Revision of Abstract Dialectical Frameworks: Preliminary Report*

Thomas Linsbichler and **Stefan Woltran**
Institute of Information Systems, TU Wien, Austria
{linsbich,woltran}@dbai.tuwien.ac.at

Abstract

Abstract Dialectical Frameworks (ADFs) enhance the capability of Dung's argumentation frameworks by modelling relations between arguments in a flexible way, thus constituting a very general formalism for abstract argumentation. Since argumentation is an inherently dynamic process, understanding how change in ADFs can be formalized is important. In this work we study AGM-style revision operators for ADFs by providing various representation results. We focus on the preferred semantics and employ tools recently developed in work on revision of Horn formulas as well as logic programs. Moreover, we present an alternative family of operators based on a variant of the postulates considering preferred interpretations of the original and admissible interpretations of the revising ADF.

1 Introduction

Within the research field of argumentation in artificial intelligence [Bench-Capon and Dunne, 2007], *abstract argumentation frameworks* (AFs) as introduced by Dung [1995] have turned out to be a suitable modelling tool for various argumentation problems. This is partly due to their conceptional simplicity, being just a directed graph where nodes represent abstract arguments and edges represent conflicts between arguments. However, this comes also with limitations in terms of expressibility, which has led to the introduction of several enhancements of Dung's AFs, incorporating support [Cayrol and Lagasquie-Schiex, 2005], preferences [Modgil, 2009], attacks on attacks [Baroni *et al.*, 2011] and other concepts (see [Brewka *et al.*, 2014] for an overview). One of the most recent and powerful generalizations of AFs constitute *abstract dialectical frameworks* (ADFs) [Brewka and Woltran, 2010; Brewka *et al.*, 2013], where the relation between arguments is modelled via *acceptance conditions* for each argument (in the form of Boolean functions), capturing various forms of attack and support. This enhanced modelling capability of ADFs has been used for preferential reasoning [Brewka *et al.*, 2013], judgment aggregation [Booth, 2015], and legal reasoning [Al-Abdulkarim *et al.*, 2016].

*Supported by FWF through projects I1102, I2854, and P25521.

Argumentation as such is a highly *dynamic process*. Therefore the evaluation of formalisms modelling argumentation problems is subject to constant changes in the model. As a consequence, there has been a tremendous amount of research on the dynamics of argumentation frameworks and in particular on the *revision* of Dung AFs (see e.g. [Falappa *et al.*, 2011]) in the last years. The prominent AGM approach for belief change [Alchourrón *et al.*, 1985; Katsuno and Mendelzon, 1991] was applied to AFs by Coste-Marquis *et al.* [2014], characterizing minimal change revision operators by so-called *representation results*. Diller *et al.* [2015] used recent insights on the expressiveness of AFs [Dunne *et al.*, 2015] as well as on how to deal with fragments of classical logic in belief revision [Delgrande and Peppas, 2015] to characterize AGM revision operators which return a single AF instead of a set of such.

In this work we study such AGM revision operators for ADFs. We obtain representation theorems characterizing *all* operators satisfying an adapted version of the AGM postulates by rankings on interpretations. The main challenge is the fact that ADFs are not able to express arbitrary sets of interpretations under these semantics (supported models being an exception in this matter). Fortunately, the exact *expressiveness* of ADF-semantics has recently been established by Pührer [2015] and Strass [2015], who gave exact characterizations for realizability under three-valued and two-valued semantics, respectively. We will extend and employ these results which will result in a different characterization for each semantics. Most semantics evaluate ADFs based on *three-valued* interpretations, generalizing labelling-based semantics of AFs [Caminada and Gabbay, 2009]. Therefore, to obtain concrete operators, we employ a distance measure for three-valued interpretations to define rankings.

We will focus on preferred and admissible semantics – preferred interpretations are defined as maximal admissible interpretations. For revision under preferred semantics we obtain a representation result by adjusting the conditions on rankings to the limited expressiveness of the semantics and adding an additional postulate inspired by [Delgrande and Peppas, 2015] preventing cycles. The approach is similar to revision of AFs [Diller *et al.*, 2015], but deals with sets of three-valued interpretations instead of two-valued extensions. Moreover, will define a three-valued version of Dalal's well-known revision operator [Dalal, 1988]. Admissible semantics, on the other hand, yield only a single operator satisfying the postulates. Since, as we will argue, both approaches have some weaknesses, we propose a novel hybrid approach which bases rankings on preferred interpretations but allows admissible interpretations of the revising ADF to be the result of the revision.

Finally, we informally discuss the representation of operators for the two-valued semantics, namely stable and supported models. Moreover, we argue that complete and grounded semantics cannot be captured by the AGM approach, i.e. there are no operators satisfying the postulates.

2 Background

We assume a fixed finite set of statements A. An *interpretation* is a mapping $v : A \to \{\mathbf{t}, \mathbf{f}, \mathbf{u}\}$ assigning one of the truth values true ($\mathbf{t}$), false ($\mathbf{f}$) or unknown ($\mathbf{u}$) to each statement. The set of statements to which v assigns a particular truth value $\mathbf{x} \in \{\mathbf{t}, \mathbf{f}, \mathbf{u}\}$ is denoted by $v^{\mathbf{x}}$. An interpretation is *two-valued* if $v^{\mathbf{u}} = \emptyset$, i.e. the truth value $\mathbf{u}$ is not assigned. Two-valued interpretations v can be extended to assign truth values $v(\varphi) \in \{\mathbf{t}, \mathbf{f}\}$ to propositional formulas φ as usual.

The three truth values are partially ordered according to their information content: we have $\mathbf{u} <_i \mathbf{t}$ and $\mathbf{u} <_i \mathbf{f}$ and no other pair in $<_i$, meaning that the classical truth values contain more information than the truth value unknown. As usual, $\leq_i$ denotes the partial order associated to the strict partial order $<_i$. The pair $(\{\mathbf{t}, \mathbf{f}, \mathbf{u}\}, \leq_i)$ forms a complete meet-semilattice with the information meet operation $\sqcap_i$. This meet can intuitively be interpreted as *consensus* and assigns $\mathbf{t} \sqcap_i \mathbf{t} = \mathbf{t}$, $\mathbf{f} \sqcap_i \mathbf{f} = \mathbf{f}$, and returns $\mathbf{u}$ otherwise. The information ordering $\leq_i$ extends in a straightforward way to interpretations v_1, v_2 over A in that $v_1 \leq_i v_2$ iff $v_1(a) \leq_i v_2(a)$ for all $a \in A$. We then say, for two interpretations v_1, v_2, that v_2 *extends* v_1 iff $v_1 \leq_i v_2$. The set $\mathcal{V}$ of all interpretations over A forms a complete meet-semilattice with respect to the information ordering $\leq_i$. The consensus meet operation $\sqcap_i$ of this semilattice is given by $(v_1 \sqcap_i v_2)(a) = v_1(a) \sqcap_i v_2(a)$ for all $a \in A$. By $\mathcal{V}^2$ we denote the set of two-valued interpretations; they are the $\leq_i$-maximal elements of the meet-semilattice $(\mathcal{V}, \leq_i)$. We denote by $[v]_2$ the set of all two-valued interpretations that extend v.

Two interpretations v_1 and v_2 are *compatible* if $v_1^{\mathbf{t}} \cap v_2^{\mathbf{f}} = v_1^{\mathbf{f}} \cap v_2^{\mathbf{t}} = \emptyset$ and *incompatible* otherwise. A set of interpretations $V \subseteq \mathcal{V}$ is compatible if each pair $v_1, v_2 \in V$ is compatible and incompatible otherwise; its *adm-closure*, $cl(V)$, contains exactly those $v \in \mathcal{V}$ such that $\forall a \in (v^{\mathbf{t}} \cup v^{\mathbf{f}}) \forall v_2 \in [v]_2 \exists v' \in V$ s.t. $v' \leq_i v_2 \wedge v'(a) = v(a)$. We will use $cl(v_1, v_2)$ as shorthand for $cl(\{v_1, v_2\})$.

We define the symmetric distance function $\triangle$ between truth values as follows: $\mathbf{t} \triangle \mathbf{f} = 1$, $\mathbf{t} \triangle \mathbf{u} = \mathbf{f} \triangle \mathbf{u} = \frac{1}{2}$, and $\mathbf{x} \triangle \mathbf{x} = 0$ for $\mathbf{x} \in \{\mathbf{t}, \mathbf{f}, \mathbf{u}\}$. Lifted to interpretations $v_1, v_2 \in \mathcal{V}$, it is defined as $v_1 \triangle v_2 = \sum_{a \in A} v_1(a) \triangle v_2(a)$.

Finally note that we will represent interpretations by sequences of truth values, assuming a total ordering on the underlying vocabulary. For instance the interpretation $\{a \mapsto \mathbf{u}, b \mapsto \mathbf{t}, c \mapsto \mathbf{f}\}$ will be abbreviated by $\mathbf{utf}$.

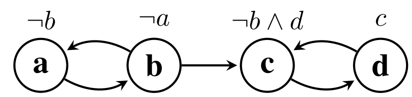

Figure 1: ADF $F = \{\langle a, \neg b \rangle, \langle b, \neg a \rangle, \langle c, \neg b \wedge d \rangle, \langle d, c \rangle\}$.

2.1 Abstract Dialectical Frameworks

An ADF F is a set of tuples $\langle s, \varphi_s \rangle$ where $s \in A$ is a statement and φ_s is a propositional formula over A, the acceptance condition of s. Note that this formalization syntactically differs from the original one [Brewka and Woltran, 2010], where an ADF is represented by a triple (A, L, C) where L is a set of links between statements and C a set of total functions $2^A \mapsto \{\mathbf{t}, \mathbf{f}\}$. It is however easy to see that these two notions are equivalent, as the set of links L is implicitly given by the atoms occurring in the acceptance conditions and the fact that the total functions C can be expressed by propositional formulas. We denote the set of all ADFs by $\mathcal{F}_A$.

The semantics of ADFs can be defined via an operator Γ_F over three-valued interpretations. Given an ADF F and an interpretation v, it is defined as

$$\Gamma_F(v)(a) = \bigsqcap_{w \in [v]_2} w(\varphi_a).$$

Intuitively, the operator returns, for each statements a, the consensus truth value of the evaluation of the acceptance formula φ_a with each two-valued interpretation extending v. The semantics can now be defined as follows:

Definition 1. Given an ADF F, an interpretation v is

- *admissible* for F iff $v \leq_i \Gamma_F(v)$,
- *complete* for F iff $v = \Gamma_F(v)$,
- *preferred* for F iff v is admissible for F and each $v' \in \mathcal{V}$ with $v <_i v'$ is not admissible for F,
- *grounded* for F iff v is complete for F and each $v' \in \mathcal{V}$ with $v' <_i v$ is not complete for F,
- a *(supported) model* of F iff v is two-valued and $v = \Gamma_F(v)$,
- a *stable model* of F iff v is a model of F and $v^t = w^t$, where w is the grounded interpretation of $F^v = \{\langle a, \varphi_a[x/\bot : v(x) = \mathbf{f}] \rangle \mid a \in v^{\mathbf{t}}\}$.

We denote the admissible, complete, preferred, and grounded interpretations for, and supported and stable models of an ADF F by *ad*(F), *co*(F), *pr*(F), *gr*(F), *mo*(F), and *st*(F), respectively. For alternative semantics we refer to [Strass, 2013; Polberg, 2014].

The semantics have been shown to be proper generalizations of AF-semantics [Brewka and Woltran, 2010; Brewka *et al.*, 2013], with both supported and stable models generalizing stable semantics of AFs, differing only in the treatment of support cycles.

Example 1. Consider the ADF F depicted in Figure 1. The admissible interpretations of F are as follows: $ad(F) = \{\mathbf{uuuu}, \mathbf{tfuu}, \mathbf{tfff}, \mathbf{tftt}, \mathbf{ftuu}, \mathbf{ftfu}, \mathbf{ftff}, \mathbf{uuff}\}$. Further observe that $co(F) = ad(F) \setminus \{\mathbf{ftuu}, \mathbf{ftfu}\}$, $pr(F) = \{\mathbf{tfff}, \mathbf{tftt}, \mathbf{ftff}\}$, $gr(F) = \{\mathbf{uuuu}\}$, $mo(F) = pr(F)$, and $st(F) = pr(F) \setminus \{\mathbf{tftt}\}$.

A set of interpretations V is *realizable* under a semantics σ if there is an ADF F with $\sigma(F) = V$. The following proposition recalls results which are either explicitly stated or immediate consequences of [Pührer, 2015] and [Strass, 2015].

Proposition 1. *A set of interpretations V is realizable under*

- *ad iff $V \neq \emptyset$ and $V = cl(V)$;*
- *pr iff $V \neq \emptyset$ and V is incompatible;*
- *mo iff $V \subseteq \mathcal{V}^2$;*
- *st iff $V \subseteq \mathcal{V}^2$ and $v_1^{\mathbf{t}} \not\subseteq v_2^{\mathbf{t}}$, $v_2^{\mathbf{t}} \not\subseteq v_1^{\mathbf{t}}$ for all $v_1, v_2 \in V$.*

For $\sigma \in \{ad, pr, mo, st\}$ it holds that $\sigma(F) \cap \sigma(G)$ is realizable under σ for arbitrary ADFs F and G, given that $\sigma(F) \cap \sigma(G) \neq \emptyset$. This does not hold for co and gr in general.

Definition 2. Given a semantics σ, the function $f_\sigma : 2^{\mathcal{V}} \mapsto \mathcal{F}_A$ maps sets of interpretations to ADFs such that $\sigma(f_\sigma(V)) = V$ if V is realizable under σ and $\sigma(f_\sigma(V)) = \mathbf{u}\ldots\mathbf{u}$ otherwise.

Note that canonical constructions for $f_\sigma(V)$ for realizable sets V can be found in [Pührer, 2015] and [Strass, 2015]. Although f_σ is not unique in general, it is assumed to be fixed for every σ throughout the paper. In particular, $f_\sigma(V) = \{\langle a, \neg a\rangle \mid a \in A\}$ for V not realizable under σ.

2.2 Belief Revision

The most prominent approach to belief revision was introduced by Alchourrón *et al.* [1985] and reformulated for propositional formulas by Katsuno and Mendelzon [1991]. They define an equivalent version of the AGM-postulates for operators $*$ mapping pairs of formulas to a revised formula.

$(R1)$ $\psi * \mu \models \mu$.

$(R2)$ If $\psi \wedge \mu$ is satisfiable, then $\psi * \mu = \psi \wedge \mu$.

$(R3)$ If μ is satisfiable, then $\psi \wedge \mu$ is also satisfiable.

$(R4)$ If $\psi_1 \equiv \psi_2$ and $\mu_1 \equiv \mu_2$, then $\psi_1 * \mu_1 \equiv \psi_2 * \mu_2$.

$(R5)$ $(\psi * \mu) \wedge \phi \models \psi * (\mu \wedge \phi)$.

$(R6)$ If $(\psi*\mu)\wedge\phi$ is satisfiable, then $\psi*(\mu\wedge\phi) \models (\psi*\mu)\wedge\phi$.

While postulates $R1$ to $R4$ are self-explanatory, note that $R5$ and $R6$ ensure that revision is performed with minimal change to the revised formula ψ.

The main result of [Katsuno and Mendelzon, 1991] is that there is a one-to-one correspondence between operators which are rational according to the AGM postulates and functions mapping each formula to a certain binary relation among interpretations. Thus, for constructing an AGM operator, a necessary and sufficient condition is the existence of such a function.

Definition 3. A *preorder $\preceq$ on $\mathcal{V}$* is a reflexive, transitive binary relation on $\mathcal{V}$. The preorder $\preceq$ is

- *total* if $v_1 \preceq v_2$ or $v_2 \preceq v_1$ for any $v_1, v_2 \in \mathcal{V}$,
- *i-max-total* if $v_1 \preceq v_2$ or $v_2 \preceq v_1$ for any $v_1, v_2 \in \mathcal{V}$ with $v_1 \not\leq_i v_2$ and $v_2 \not\leq_i v_1$.

Moreover, for $v_1, v_2 \in \mathcal{V}$, $v_1 \prec v_2$ denotes the strict part of $\preceq$, i.e. $v_1 \preceq v_2$ and $v_2 \not\preceq v_1$. We write $v_1 \approx v_2$ in case $v_1 \preceq v_2$ and $v_2 \preceq v_1$.

Given a preorder, the construction of the corresponding operator is then based on the following selection function:

$$\min(V, \preceq) = \{v_1 \in V \mid \nexists v_2 \in V : v_2 \prec v_1\}.$$

3 Revising ADFs

In this section we apply the AGM approach to the revision of ADFs by studying operators $* : \mathcal{F}_A \times \mathcal{F}_A \mapsto \mathcal{F}_A$. It is inspired by the approach by Diller *et al.* [2015] to revision of AFs. We begin by reformulating the postulates for our setting, parameterized by the used semantics.

$(A1_\sigma)$ $\sigma(F * G) \subseteq \sigma(G)$.

$(A2_\sigma)$ If $\sigma(F) \cap \sigma(G) \neq \emptyset$, then $\sigma(F * G) = \sigma(F) \cap \sigma(G)$.

$(A3_\sigma)$ If $\sigma(G) \neq \emptyset$, then $\sigma(F * G) \neq \emptyset$.

$(A4_\sigma)$ If $\sigma(G) = \sigma(H)$, then $\sigma(F * G) = \sigma(F * H)$.

$(A5_\sigma)$ $\sigma(F * G) \cap \sigma(H) \subseteq \sigma(F * f_\sigma(\sigma(G) \cap \sigma(H)))$.

$(A6_\sigma)$ If $\sigma(F * G) \cap \sigma(H) \neq \emptyset$, then $\sigma(F * f_\sigma(\sigma(G) \cap \sigma(H))) \subseteq \sigma(F * G) \cap \sigma(H)$.

Next we define two types of rankings which will be the counterpart to the postulates in the representation results.

Definition 4. Given a semantics σ and an ADF F, a preorder $\preceq_F$ is a *(i-max-)faithful ranking* for F if it is (i-max-)total and for all (incompatible) interpretations $v_1, v_2 \in \mathcal{V}$ it holds that

(i) if $v_1, v_2 \in \sigma(F)$ then $v_1 \approx_F v_2$, and

(ii) if $v_1 \in \sigma(F)$ and $v_2 \notin \sigma(F)$ then $v_1 \prec_F v_2$.

A function mapping each ADF to a (i-max-)faithful ranking is called *(i-max-)faithful assignment*.

3.1 Revision under Preferred Semantics

In the remainder of this section we will focus on the preferred semantics. To fulfill the postulates, a revision operator will have to result in an ADF having certain preferred interpretations. However, as can be already seen by Proposition 1, preferred semantics (and all the others) underlie certain limits in terms of expressiveness. That is, certain desired outcomes may not be realizable. We first give sufficient conditions for realizability we will make use of in the following.

Proposition 2. *A set of interpretations $V \subseteq \mathcal{V}$ is realizable under pr if one of the following holds:*

1. *$V \subseteq pr(F)$ and $V \neq \emptyset$ for some $F \in \mathcal{F}_A$;*
2. *$V = \{v_1, v_2\}$ and v_1 and v_2 are incompatible; or*
3. *$V = \{v\}$.*

The following example shows that the standard set of postulates is not enough to get a correspondence to preorders on interpretations.

Example 2. Consider an arbitrary ADF F and the binary relation $\preceq$ having $pr(F)$ as least elements, containing the cycle $\mathbf{uft} \prec \mathbf{ttf} \prec \mathbf{fut} \prec \mathbf{tuf} \prec \mathbf{uft}$ and being a linear order otherwise. Note that $\preceq$ is not transitive and therefore only a pseudo-preorder. However, the revision operator $*$ induced by $\preceq$ can be shown to satisfy all postulates $A1_{pr}-A6_{pr}$. Moreover, every binary relation $\preceq'$ inducing the same operator $*$ must contain this cycle. Consider the pair of interpretations $\mathbf{uft}$ and $\mathbf{ttf}$. They are incompatible, hence realizable (cf. Proposition 1). Therefore the revision of F by $f_{pr}(\{\mathbf{uft}, \mathbf{ttf}\})$ must have $\mathbf{uft}$ as single preferred interpretation, hence $\mathbf{uft} \prec' \mathbf{ttf}$. This holds for every neighboring pair of the cycle, hence $\preceq'$ must contain the same cycle.

The following postulate, which is adapted from Delgrande and Peppas [2015], closes this gap. Note that it is redundant in classical AGM revision.

($Acyc_\sigma$) If for $1 \leq i < n$, $\sigma(F * G_{i+1}) \cap \sigma(G_i) \neq \emptyset$ and $\sigma(F * G_1) \cap \sigma(G_n) \neq \emptyset$ then $\sigma(F * G_n) \cap \sigma(G_1) \neq \emptyset$.

We are now ready to give our first representation result.

Theorem 1. *Let F be an* ADF *and $\preceq_F$ an i-max-faithful ranking for F. Define an operator $* : \mathcal{F}_A \times \mathcal{F}_A \mapsto \mathcal{F}_A$ by*

$$F * G = f_{pr}(\min(pr(G), \preceq_F)).$$

Then $*$ satisfies postulates $A1_{pr} - A6_{pr}$ and $Acyc_{pr}$.

Proof. By definition of f_{pr} and the fact that any non-empty $V \subseteq pr(G)$ is realizable under pr (cf. Proposition 2.1) it holds that $pr(f_{pr}(\min(pr(G), \preceq_F))) = \min(pr(G), \preceq_F)$, hence $pr(F * G) = \min(pr(G), \preceq_F)$. Thus $A1_{pr}$ and $A4_{pr}$ follow.

For $A2_{pr}$, assume $pr(F) \cap pr(G) \neq \emptyset$. Since $\preceq_F$ is i-max-faithful and all $v_1, v_2 \in pr(G)$ are pairwise incompatible we get that $\min(pr(G), \preceq_F) = pr(F) \cap pr(G)$ and hence $pr(F * G) = pr(F) \cap pr(G)$.

As $\preceq_F$ is transitive (by being a preorder) and A is finite, $\min(pr(G), \preceq_F) \neq \emptyset$, hence $A3_{pr}$ holds.

For $A5_{pr}$ and $A6_{pr}$ we consider the non-trivial case where $pr(F * G) \cap pr(H) \neq \emptyset$. Recalling that $pr(F) \cap pr(G)$ is realizable under pr (cf. Proposition 1), we have to show that $\min(pr(G), \preceq_F) \cap pr(H) = \min(pr(G) \cap pr(H), \preceq_F)$. Towards a contradiction assume there is some $v \in \min(pr(G), \preceq_F) \cap pr(H)$ such that $v \notin \min(pr(G) \cap pr(H), \preceq_F)$. As then $v \in pr(G)$ and $v \in pr(H)$ there must be some $v' \in pr(G) \cap pr(H)$ with $v' \prec_F v$, contradicting $v \in \min(pr(G), \preceq_F)$. On the other hand assume, again to the contrary, that there is some $v \in \min(pr(G) \cap pr(H), \preceq_F)$ such that $v \notin \min(pr(G), \preceq_F) \cap pr(H)$. From $v \in pr(H)$ we get $v \notin \min(pr(G), \preceq_F)$. As by assumption $pr(F * G) \cap prf(H) \neq \emptyset$, let $v' \in \min(pr(G), \preceq_F)$ and $v' \in pr(H)$. Then also $v' \in pr(G) \cap pr(H)$. Since $v, v' \in pr(H)$, v and v' are incompatible, $\preceq_F$ is i-max-total and $v \in \min(pr(G) \cap pr(H), \preceq_F)$ by assumption, we get $v \preceq_F v'$. Thus $v \in \min(pr(G), \preceq_F)$ because $v' \in \min(pr(G), \preceq_F)$, a contradiction.

For $Acyc_{pr}$ consider a sequence of ADFs $G_1, \ldots, G_n$ such that $pr(F * G_{i+1}) \cap pr(G_i) \neq \emptyset$ for $1 \leq i < n$ and $pr(F * G_1) \cap pr(G_n) \neq \emptyset$. Let $1 \leq i < n$. By definition of $*$ we have $pr(f_{pr}(\min(pr(G_{i+1}), \preceq_F))) \cap pr(G_i) \neq \emptyset$. Then, by Proposition 2, $\min(pr(G_{i+1}), \preceq_F) \cap pr(G_i) \neq \emptyset$ follows. Hence there is some $v'_i \in pr(G_i)$ such that $v'_i \preceq_F v_{i+1}$ for all $v_{i+1} \in pr(G_{i+1})$. From transitivity of $\preceq_F$ we infer that there is a $v'_1 \in pr(G_1)$ such that $v'_1 \preceq_F v_n$ for all $v_n \in pr(G_n)$. From $pr(F * G_1) \cap pr(G_n) \neq \emptyset$ it follows that there is some $v''_1 \in \min(G_1, \preceq_F)$ (hence also $v''_1 \in pr(G_1)$ and $v''_1 \preceq_F v'_1$) with $v''_1 \in pr(G_n)$. We have $v''_1 \preceq_F v'_1 \preceq_F v_n$ (for each $v_n \in pr(G_n)$), hence $v''_1 \in \min(pr(G_n), \preceq_F)$. This together with $v''_1 \in pr(G_0)$ means that $pr(F * G_n) \cap pr(G_1) \neq \emptyset$, which was to show. □

Theorem 2. *Let $* : \mathcal{F}_A \times \mathcal{F}_A \mapsto \mathcal{F}_A$ be a revision operator satisfying postulates $A1_{pr} - A6_{pr}$ and $Acyc_{pr}$. Then there is an assignment mapping each* ADF *F to an i-max-faithful ranking $\preceq$ such that $pr(F * G) = \min(pr(G), \preceq)$ for every* ADF *G.*

Proof. Assume an arbitrary ADF F. We will define $\preceq$ and show that it is an i-max-faithful ranking for F and that $pr(F * G) = \min(pr(G), \preceq)$.

First let $\preceq'$ be the relation on $\mathcal{V}$ such that for each $v \in \mathcal{V}$, $v \approx' v$, and for any incompatible interpretations $v_1, v_2 \in \mathcal{V}$,

$$v_1 \preceq' v_2 \Leftrightarrow v_1 \in pr(F * f_{pr}(\{v_1, v_2\})).$$

The relation $\preceq$ is defined as the transitive closure of $\preceq'$:

$$v \preceq v' \Leftrightarrow \exists w_1, \ldots, w_n : v \preceq' w_1 \preceq' \cdots \preceq' w_n \preceq' v'.$$

First, $\preceq$ is clearly reflexive and transitive, making it a preorder on $\mathcal{V}$. Moreover, for incompatible interpretations $v_1, v_2 \in \mathcal{V}$ we know from Proposition 2.2 that $\{v_1, v_2\}$ is realizable under pr, hence $pr(f_{pr}(\{v_1, v_2\})) = \{v_1, v_2\}$. By $A1_{pr}$ and $A3_{pr}$ we therefore get that either $v_1 \preceq' v_2$ or $v_2 \preceq' v_1$, and, consequently, also $v_1 \preceq v_2$ or $v_2 \preceq v_1$, hence $\preceq$ is i-max-total.

We proceed by showing that $\preceq$ is i-max-faithful. To show (i), let $v_1, v_2 \in pr(F)$ and note that $pr(f_{pr}(\{v_1, v_2\})) = \{v_1, v_2\}$. Hence, by $A2_{pr}$, we get $pr(F * f_{pr}(\{v_1, v_2\})) = \{v_1, v_2\}$. Therefore, by definition of $\preceq$, $v_1 \preceq v_2$ and $v_2 \preceq v_1$, i.e. $v_1 \approx v_2$. For (ii), we begin with the intermediate statement

$$\text{for } v_1 \ldots v_n \in \mathcal{V} : v_1 \preceq' \cdots \preceq' v_n \preceq' v_1 \Rightarrow v_1 \preceq' v_n \quad (1)$$

For $n \leq 2$ the statement is immediate. Assume $n > 2$. By definition of $\preceq'$ we first get that v_i and v_{i+1} for $i \in \{1, \ldots, n-1\}$ as well as v_n and v_1 are incompatible, hence $f_{pr}(\{v_i, v_{i+1}\}) = \{v_i, v_{i+1}\}$ and $f_{pr}(\{v_n, v_1\}) = \{v_n, v_1\}$ by Proposition 2. Moreover, we get $v_i \in pr(F * f_{pr}(\{v_i, v_{i+1}\}))$ for $i \in \{1, \ldots, n-1\}$ and $v_n \in pr(F * f_{pr}(\{v_n, v_1\}))$. It follows that $v_1 \in pr(F * f_{pr}(\{v_1, v_2\})) \cap \{v_n, v_1\}$, $v_i \in pr(F * f_{pr}(\{v_i, v_{i+1}\})) \cap \{v_{i-1}, v_i\}$ for $i \in \{2, \ldots, n-1\}$, and $v_n \in pr(F * f_{pr}(\{v_n, v_1\})) \cap \{v_{n-1}, v_n\}$. Considering $Acyc$ we get $pr(F * f_{pr}(\{v_n, v_1\})) \cap \{v_1, v_2\} \neq \emptyset$, meaning further by $A5_{pr}$ and $A6_{pr}$ that $pr(F * f_{pr}(\{v_n, v_1\})) \cap \{v_1, v_2\} = pr(F * f_{pr}(\{v_n, v_1\} \cap \{v_1, v_2\})) = pr(F * f_{pr}(\{v_1\}))$. By $pr(f_{pr}(\{v_1\})) = \{v_1\}$ (cf. Proposition 2.3), $A1_{pr}$ and $A3_{pr}$, we follow that $v_1 \in pr(F * f_{pr}(\{v_n, v_1\}))$, meaning that $v_1 \preceq' v_n$, concluding the proof for (1). We proceed by showing the statement

$$\text{for } v_1, v_2 \in \mathcal{V} : v_1 \prec' v_2 \Rightarrow v_1 \prec v_2 \quad (2)$$

$v_1 \preceq v_2$ is clear by definition. Assume, towards a contradiction, that $v_2 \preceq v_1$. Then $\exists w_1, \ldots, w_n$ such that $v_1 \preceq'$

$w_1 \preceq' \cdots \preceq' w_n \preceq' v_2$. As by assumption $v_1 \preceq' v_2$ we follow by (1) that $v_2 \preceq' v_1$, a contradiction to $v_1 \prec' v_2$, showing (2). Now let v_1 and v_2 be incompatible interpretations such that $v_1 \in pr(F)$ and $v_2 \notin pr(F)$. By $A2_{pr}$ we get $pr(F * f_{pr}(\{v_1, v_2\})) = pr(F) \cap \{v_1, v_2\} = \{v_1\}$, implying $v_1 \preceq' v_2$. Therefore, by (2), also $v_1 \preceq v_2$, showing (ii) and, consequently, that $\preceq$ is i-max-faithful.

Before showing that $*$ is indeed simulated by $\preceq$, we prove

$$\text{for } v_1, v_2 \in \mathcal{V} \text{ s.t. } v_1 \preceq' v_2, G \in \mathcal{F}_A : \quad v_1 \in pr(G) \wedge v_2 \in pr(F * G) \Rightarrow v_1 \in pr(F * G) \tag{3}$$

Let $G \in \mathcal{F}_A$ such that $v_1 \in pr(G)$ and $v_2 \in pr(F * G)$. First note that, by $*$ fulfilling $A1_{pr}$, $v_2 \in pr(G)$, meaning that v_1 and v_2 are incompatible and therefore $pr(f_{pr}(\{v_1, v_2\})) = \{v_1, v_2\}$. From $A5_{pr}$ and $A6_{pr}$ we then get that $pr(F * G) \cap \{v_1, v_2\} = pr(F * f_{pr}(pr(G) \cap \{v_1, v_2\})) = pr(F * f_{pr}(\{v_1, v_2\}))$. By the assumption that $v_1 \preceq' v_2$ it holds that $v_1 \in pr(F * f_{pr}(\{v_1, v_2\}))$, hence (3) follows. The last intermediate step is to show that

$$\text{for } G \in \mathcal{F}_A : \min(pr(G), \preceq) = \min(pr(G), \preceq') \tag{4}$$

Consider some $G \in \mathcal{F}_A$. ($\subseteq$) Let $v_1 \in \min(pr(G), \preceq)$ and suppose there exists an $v_2 \in pr(G)$ with $v_2 \prec' v_1$. This means, by (2), that also $v_2 \preceq v_1$, a contradiction. Hence $v_2 \not\prec' v_1$ for all $v_2 \in pr(G)$, i.e. $v_1 \in \min(pr(G), \preceq')$. ($\supseteq$) Let $v_1 \in \min(pr(G), \preceq')$ and $v_2 \in pr(G)$. We show that $v_1 \preceq' v_2$, since then $v_1 \preceq v_2$ and, consequently, $v_1 \in \min(pr(G), \preceq)$ follows by definition of $\preceq$. If $v_1 = v_2$ we have $v_1 \preceq' v_2$ by definition of $\preceq'$. If $v_1 \neq v_2$ observe that, by $v_1, v_2 \in pr(G)$, v_1 and v_2 are incompatible, hence at least one of $v_1 \preceq' v_2$ and $v_2 \preceq' v_1$ must hold. By $v_1 \in \min(pr(G), \preceq')$ it cannot hold that $v_2 \prec' v_1$, hence $v_1 \preceq' v_2$.

We are now ready to show that, for any ADF G, $pr(F*G) = \min(pr(G), \preceq)$. Considering (4) we just have to show that

$$\text{for } G \in \mathcal{F}_A : pr(F * G) = \min(pr(G), \preceq') \tag{5}$$

($\subseteq$) Let $v \in pr(F * G)$ and keep in mind that, by $A1_{pr}$, also $v \in pr(G)$. We show for each $w \in pr(G)$ that $v \preceq' w$. Consider an arbitrary $w \in pr(G)$. Note that by $v, w \in pr(G)$ we have that $pr(f_{pr}(\{v, w\})) = \{v, w\}$. From $A5_{pr}$ and $A6_{pr}$ we get $pr(F * G) \cap \{v, w\} = pr(F * f_{pr}(pr(G) \cap \{v, w\})) = pr(F * f_{pr}(\{v, w\}))$. As by assumption $v \in pr(F * G)$ we get $v \preceq' w$ by definition of $\preceq'$. ($\supseteq$) Towards a contradiction, assume some $v \in \min(pr(G), \preceq')$ such that $v \notin pr(F * G)$ (again note that also $v \in pr(G)$ by $A1_{pr}$). By $A3_{pr}$ and the fact that $pr(G) \neq \emptyset$ there is some $w \in pr(F*G)$. From (3) we infer that $v \not\preceq' w$. But the by assumption also $w \not\prec' v$. Since v and w must be incompatible by $v, w \in pr(G)$ this means $pr(F * f_{pr}(\{v, w\})) \cap \{v, w\} = \emptyset$ and by $*$ fulfilling $A1_{pr}$ even $pr(F * f_{pr}(\{v, w\})) = \emptyset$, a contradiction to $*$ satisfying $A3_{pr}$. □

With Theorems 1 and 2 we have obtained a one-to-one correspondence between i-max-faithful rankings and revision operators satisfying all postulates. In particular, we can use standard revision operators from the literature which work on faithful rankings (as each faithful ranking is also i-max-faithful) to get concrete revision operators. To exemplify the obtained result, we consider Dalal's operator (1988), customized to the three-valued setting (using the same distance measure as, for instance, in [Arieli, 2008]):

Definition 5. Given an ADF F and semantics σ, the ranking $\preceq^\sigma_F$ based on three-valued distance is defined as

$$v_1 \preceq^\sigma_F v_2 \Leftrightarrow \min_{v \in \sigma(F)} (v \triangle v_1) \leq \min_{v \in \sigma(F)} (v \triangle v_2).$$

for each $v_1, v_2 \in \mathcal{V}$. The operator $*^D_\sigma$ induced by $\preceq^\sigma_F$ returns $F *^D_\sigma G = f_{pr}(\min(\sigma(G), \preceq^\sigma_F))$ for each $G \in \mathcal{F}_A$.

It is easy to see that $\preceq^\sigma_F$ is i-max-faithful, as the minimal distance to $\sigma(F)$ is 0 for interpretations $v \in \sigma(F)$ and greater than 0 for interpretations $v \notin \sigma(F)$. Hence, by Theorem 1, $*^D_\sigma$ satisfies all postulates.

Example 3. Consider the ADF $F = \{\langle a, a\rangle, \langle b, a\rangle, \langle c, \neg a \wedge b\rangle\}$, having $pr(F) = \{\mathbf{ttf}, \mathbf{fff}\}$. First note that the minimal elements of $\preceq^{pr}_F$ coincide with $pr(F)$, i.e. $\mathbf{ttf} \approx^{pr}_F \mathbf{fff} \prec^{pr}_F$ *others*. Now consider the revision by the ADF G having $pr(G) = \{\mathbf{tft}, \mathbf{ttu}, \mathbf{ffu}\}$ and observe that $\mathbf{ttu} \approx^{pr}_F \mathbf{ffu} \prec^{pr}_F \mathbf{tft}$ ($\mathbf{ttu}$ and $\mathbf{ffu}$ have minimal distance to $pr(F)$ of $\frac{1}{2}$, while $\mathbf{tft}$ has 2). Therefore we get $F *^D_{pr} G = f_{pr}(\{\mathbf{ttu}, \mathbf{ffu}\})$. On the other hand consider the ADF $G' = \{\langle a, \top\rangle, \langle b, \neg a\rangle, \langle c, \neg b\rangle\}$, having $pr(G') = \{\mathbf{tft}\}$. The revision of F by G' obviously results in an ADF also having $\mathbf{tft}$ – minimal distance 2 to $pr(F)$ – as only preferred interpretation. Inspecting the set of admissible interpretations of G', which can be seen as reasonable (but not maximal) positions in the revising ADF, $ad(G') = \{\mathbf{tft}, \mathbf{tfu}, \mathbf{tuu}, \mathbf{uuu}\}$, we observe that it contains elements which are closer to $pr(F)$ than $\mathbf{tft}$. In particular, the interpretation $\mathbf{tuu}$ has distance 1 to $pr(F)$ and is even admissible in F.

3.2 Revision under Admissible Semantics

Example 3 suggests to take the admissible interpretations into account when revising with respect to the preferred interpretations. A quite radical step would be to just revise with respect to admissible interpretations instead. By the fact that $ad(F_1) \cap ad(F_2) \neq \emptyset$ for all ADFs $F_1, F_2 \in \mathcal{F}_A$ we get only one operator satisfying postulate $A2_{ad}$ and the following result immediately follows:

Theorem 3. *An operator $* : \mathcal{F}_A \times \mathcal{F}_A \mapsto \mathcal{F}_A$ fulfills $A1_{ad}$ – $A6_{ad}$ iff $*$ is defined as $F * G = f_{ad}(ad(F) \cap ad(G))$.*

It is important to note that admissible semantics is closed under intersection (cf. Proposition 1), therefore $f_{ad}(ad(F) \cap ad(G))$ always realizes $ad(F) \cap ad(G)$.

Example 4. Again consider the ADFs F and G' from Example 3 and note that $ad(F) = \{\mathbf{ttf}, \mathbf{fff}, \mathbf{ttu}, \mathbf{tuf}, \mathbf{ffu}, \mathbf{tuu}, \mathbf{fuu}, \mathbf{uuu}\}$ and $ad(G') = \{\mathbf{tft}, \mathbf{tfu}, \mathbf{tuu}, \mathbf{uuu}\}$. Moreover, let $*_{ad}$ be the operator from Theorem 3. As expected, we get $F *_{ad} G' = f_{ad}(\{\mathbf{tuu}, \mathbf{uuu}\})$, i.e. the resulting ADF has $\mathbf{tuu}$ as single preferred interpretation, which was somehow seen as one of the more desired scenarios in Example 3.

But now consider the ADF G'' having $ad(G'') = \{\mathbf{utf}, \mathbf{uuu}\}$ and observe that $F *_{ad} G'' = f_{ad}(\{\mathbf{uuu}\})$. From the perspective of the preferred interpretations of F (being $\{\mathbf{ttf}, \mathbf{fff}\}$) this might not be desired, as $\mathbf{utf}$ is admissible in G'' and has a distance of only $\frac{1}{2}$ to $pr(F)$, while the result of the revision has distance $\frac{3}{2}$.

3.3 Hybrid Approach

Due to the problem illustrated in Example 4 we are interested in operators selecting out of the admissible interpretations of the revising ADF (in a sense accepting all reasonable positions as valid outcomes of the revision), but basing the amount of change on the preferred interpretations of the original ADF. To this end we reformulate the postulates to this setting:

($P1$) $pr(F \star G) \subseteq ad(G)$.

($P2$) If $pr(F) \cap ad(G) \neq \emptyset$, then $pr(F \star G) = pr(F) \cap ad(G)$.

($P3$) If $ad(G) \neq \emptyset$, then $pr(F \star G) \neq \emptyset$.

($P4$) If $ad(G) = ad(H)$, then $pr(F \star G) = pr(F \star H)$.

($P5$) $pr(F \star G) \cap ad(H) \subseteq pr(F \star f_{ad}(ad(G) \cap ad(H)))$.

($P6$) If $pr(F \star G) \cap ad(H) \neq \emptyset$, then $pr(F \star f_{ad}(ad(G) \cap ad(H))) \subseteq pr(F \star G) \cap ad(H)$.

(*Acyc*) If for $1 \leq i < n$, $pr(F * G_{i+1}) \cap ad(G_i) \neq \emptyset$ and $pr(F * G_1) \cap ad(G_n) \neq \emptyset$ then $pr(F * G_n) \cap ad(G_1) \neq \emptyset$.

As admissible semantics may give pairwise compatible interpretations, we will not restrict ourselves to i-max-faithful rankings for the representation result. However, we face another challenge, as illustrated in the following example.

Example 5. Consider the ranking $\mathbf{ff} \prec \mathit{others} \prec \mathbf{tu} \approx \mathbf{ut} \prec \mathbf{tt} \prec \mathbf{uu}$ and the ADFs $F = \{\langle a, \bot\rangle, \langle b, \bot\rangle\}$, $G = \{\langle a, \top\rangle, \langle b, \top\rangle\}$, and $H = \{\langle a, \neg a \vee b\rangle, \langle b, a \vee \neg b\rangle\}$. We have $pr(F) = \{\mathbf{ff}\}$, $ad(G) = \{\mathbf{uu}, \mathbf{ut}, \mathbf{tu}, \mathbf{tt}\}$, and $ad(H) = \{\mathbf{uu}, \mathbf{tt}\}$. It can be seen that $\preceq$ is a faithful ranking for F. However, the revision operator $\star$ induced by $\preceq$ gives us $F \star G = f_{pr}(\{\mathbf{ut}, \mathbf{tu}\})$ and we further get

- $pr(F \star G) \cap ad(H) = \{\mathbf{uu}\}$, but
- $pr(F \star f_{ad}(ad(G) \cap ad(H))) = \{\mathbf{tt}\}$.

Therefore $\star$ violates $P5$. The problem is somehow hidden in the fact that $\mathbf{ut}$ and $\mathbf{tu}$ are compatible. That is, the set of interpretations $\{\mathbf{ut}, \mathbf{tu}\}$ cannot be realized under preferred semantics, hence $pr(f_{pr}(\{\mathbf{ut}, \mathbf{tu}\})) = \{\mathbf{uu}\}$.

To overcome this issue we introduce the concept of compliance, generalizing similar notions from [Delgrande *et al.*, 2013; Delgrande and Peppas, 2015; Diller *et al.*, 2015].

Definition 6. A preorder $\preceq$ is σ-τ-compliant if, for every ADF $F \in \mathcal{F}_A$, $\min(\tau(F), \preceq)$ is realizable under σ.

In general, this condition depends on the concrete capabilities in terms of realizability of σ and τ. Fortunately, we can capture *pr-ad*-compliance with conditions on the ranking.

Proposition 3. *A preorder $\preceq$ is pr-ad-compliant iff: if $v_1, v_2 \in \mathcal{V}$ are compatible and $v_1 \approx v_2$ then $\exists v_3 \in cl(v_1, v_2) : v_3 \prec v_1, v_2$.*

We will make use of the following properties of the adm-closure in the following results.

Lemma 1. *For each $V, V_1, V_2 \subseteq \mathcal{V}$ and $v, v' \in \mathcal{V}$ it holds:*

1. $cl(V) = cl(cl(V))$ *(*idempotence*)*
2. $V_1 \subseteq V_2 \Rightarrow cl(V_1) \subseteq cl(V_2)$ *(*monotonicity*)*
3. $\forall v'' \in cl(v, v') : cl(v, v'') \subseteq cl(v, v')$.

Proof. Note that $V \subseteq cl(V)$ for any $V \subseteq \mathcal{V}$ is clear by definition. (1) $cl(V) \subseteq cl(cl(V))$ follows from the initial observation. Assume there is some $v \in cl(cl(V))$ with $v \notin cl(V)$. The latter means that $\exists a \in (v^{\mathbf{t}} \cup v^{\mathbf{f}}) \exists v_2 \in [v]_2$ s.t. $\nexists v' \in V : v' \leq_i v_2 \wedge v'(a) = v(a)$. Now for this particular a and v_2 it holds, by $v \in cl(cl(V))$, that $\exists w \in cl(V) : w \leq_i v_2 \wedge w(a) = v(a)$. Hence $\exists w' \in V : w' \leq_i v_2 \wedge w'(a) = w(a)$. We have $w'(a) = w(a) = v(a)$, a contradiction.
(2) Let $v \in cl(V_1)$ and consider some $a \in v^{\mathbf{t}} \cup v^{\mathbf{f}}$ and $v_2 \in [v]_2$. There is some $v' \in V_1$ s.t. $v' \leq_i v_2$ and $v(a) = v'(a)$. As $V_1 \subseteq V_2$ by assumption, also $v \in V_2$, hence $v \in cl(V_2)$.
(3) Consider some $v'' \in cl(v, v')$, i.e. $\forall a \in v''^{\mathbf{t}} \cup v''^{\mathbf{f}} \forall v_2 \in [v'']_2 (v \leq_i v_2 \wedge v(a) = v''(a)) \vee (v' \leq_i v_2 \wedge v'(a) = v''(a))$. Assume there is some $w \in cl(v, v'')$ and $w \notin cl(v, v')$. The latter means that $\exists a \in w^{\mathbf{t}} \cup w^{\mathbf{f}} \exists w_2 \in [w]_2$ s.t. $\neg(v \leq_i w_2 \wedge v(a) = w(a)) \wedge \neg(v' \leq_i w_2 \wedge v'(a) = w(a))$. Hence, by $w \in cl(v, v'')$, we get for this particular a and w_2 that $v'' \leq_i w_2$ and $v''(a) = w(a)$. From $a \in w^{\mathbf{t}} \cup w^{\mathbf{f}}$ and $v''(a) = w(a)$ it follows that $a \in v''^{\mathbf{t}} \cup v''^{\mathbf{f}}$ and from $v'' \leq_i w_2$ we get $w_2 \in [v'']_2$. Therefore, from $v'' \in cl(v, v')$ and $\neg(v' \leq_i w_2 \wedge v'(a) = w(a))$, we get $v \leq_i w_2$ and $v(a) = v''(a)$ and, consequently, $v(a) = w(a)$, a contradiction. ☐

We now show the representation result for our hybrid operators which work on the admissible interpretations of the revising ADF but basing the distance measure on the preferred interpretations of the original ADF. The first direction follows similar to Theorem 1 with the help of *pr-ad*-compliance.

Theorem 4. *Let F be an* ADF *and $\preceq_F$ a pr-ad-compliant, faithful ranking for F. Define operator $\star : \mathcal{F}_A \times \mathcal{F}_A \mapsto \mathcal{F}_A$ by $F \star G = f_{pr}(\min(ad(G), \preceq_F))$. Then $\star$ satisfies postulates $P1 - P6$ and Acyc.*

Theorem 5. *Let $\star$ be a revision operator satisfying $P1 - P6$ and Acyc. Then there is an assignment mapping each* ADF *F to a faithful ranking $\preceq$ for F that is pr-ad-compliant and $pr(F \star G) = \min(ad(G), \preceq)$ for every* ADF *G.*

Proof. Given a revision operator $\star$ satisfying $P1 - P6$ and *Acyc*, let F be an arbitrary ADF. We will gradually define the ranking $\preceq$ and show that it is faithful and *pr-ad*-compliant and it indeed simulates $\star$. First, we define $\preceq'$ as

$$v_1 \preceq' v_2 \Leftrightarrow v_1 \in pr(F \star f_{ad}(cl(v_1, v_2)))$$

for each $v_1, v_2 \in \mathcal{V}$. Note that $\preceq'$ is reflexive, but neither transitive nor total. This is because there might be interpretations $v_1, v_2 \in \mathcal{V}$ for which $pr(F \star f_{ad}(cl(v_1, v_2))) \cap \{v_1, v_2\} = \emptyset$ due to $cl(v_1, v_2) \supset \{v_1, v_2\}$. After showing three properties of $\preceq'$ we will extend it first to the transitive $\preceq^t$ and then to the desired ranking $\preceq$.

$$\begin{aligned} &\text{for } v_1, v_2 \in \mathcal{V} \text{ s.t. } v_1 \preceq' v_2, G \in \mathcal{F}_A : \\ &\quad v_1 \in ad(G) \wedge v_2 \in pr(F \star G) \Rightarrow v_1 \in pr(F \star G) \end{aligned} \tag{6}$$

Let $G \in \mathcal{F}_A$, $v_1 \in ad(G)$, $v_2 \in pr(F \star G)$ with $v_1 \preceq' v_2$. First, we get $v_2 \in ad(G)$ from $P1$. Moreover, from $P5$ and $P6$ we get $pr(F \star G) \cap cl(v_1, v_2) = pr(F \star f_{ad}(ad(G) \cap cl(v_1, v_2))$. As both $v_1, v_2 \in ad(G)$ we get that $cl(v_1, v_2) \subseteq ad(G)$ from Lemma 1.2, hence $pr(F \star G) \cap cl(v_1, v_2) = pr(F \star f_{ad}(cl(v_1, v_2))$. Now as $v_1 \preceq' v_2$ by assumption it must hold that $v_1 \in pr(F \star f_{ad}(v_1, v_2))$, hence $v_1 \in pr(F \star G)$.

We proceed with

$$\text{for } G \in \mathcal{F}_A : \min(ad(G), \preceq') = pr(F \star G) \qquad (7)$$

($\subseteq$): To the contrary, assume some $v_1 \in \min(ad(G), \preceq')$ with $v_1 \notin pr(F \star G)$. From $P3$ we know $pr(F \star G) \neq \emptyset$, so assume an arbitrary $v_2 \in pr(F \star G)$. From (6) we follow that $v_1 \not\preceq' v_2$ and, consequently, from $v_1 \in \min(ad(G, \preceq')$ also $v_2 \not\preceq' v_1$. By the definition of $\preceq'$ and considering $P3$ there must then be some $v_3 \in pr(F \star f_{ad}(cl(v_1, v_2)))$. From $P1$ it follows that $v_3 \in f_{ad}(cl(v_1, v_2))$, i.e. $v_3 \in cl(v_1, v_2)$. Then from $P5$ and $P6$ we get $pr(F \star f_{ad}(cl(v_1, v_2))) \cap cl(v_1, v_3) = pr(F \star f_{ad}(cl(v_1, v_2) \cap cl(v_1, v_3)))$. From Lemma 1.3 it follows that $cl(v_1, v_3) \subseteq cl(v_1, v_2)$, hence $pr(F \star f_{ad}(cl(v_1, v_2))) \cap cl(v_1, v_3) = pr(F \star f_{ad}(cl(v_1, v_3)))$. Therefore $v_1 \notin pr(F \star f_{ad}(cl(v_1, v_3)))$ and $v_3 \in pr(F \star f_{ad}(cl(v_1, v_3)))$, hence $v_3 \prec' v_1$. Finally, note that $cl(v_1, v_2) \subseteq ad(G)$, hence $v_3 \in ad(G)$ contradicting $v_1 \in \min(ad(G, \preceq')$. ($\supseteq$): Let $v_1 \in pr(F \star G)$ and consider an arbitrary $v_2 \in ad(G)$. Observing $v_1 \in ad(G)$ by $P1$ we get $pr(F \star G) \cap cl(v_1, v_2) = pr(F \star f_{ad}(ad(G) \cap cl(v_1, v_2)))$ by $P5$ and $P6$. Moreover, $cl(v_1, v_2) \subseteq ad(G)$ by Lemma 1.2, hence $pr(F \star G) \cap cl(v_1, v_2) = pr(F \star f_{ad}(cl(v_1, v_2)))$ and, consequently, $v_1 \in pr(F \star f_{ad}(cl(v_1, v_2)))$, meaning $v_1 \preceq' v_2$. Therefore, recalling that v_2 was chosen arbitrarily, $v_1 \in \min(ad(G), \preceq')$.

The following can be shown similarly as (1).

$$\text{for } v_1, \ldots, v_n \in \mathcal{V} : v_1 \preceq' \cdots \preceq' v_n \preceq' v_1 \Rightarrow v_1 \preceq' v_n \qquad (8)$$

Now we define $\preceq^t$ to be the transitive closure of $\preceq'$. As a consequence of (8) we infer

$$\text{for } v_1, v_2 \in \mathcal{V} : v_1 \prec' v_2 \Rightarrow v_1 \prec^t v_2 \qquad (9)$$

Defining, for any set of interpretations V, $\max(V, \preceq^t)$ as the set $\{v_1 \in V \mid \nexists v_2 \in V : v_1 \prec^t v_2\}$ we get, by (8) and the fact that $\mathcal{V}$ is finite, that

$$\text{for } V \subseteq \mathcal{V} : V \neq \emptyset \Rightarrow \max(V, \preceq^t) \neq \emptyset \qquad (10)$$

We are now ready to define $\preceq$. To this end consider the sequence of sets of interpretations $V_0, V_1, \ldots$ defined as

$$V_0 = \max(\mathcal{V}, \preceq^t),$$
$$V_1 = \max\left(\mathcal{V} \setminus V_0, \preceq^t\right),$$
$$V_i = \max(\mathcal{V} \setminus \bigcup_{0 \leq j < i} V_j, \preceq^t) \text{ for } i > 1.$$

Since $\mathcal{V}$ is finite we conclude from (10) that the sequence will reach the empty set of interpretations at some point and each of the following elements will also be empty. The sequence $V_1, \ldots, V_m$ of non-empty sets of interpretation then forms a partition of $\mathcal{V}$. Based on this we define $\preceq$ as

$$v_1 \preceq v_2 \Leftrightarrow \exists V_i, V_j \text{ s.t. } v_1 \in V_i, v_2 \in V_j, i \geq j$$

for each $v_1, v_2 \in \mathcal{V}$. It is easy to see that $\preceq$ is total, reflexive, and transitive. Its minimal elements coincide with $\preceq'$:

$$\text{for } G \in \mathcal{F}_A : \min(ad(G), \preceq) = \min(ad(G), \preceq') \qquad (11)$$

Let V_k be the last set in the sequence $V_0, \ldots, V_m$ such that $V_k \cap ad(G) \neq \emptyset$. By definition of $\preceq$, $\min(ad(G), \preceq) = V_k \cap ad(G)$. Hence we have to show that $V_k \cap ad(G) = \min(ad(G), \preceq')$. ($\subseteq$): Assume there is some $v \in V_k \cap ad(G)$ such that $v \notin \min(ad(G), \preceq')$. From the latter it follows that $\exists v_0 \in ad(G) : v_0 \prec' v$. From (9) we get $v_0 \prec^t v$, hence $v_0 \notin \max(V_k, \preceq^t)$. As V_k is the last set with $V_k \cap ad(G) \neq \emptyset$ it must hold that $v_0 \in V_j$ with $j < k$, i.e. $v_0 \in \max(\mathcal{V} \setminus \bigcup_{0 \leq i < j} V_i, \preceq^t)$. Therefore, recalling $v_0 \prec^t v$, $v \notin \mathcal{V} \setminus \bigcup_{0 \leq i < j} V_i$, contradicting $v \in V_k$ and $j < k$. ($\supseteq$): Assume there is some $v_0 \in \min(ad(G), \preceq')$ such that $v_0 \notin V_k \cap ad(G)$. That means $v_0 \in ad(G)$ and $v_0 \notin V_k$ and further that $v_0 \in V_j$ for $j < k$. Now let $v_1 \in V_k \cap ad(G)$. As $j < k$ hence $v_1 \in \mathcal{V} \setminus \bigcup_{0 \leq i < j} V_i$. Since v_0 is maximal wrt. $\preceq^t$ in this set, $v_0 \not\prec^t v_1$ and further by (9) $v_0 \not\prec' v_1$. It holds that $v_0 \in pr(F \star f_{ad}(cl(v_0, v_1)))$ and therefore $v_0 \preceq' v_1$ though. We show this by assuming, towards a contradiction, that $v_0 \notin pr(F \star f_{ad}(cl(v_0, v_1)))$. Hence $v_0 \not\preceq' v_1$. As $v_0 \in \min(ad(G), \preceq')$ by assumption, then also $v_1 \not\preceq' v_0$. By $P3$ there has to be some $v_2 \in pr(F \star f_{ad}(cl(v_0, v_1)))$. As also $v_2 \in cl(v_0, v_2)$ we get by $P5$ and $P6$ that $v_2 \in pr(F \star f_{ad}(cl(v_0, v_1) \cap cl(v_0, v_2)))$. From Lemma 1.2 we infer that $cl(v_0, v_2) \subseteq cl(v_0, v_1)$, hence $v_2 \in pr(F \star f_{ad}(v_0, v_2))$, meaning that $v_2 \preceq' v_0$. Moreover, $v_0 \notin pr(F \star f_{ad}(v_0, v_2))$, hence even $v_2 \prec' v_0$. As $v_2 \in ad(G)$ from $v_0, v_1 \in ad(G)$ and $cl(v_0, v_1) \subseteq cl(ad(G)) = ad(G)$, we get a contradiction to $v_0 \in \min(ad(G), \preceq')$. Now consider an arbitrary $v_3 \in \mathcal{V} \setminus \bigcup_{0 \leq i < j} V_i$ such that $v_1 \preceq^t v_3$. From $v_0 \preceq' v_1 \preceq^t v_3$ we get $v_0 \preceq^t v_3$. But since $v_0 \in \max(\mathcal{V} \setminus \bigcup_{0 \leq i < j} V_i, \preceq^t)$ it must also hold that $v_3 \preceq^t v_0$, meaning, together with $v_0 \preceq' v_1$, that $v_3 \preceq^t v_1$. As v_3 was chosen arbitrarily we have that $v_1 \in \max(\mathcal{V} \setminus \bigcup_{0 \leq i < j} V_i, \preceq^t)$, i.e. $v_1 \in V_j$, a contradiction to $v_1 \in V_k$ and $j < k$.

The fact that $\preceq$ indeed simulates $\star$ is now obtained from (7) and (11): We get that $pr(F \star G) = \min(ad(G), \preceq)$ for each ADF G. This also makes $\preceq$ *pr-ad*-compliant. To show that $\preceq$ is faithful for F assume $pr(F) \neq \emptyset$ (otherwise faithfulness is trivial). By $P2$ it holds that $pr(F \star f_{ad}(\mathcal{V})) = pr(F)$, hence $pr(F) = \min(\mathcal{V}, \preceq)$, meaning that (i) $v_1 \approx v_2$ for $v_1, v_2 \in pr(F)$ and (ii) $v_1 \prec v_2$ for $v_1 \in pr(F)$ and $v_2 \notin pr(F)$. $\square$

With the insights from Theorems 4 and 5 we obtain concrete operators from faithful and *pr-ad*-compliant rankings. For instance, a valid operator is induced from the ranking $\preceq_F$ where $pr(F)$ are the minimal elements and all other interpretations form a $\prec_F$-chain. The three-valued version of Dalal's operator (cf. Definition 5) is not directly applicable here, as $\preceq_F^{pr}$ does not yield a *pr-ad*-compliant ranking for every ADF:

Example 6. Consider the ADF $F = \{\langle a, a \wedge b\rangle, \langle b, a \wedge b\rangle\}$ having $pr(F) = \{\mathbf{tt}, \mathbf{ff}\}$. It yields the ranking $\mathbf{tt} \approx_F^{pr} \mathbf{ff} \prec_F^{pr} \mathbf{tu} \approx_F^{pr} \mathbf{ut} \approx_F^{pr} \mathbf{uf} \approx_F^{pr} \mathbf{fu} \prec_F^{pr} \mathbf{tf} \approx_F^{pr} \mathbf{ft} \approx_F^{pr} \mathbf{uu}$. Now consider the compatible interpretations $\mathbf{tu}$ and $\mathbf{uf}$ and observe that all $v \in cl(\mathbf{tu}, \mathbf{uf}) = \{\mathbf{uu}, \mathbf{tu}, \mathbf{ut}, \mathbf{tf}\}$ have $v \not\prec_F^{pr} \mathbf{tu}, \mathbf{uf}$. Therefore, according to Proposition 3, $\preceq_F^{pr}$ is not *pr-ad*-compliant. In practice, this means that $F *_{pr}^{D} G$, where $ad(G) = \{\mathbf{uu}, \mathbf{tu}, \mathbf{uf}, \mathbf{tf}\}$, would yield $f_{pr}(\{\mathbf{tu}, \mathbf{uf}\})$; but as $\{\mathbf{tu}, \mathbf{uf}\}$ is not realizable under *pr* we do net get the preferred interpretations prescribed by the postulates.

A refinement of the distance measure in order to result in *pr-ad*-compliant rankings is subject to future work.

4 Discussion

Summary. We have characterized operators for the revision of ADFs under preferred semantics. Using recent insights on realizability we showed that rankings giving rise to concrete operators underlie milder conditions than in classical AGM revision (i-max-faithful versus faithful). We have exemplified these results by a three-valued version of Dalal's operator. While admissible semantics yield a single rational operator, we have proposed an alternative family of revision operators combining admissible and preferred semantics. Their representation by rankings is based on *pr-ad*-compliance.

Other semantics. First consider complete semantics and recall that there might be ADFs F and G such that their common complete interpretations might not be realizable under complete semantics. Therefore, when revising F by G, it is impossible to satisfy $A2_{co}$ since it would require the resulting ADF to have exactly $co(F) \cap co(G)$ as complete interpretations. The same applies to grounded semantics. Therefore it holds that for $\sigma \in \{co, gr\}$ there is no operator $* : \mathcal{F}_A \times \mathcal{F}_A \mapsto \mathcal{F}_A$ satisfying $A1_\sigma - A6_\sigma$.

For supported models, on the other hand, we observe that they have the same expressiveness as propositional logic, therefore results from classical AGM revision carry over.

Finally, stable models have similar sufficient conditions for realizability as preferred semantics, namely that a set of interpretations V is realizable if (1) $V \subseteq st(F)$ for some $F \in \mathcal{F}_A$, (2) $V = \{v_1, v_2\}$ for $v_1, v_2 \in \mathcal{V}^2$ with $v_1^{\mathbf{t}}$ and $v_2^{\mathbf{t}}$ being $\subseteq$-incomparable, or (3) $V = \{v\}$ with $v \in \mathcal{V}^2$. Therefore we expect to get similar representation results as for preferred semantics, just with slightly different conditions on the ranking.

Future work. While in this work we only dealt with the semantic outcome of operators, we also plan to study syntactic aspects of revision. Moreover, we want to study the computational complexity of Dalal's operator under preferred semantics, given that the complexity of reasoning tasks in ADFs is studied comprehensively [Strass and Wallner, 2015; Gaggl *et al.*, 2015]. Finally, we want to see how gained insights carry over to the revision of AFs: operators combining preferred and admissible semantics as well as revision under three-valued semantics [Caminada and Gabbay, 2009].

References

Latifa Al-Abdulkarim, Katie Atkinson, and Trevor J. M. Bench-Capon. A methodology for designing systems to reason with legal cases using abstract dialectical frameworks. *Artif. Intell. Law*, 24(1):1–49, 2016.

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.

Ofer Arieli. Distance-based paraconsistent logics. *Int. J. Approx. Reasoning*, 48(3):766–783, 2008.

Pietro Baroni, Federico Cerutti, Massimiliano Giacomin, and Giovanni Guida. AFRA: Argumentation framework with recursive attacks. *Int. J. Approx. Reasoning*, 52(1):19–37, 2011.

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

Richard Booth. Judgment aggregation in abstract dialectical frameworks. In *Advances in KR, Logic Programming, and Abstract Argumentation*, volume 9060 of *LNCS*, pages 296–308, 2015.

Gerhard Brewka and Stefan Woltran. Abstract dialectical frameworks. In *Proc. KR*, pages 102–111, 2010.

Gerhard Brewka, Stefan Ellmauthaler, Hannes Strass, Johannes P. Wallner, and Stefan Woltran. Abstract dialectical frameworks revisited. In *Proc. IJCAI*, pages 803–809, 2013.

Gerhard Brewka, Sylwia Polberg, and Stefan Woltran. Generalizations of Dung frameworks and their role in formal argumentation. *IEEE Intelligent Systems*, 29(1):30–38, 2014.

Martin Caminada and Dov M. Gabbay. A logical account of formal argumentation. *Studia Logica*, 93(2):109–145, 2009.

Claudette Cayrol and Marie-Christine Lagasquie-Schiex. On the acceptability of arguments in bipolar argumentation frameworks. In *Proc. ECSQARU*, pages 378–389, 2005.

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

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

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

James P. Delgrande, Pavlos Peppas, and Stefan Woltran. AGM-Style Belief Revision of Logic Programs under Answer Set Semantics. In *Proc. LPNMR*, pages 264–276, 2013.

Martin Diller, Adrian Haret, Thomas Linsbichler, Stefan Rümmele, and Stefan Woltran. An extension-based approach to belief revision in abstract argumentation. In *Proc. IJCAI*, pages 2926–2932, 2015.

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):321–358, 1995.

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

Marcelo A. Falappa, Alejandro J. García, Gabriele Kern-Isberner, and Guillermo R. Simari. On the evolving relation between belief revision and argumentation. *Knowledge Eng. Review*, 26(1):35–43, 2011.

Sarah A. Gaggl, Sebastian Rudolph, and Hannes Strass. On the computational complexity of naive-based semantics for abstract dialectical frameworks. In *Proc. IJCAI*, pages 2985–2991, 2015.

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

Sanjay Modgil. Reasoning about preferences in argumentation frameworks. *Artif. Intell.*, 173(9-10):901–934, 2009.

Sylwia Polberg. Extension-based semantics of abstract dialectical frameworks. In *Proc. STAIRS*, pages 240–249, 2014.

Jörg Pührer. Realizability of three-valued semantics for abstract dialectical frameworks. In *Proc. IJCAI*, pages 3171–3177, 2015.

Hannes Strass and Johannes P. Wallner. Analyzing the computational complexity of abstract dialectical frameworks via approximation fixpoint theory. *Artif. Intell.*, 226:34–74, 2015.

Hannes Strass. Approximating operators and semantics for abstract dialectical frameworks. *Artif. Intell.*, 205:39–70, 2013.

Hannes Strass. Expressiveness of two-valued semantics for abstract dialectical frameworks. *J. Artif. Intell. Res.*, 54:193–231, 2015.