G_1) \cap \sigma(G_2)$ such that $E_1' \preceq_F E_2$ for all $E_2 \in \sigma(G_2)$, $\ldots$, and from $(F * G_n) \cap \sigma(G_{n-1}) \neq \emptyset$ that there is an extension $E_{n-1}' \in \sigma(G_{n-1}) \cap \sigma(G_n)$ 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 $E_n' \in \sigma(G_n) \cap \sigma(G_0)$ 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 $\sigma(F *_\sigma G_n) \cap \sigma(G_0) \neq \emptyset$. $\square$

Again, we will show intermediate results as lemmas within the proof of the following theorem.

Theorem 6. *If $*_\sigma \colon AF_\mathfrak{A} \times AF_\mathfrak{A} \to AF_\mathfrak{A}$ 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 $*_\sigma \colon AF_\mathfrak{A} \times AF_\mathfrak{A} \to AF_\mathfrak{A}$ 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 $E, E' \in 2^\mathfrak{A}$:

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

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:

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

The remainder of the proof shows that $\preceq_F$ is the desired I-faithful ranking. First, notice that if $E_1 \preceq'_F E_2$ then $E_1 \preceq_F E_2$. Hence $\preceq_F$ is reflexive and, by construction, it is transitive, which makes it a 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 the 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 $E_1, \ldots, E_n$ are pairwise distinct extensions with $E_1 \preceq'_F E_2 \preceq'_F \cdots \preceq'_F E_n \preceq'_F E_1$, then $E_1 \preceq'_F E_n$.*

Proof. If $n = 2$ the conclusion follows immediately. In the following we assume that $n > 2$. From the hypothesis we have that $E_i \in \sigma(F *_\sigma f_\sigma(\{E_i, E_{i+1}\}))$, for $i \in \{1, n-1\}$, and $E_n \in \sigma(F *_\sigma f_\sigma(\{E_n, E_1\}))$. It follows that $E_1 \in \sigma(F *_\sigma f_\sigma(\{E_1, E_2\})) \cap \{E_n, E_1\}$, $E_i \in \sigma(F *_\sigma f_\sigma(\{E_i, E_{i+1}\})) \cap \{E_{i-1}, E_i\}$, for $i \in \{2, \ldots, n-1\}$, 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 $E_1 \preceq'_F E_n$. □

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 $E_1, \ldots, E_n$ 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. □

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

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. $\square$

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 $F \in AF_{\mathfrak{A}}$. In particular, since revising by an AF G produces, as it were by design, a non-empty subset of $\sigma(G)$, we are guaranteed to obtain a σ-realizable result on every occasion. 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 [25]) can be used as a revision operator of AFs by AFs.

Example 17. Consider an AF F as in Example 10, with $stb(F) = \{\{a, b, c\}\}$, for instance $F = (\{a, b, c\}, \emptyset)$. The corresponding ranking obtained with Hamming distance, $\{a, b, c\} \prec_F^D \{a, b\} \approx_F^D \{a, c\} \approx_F^D \{b, c\} \prec_F^D \{a\} \approx_F^D \{b\} \approx_F^D \{c\} \prec_F^D \emptyset$, 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 10). 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 $\min(\sigma(G), \preceq_F^D)$. 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.

The complete semantics is not proper I-maximal and is therefore not captured by the representation result given by Theorems 5 and 6. As it turns out, there is no way of tuning the requirements on the ranking side in order to get a similar result, since postulate A$*$2 cannot be satisfied by operators for revision under complete semantics. More generally, this applies to any semantics which is not closed under intersection in terms of expressiveness:

Proposition 5. *Let τ be a semantics such that there are some AFs F and G with $\tau(F) \cap \tau(G) \neq \emptyset$ and $\tau(F) \cap \tau(G) \notin \Sigma_\tau$. Then there is no operator $*_\tau : AF_{\mathfrak{A}} \times AF_{\mathfrak{A}} \to AF_{\mathfrak{A}}$ satisfying A$*$2.*

Proof. Let F and G be AFs with $\tau(F) \cap \tau(G) \neq \emptyset$ and $\tau(F) \cap \tau(G) \notin \Sigma_\tau$. For $*_\tau$ to satisfy A$*$2 it would be required to give $\tau(F *_\tau G) = \tau(F) \cap \tau(G)$, which is not possible. $\square$

The fact that complete semantics lacks this form of closure was shown in [36, Theorem 4].

Corollary 2. *There is no operator $*_{com} : AF_{\mathfrak{A}} \times AF_{\mathfrak{A}} \to AF_{\mathfrak{A}}$ for complete semantics satisfying A$*$2.*

For semantics which are not proper I-maximal but closed under intersection in terms of expressiveness estabilshing a representation result is subject to future work. Assuming that there exist operators adhering to the full set of revision postulates, appropriate restrictions to the rankings have to be figured out.

For the special case of unique-status semantics, the only operator which can be rational according to the postulates, is the one always returning an AF having exactly the same extension as the revising AF (due to A$*$1).

5. Complexity

Next, we study the complexity of Dalal's operator and its refinement in the argumentation setting.[9] 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 [47]. 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 subproblem from $\mathcal{C}$ in one computation step. The class Σ_k^{P} (and Δ_k^{P}) contains the problems that can be decided in polynomial time by a non-deterministic (deterministic) Turing machine with unrestricted access to a $\Sigma_{k-1}^{\mathsf{P}}$ 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 $\{\overline{x_1}, \dots, \overline{x_n}\}$. Moreover, if $z = \overline{x}$ then $\overline{z} = x$ (or, in other words, $\overline{\overline{x}} = x$). Finally, for a formula φ (resp. a clause c, a monom d) and a set of atoms S, we write $S \models \varphi$ (resp. $S \models c$, $S \models d$) if the interpretation where all atoms in S are assigned to true and all other atoms are assigned to false satisfies φ (resp. c, d).

The classes $\Theta_{k+1}^{\mathsf{P}}$ (for $k \geq 1$) have the following complete problems [38, 67, 64], which we will make use of in the subsequent hardness proofs:

Given an AF F and a set of arguments E, deciding whether $E \in stb(F)$ is in P and deciding whether $E \in prf(F)$ is coNP-complete [31].

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

[9]In what follows, we consider a more general setting by giving up the restriction that the domain of arguments $\mathfrak{A}$ is finite, but still assume AFs to be finite.

GIVEN:	k-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 $\max\{1 \leq i \leq m \mid \Phi_i \text{ is true}\}$ is odd.

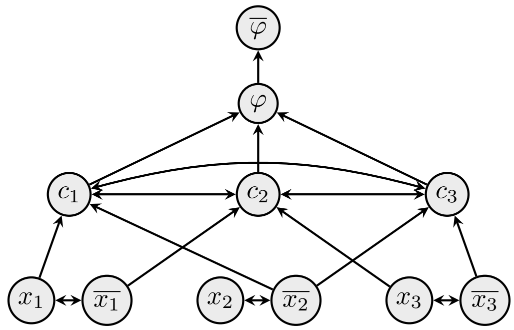

Figure 13: AF F_φ for $\varphi(X) = (x_1 \vee \neg x_2) \wedge (\neg x_1 \vee x_3) \wedge (\neg x_2 \vee \neg x_3)$.

adapted from reductions used in proofs by Dimopoulos and Torres [31] and Dunne and Bench-Capon [35].

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 13 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 $E \in stb(F_\varphi)$ such that $\overline{\varphi} \notin E$;*
2. *for each $E, E' \in stb(F_\varphi)$ such that $\overline{\varphi} \notin E$ and $\overline{\varphi} \in E'$ it holds that $|E|+1 = |E'|$;*
3. *for each $E \in stb(F_\varphi)$ such that $\overline{\varphi} \notin E$ and each $E' \in stb(F_\varphi - (C \cup \{\overline{\varphi}\}))$ it holds that $|E| = |E'|$.*

Proof. We begin with the observation that every stable extension of F_φ 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. Note that this is also the case for the AF $F_\varphi - (C \cup \{\overline{\varphi}\})$.

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 $S \cup (\overline{X \setminus S}) \cup \{\varphi\} \in stb(F_\varphi)$. ($\Leftarrow$): Let $E \in stb(F_\varphi)$ 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. $\square$

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. $T \subseteq A_G$) for each $S \in \mathit{stb}(F)$ (resp. $T \in \mathit{stb}(G)$). 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 $S \in \mathit{stb}(F)$, $T \in \mathit{stb}(G)$ 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 $S \in \mathit{stb}(F)$ such that $d_H(S, E) = z$. If such an S does exist, then $E \in \mathit{stb}(F *^D_{stb} G)$, 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 $\varphi_1(X_1), \ldots, \varphi_m(X_m)$ 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:

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

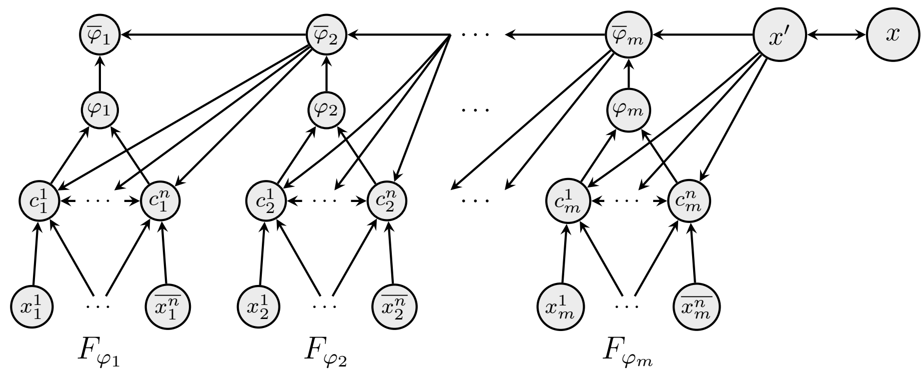

Figure 14: Illustration of the AF obtained from the reduction in the proof of Theorem 7.

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 14. Moreover, we define $G = (\{x, x'\}, \{(x, x'), (x', x)\})$ and $E = \{x\}$.

Due to the splitting Property [6], 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 $F'_{\varphi_{i-1}} = F_{\varphi_{i-1}} - (C_{i-1} \cup \{\overline{\varphi}_{i-1}\})$. 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 $E \in stb(F_{\varphi_i})$ 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.

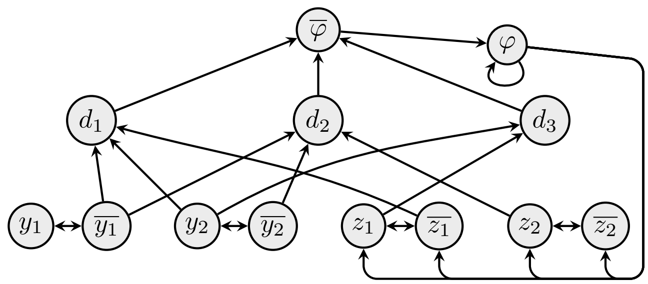

Figure 15: AF F_Φ for the QBF $\Phi = \exists y_1, y_2 \forall z_1, z_2 : (y_1 \wedge \neg y_2 \wedge z_1) \vee (y_1 \wedge y_2 \wedge \neg z_2) \vee (\neg y_2 \wedge \neg z_1)$.

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 15. 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. Hence, it is 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 $S \cup (\overline{Y \setminus S}) \cup T \cup (\overline{Z \setminus T})$. $\square$

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 $E \in prf(F_\Phi - \{\overline{\varphi}\})$ it 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 argument $d \in D$ (representing a monom $d \in D$) included in $E' \setminus E$. Due to the non-triviality of the monom 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 an element of an admissible set of F'_Φ as it is attacked by the self-attacking, and otherwise unattacked, argument φ. Moreover, as φ does not contain monoms which are trivial for Y, each argument $d \in D$ is attacked by at least one $z \in Z \cup \overline{Z}$. Besides φ, the only attacker of this z is $\overline{z}$, which, as just shown, cannot be part of an admissible set of F'_Φ. Hence it follows that no $d \in D$ can be part of an admissible set of F'_Φ. 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, with the exception of the attack from itself, unattacked in F'_Φ. Hence none of the arguments $Z \cup \overline{Z}$ can be an element of an admissible set of F'_Φ.

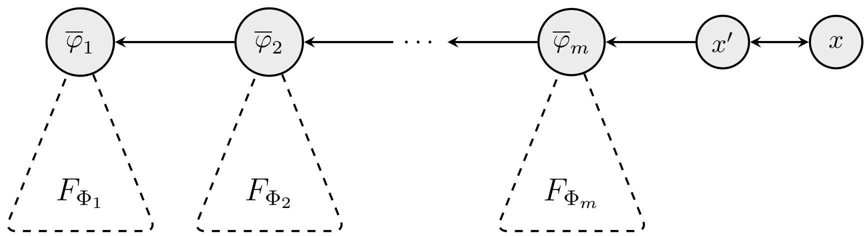

Figure 16: Illustration of the AF obtained from the reduction in the proof of Theorem 8.

Moreover, as the formula φ does not contain monoms which are trivial for Y, each argument $d \in D$ is attacked by at least one $z \in Z \cup \overline{Z}$. Besides φ, the only attacker of this z is $\overline{z}$, which, as just shown, cannot be part of an admissible set of F'_Φ. Hence no argument $d \in D$ can be part of an admissible set of F'_Φ. 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 $E \in \mathit{prf}(F *^D_{\mathit{prf}} G)$ is Θ_3^{P}-complete.*

Proof. To show membership in Θ_3^{P} we sketch an algorithm that decides $E \in \mathit{prf}(F *^D_{\mathit{prf}} G)$ in polynomial time with $\mathcal{O}(\log m)$ calls to a Σ_2^{P} oracle, where $m = |A_F| + |A_G|$. First, we check whether $E \in \mathit{prf}(G)$ (in coNP); if no we return with a negative answer. Second, the minimal distance $z = \min\{d_{\mathit{prf}}(T, F) \mid T \in \mathit{prf}(G)\}$ is determined. Since $S \subseteq A_F$ (resp. $T \subseteq A_G$) for each $S \in \mathit{prf}(F)$ (resp. $T \in \mathit{prf}(G)$), 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_2^{P} procedure: Guess $S \subseteq A_F$, $T \subseteq A_G$ and check (in coNP) whether $S \in \mathit{prf}(F)$, $T \in \mathit{prf}(G)$ and $d_H(S, T) < z$. Having obtained z, we finally call the oracle once again to check whether there is an $S \in \mathit{prf}(F)$ with $d_H(S, E) = z$. If such an S does exist then $E \in \mathit{prf}(F *^D_{\mathit{prf}} G)$, otherwise not.

To show Θ_3^{P}-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:

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

Figure 16 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 $E \in \mathit{prf}(F *^D_{\mathit{prf}} G)$ if and only if k is odd.

Due to the splitting property of preferred semantics [6], the preferred extensions of F are composed as $\mathit{prf}(F) = \{\{\alpha\} \cup \bigcup_{1 \leq i \leq m} E_i \mid \alpha \in \{x, x'\}, E_i \in \mathit{prf}(F'_{\Phi_i})\}$, 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 3.3, 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 $E \in \mathit{stb}(F \star^{D,\mathrm{ir}}_{\mathit{stb}} \varphi)$ is Θ^{P}_2-hard;*
- *deciding whether $E \in \mathit{prf}(F \star^{D,\mathrm{ir}}_{\mathit{prf}} \varphi)$ is Θ^{P}_3-hard.*

Proof. Let $\sigma \in \{stb, prf\}$. 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 $E \in \sigma(F *_{\sigma}^{D} G)$. Inspecting the hardness proofs of Theorems 7 and 8, we see that this problem is Θ_2^{P}-hard for $\sigma = stb$ and Θ_3^{P}-hard for $\sigma = prf$ even for this fixed G.[10] Hence the reduction will give the desired result.

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:

$$F' = F \cup (\{y_1, \ldots, y_{\frac{n}{2}}\}, \emptyset), \text{ and}$$

$$\varphi = \phi(\sigma(G)) \wedge \bigwedge_{1 \le i \le n} \left(\neg \left(\bigwedge_{a \in E} a \wedge \bigwedge_{a' \in (A_G \setminus E)} \neg a' \right) \leftrightarrow y_i \right),$$

with $\{y_1, \ldots, y_n\}$ being newly introduced arguments. We show that $E \in \sigma(F *_{\sigma}^{D} G)$ if and only if $E \in \sigma(F' \star_{\sigma}^{D,\,\mathrm{ir}} \varphi)$. 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 $Y \subseteq \{y_1, \ldots, y_n\}$ $(Y \neq \emptyset)$, and if $S \neq E$ then $S \cup \{y_1, \ldots, y_n\} \in [\varphi]$ and $S \cup Y \notin [\varphi]$ for any $Y \subset \{y_1, \ldots, y_n\}$. Therefore we derive the following:

- denoting $S' = S \cup \{y_1, \ldots, y_n\}$ for every $S \in (\sigma(G) \setminus \{E\})$ and denoting $E' = E$, it holds that $[\varphi] = \{S' \mid S \in \sigma(G)\}$;
- denoting $T' = T \cup \{y_1, \ldots, y_{\frac{n}{2}}\}$ 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 $S_1 \preceq_F^D S_2$ if and only if $S_1' \preceq_{F'}^D S_2'$.

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

On the other hand assume $E \in \sigma(F *_{\sigma}^{D} G)$. That means that for all $S \in \sigma(G)$ it holds that $E \preceq_F^D S$. For those $S \in \sigma(G)$ with $E \prec_F^D S$ we get $E' \prec_{F'}^{D,\mathrm{ir}} S'$ as before. Consider an $S \in \sigma(G)$ with $E \approx_F^D S$. 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, $E \preceq_{F'}^{D,\mathrm{ir}} S'$. Since this holds for every $S \in \sigma(G)$ we conclude that $E \in \sigma(F' \star_{\sigma}^{D,\,\mathrm{ir}} \varphi)$. $\square$

As an upper bound for the complexity, we show membership in Δ_2^{P} for revision with respect to stable semantics and membership in Δ_3^{P} for preferred semantics.

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

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

- *deciding whether $E \in stb(F \star_{stb}^{D,\mathrm{ir}} \varphi)$ is in Δ_2^{P};*
- *deciding whether $E \in prf(F \star_{prf}^{D,\mathrm{ir}} \varphi)$ is in Δ_3^{P}.*

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 $E \in stb(F \star_{stb}^{D,\mathrm{ir}} \varphi)$ in polynomial time with access to an NP oracle. Let $m = |A_F| + |var(\varphi)|$.[11] 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 $z = \min\{d_{stb}(T, F) \mid T \in [\varphi]\}$ 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 $a_i \in A_F \cup var(\varphi)$. 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 $S \subseteq A_F$, $T \subseteq var(\varphi)$, and check whether $S \in stb(F)$, $T \in [\varphi]$, $d_H(S, T) = z$, $|T| \leq e$ and $\overrightarrow{T}^{\#} \preceq_{\mathrm{lex}} p$. 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 $S \in stb(F)$ such that $d_H(S, E) = z$; if these checks turn out positive, $E \in stb(F \star_{stb}^{D,\mathrm{ir}} \varphi)$, otherwise not.

The proof for Δ_3^{P}-membership of deciding whether $E \in prf(F \star_{prf}^{D,\mathrm{ir}} \varphi)$ uses the same polynomial time procedure, now with access to a Σ_2^{P} 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_2^{P} oracle. $\square$

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

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

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 [11, 13, 14, 18, 19, 46].

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 [7] 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. [44] 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. [15] give an AGM-like characterization of revision of AFs when certain logical “constraints” expressing beliefs regarding the labellings of the AFs are “strengthened” to reflect 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. [21], where revision functions are defined following a two step process: first a counterpart 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. [21] 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 10, 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 18. Consider the AF F depicted in Figure 17 (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 $\star\colon AF_{\mathfrak{A}} \times \mathsf{P}_{\mathfrak{A}} \to AF_{\mathfrak{A}}$ 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

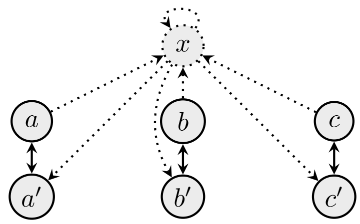

Figure 17: AFs discussed in Example 18.

Figure 17 including the dotted part.

As the previous example shows, the choice of Coste-Marquis et al. [21] 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. [9] 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. [32], where enforcement is encoded in the framework of Dynamic Logic of Propositional Assignments (DL-PA). In the same direction, work by Nouioua and Würbel [57] 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. [22] 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. [23] 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 [8] 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 [51] 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 [29]. 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 [27] (see also preceding work [12]) 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 [52] and Snaith and Reed [62] present models of dynamics in *structured* (as opposed to abstract) argumentation. The former offers a model building on results by Moguillansky [51] (see also previous work from this group [53, 59, 54, 55], as well as the quantitative approach presented in [61]), 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 [36] and belief change [28]. We have considered two different approaches to revision of AFs; our results are summarized in Table 4. 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 analyzed 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

Revision Type	Semantics	Postulates	Rankings	Impossibility
by formula	all	P$\star$1–P$\star$6	faithful, σ-compliant (Defs. 3, 4)	not $\subseteq$-closed (Prop. 2)
by AF	proper I-max. (Def. 1)	A$\ast$1–A$\ast$6, Acyc	I-faithful (Def. 8)	not $\cap$-closed (Prop. 5)

Table 4: Summary of the respresentation results provided by Theorems 2 and 3 and Theorems 5 and 6, respectively, including the semantics they apply to, the postulates we require the operator to adhere to, and the restrictions on rankings needed. The last column recalls sufficient conditions for non-existence of operators for the respective approaches.

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 [49]. 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. Then, in order to apply our approach to concrete instantiations AFs, one has to consider the fact that these might yield only fragments of the class of all AFs. We envisage that the notion of compliance allows to adapt to these fragments as soon as we have a characterization of their signature.

On a more general level, in this work we have remained faithful to the traditional view of revision where the input to the revision process must be fully accommodated in the result; it would also be worthwile to explore non-prioritized revision [43, 40, 41], especially in the revison-by-AF scenario. Also, we want to analyze 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.

Acknowledgement

Supported by the Austrian Science Fund (FWF) under grants I1102, I2854, P25518, P25521, P30168, W1255, and Y698.

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] Katie Atkinson, Pietro Baroni, Massimiliano Giacomin, Anthony Hunter, Henry Prakken, Chris Reed, Guillermo R. Simari, Matthias Thimm, and Serena Villata. Towards artificial argumentation. *AI Magazine*, 38(3):25–36, 2017.

[3] Salvador Barberà, Walter Bossert, and Prasanta K. Pattanaik. Ranking sets of objects. In Salvador Barberà, Peter J. Hammond, and Christian Seidl, editors, *Handbook of Utility Theory: Volume 2 Extensions*, pages 893–977. Springer US, Boston, MA, 2004.

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

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

[6] 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.

[7] 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.

[8] 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.

[9] 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.

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

[11] 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.

[12] 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.

[13] 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.

[14] 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.

[15] 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.

[16] Martin Caminada. Comparing two unique extension semantics for formal argumentation: Ideal and eager. In *Proceedings of the 19th Belgian-Dutch Conference on Artificial Intelligence (BNAIC 2007)*, pages 81–87, 2007.

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

[18] 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.

[19] 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.

[20] Sylvie Coste-Marquis, Caroline Devred, Sébastien Konieczny, Marie-Christine Lagasquie-Schiex, and Pierre Marquis. On the merging of Dung's argumentation systems. *Artif. Intell.*, 171(10-15):730–753, 2007.

[21] 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.

[22] 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.

[23] 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.

[24] 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.

[25] 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.

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

[27] 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.

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

[29] 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.

[30] 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.

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

[32] 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.

[33] 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.

[34] Phan Minh Dung, Paolo Mancarella, and Francesca Toni. Computing ideal sceptical argumentation. *Artif. Intell.*, 171(10-15):642–674, 2007.

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

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

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

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

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

[40] Sven Ove Hansson. Semi-Revision (invited paper). *Journal of Applied Non-Classical Logics*, 7(2), 1997.

[41] Sven Ove Hansson, Eduardo L. Fermé, John Cantwell, and Marcelo A. Falappa. Credibility Limited Revision. *J. Symb. Log.*, 66(4):1581–1596, 2001.

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

[43] Sébastien Konieczny and Ramón Pino Pérez. Logic Based Merging. *J. Philosophical Logic*, 40(2):239–270, 2011.

[44] 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.

[45] 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.

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

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

[48] Thomas Linsbichler. *Advances in Abstract Argumentation – Expressiveness and Dynamics*. PhD thesis, TU Wien, 2017.

[49] 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.

[50] Sanjay Modgil, Francesca Toni, Floris Bex, Ivan Bratko, Carlos I. Chesñevar, Wolfgang Dvořák, Marcelo A. Falappa, Xiuyi Fan, Sarah A. Gaggl, Alejandro J. García, María P. González, Thomas F. Gordon, João Leite, Martin Možina, Chris Reed, Guillermo R. Simari, Stefan Szeider, Paolo Torroni, and Stefan Woltran. The Added Value of Argumentation. In Sascha Ossowski, editor, *Agreement Technologies*, volume 8 of *Law, Governance and Technology Series*, pages 357–403. Springer Netherlands, 2013.

[51] 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.

[52] 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.

[53] 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.

[54] 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.

[55] Martín O. Moguillansky, Nicolás D. Rotstein, Marcelo A. Falappa, Alejandro Javier García, and Guillermo Ricardo Simari. Dynamics of knowledge in *DeLP* through Argument Theory Change. *TPLP*, 13(6):893–957, 2013.

[56] 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.

[57] 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.

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

[59] 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.

[60] 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.

[61] Gerardo I. Simari, Paulo Shakarian, and Marcelo A. Falappa. A quantitative approach to belief revision in structured probabilistic argumentation. *Ann. Math. Artif. Intell.*, 76(3-4):375–408, 2016.

[62] Mark Snaith and Chris Reed. Argument revision. *J. Log. Comput.*, 27(7): 2089–2134, 2017.

[63] 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.

[64] 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.

[65] Phillip E. Tetlock and Dan Gardner. *Superforecasting: The Art and Science of Prediction*. Crown/Archetype, 2015. ISBN 9780804136709. URL `https://books.google.at/books?id=hC_qBQAAQBAJ`.

[66] 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.

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