

Probabilistic Bipolar Abstract Argumentation Frameworks: Complexity Results

Bettina Fazzinga¹ and Sergio Flesca² and Filippo Furfaro²

¹ ICAR - CNR, Via P. Bucci, 87036 Rende (CS) - Italy

² DIMES - Università della Calabria, Via P. Bucci, 87036 Rende (CS) - Italy
 fazzinga@icar.cnr.it, flesca@dimes.unical.it, furfaro@dimes.unical.it

Abstract

Probabilistic Bipolar Abstract Argumentation Frameworks (prBAFs) are considered, and the complexity of the fundamental problem of computing extensions' probabilities is addressed. The most popular semantics of supports and extensions are considered, as well as different paradigms for defining the probabilistic encoding of the uncertainty. Interestingly, the presence of supports, which does not alter the complexity of verifying extensions in the deterministic case, is shown to introduce a new source of complexity in some probabilistic settings, for which tractable cases are also identified.

1 Introduction

An *abstract argumentation framework* (AAF) represents a dispute as an *argumentation graph* $\langle A, D \rangle$, where A is the set of nodes (called *arguments*) and D is the set of edges (called *defeats* or *attacks*). Herein, an argument is an abstract entity that may attack and/or be attacked by other arguments, where “ a attacks b ” means that argument a rebuts/weakens b . Given this, reasoning on the possible strategies for winning the dispute typically requires looking into the *extensions* of the AAF, i.e. a set of arguments that satisfies some properties certifying its “strength”. Different semantics for AAFs have proven reasonable, such as *admissible* (*ad*), *stable* (*st*), *preferred* (*pr*), *complete* (*co*), *grounded* (*gr*), *ideal* (*id*) [Dung, 1995; Dung *et al.*, 2007; Baroni and Giacomin, 2009], and the complexity of the problem EXT of verifying whether a set is an extension has been studied under each of these semantics [Dunne and Wooldridge, 2009; Dunne, 2009].

Since the introduction of AAFs in [Dung, 1995], many variants have been proposed, with the aim of modeling disputes more accurately. Among these, *Bipolar Abstract Argumentation Frameworks* (BAFs) allow *supports*, besides attacks, to be specified between arguments. Specifically, two alternative formal semantics of support have been introduced: in [Cayrol and Lagasque-Schieux, 2005], the support is a generic “inverse” of the notion of attack (“*abstract semantics*”: “ a supports b ” means that there is a positive interaction between a and b – from a to b), while, in [Boella *et al.*,




Figure 1: A bipolar abstract argumentation framework

2010], it is viewed as a “deductive” correlation between arguments (“*deductive semantics*”: if a supports b , the acceptance of a implies the acceptance of b). The various extensions’ semantics defined for AAFs have been shown to have a natural counterpart over BAFs, after noticing that combining attacks with supports (of any semantics) generates “implicit” attacks.

Example 1 *The graph in Figure 1 is a BAF with six arguments a, b, c, d, e, f . The dashed and standard arrows denote supports and attacks, respectively. The co-existence of supports and attacks entails the existence of implicit attacks. For instance, under both the abstract and deductive semantics, the fact that a strengthens b and b attacks c implicitly says that a attacks c . This kind of implicit attack is often called “supported attack”. If the deductive semantics is adopted, there are other forms of implicit attacks. For instance, since a supports b and e attacks b , there is an implicit attack from e to a . Otherwise, a would be acceptable while b would be not, thus contradicting the deductive support from a to b .*

Other variants of AAFs are those addressing the representation of uncertainty. In this regard, *probabilistic AAFs* (prAAFs) are a popular paradigm, and in particular those following the *constellation approach*. Here, the dispute is modeled as a set of possible scenarios, each consisting of a standard AAF (called *possible AAF*) associated with a probability of representing all and only the arguments and attacks actually occurring in the dispute. In particular, two main paradigms have been adopted for specifying the probability distribution function (pdf), called EX and IND. In the general case, the *extensive* form EX is used, where the each possible AAF must be explicitly specified along with its probability. Otherwise, when *independence* between arguments/attacks is assumed, the form IND can be used, where the probabilities of the possible scenarios are represented implicitly by specifying the marginal probabilities of the arguments and attacks. For both EX and IND, the complexity of the probabilistic counterpart P-EXT of EXT (asking for the probability that a set of arguments is an extension) has been characterized in [Fazzinga *et al.*, 2016].

In this paper, we consider the *probabilistic Bipolar Argu-*

mentation Framework (prBAF), where the bipolarity of BAFs is combined with the probabilistic modeling of the uncertainty of prAAFs, and characterize the complexity of reasoning on extensions. In our analysis, we start by observing that the introduction of supports in the deterministic case does not increase the complexity of verifying extensions, meaning that EXT over BAFs (for both supports’ semantics) has the same complexity as over AAFs. Then, we show that allowing supports can introduce a new source of complexity when probabilities are considered: under the admissible and stable semantics, and under the paradigm IND, for which P-EXT is polynomial over prAAFs, the problem P-EXT becomes intractable over prBAFs. For these cases, we give an insight on this behavior by showing sufficient conditions restoring the polynomiality of P-EXT. Table 1 summarizes our results.

<i>sem</i>	Literature			Our results		
	EXT (AAF)	P-EXT (prAAF)		EXT (BAF)	P-EXT (prBAF)	
		IND	EX		IND	EX
admissible	P	FP	FP	P	$FP^{\#P}$	FP
stable	P	FP	FP	P	$FP^{\#P}$	FP
complete	P	$FP^{\#P}$	FP	P	$FP^{\#P}$	FP
grounded	P	$FP^{\#P}$	FP	P	$FP^{\#P}$	FP
preferred	$coNP$	$FP^{\#P}$	$FP^{ NP}$	$coNP$	$FP^{\#P}$	$FP^{ NP}$
ideal	in θ_2^p , $coNP-h$	$FP^{\#P}$	$FP^{ NP}$	in θ_2^p , $coNP-h$	$FP^{\#P}$	$FP^{ NP}$

Table 1: Complexity of EXT and P-EXT. Here: FP (resp., FP^X) is the class of functions computable by a polynomial-time TM (resp., with invocations to an oracle for the class X); $FP^{|X}$ is similar to FP^X , but the oracle invocations can be done parallelly; θ_2^p is $P^{|NP}$; $\#P$ is the class of functions counting the solutions of an NP problem

Related Work. [Cayrol and Lagasque-Schiex, 2005] first introduced BAFs, where supports have the general “abstract” semantics of positive interactions between arguments. Later, three more specific interpretations for supports have been proposed: [Boella *et al.*, 2010], [Nouioua and Risch, 2010] and [Oren and Norman, 2008] introduced the *deductive*, *necessary*, and *evidential* semantics for the support relation, respectively. In this paper, we focus on the abstract and deductive semantics, but our results also hold for necessary supports (as shown in [Cayrol and Lagasque-Schiex, 2013], they are dual to deductive ones). [Cayrol and Lagasque-Schiex, 2013] reviews the four different semantics for supports, and discusses the similarities and differences among these interpretations. [Cohen *et al.*, 2012] introduces a more general framework that incorporates attacks, supports and a preference relation. In [Martínez *et al.*, 2006], subarguments in AAF have been introduced, that in [Cohen *et al.*, 2014] have been shown to be closely related with the necessary support. Other related works are [Brewka and Woltran, 2010; Verheij, 2009] where, although supports are not mentioned, similar dependencies have been considered. A detailed survey over BAFs can be found in [Cohen *et al.*, 2014].

As regards uncertainty in AAFs, the approaches based on probability theory can be classified in two categories: those adopting the classical *constellations* approach [Hunter, 2014;

Rienstra, 2012; Doder and Woltran, 2014; Dondio, 2014; Hunter, 2012; Li *et al.*, 2011; Fazzinga *et al.*, 2013; Fazzinga *et al.*, 2015] and those adopting the recent *epistemic* one [Thimm, 2012; Hunter and Thimm, 2014b; Hunter and Thimm, 2014a]. The former category has the two subcategories EX [Rienstra, 2012; Dondio, 2014] and IND [Doder and Woltran, 2014; Li *et al.*, 2011; Fazzinga *et al.*, 2013; Fazzinga *et al.*, 2015], described in the paper. The interested reader can find a more detailed comparative description of the two categories in [Hunter, 2013]. Furthermore, many proposals have been made where uncertainty is represented by exploiting weights or preferences on arguments and/or defeats [Bench-Capon, 2003; Amgoud and Vesic, 2011; Modgil, 2009; Dunne *et al.*, 2011; Coste-Marquis *et al.*, 2012]. Although these approaches have proved effective in different contexts, there is no common agreement on what kind of approach should be used in general. In this regard, [Hunter, 2012; Hunter, 2013] observed that the probability-based approaches may take advantage from relying on a well-established and well-founded theory, whereas the approaches based on weights or preferences do not. Finally, recent works [Proietti, 2017; Polberg and Hunter, 2018] investigate the need for extending BAFs with probabilities.

2 Preliminaries

We now review *Bipolar Abstract Argumentation Frameworks* (BAFs) and the concepts of support, attack and defense, along with the most popular extensions’ semantics over BAFs.

2.1 Bipolar Abstract Argumentation Frameworks

Definition 1 [BAF] A bipolar abstract argumentation framework (BAF) is a tuple $\mathcal{F} = \langle \mathcal{A}, \mathcal{R}_a, \mathcal{R}_s \rangle$, where \mathcal{A} is a set of arguments, $\mathcal{R}_a \subseteq \mathcal{A} \times \mathcal{A}$ is a defeat/attack relation and $\mathcal{R}_s \subseteq \mathcal{A} \times \mathcal{A}$ is a support relation.

In the first proposal of BAF [Cayrol and Lagasque-Schiex, 2005], supports are given an *abstract semantics*, that is the opposite of the traditional semantics of attack, inherited from classical AAFs. This was shown to make the combination of supports and attacks imply the so-called *supported attacks*¹.

Definition 2 [Supported attack] Let $\mathcal{F} = \langle \mathcal{A}, \mathcal{R}_a, \mathcal{R}_s \rangle$ be a BAF, and $a, b \in \mathcal{A}$. There is a supported attack from a to b iff there is a sequence $a_1 \mathcal{R}_1 \dots \mathcal{R}_{n-1} a_n$, with $n \geq 2$, where $a_1 = a$, $a_n = b$, $\forall i \in [1..n-2] \mathcal{R}_i = \mathcal{R}_s$, and $\mathcal{R}_{n-1} = \mathcal{R}_a$.

Example 2 In the BAF in Figure 1, there is a supported attack from a to c . Also the three direct attacks (e, b) , (b, c) and (f, d) are special cases of supported attack.

Besides the abstract, other semantics have been proposed for supports (see *Related Work*). In particular, we consider the well-established *deductive semantics*, first proposed in [Boella *et al.*, 2010]. Here, “ a supports b ” is interpreted as a strong correlation between a and b , meaning that if a is acceptable, then b is acceptable too. As observed in [Cayrol and Lagasque-Schiex, 2013], under this semantics, a new form of implicit attack, called *d-attack*, must be considered.

¹[Cayrol and Lagasque-Schiex, 2005] discussed also *indirect* implicit attacks. W.l.o.g., as in [Cayrol and Lagasque-Schiex, 2009], we disregard them, as they would not affect our results.

Definition 3 [d-attack] Given a BAF $\mathcal{F} = \langle \mathcal{A}, \mathcal{R}_a, \mathcal{R}_s \rangle$ and $a, b \in \mathcal{A}$, there is a d-attack from a to b iff

- $a\mathcal{R}_a b$, or
- there is an argument a' such that there is a path from a to a' consisting of only support edges, and a' attacks b , or
- there is an argument a' such that there is a path from b to a' consisting of only support edges, and a attacks a' .

Example below shows that d-attacks include supported attacks, but can be also of the form of *supermediated attacks* (described by the last point in Definition 3).

Example 3 Under the deductive semantics for supports, in the BAF of Figure 1, it is easy to see that the supported attacks reported in Example 2 are d-attacks. Further d-attacks are the supermediated attacks from e to a , and from f to c .

In order to analyze what changes when moving from one semantics of supports to the other, we partition BAFs into two classes: *s-BAFs* and *d-BAFs*, where only supported attacks and d-attacks are considered, respectively. From now on, we assume the presence of a BAF $\mathcal{F} = \langle \mathcal{A}, \mathcal{R}_a, \mathcal{R}_s \rangle$, and, when needed, we will specify whether \mathcal{F} is an s- or a d-BAF.

Definition 4 [Set-support] A set $S \subseteq \mathcal{A}$ set-supports an argument $a \in \mathcal{A}$ iff there is an argument $a' \in S$ such that there is a path from a' to a consisting of only support edges.

Definition 5 [Set-attack] Let $\mathcal{F} = \langle \mathcal{A}, \mathcal{R}_a, \mathcal{R}_s \rangle$ be an s-BAF (resp., d-BAF). A set $S \subseteq \mathcal{A}$ set-attacks $a \in \mathcal{A}$ iff there is a supported attack (resp., d-attack) from some $b \in S$ to a .

Definition 6 [Set-defense] A set $S \subseteq \mathcal{A}$ set-defends an argument $a \in \mathcal{A}$ iff, $\forall b \in \mathcal{A}$, if $\{b\}$ set-attacks a then $\exists c \in S$ such that $\{c\}$ set-attacks b .

Example 4 Consider the BAF \mathcal{F} in Figure 1. Independently from supports' semantics, $\{a, e\}$ both set-supports and set-attacks b , and also set-attacks c . If \mathcal{F} is an s-BAF, then $\{a, e\}$ does not set-defend c , since there is a supported attack from a to c , and no attack from e to a . Observe that $\{a, e\}$ does not set-defend c if \mathcal{F} is a d-BAF either.

2.2 Semantics

We first recall the notions of conflict-freeness and safety.

Definition 7 [Conflict-free and safe sets of arguments] A set of arguments $S \subseteq \mathcal{A}$ is:

- conflict-free iff $\nexists a, b \in S$ such that $\{a\}$ set-attacks b ;
- safe iff $\nexists b \in \mathcal{A}$ such that S set-attacks b and either S set-supports b or $b \in S$.

Example 5 If the BAF \mathcal{F} in Figure 1 is an s-BAF, both $\{a, e\}$, and $\{f, c\}$ are conflict-free but not safe, while both $\{a, b, f\}$ and $\{a, b, d\}$ are conflict-free and safe. If \mathcal{F} is a d-BAF, both $\{a, e\}$, and $\{f, c\}$ are not conflict-free, while both $\{a, b, f\}$ and $\{a, b, d\}$ are still conflict-free and safe.

All the most popular semantics of extensions of “standard” AAFs have been extended to the case of BAFs [Cayrol and Lagasque-Schiex, 2005]. We start with the stable semantics.

Definition 8 [Stable extension] A set of arguments $S \subseteq \mathcal{A}$ is a stable extension iff S is conflict-free and $\forall a \in \mathcal{A} \setminus S$ it holds that S set-attacks a .

Due to the presence of supports and the fact that, in BAFs, conflict-freeness and safety do not coincide, for some AAF's semantics, different variants are considered when moving to BAFs. This is the case of the admissible semantics.

Definition 9 [Admissible extension] A set $S \subseteq \mathcal{A}$ is

- a d-admissible extension iff S is conflict-free and set-defends all of its arguments;
- an s-admissible extension iff S is safe and set-defends all of its arguments;
- a c-admissible extension iff S is conflict-free, closed for \mathcal{R}_s and set-defends all of its arguments.

In turn, the other semantics subsuming the admissible one are defined as follows. A set $S \subseteq \mathcal{A}$ is said to be:

- a *d-complete* (resp. *s-complete*, *c-complete*) extension iff S is d-admissible (resp., *s-admissible*, *c-admissible*) and S contains all the arguments set-defended by S ;
- a *d-grounded* (resp. *s-grounded*, *c-grounded*) extension iff S is a minimal (w.r.t. \subseteq) d-complete (resp. *s-complete*, *c-complete*) extension;
- a *d-preferred* (resp. *s-preferred*, *c-preferred*) extension iff S is a maximal (w.r.t. \subseteq) d-complete (resp. *s-complete*, *c-complete*) extension;
- a *d-ideal* (resp. *s-ideal*, *c-ideal*) extension iff S is a maximal (w.r.t. \subseteq) d-admissible (resp. *s-admissible*, *c-admissible*) extension and S is contained in every d-preferred (resp. *s-preferred*, *c-preferred*) extension.

We denote the set $\{d\text{-ad}, s\text{-ad}, c\text{-ad}, st, d\text{-co}, s\text{-co}, c\text{-co}, d\text{-gr}, c\text{-gr}, c\text{-gr}, d\text{-pr}, s\text{-pr}, c\text{-pr}, d\text{-id}, s\text{-id}, c\text{-id}\}$ consisting of the above semantics as SEM (herein, *st* means stable, *d-ad* d-admissible, *s-ad* s-admissible, and so on).

Example 6 Consider the BAF in Figure 1. $\{a, b, f\}$, although conflict-free and safe, is not a d-ad extension for both s-BAF and d-BAF (since b is not set-defended). Furthermore, for the s-BAF case, $\{a, f\}$ is a s-ad, s-gr and s-pr extension, $\{a, e, f\}$ is a st, d-ad, d-gr, d-pr, and d-id extension, $\{f\}$ is an s-id extension, $\{e, f\}$ is a c-pr, c-gr and c-id extension.

For the d-BAF case we have: $\{e, f\}$ is the unique stable extension, that is also c-preferred, c-grounded and c-ideal.

The fundamental problem of verifying whether a set S of arguments is an extension over a given BAF under a semantics $sem \in SEM$ will be denoted as $EXT^{sem}(S)$. Given a BAF $\mathcal{F} = \langle \mathcal{A}, \mathcal{R}_a, \mathcal{R}_s \rangle$, a set $S \subseteq \mathcal{A}$, and a semantics $sem \in SEM$, we define the boolean function $ext(\mathcal{F}, sem, S)$ returning *true* iff S is an extension under sem .

3 Probabilistic BAFs (prBAFs)

We now consider the extension of BAFs where uncertainty is addressed and modeled probabilistically as in “traditional” probabilistic AAFs - prAAF. A probabilistic BAF (prBAF) \mathcal{F} is a tuple $\mathcal{F} = \langle \mathcal{A}, \mathcal{R}_a, \mathcal{R}_s, \mathcal{P} \rangle$, where $\mathcal{F} = \langle \mathcal{A}, \mathcal{R}_a, \mathcal{R}_s \rangle$ is a BAF and \mathcal{P} is a probability distribution function (pdf) over the set $PS = \{\alpha = \langle \mathcal{A}', \mathcal{R}'_a, \mathcal{R}'_s \rangle \mid \mathcal{A}' \subseteq \mathcal{A} \wedge \mathcal{R}'_a \subseteq$

$(\mathcal{A}' \times \mathcal{A}') \cap \mathcal{R}_a \wedge \mathcal{R}'_s \subseteq (\mathcal{A}' \times \mathcal{A}') \cap \mathcal{R}_s$. That is, the elements in $PS(\mathcal{F})$, called *possible scenarios* or *possible BAFs*, are the alternative cases of dispute that may occur, and each of them is encoded by a BAF. Then, \mathcal{P} assigns to each possible BAF the probability that it describes the actual dispute.

As happens with prAAFs, we consider two encodings of the pdf over the possible BAFs, namely EX and IND. That is, a prBAF \mathcal{F} of form EX is a tuple $\langle \mathcal{A}, \mathcal{R}_a, \mathcal{R}_s, \vec{\alpha}, \vec{P} \rangle$, where \mathcal{A} , \mathcal{R}_a and \mathcal{R}_s are sets of arguments, attacks and supports, respectively, while $\vec{\alpha} \subseteq PS(\mathcal{F})$ is the sequence $\vec{\alpha} = \alpha_1, \dots, \alpha_m$ of the possible BAFs that are assigned non-zero probability and $\vec{P} = P(\alpha_1), \dots, P(\alpha_m)$ are their probabilities. The size of a prBAF $\langle \mathcal{A}, \mathcal{R}_a, \mathcal{R}_s, \vec{\alpha}, \vec{P} \rangle$ of form EX is thus $O((|\mathcal{A}| + |\mathcal{R}_a| + |\mathcal{R}_s|) \cdot (|\vec{\alpha}| + |\vec{P}|))$. Obviously, the form EX allows any pdf (encoding any correlation between arguments/attacks/supports) to be represented. Otherwise, when independence is assumed (thus, no correlation must be encoded), the more compact paradigm IND can be used. A prBAF of type IND is a tuple $\langle \mathcal{A}, \mathcal{R}_a, \mathcal{R}_s, \mathcal{P}_A, \mathcal{P}_R \rangle$ where $A = \{a_1, \dots, a_m\}$, $\mathcal{R}_a = \{\delta_1, \dots, \delta_n\}$ and $\mathcal{R}_s = \{\sigma_1, \dots, \sigma_k\}$ are the sets of arguments, attacks and supports, respectively, and $\mathcal{P}_A = \{P(a_1), \dots, P(a_m)\}$, $\mathcal{P}_R = \{P(\delta_1), \dots, P(\delta_n), P(\sigma_1), \dots, P(\sigma_k)\}$ are their marginal probabilities. The pdf P over the possible scenarios that is implied by the independence assumption and the marginal probabilities $\mathcal{P}_A, \mathcal{P}_R$ is as follows. For each possible BAF $\alpha' = \langle \mathcal{A}', \mathcal{R}'_a, \mathcal{R}'_s \rangle$, the probability $P(\alpha')$ is:

$$P(\alpha') = \prod_{a \in \mathcal{A}'} P(a) \times \prod_{a \in \mathcal{A} \setminus \mathcal{A}'} (1 - P(a)) \times \prod_{\delta \in \mathcal{R}'_a} P(\delta) \times \prod_{\delta \in (\mathcal{R}_s \cap (\mathcal{A}' \times \mathcal{A}')) \setminus \mathcal{R}'_a} (1 - P(\delta)) \times \prod_{\sigma \in \mathcal{R}'_s} P(\sigma) \times \prod_{\sigma \in (\mathcal{R}_s \cap (\mathcal{A}' \times \mathcal{A}')) \setminus \mathcal{R}'_s} (1 - P(\sigma)). \quad (1)$$

The size of a prBAF of type IND is $O(|\mathcal{A}| + |\mathcal{R}_a| + |\mathcal{R}_s| + |P_A| + |P_R|)$. For both EX and IND, probabilities are assumed to be rationals, whose sizes contribute to $|\vec{P}|$, $|P_A|$ and $|P_R|$.

Example 7 Consider a prBAF \mathcal{F}' of form EX, where $\mathcal{A}, \mathcal{R}_a$ and \mathcal{R}_s are those of Figure 1, and $\vec{\alpha} = \alpha_1, \alpha_2, \alpha_3$ and $\vec{P} = P(\alpha_1), P(\alpha_2), P(\alpha_3)$ are: $\alpha_1 = \langle \mathcal{A}, \mathcal{R}_a, \mathcal{R}_s \rangle$, $\alpha_2 = \langle \mathcal{A}, \mathcal{R}_a \setminus \{(e, b)\}, \mathcal{R}_s \rangle$, $\alpha_3 = \langle \mathcal{A} \setminus \{e\}, \mathcal{R}_a \setminus \{(e, b)\}, \mathcal{R}_s \rangle$, and $P(\alpha_1) = 0.6$, $P(\alpha_2) = 0.2$ and $P(\alpha_3) = 0.2$.

Consider now a prBAF \mathcal{F}'' of form IND, where $\mathcal{A}, \mathcal{R}_a$ and \mathcal{R}_s are those of Figure 1, and \mathcal{P}_A and \mathcal{P}_R are the following: $P(a) = P(b) = P(c) = P(d) = P(f) = 1$, $P(e) = 0.5$, $P(e, b) = 0.5$ and the probabilities of the other supports and attacks are equal to 1. We have three possible scenarios, α_1, α_2 and α_3 , that are the same of the case of \mathcal{F}' , but with probabilities: $P(\alpha_1) = 0.25$, $P(\alpha_2) = 0.25$, $P(\alpha_3) = 0.5$.

In what follows, given a prBAF $\mathcal{F} = \langle \mathcal{A}, \mathcal{R}_a, \mathcal{R}_s, \mathcal{P} \rangle$ of any kind (thus, independently from the way \mathcal{P} is encoded), we denote as $\mathcal{F}.\vec{\alpha} = \alpha_1, \dots, \alpha_m$ the possible BAFs that are assigned non-zero probability by \mathcal{P} , and as $\mathcal{F}.\vec{P} = P(\alpha_1), \dots, P(\alpha_m)$ their probabilities. For instance, if \mathcal{F} is encoded as a prBAF of form EX, then $\mathcal{F}.\vec{\alpha}$ and $\mathcal{F}.\vec{P}$ are exactly the terms $\vec{\alpha}$ and \vec{P} in the tuple $\langle \mathcal{A}, \mathcal{R}_a, \mathcal{R}_s, \vec{\alpha}, \vec{P} \rangle$ encoding \mathcal{F} . Analogously, if \mathcal{F} is of form IND, then $\mathcal{F}.\vec{\alpha}$ is the set of all the possible BAFs having non-zero probability

definable over $\mathcal{A}, \mathcal{R}_a$ and \mathcal{R}_s , and $\mathcal{F}.\vec{P}$ is the pdf defined in Equation 1. The probabilistic versions of the two sub-classes s-BAF and d-BAF will be called *s-prBAF* and *d-prBAF*, respectively. Obviously, they can be of form IND or EX.

When switching to the probabilistic setting, the decision problem $\text{EXT}^{sem}(S)$ makes no sense, since a number of different scenarios are possible, and a set of arguments can be an extension in some scenarios, but not in others. Thus, the most natural “translation” of the problem of examining the “reasonability” of a set of arguments S becomes the functional problem $\text{P-EXT}^{sem}(S)$ of evaluating the probability that S is an extension, according to the following definition.

Definition 10 (P-EXT^{sem}(S) and P^{sem}(S)) Given a prBAF \mathcal{F} , a set S of arguments, and a semantics $sem \in SEM$, $\text{P-EXT}^{sem}(S)$ is the problem of computing the probability $P_{\mathcal{F}}^{sem}(S)$ that S is an extension under sem , i.e.

$$P_{\mathcal{F}}^{sem}(S) = \sum_{\alpha \in \mathcal{F}.\vec{\alpha} \wedge ext(\alpha, sem, S)} \mathcal{F}.P(\alpha) \quad (2)$$

In the following, we will denote as $\text{P-EXT}_{EX}^{sem}(S)$ (resp., $\text{P-EXT}_{IND}^{sem}(S)$) the problem $\text{P-EXT}^{sem}(S)$ restricted to the case that the input prBAF is of form EX (resp., IND).

Example 8 Continuing examples 6 and 7, we now compute the probability that $S = \{a, e\}$ is *d-admissible* in both the *s-and d-prBAF* cases, and for both the forms \mathcal{F}' and \mathcal{F}'' .

Case s-prBAF: S is *d-admissible* in both α_1 and α_2 (as e is missing in α_3), thus $P_{\mathcal{F}'}^{d-ad}(S) = P(\alpha_1) + P(\alpha_2) = 0.8$ and $P_{\mathcal{F}''}^{d-ad}(S) = P(\alpha_1) + P(\alpha_2) = 0.5$.

Case d-prBAF: S is *d-admissible* only in α_2 , as in α_1 e *d-attacks* a and in α_3 e is missing, thus we have $P_{\mathcal{F}'}^{d-ad}(S) = P(\alpha_2) = 0.2$ and $P_{\mathcal{F}''}^{d-ad}(S) = P(\alpha_2) = 0.25$.

4 Complexity Results

We now provide our main contribution: we characterize the complexity of $\text{P-EXT}^{sem}(S)$ over prBAFs for all the combinations of semantics of extensions, of supports (s- and d-BAFs), and probabilistic paradigms (EX and IND). Table 1 summarizes our results: the variants d-, s-, and c- of the same semantics are grouped in the same row (for instance, “*admissible*” groups d-, s-, and c-*admissible*), as we will show that moving from one variant to the other does not affect the complexity. Before considering the probabilistic setting, we address the deterministic one, by discussing the complexity of $\text{EXT}^{sem}(S)$ over BAFs.

Proposition 1 For both s-BAFs and d-BAFs, $\text{EXT}^{sem}(S)$ is in P for $sem \in \{d-ad, s-ad, c-ad, st, d-co, s-co, c-co, d-gr, c-gr, c-gr\}$, is coNP-complete for $sem \in \{d-pr, s-pr, c-pr\}$ and is in θ_2^P and coNP-hard for $sem \in \{d-id, s-id, c-id\}$.

Proposition 1 follows from the fact that the conditions that must be checked to verify an extension in the presence of supports (for the three variants d-, s-, and c-) require only polynomial computations on top of those needed in the absence of supports. This implies that the complexity of $\text{EXT}^{sem}(S)$ over BAFs is the same as over AAFs (second column of Table 1). We now characterize the complexity of $\text{P-EXT}_{EX}^{sem}(S)$.

Theorem 1 For both s -prBAFs and d -prBAFs, $P\text{-EXT}_{\text{EX}}^{\text{sem}}(S)$ is in P for $\text{sem} \in \{d\text{-ad}, s\text{-ad}, c\text{-ad}, st, d\text{-co}, s\text{-co}, c\text{-co}, d\text{-gr}, c\text{-gr}, c\text{-gr}\}$, and $FP^{\parallel NP}$ -complete for $\text{sem} \in \{d\text{-pr}, s\text{-pr}, c\text{-pr}, d\text{-id}, s\text{-id}, c\text{-id}\}$.

(Proof.) Evaluating the probability that S is an extension requires checking whether S is an extension over all the possible BAFs. Under EX, the number of possible scenarios is linear in the input size. Hence, $P\text{-EXT}_{\text{EX}}^{\text{sem}}(S)$ is in P for all the semantics that are polynomial in the deterministic case. For the other semantics, membership in $FP^{\parallel NP}$ follows from the fact that the check can be done in coNP (for the group “preferred”) or with parallel invocations to NP oracles (for the group “ideal”). Hardness for $FP^{\parallel NP}$ follows from the hardness for prAAF proved in [Fazzinga et al., 2016]. \square

We now characterize the complexity of $P\text{-EXT}_{\text{IND}}^{\text{sem}}(S)$.

Theorem 2 For any $\text{sem} \in \text{SEM}$, and for both s - and d -prBAFs, $P\text{-EXT}_{\text{IND}}^{\text{sem}}(S)$ is $FP^{\#P}$ -complete.

(Proof.) The membership can be shown applying minor modifications to the TM used in Theorem 3.14 in [Fazzinga et al., 2015], where the membership in $FP^{\#P}$ of the correspondent problem for prAAFs (thus, with no supports) was shown. For the hardness, we have to reason by cases. The case that sem is not “stable” or a variant of “admissible” is trivial, since it follows from the fact that the problem is $FP^{\#P}$ -hard even when no supports occur [Fazzinga et al., 2015]. Otherwise, if $\text{sem} \in \{d\text{-ad}, s\text{-ad}, c\text{-ad}, st\}$, we consider two cases.

Case 1: the prBAF is an s-prBAF. We first prove the hardness for $\text{sem} \in \{d\text{-ad}, s\text{-ad}, c\text{-ad}\}$. It suffices to show a reduction from a $\#P$ -complete problem, since, for functional problems, $\#P$ - and $FP^{\#P}$ -hardnesses coincide. We show a Cook reduction from $\#BP2DNF$, the $\#P$ -complete problem of counting the satisfying assignments of a bipartite positive 2-DNF formula. Specifically, given two disjoint sets of propositional variables $\mathcal{X} = \{X_1, \dots, X_n\}$, $\mathcal{Y} = \{Y_1, \dots, Y_m\}$, and a BP2DNF formula $\phi = C_1 \vee \dots \vee C_k$, where every clause C_i is of the form $X \wedge Y$, with $X \in \mathcal{X}$ and $Y \in \mathcal{Y}$, we consider the s -prBAF $\mathcal{F}(\phi) = \langle \mathcal{A}, \mathcal{R}_a, \mathcal{R}_s, \mathcal{P}_A, \mathcal{P}_R \rangle$ where:

- $\mathcal{A} = \{a, b, c, d\} \cup \{x_i \mid X_i \in \mathcal{X}\} \cup \{y_i \mid Y_i \in \mathcal{Y}\}$;
- $\mathcal{R}_a = \{\langle c, x_i \rangle \mid X_i \in \mathcal{X}\} \cup \{\langle c, y_i \rangle \mid Y_i \in \mathcal{Y}\} \cup \{\langle c, b \rangle, \langle b, d \rangle\}$;
- $\mathcal{R}_s = \{\langle a, x_i \rangle \mid X_i \in \mathcal{X}\} \cup \{\langle y_j, b \rangle \mid Y_j \in \mathcal{Y}\} \cup \{\langle x_i, y_j \rangle \mid X_i \wedge Y_j \text{ is a clause of } \phi\}$;
- \mathcal{P}_A assigns probability 1 to every argument;
- \mathcal{P}_R assigns probability 1 to every defeat and to every support $\langle x_i, y_j \rangle$ such that $X_i \wedge Y_j$ is in ϕ , and probability $\frac{1}{2}$ to the other supports (i.e., those of form $\langle a, x_i \rangle$ or $\langle y_j, b \rangle$).

Let β be the bijection from the truth assignments for $X_1, \dots, X_n, Y_1, \dots, Y_m$ to the possible BAFs of $\mathcal{F}(\phi)$ with non-zero probability, such that, for each truth assignment t , $\beta(t) = \langle \mathcal{A}', \mathcal{R}'_a, \mathcal{R}'_s \rangle$ is the possible BAF where: i) $\langle a, x_i \rangle \in \mathcal{R}'_s$ iff $t(X_i) = \text{true}$ and ii) $\langle y_j, b \rangle \in \mathcal{R}'_s$ iff $t(Y_j) = \text{true}$.

Consider the set of arguments $\{c, d\}$. It is easy to see that, for each truth assignment t , the set $\{c, d\}$ is d - s - c -admissible in $\beta(t)$ iff t does not satisfy ϕ (otherwise, there would be a supported attack from a to d , and a is not attacked by c or d).

Hence, any instance ϕ of $\#BP2DNF$ can be reduced to an instance of $P\text{-EXT}_{\text{IND}}^{\text{sem}}(S)$ by first constructing the prBAF

$\mathcal{F}(\phi)$, and next returning $2^{n+m} \cdot (1 - P_{\mathcal{F}(\phi)}^{\text{sem}}(\{c, d\}))$ as the number of satisfying assignments of ϕ .

Reasoning analogously, the statement for the stable semantics can be proved by taking the same construction and returning $2^{n+m} \cdot (1 - P_{\mathcal{F}(\phi)}^{\text{st}}(\{a, c, d\}))$ (instead of $2^{n+m} \cdot (1 - P_{\mathcal{F}(\phi)}^{\text{st}}(\{c, d\}))$) as the number of satisfying assignments of ϕ .

Case 2: the prBAF is an d-prBAF. We show a Cook reduction from the $\#P$ -complete problem $\#P2CNF$ of counting the satisfying assignments of a positive 2-CNF formula. Specifically, given a positive 2-CNF formula $\phi = C_1 \wedge \dots \wedge C_k$ over the set of propositional variables $\mathcal{X} = \{X_1, \dots, X_n\}$ where each clause C_h is of the form $X_i \vee X_j$ with $X \in \mathcal{X}$, we consider the d -prBAF $\mathcal{F}(\phi) = \langle \mathcal{A}, \mathcal{R}_a, \mathcal{R}_s, \mathcal{P}_A, \mathcal{P}_R \rangle$ where:

- $\mathcal{A} = \{a\} \cup \{x_i \mid X_i \in \mathcal{X}\} \cup \{c_h \mid C_h \in \{C_1, \dots, C_k\}\}$;
- $\mathcal{R}_a = \{\langle a, x_i \rangle \mid X_i \in \mathcal{X}\} \cup \{\langle c_h, a \rangle \mid C_h \in \{C_1, \dots, C_k\}\}$;
- $\mathcal{R}_s = \{\langle c_h, x_i \rangle \mid X_i \text{ occurs in the clause } C_h\}$;
- \mathcal{P}_A assigns probability 1.0 to the arguments a, c_1, \dots, c_k and probability $\frac{1}{2}$ to the arguments x_1, \dots, x_n ;
- \mathcal{P}_R assigns probability 1.0 to every attack and support.

Let β be the bijection from the truth assignments for X_1, \dots, X_n to the possible BAFs of $\mathcal{F}(\phi)$ with non-zero probability such that, for each truth assignment t , we have that $\beta(t) = \langle \mathcal{A}', \mathcal{R}'_a, \mathcal{R}'_s \rangle$ is the possible BAF such that $x_i \in \mathcal{A}'$ iff $t(X_i) = \text{true}$. Consider the set of arguments $\{a\}$. It is easy to see that, for each truth assignment t , $\{a\}$ is d - s - c -admissible or stable in $\beta(t)$ iff t satisfies ϕ . This derives from the fact that $\beta(t)$ contains an attack from a to x_i iff $t(X_i) = \text{true}$. In turn, the existence of $\langle a, x_i \rangle$ raises a supermediated attack from a to every c_j such that X_i occurs in C_j , thus defending a from c_j . Hence, any instance ϕ of $\#P2CNF$ can be reduced to an instance of $P\text{-EXT}_{\text{IND}}^{\text{sem}}(S)$ by first constructing the prBAF $\mathcal{F}(\phi)$, and next returning $2^{n+m} \cdot P_{\mathcal{F}(\phi)}^{\text{sem}}(\{a\})$ as the number of satisfying assignments of ϕ . \square

4.1 Tractable Cases for $P\text{-EXT}_{\text{IND}}^{\text{sem}}(S)$

The results above say that the combination of bipolarity and probabilities makes reasoning over extensions harder only for the form IND under the s - d - c -admissible and stable semantics, for which $P\text{-EXT}_{\text{IND}}^{\text{sem}}(S)$ is polynomial in the absence of supports (see Table 1). Thus, finding conditions restoring the polynomiality is worth investigating, as this gives an insight on the new source of complexity introduced by combining probabilities and supports. With this aim, we focus on the form IND and show a restricted form of prBAFs (namely, C -bounded prBAFs) for which $P\text{-EXT}_{\text{IND}}^{\text{sem}}(S)$ is polynomial. At the end of this section, we will show that the relevance of this restriction goes beyond the definition of a tractable class: our proof of tractability is based on introducing a general paradigm for computing extensions’ probability that can enhance the efficiency even when the restriction does not hold.

Formally, given a prBAF $\mathcal{F} = \langle \mathcal{A}, \mathcal{R}_a, \mathcal{R}_s, \mathcal{P}_A, \mathcal{P}_R \rangle$, consider the sets: i) $\mathcal{F}.\mathcal{A}_e = \{a \mid \exists \langle a, b \rangle \in \mathcal{R}_s \vee \exists \langle b, a \rangle \in \mathcal{R}_s\}$ (called set of *supp-arguments*, as they are those involved in supports), and ii) $\mathcal{F}.\mathcal{R}_e = \{\langle a, b \rangle \in (\mathcal{R}_a \cup \mathcal{R}_s) \mid (a \in \mathcal{A}_e \vee b \in \mathcal{A}_e)\}$ (called set of *supp-attacks* and *supports*, as they

are attacks/supports incident to supp-arguments). Given this, for any constant C , \mathcal{F} is said to be C -bounded if $|\mathcal{F}.\mathcal{R}_e| \leq C$ (meaning that the number of supports and of attacks to/from arguments involved in supports is bounded by C).

Our strategy is based on the notions of *contraction* and *completion*. A *contraction* for a prBAF $\mathcal{F} = \langle \mathcal{A}, \mathcal{R}_a, \mathcal{R}_s, \mathcal{P}_A, \mathcal{P}_R \rangle$ is a prBAF $\mathcal{F}^* = \langle \mathcal{A}^*, \mathcal{R}_a^*, \mathcal{R}_s^*, \mathcal{P}_A^*, \mathcal{P}_R^* \rangle$ where:

- $\mathcal{A}^* \subseteq \mathcal{A}$ and $\mathcal{A} \setminus \mathcal{A}^* \subseteq \mathcal{F}.\mathcal{A}_e$;
- $\mathcal{R}_a^* \subseteq \mathcal{R}_a \cap (\mathcal{A}^* \times \mathcal{A}^*)$ and $\mathcal{R}_a \setminus \mathcal{R}_a^* \subseteq \mathcal{F}.\mathcal{R}_e$;
- $\mathcal{R}_s^* \subseteq \mathcal{R}_s \cap (\mathcal{A}^* \times \mathcal{A}^*)$;
- $\mathcal{P}_A^*(a) = 1$ if $a \in \mathcal{F}.\mathcal{A}_e$ and $\mathcal{P}_A^*(a) = \mathcal{P}_A(a)$ otherwise;
- $\mathcal{P}_R^*((a, b)) = 1$ if $a \in \mathcal{F}.\mathcal{A}_e \vee b \in \mathcal{F}.\mathcal{A}_e$, and $\mathcal{P}_R^*((a, b)) = \mathcal{P}_R((a, b))$ otherwise.

Basically, \mathcal{F}^* 's supports are a subset of \mathcal{F} 's, and arguments (resp., attacks) are a subset of \mathcal{F} 's containing at least the non supp-arguments (resp., the non supp-attacks). Then, the probabilities are copied from those specified in \mathcal{F} , except for those over supp-arguments and attacks, that are overwritten with 1.

$\mathcal{E}(\mathcal{F})$ will denote the set of possible contractions of \mathcal{F} . For $\mathcal{F}^* \in \mathcal{E}(\mathcal{F})$, the probability of \mathcal{F}^* given \mathcal{F} is: $P(\mathcal{F}^*|\mathcal{F}) = \prod_{a \in \mathcal{A}_e} P_{\mathcal{F}^*}(a) \times \prod_{\delta \in \mathcal{R}_e \cap (\mathcal{A}^* \times \mathcal{A}^*)} P_{\mathcal{F}^*}(\delta)$, where:

- if $a \in \mathcal{A}^*$, $P_{\mathcal{F}^*}(a) = \mathcal{P}_A(a)$; else, $P_{\mathcal{F}^*}(a) = 1 - \mathcal{P}_A(a)$;
- if $\delta \in \mathcal{R}_a^* \cup \mathcal{R}_s^*$, $P_{\mathcal{F}^*}(\delta) = \mathcal{P}_R(\delta)$; else, $P_{\mathcal{F}^*}(\delta) = 1 - \mathcal{P}_R(\delta)$.

As regards *completions*, their definition uses the function $\text{cert}(\mathcal{F})$, returning the BAF consisting of all and only the certain arguments/attacks/supports of \mathcal{F} (i.e., those having probability 1). Thus, the *completion* of \mathcal{F} is the prBAF $\text{compl}(\mathcal{F}) = \langle \mathcal{A}', \mathcal{R}'_a, \mathcal{R}'_s, \mathcal{P}'_A, \mathcal{P}'_R \rangle$ where:

- $\mathcal{A}' = \mathcal{A}$ and $\mathcal{R}'_s = \mathcal{R}_s$;
- $\mathcal{R}'_a = \mathcal{R}_a \cup R'$, where R' consists of the s- or the d- attacks of $\text{cert}(\mathcal{F})$, depending on whether \mathcal{F} is an s- or a d- prBAF;
- $\forall a \in \mathcal{A}$, $\mathcal{P}'_A(a) = \mathcal{P}_A(a)$;
- $\forall \delta \in \mathcal{R}'_a$, if $\delta \in R'$ then $\mathcal{P}'_R(\delta) = 1$, else $\mathcal{P}'_R(\delta) = \mathcal{P}_R(\delta)$.

Next lemma allows for decomposing the evaluation of $P_{\mathcal{F}}^{\text{sem}}(S)$ into evaluating $P_{\mathcal{F}^*}^{\text{sem}}(S)$ over each contraction \mathcal{F}^* .

Lemma 1 *Let \mathcal{F} be a prBAF and S a set of its arguments. For $\text{sem} \in \{d\text{-ad}, s\text{-ad}, c\text{-ad}, st\}$, it holds that $P_{\mathcal{F}}^{\text{sem}}(S) = \sum_{\mathcal{F}^* \in \mathcal{E}(\mathcal{F})} P(\mathcal{F}^*|\mathcal{F}) \times P_{\mathcal{F}^*}^{\text{sem}}(S)$.*

The following lemma provides a method for computing $P_{\mathcal{F}^*}^{\text{sem}}(S)$ for each $\mathcal{F}^* \in \mathcal{E}(\mathcal{F})$, as it states that $P_{\mathcal{F}^*}^{\text{sem}}(S)$ can be computed by taking the prAAF $\overline{\mathcal{F}}$ obtained by removing the supports from the completion of \mathcal{F}^* , and then using over $\overline{\mathcal{F}}$ any state-of-the-art algorithm for computing the extensions' probabilities over "traditional" prAAFs (that is polynomial for the semantics ad, st).

Lemma 2 *Let \mathcal{F} be a prBAF, \mathcal{F}^* a contraction for \mathcal{F} , $\overline{\mathcal{F}}$ the prAAF obtained from $\text{compl}(\mathcal{F}^*)$ by removing the supports, and S a set of arguments of \mathcal{F}^* . Then:*

- For $\text{sem} \in \{d\text{-ad}, st\}$ $P_{\mathcal{F}^*}^{\text{sem}}(S) = P_{\overline{\mathcal{F}}}^{\text{sem}'}(S)$, where $\text{sem}' \in \{ad, st\}$ respectively;
- $P_{\mathcal{F}^*}^{s\text{-ad}}(S) = 0$ if S is not safe over $\text{cert}(\text{compl}(\mathcal{F}^*))$; otherwise, $P_{\mathcal{F}^*}^{s\text{-ad}}(S) = P_{\overline{\mathcal{F}}}^{ad}(S)$;

– $P_{\mathcal{F}^*}^{c\text{-ad}}(S) = 0$ if S is not closed for \mathcal{R}_s over $\text{cert}(\text{compl}(\mathcal{F}^*))$; otherwise, $P_{\mathcal{F}^*}^{c\text{-ad}}(S) = P_{\overline{\mathcal{F}}}^{ad}(S)$.

Observe that lemmas 1 and 2 do not require that \mathcal{F} is C-bounded: they state general properties that, on C-bounded prBAFs, can be exploited to compute $\text{P-EXT}_{\text{IND}}^{\text{sem}}(S)$ efficiently (for what shown in the proof of Theorem 3).

Theorem 3 *For $\text{sem} \in \{d\text{-ad}, s\text{-ad}, c\text{-ad}, st\}$, $\text{P-EXT}_{\text{IND}}^{\text{sem}}(S)$ is in FP for both C-bounded s- and d- prBAFs.*

Proof. Lemma 1 ensures that $P_{\mathcal{F}}^{\text{sem}}(S)$ can be computed as $\sum_{\mathcal{F}^* \in \mathcal{E}(\mathcal{F})} P(\mathcal{F}^*|\mathcal{F}) \times P_{\mathcal{F}^*}^{\text{sem}}(S)$, where for each $\mathcal{F}^* \in \mathcal{E}(\mathcal{F})$. Lemma 2 and the fact that computing $P_{\overline{\mathcal{F}}}^{\text{sem}'}(S)$ over a prAAF is feasible in polynomial time for $\text{sem}' \in \{ad, st\}$ [Fazzinga et al., 2015] in turn imply that $P_{\mathcal{F}}^{\text{sem}}(S)$ can be computed in polynomial time w.r.t. the size of \mathcal{F} . It is easy to see that, since \mathcal{F} is C-bounded, $|\mathcal{E}(\mathcal{F})| \leq 2^C$, as the number of contractions is bounded by the number of subsets of $\mathcal{F}.\mathcal{R}_e$. Thus, the sum in Lemma 1 consists of a constant number of polynomial time computable addends. \square

Remark. The strategy used to prove the tractability of the C-bounded class can be applied also to non C-bounded prBAFs, where, even if it is not guaranteed to work in polynomial time, it is likely to perform better than an 'exhaustive' evaluation, as it requires only contractions to be enumerated.

5 Conclusions

We have characterized the computational complexity of evaluating extensions' probabilities over probabilistic BAFs. The most popular semantics of extensions and of supports have been considered, along with different probabilistic paradigms for encoding uncertainty. Interestingly, we have shown that, considering bipolarity does not do not increase the computational complexity, except for the case of admissible and stable semantics under the assumption of independence between the terms of the dispute. For this case, we have shown that the complexity explodes from FP to $FP^{\#P}$ -complete, and islands of tractability have been detected.

References

- [Amgoud and Vesic, 2011] Leila Amgoud and Srdjan Vesic. A new approach for preference-based argumentation frameworks. *A. Math. Artif. Intell.*, 63(2):149–183, 2011.
- [Baroni and Giacomin, 2009] Pietro Baroni and Massimiliano Giacomin. Semantics of abstract argument systems. In *Argum. in Artificial Intelligence*, pages 25–44. 2009.
- [Bench-Capon, 2003] Trevor J. M. Bench-Capon. Persuasion in practical argument using value-based argumentation frameworks. *J. Log. Comput.*, 13(3):429–448, 2003.
- [Boella et al., 2010] Guido Boella, Dov M. Gabbay, Leendert W. N. van der Torre, and Serena Villata. Support in abstract argumentation. In *Proc. of COMMA 2010*, pages 111–122, 2010.
- [Brewka and Woltran, 2010] Gerhard Brewka and Stefan Woltran. Abstract dialectical frameworks. In *Proc. of KR*, pages 102–111, 2010.

- [Cayrol and Lagasquie-Schiex, 2005] Claudette Cayrol and Marie-Christine Lagasquie-Schiex. On the acceptability of arguments in bipolar argumentation frameworks. In *Proc. of ECSQARU*, pages 378–389, 2005.
- [Cayrol and Lagasquie-Schiex, 2009] Claudette Cayrol and Marie-Cristine Lagasquie-Schiex. Bipolar abstract argumentation systems. In *Argum. in Artif. Intell.*, pages 65–84. 2009.
- [Cayrol and Lagasquie-Schiex, 2013] Claudette Cayrol and Marie-Christine Lagasquie-Schiex. Bipolarity in argumentation graphs: Towards a better understanding. *Int. J. Approx. Reasoning*, 54(7):876–899, 2013.
- [Cohen *et al.*, 2012] Andrea Cohen, Alejandro Javier García, and Guillermo Ricardo Simari. Backing and undercutting in abstract argumentation frameworks. In *Proc. of FoIKS*, pages 107–123, 2012.
- [Cohen *et al.*, 2014] Andrea Cohen, Sebastian Gottifredi, Alejandro Javier García, and Guillermo Ricardo Simari. A survey of different approaches to support in argumentation systems. *Know. Eng. Review*, 29(5):513–550, 2014.
- [Coste-Marquis *et al.*, 2012] Sylvie Coste-Marquis, Sébastien Konieczny, Pierre Marquis, and Mohand Akli Ouali. Weighted attacks in argumentation frameworks. In *Proc. of KR*, pages 593–597, 2012.
- [Doder and Woltran, 2014] Dragan Doder and Stefan Woltran. Probabilistic argumentation frameworks - A logical approach. In *Proc. of SUM*, pages 134–147, 2014.
- [Dondio, 2014] Pierpaolo Dondio. Toward a computational analysis of probabilistic argumentation frameworks. *Cybernetics and Systems*, 45(3):254–278, 2014.
- [Dung *et al.*, 2007] Phan Minh Dung, Paolo Mancarella, and Francesca Toni. Computing ideal sceptical argumentation. *Artif. Intell.*, 171(10-15):642–674, 2007.
- [Dung, 1995] Phan Minh Dung. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. *Artif. Intell.*, 77(2):321–358, 1995.
- [Dunne and Wooldridge, 2009] Paul E. Dunne and Michael Wooldridge. Complexity of abstract argumentation. In *Argum. in Artificial Intelligence*, pages 85–104. 2009.
- [Dunne *et al.*, 2011] Paul E. Dunne, Anthony Hunter, Peter McBurney, Simon Parsons, and Michael Wooldridge. Weighted argument systems: Basic definitions, algorithms, and complexity results. *Artif. Intell.*, 175(2), 2011.
- [Dunne, 2009] Paul E. Dunne. The computational complexity of ideal semantics. *Artif. Intell.*, 173(18), 2009.
- [Fazzinga *et al.*, 2013] Bettina Fazzinga, Sergio Flesca, and Francesco Parisi. On the complexity of probabilistic abstract argumentation. In *Proc. of IJCAI*, pages 898–904, 2013.
- [Fazzinga *et al.*, 2015] Bettina Fazzinga, Sergio Flesca, and Francesco Parisi. On the complexity of probabilistic abstract argumentation frameworks. *ACM Trans. Comput. Log. (TOCL)*, 16(3):22, 2015.
- [Fazzinga *et al.*, 2016] Bettina Fazzinga, Sergio Flesca, and Filippo Furfaro. Computing extensions’ probabilities in probabilistic abstract argumentation: Beyond independence. In *Proc. of ECAI*, pages 1588–1589, 2016.
- [Hunter and Thimm, 2014a] Anthony Hunter and Matthias Thimm. Probabilistic argumentation with epistemic extensions. In *Proc. of DARE@ECAI*, 2014.
- [Hunter and Thimm, 2014b] Anthony Hunter and Matthias Thimm. Probabilistic argumentation with incomplete information. In *Proc. of ECAI*, pages 1033–1034, 2014.
- [Hunter, 2012] Anthony Hunter. Some foundations for probabilistic abstract argumentation. In *Proc. of COMMA*, pages 117–128, 2012.
- [Hunter, 2013] Anthony Hunter. A probabilistic approach to modelling uncertain logical arguments. *Int. J. Approx. Reasoning*, 54(1):47–81, 2013.
- [Hunter, 2014] Anthony Hunter. Probabilistic qualification of attack in abstract argumentation. *Int. J. Approx. Reasoning*, 55(2):607–638, 2014.
- [Li *et al.*, 2011] Hengfei Li, Nir Oren, and Timothy J. Norman. Probabilistic argumentation frameworks. In *Proc. of TAFE*, 2011.
- [Martínez *et al.*, 2006] Diego C. Martínez, Alejandro Javier García, and Guillermo Ricardo Simari. On acceptability in abstract argumentation frameworks with an extended defeat relation. In *Proc. of COMMA*, pages 273–278, 2006.
- [Modgil, 2009] Sanjay Modgil. Reasoning about preferences in argumentation frameworks. *Artif. Intell.*, 173(9-10):901–934, 2009.
- [Nouioua and Risch, 2010] Farid Nouioua and Vincent Risch. Bipolar argumentation frameworks with specialized supports. In *Proc. of ICTAI*, pages 215–218, 2010.
- [Oren and Norman, 2008] Nir Oren and Timothy J. Norman. Semantics for evidence-based argumentation. In *Proc. of COMMA*, pages 276–284, 2008.
- [Polberg and Hunter, 2018] Sylwia Polberg and Antony Hunter. Empirical evaluation of abstract argumentation: Supporting the need for bipolar and probabilistic approaches. *Int. J. Approx. Reasoning*, 93:487–543, 2018.
- [Proietti, 2017] Carlo Proietti. Polarization and bipolar probabilistic argumentation frameworks. In *Proc. of Work. AI³*, pages 22–27, 2017.
- [Rienstra, 2012] Tjitze Rienstra. Towards a probabilistic dung-style argumentation system. In *Proc. of AT*, pages 138–152, 2012.
- [Thimm, 2012] Matthias Thimm. A probabilistic semantics for abstract argumentation. In *Proc. of ECAI*, pages 750–755, 2012.
- [Verheij, 2009] Bart Verheij. The toulmin argument model in artificial intelligence. In *Argum. in Artificial Intelligence*, pages 219–238. 2009.