Verification in Attack-Incomplete Argumentation Frameworks

Dorothea Baumeister, Daniel Neugebauer(✉), and Jörg Rothe

Heinrich-Heine-Universität Düsseldorf, 40225 Düsseldorf, Germany
{baumeister,neugebauer,rothe}@cs.uni-duesseldorf.de

Abstract. We tackle the problem of expressing incomplete knowledge about the attack relation in abstract argumentation frameworks. In applications, incomplete argumentation frameworks may arise as intermediate states in an elicitation process, when merging different beliefs about an argumentation framework's state, or in cases where the complete information cannot be fully obtained. To this end, we employ a model introduced by Cayrol et al. [10] and analyze the question of whether certain justification criteria are possibly (or necessarily) fulfilled, i.e., whether they are fulfilled in some (or in every) completion of the incomplete argumentation framework. We formally extend the definition of existing criteria to these incomplete argumentation frameworks and provide characterization and complexity results for variants of the verification problem.

1 Introduction

Argumentation frameworks are used to model discussions and deliberations among agents, be it human beings or software agents. The aim is to find sets of arguments that can be considered "justified" by satisfying certain properties. In a pathbreaking paper, Dung [17] introduced a formal model to describe argumentation frameworks and their semantics, which abstracts from the content of arguments and regards their interaction only. More background on abstract argumentation in artificial intelligence can be found in the book by Rahwan and Simari [30].

We revisit a generalized model for abstract argumentation frameworks originally proposed by Cayrol et al. [10] who extend the classical model to an attack-incomplete setting. In *attack-incomplete argumentation frameworks*, all arguments are known, but the set of all possible attacks between them is partitioned into attacks that are either known to definitely exist, or known to definitely never exist, or currently unknown to exist but that may potentially arise in the future. We study central properties and semantics of argumentation frameworks, such as conflict-freeness, admissibility, stability, preferredness, completeness, and groundedness [17], which we extend to the attack-incomplete setting by asking whether they are *possibly* or *necessarily* fulfilled. As our technical contribution, we provide characterization and complexity results for variants of the standard verification problem in attack-incomplete argumentation frameworks.

© Springer International Publishing Switzerland 2015
T. Walsh (Ed.): ADT 2015, LNAI 9346, pp. 341–358, 2015.
DOI: 10.1007/978-3-319-23114-3_21

Related Work and Motivation: Our work is motivated by the “Online Participation” project, an interdisciplinary graduate college of HHU Düsseldorf and other institutions[1] in which researchers from economics, communication theory, political sciences, social sciences, law, and computer science are participating. A central goal in this project is to build an internet platform that can be used for online discussions and deliberations. While these—as mentioned above—can be modeled abstractly by argumentation frameworks, a major drawback of the classical model due to Dung [17] is that it assumes complete knowledge of the arguments and the attack relation, that is, the *process* of arguing is assumed to have been completed already. However, such complete information is rarely available in practical applications; rather, one would like to model such an online discussion *dynamically*, evolving over time.

First ideas regarding dynamic changes in argumentation frameworks applying the theory of belief revision are due to Cayrol et al. [11], who also survey the literature on the dynamics of abstract argumentation frameworks [12]. They limit themselves to the addition or deletion of one argument, together with a respective change in the attack relation. Their work focuses on a classification of how and why those changes can alter the set of extensions of the given argumentation framework. Boella et al. [6] define general principles for the abstraction of arguments and attacks for the grounded semantics mainly. Liao et al. [26] investigate the question of how one can efficiently compute the status of an argument (i.e., whether it is accepted, rejected, or undecided) upon changing the arguments and attacks. Coste-Marquis et al. [14] study how belief revision postulates can be applied to argumentation systems.

Also, the concept of incomplete knowledge in abstract argumentation has recently received some attention. In *probabilistic argumentation frameworks* (see, for example, the work of Li et al. [25], Rienstra [31], Fazzinga et al. [19,20], Hunter [22], and Doder and Woltran [16]), arguments and/or attacks have an associated probability, which represents an agent’s degree of belief that the argument or attack is in force, or their reluctance to disregard the argument or attack. This can be considered as a quantified model of uncertainty that allows to derive the probability of certain criteria to hold. Baumeister et al. [5] study a model of *argument*-incomplete argumentation frameworks.

Cayrol et al. [10] propose argumentation frameworks with an additional “ignorance relation” among arguments that contains the attacks for which there is uncertainty. We adopt their extended framework model, but take a different perspective: In their work [10], new semantics for attack-incomplete argumentation frameworks are defined, which puts a lot of focus on the incomplete framework itself, rather than on its completions. Opposed to that, we analyze whether standard semantics apply in some (or all) completions of an incomplete framework. This is a natural question arising when dealing with incomplete

[1] Besides four faculties of HHU Düsseldorf and the Fachhochschule für öffentliche Verwaltung NRW, the practice partners of this project include registered societies, limited liability companies, and the municipal councils of Köln, Bonn, and Münster, among others. We refer to the website http://www.fortschrittskolleg.de for more details.

knowledge and has already been considered for similar notions of uncertainty in various areas. In the related field of computational social choice (see, e.g., the book chapter by Brandt et al. [9]), and especially so in voting, classical complete-information settings have been extended to allow for incomplete information as well. The book chapters by Boutilier and Rosenschein [7] and Baumeister and Rothe [1] survey the known results on incomplete information and communication in voting, in particular covering the concepts of *possible* and *necessary winners* in elections that have been introduced by Konczak and Lang [23] and studied in terms of computational complexity both for the original problems (see, e.g., [23,34]) and for a number of variants, such as possible winners when new alternatives are added [13], when there is uncertainty about which voting rule is used [2], and when there is uncertainty about the voters' weights in weighted elections [3].[2] The notions of possible and necessary winners have also been transferred to other fields where information may be incomplete, including fair division [8], algorithmic game theory [24], and judgment aggregation [4].

In Sect. 2, we describe the classical model of abstract argumentation frameworks, and we provide the needed notions from complexity theory. In Sect. 3, we introduce attack-incomplete argumentation frameworks and in Sect. 4 we present our results. In Sect. 5, we give our conclusions and state some open questions.

2 Preliminaries

In this section, we introduce the classical argumentation framework model and the notation used in this paper and provide some basic notions of complexity theory. Our models are based on the seminal work of Dung [17] who introduced an abstract model for argumentation frameworks; while using his notions and concepts, we adopt some notation from the book chapter by Dunne and Wooldridge [18].

An *argumentation framework* is a pair $AF = \langle \mathscr{A}, \mathscr{R} \rangle$ that contains a set $\mathscr{A}$ of n *arguments* and a binary *attack relation* $\mathscr{R} \subseteq \mathscr{A} \times \mathscr{A}$ of up to n^2 pairs of arguments. We say that a *attacks* b if $(a, b) \in \mathscr{R}$. Given an argumentation framework $AF = \langle \mathscr{A}, \mathscr{R} \rangle$, the set of *attackers* of a set B of arguments is $\{a \in \mathscr{A} \mid \exists b \in B : (a, b) \in \mathscr{R}\}$. We say that a set D of arguments *defends* a set B of arguments if for each attacker a of B, there is an argument $d \in D$ with $(d, a) \in \mathscr{R}$. Accordingly, D does not defend B if there is an attacker of B that is not attacked by any $d \in D$.

Every argumentation framework can be illustrated as a directed graph $G = (V, E)$ by identifying $V = \mathscr{A}$ and $E = \mathscr{R}$ (see Example 1 and Fig. 1).

Example 1. A very basic argumentation framework is $AF = \langle \mathscr{A}, \mathscr{R} \rangle$ with the argument set $\mathscr{A} = \{a, b, c, d\}$ and the attacks $\mathscr{R} = \{(a, b), (a, c), (a, d), (b, d), (c, c), (d, a), (d, b)\}$ (see Fig. 1 for its graph representation).

[2] Other models of incomplete-information settings in voting include dynamic social choice with evolving preferences [29] and online manipulation in sequential elections [21].

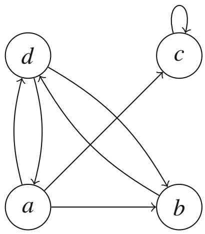

Fig. 1. Graph representation of the argumentation framework in Example 1

We now formally define the properties of sets of arguments in argumentation frameworks that were introduced in Dung's initial work and that are central to this paper:

1. The most basic property is conflict-freeness, which simply forbids attacks within a subset of the arguments. Formally, a subset $S \subseteq \mathscr{A}$ is *conflict-free* if there are no arguments a and b in S such that $(a, b) \in \mathscr{R}$.
2. An argument $a \in \mathscr{A}$ is *acceptable with respect to* $S \subseteq \mathscr{A}$ if S defends a, i.e., if for all $b \in \mathscr{A}$ with $(b, a) \in \mathscr{R}$, there is at least one $c \in S$ with $(c, b) \in \mathscr{R}$.
3. Further, a conflict-free set S of arguments is called *admissible* if every argument $a \in S$ is acceptable with respect to S.

Dung defines several *semantics* based on these properties in his original work, namely the *preferred*, *stable*, *complete*, and *grounded semantics*. Subsequent papers on argumentation frameworks proposed a variety of further semantics, but we will only be concerned with the semantics mentioned above.[3]

1. A set $S \subseteq \mathscr{A}$ is *preferred* if S is a maximal (with respect to set inclusion) admissible set.
2. A conflict-free set $S \subseteq \mathscr{A}$ is *stable* if it attacks all other arguments, i.e., if for every argument $b \in \mathscr{A} \setminus S$, there exists an $a \in S$ with $(a, b) \in \mathscr{R}$.
3. The *complete* semantics is defined via the *characteristic function* of an argumentation framework AF, which is $F_{AF} : 2^{\mathscr{A}} \to 2^{\mathscr{A}}$ with
$$F_{AF}(S) = \{a \in \mathscr{A} \mid a \text{ is acceptable with respect to } S\}.$$
The characteristic function is monotonic with respect to set inclusion and there is always an $i \in \mathbb{N}$ for which the i-fold composition of F_{AF} has a fixed point. A set $S \subseteq \mathscr{A}$ is *complete* if it is a fixed point of F_{AF}, or equivalently, if every $a \in \mathscr{A}$ that is acceptable with respect to S is contained in S.
4. The (unique) *grounded set* of an argumentation framework AF is the least (with respect to set inclusion) fixed point of F_{AF}, i.e., the complete set obtained when starting with the empty set.

[3] In addition to these semantics we also use conflict-freeness and admissibility as criteria. While these are generally not considered to be semantics, we will not always explicitly distinguish between semantics and basic properties for the sake of conciseness.

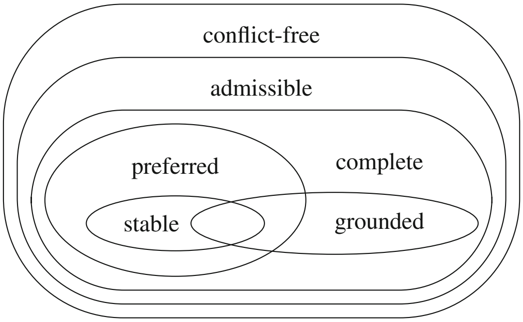

Fig. 2. Relations among the various criteria and semantics for sets of arguments

Dung investigates how these properties are correlated, and provides many results for this that we can make use of. Figure 2 displays all relations among the various criteria and semantics that we use. If an area labeled with criterion $\mathbf{s}$ is fully included in an area labeled with criterion $\mathbf{s}'$, this indicates that in all argumentation frameworks all sets of arguments that fulfill $\mathbf{s}$ also fulfill $\mathbf{s}'$. The converse is not necessarily true, i.e., all displayed set inclusions are strict. Further, none of the areas are disjoint, so one and the same set of arguments might fulfill all criteria/semantics simultaneously.

Given an argumentation framework AF and a semantics $\mathbf{s}$, a set S of arguments that fulfills the conditions imposed by $\mathbf{s}$ in AF is also called an $\mathbf{s}$ *extension of* AF. If it is clear from the context, we omit stating explicitly the argumentation framework that the subset is an extension of.

Dunne and Wooldridge [18] give an overview of a number of decision problems, each defined for various semantics. We will focus on only one of them, namely, the verification problem.

s-VERIFICATION	
Given:	An argumentation framework $\langle \mathscr{A}, \mathscr{R} \rangle$ and a subset $S \subseteq \mathscr{A}$.
Question:	Is S an $\mathbf{s}$ extension?

Here, the letter $\mathbf{s}$ is a placeholder for a specific semantics. For better readability, we will sometimes use CF as a shorthand for *conflict-freeness*, AD for *admissibility*, PR for *preferredness*, ST for *stability*, CP for *completeness*, and GR for *groundedness*.

Other previously considered decision problems are, for example, **s**-EXISTENCE, **s**-CREDULOUS-ACCEPTANCE, and **s**-SKEPTICAL-ACCEPTANCE (for a definition, see, for example, [18]). Many of these problems are hard to decide: They are complete for the complexity classes NP, coNP, or even Π_2^p. By contrast, **s**-VERIFICATION is easy for most semantics $\mathbf{s}$ studied here, which follows immediately from the work of Dung [17], with the only exception being pr-Verification, which is known to be coNP-complete [15].

We assume the reader to be familiar with the basic notions of complexity theory, such as the complexity classes P, NP, and coNP mentioned above and with the notions of hardness and completeness (based on the polynomial-time many-one reducibility, $\leq_{\mathrm{m}}^{\mathrm{p}}$). $\Sigma_2^p = \mathrm{NP}^{\mathrm{NP}}$ and $\Pi_2^p = \mathrm{coNP}^{\mathrm{NP}}$ are the second level of the polynomial hierarchy, which has been introduced by Meyer and Stockmeyer [27,33]. It holds that $\mathrm{P} \subseteq \mathrm{NP} \subseteq \Sigma_2^p \cup \Pi_2^p$ and $\mathrm{P} \subseteq \mathrm{coNP} \subseteq \Sigma_2^p \cup \Pi_2^p$, and none of these inclusions is known to be strict. For further details, see, e.g., [28,32].

3 Attack-Incomplete Argumentation Frameworks

We will now consider argumentation frameworks with incomplete knowledge about the attack relation, where a set of n arguments is fixed and only a subset of all n^2 possible attacks is known to either definitely exist or to definitely not exist—the state of the remaining attacks is currently unknown. We call this an *attack-incomplete argumentation framework.*

3.1 Model and Formal Definitions

An extension of standard argumentation frameworks to attack-incomplete argumentation frameworks was proposed by Cayrol et al. [10], which allows to distinguish between definite attacks, impossible attacks, and possible attacks. We apply their extended model using a slightly different notation.

Definition 1. *An* attack-incomplete argumentation framework *is a triple $\langle \mathscr{A}, \mathscr{R}^+, \mathscr{R}^- \rangle$, where $\mathscr{A}$ is a nonempty set of arguments and $\mathscr{R}^+$ and $\mathscr{R}^-$ are disjoint subsets of $\mathscr{A} \times \mathscr{A}$. $\mathscr{R}^+$ denotes the set of all ordered pairs of arguments between which an attack is known to definitely exist, while $\mathscr{R}^-$ denotes the set of all ordered pairs of arguments between which an attack is known to never exist. The set of possible attacks $(\mathscr{A} \times \mathscr{A}) \smallsetminus (\mathscr{R}^+ \cup \mathscr{R}^-)$, which is implicitly given through $\mathscr{R}^+$ and $\mathscr{R}^-$, is denoted as $\mathscr{R}^?$.*

Let $AF = \langle \mathscr{A}, \mathscr{R}^+, \mathscr{R}^- \rangle$ be a given attack-incomplete argumentation framework. An argumentation framework $AF^* = \langle \mathscr{A}, \mathscr{R}^* \rangle$ with $\mathscr{R}^+ \subseteq \mathscr{R}^* \subseteq \mathscr{R}^+ \cup \mathscr{R}^?$ is called a *completion of* AF. Every attack-incomplete argumentation framework obviously has $2^{\|\mathscr{R}^?\|}$ different completions. In particular, we call the completion that discards all possible attacks ($\mathscr{R}^* = \mathscr{R}^+$) the *minimal completion of* AF, and the completion that includes all possible attacks ($\mathscr{R}^* = \mathscr{R}^+ \cup \mathscr{R}^?$) is called the *maximal completion of* AF.

We now extend the notions for classical argumentation frameworks that we described in Sect. 2 to attack-incomplete argumentation frameworks, distinguishing between properties holding either *possibly* or *necessarily.* Generally, a property holds *possibly* for an attack-incomplete argumentation framework AF if

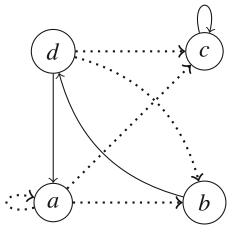

Fig. 3. An attack-incomplete argumentation framework

there *exists* a completion AF^* of AF for which the property holds, and a property holds *necessarily* if it holds *for all* completions of AF.[4] Note that if a property holds necessarily, it also holds possibly; even if $\mathscr{R}^? = \emptyset$, there exists $2^0 = 1$ completion, which happens to be both minimal and maximal in this case.

Example 2. Figure 3 gives the graph representation of an attack-incomplete argumentation framework $AF = \langle \mathscr{A}, \mathscr{R}^+, \mathscr{R}^- \rangle$ with the set $\mathscr{A} = \{a, b, c, d\}$ of arguments, where definite attacks $\mathscr{R}^+ = \{(b, d), (c, c), (d, a)\}$ are drawn as solid arcs, possible attacks $\mathscr{R}^? = \{(a, a), (a, b), (a, c), (d, b), (d, c)\}$ as dotted arcs, and attacks that are known to never exist (i.e., $\mathscr{R}^- = \{(a, d), (b, a), (b, b), (b, c), (c, a), (c, b), (c, d), (d, d)\}$) are not displayed.

For this example, it holds that the sets $\emptyset$, $\{b\}$, and $\{d\}$ are the only necessarily conflict-free extensions of AF, while the sets $\emptyset$, $\{a\}$, $\{b\}$, $\{d\}$, and $\{a, b\}$ are the only possibly conflict-free extensions of AF.

Deciding whether a given property holds possibly (respectively, necessarily) adds an existential (respectively, universal) quantifier over an exponential space to the standard problem, potentially making it intractable or increasing its level of intractability. However, some problems remain easy to solve. This is obvious, for example, for the possible and necessary attack between two given arguments: An argument $a \in \mathscr{A}$ is possibly attacked (respectively, necessarily attacked) by an argument $b \in \mathscr{A}$ if and only if $(b, a) \notin \mathscr{R}^-$ (respectively, $(b, a) \in \mathscr{R}^+$), which can clearly be verified in polynomial time. In Sect. 4, we will present our results on the complexity of deciding whether a set of arguments is a possible or necessary **s**-extension for all considered semantics **s**.

3.2 Comparison with the Model of Cayrol et al. [10]

Although we use the notion of attack-incomplete argumentation framework due to Cayrol et al. [10] (called *Partial Argumentation Framework (PAF)* in their

[4] Unlike the concepts of *credulous* and *skeptical* acceptance in the related literature, which denote membership of arguments in, respectively, *some* and *all extensions* of a specific argumentation framework, our notions of properties holding *possibly* and *necessarily* describe criteria holding in, respectively, *some* and *all argumentation frameworks* (i.e., *completions*), and are therefore settled one level of abstraction higher.

Table 1. Comparison of properties of sets of arguments

Property from [10]		Our property
S is R-conflict-free	$\Longleftrightarrow$	S is possibly conflict-free
S is RI-conflict-free	$\Longleftrightarrow$	S is necessarily conflict-free
a is R-acceptable w.r.t. S	$\not\Longleftrightarrow$	a is possibly acceptable w.r.t. S
a is RI-acceptable w.r.t. S	$\Longleftrightarrow$	a is necessarily acceptable w.r.t. S
S is R-admissible	$\not\Longleftrightarrow$	S is possibly admissible
S is RI-admissible	$\Longleftrightarrow$	S is necessarily admissible
S is R-preferred	$\not\Longleftrightarrow$	S is possibly preferred
S is RI-preferred	$\not\Longleftrightarrow$	S is possibly preferred

work), we do not take the same perspective on properties and semantics: While they define new semantics for the PAFs *themselves*, we will analyze whether the conditions of standard semantics are fulfilled in some or all *completions* of it. This avoids the strange case where an incomplete framework satisfies some property, despite none of its completions satisfying this property; for example, in the model by Cayrol et al. it may be the case that a set S of arguments is the only RI-preferred extension[5] of an attack-incomplete argumentation framework, even though it is not a preferred extension for any of the framework's completions.

While the formal conditions imposed by both approaches coincide in some cases, they are generally different. Table 1 gives an overview of all criteria and semantics introduced by Cayrol et al.[6] and their counterparts in our model, and indicates whether or not they are equivalent. A formal proof of why equivalence does or does not hold in each individual case is omitted due to space constraints.

4 Possible and Necessary Verification

The problem **s**-Verification for standard argumentation frameworks naturally yields two problems for attack-incomplete argumentation frameworks, **s**-Att-Inc-Possible-Verification and **s**-Att-Inc-Necessary-Verification, for each semantics **s**.

s-Att-Inc-Possible-Verification (**s**-AttIncPV)	
Given:	An attack-incomplete argumentation framework $AF = \langle \mathscr{A}, \mathscr{R}^+, \mathscr{R}^- \rangle$ and a set $S \subseteq \mathscr{A}$.
Question:	Is there a completion AF^* of AF such that S is an **s** extension in AF^*?

[5] A set of arguments is RI*-preferred* if it is maximal among all necessarily admissible sets, where $R \widehat{=} \mathscr{R}^+$ and $I \widehat{=} \mathscr{R}^?$ in our notation.

[6] For formal definitions of these criteria, see their work [10].

s-ATT-INC-NECESSARY-VERIFICATION (**s**-ATTINCNV)	
Given:	An attack-incomplete argumentation framework $AF = \langle \mathscr{A}, \mathscr{R}^+, \mathscr{R}^- \rangle$ and a set $S \subseteq \mathscr{A}$.
Question:	For all completions AF^* of AF, is S an **s** extension in AF^*?

As already mentioned, the original problem can be solved efficiently for the admissible, stable, complete, and grounded semantics. We prove that both new problems can still be solved efficiently for these semantics, even though the number of completions is exponential in the number of possible attacks. We define best-case and worst-case completions for the different semantics, a given attack-incomplete argumentation framework AF, and a given set S of arguments. Intuitively, a best-case completion includes all attacks that are beneficial for S with respect to the considered semantics, whereas a worst-case completion includes those attacks that harm the conditions imposed by the semantics. We prove that these completions are critical completions for the respective decision problem, i.e., the answer to the VERIFICATION variant corresponding to the attack-incomplete framework is the same as that to VERIFICATION for the respective completion.

4.1 Verifying Conflict-Freeness, Admissibility, and Stability

For conflict-freeness, admissibility, or stability of a set S of arguments, all attacks against elements of S are never beneficial and possibly harmful, and all attacks against arguments outside of S are never harmful and possibly beneficial. Thus the simple and straightforward "optimistic" and "pessimistic" completions, defined as follows, can serve as critical completions for these three criteria.

Definition 2. *Let $AF = \langle \mathscr{A}, \mathscr{R}^+, \mathscr{R}^- \rangle$ be an attack-incomplete argumentation framework and let $S \subseteq \mathscr{A}$. The* optimistic completion of AF *for* S *is* $AF_S^{\text{opt}} = \langle \mathscr{A}, \mathscr{R}_S^{\text{opt}} \rangle$ *with* $\mathscr{R}_S^{\text{opt}} = \mathscr{R}^+ \cup \{(a,b) \in \mathscr{R}^? \mid b \notin S\}$*. The* pessimistic completion of AF *for* S *is* $AF_S^{\text{pes}} = \langle \mathscr{A}, \mathscr{R}_S^{\text{pes}} \rangle$ *with* $\mathscr{R}_S^{\text{pes}} = \mathscr{R}^+ \cup \{(a,b) \in \mathscr{R}^? \mid b \in S\}$.

Example 3. Figure 4 displays the optimistic and the pessimistic completion for $S = \{a, b\}$ in the argumentation framework from Example 2: Possible attacks that are added to the set of definite attacks in the respective completion are drawn as boldfaced arcs, possible attacks that are not added to the set of definite attacks in the respective completion are omitted in Fig. 4(b) and (c), and the arguments in S are displayed by boldfaced circles.

Lemma 1. *Let $AF = \langle \mathscr{A}, \mathscr{R}^+, \mathscr{R}^- \rangle$ be an attack-incomplete argumentation framework, let $S \subseteq \mathscr{A}$, and let AF_S^{opt} be the optimistic completion of AF for S.*

1. *S is possibly conflict-free in AF if and only if S is a conflict-free extension of AF_S^{opt}.*

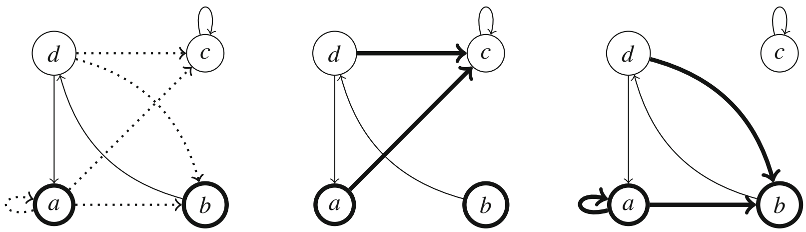

(a) Attack-incomplete AF (b) Optimistic completion (c) Pessimistic completion

Fig. 4. Optimistic and pessimistic completions for $S = \{a, b\}$

2. *$a \in S$ is possibly acceptable with respect to S in AF if and only if a is acceptable with respect to S in AF_S^{opt}.*
3. *S is possibly admissible in AF if and only if S is an admissible extension of AF_S^{opt}.*
4. *S is possibly stable in AF if and only if S is a stable extension of AF_S^{opt}.*

Proof. The converse is trivial in all cases: If S fulfills a given criterion in AF_S^{opt}, this immediately yields that S possibly fulfills the criterion in AF. We now prove the other direction of the equivalence individually for each criterion:

1. If a set S of arguments is not conflict-free in AF_S^{opt}, then there must be an attack between elements of S in $\mathscr{R}_S^{\text{opt}}$, which must be already in $\mathscr{R}^+$ due to how $\mathscr{R}_S^{\text{opt}}$ is constructed, and which therefore exists in every completion of AF. Thus S is not a possibly conflict-free set in AF.
2. If there is some $a \in S$ that is not acceptable with respect to S in AF_S^{opt}, then it is attacked by some $b \in \mathscr{A}$ in $\mathscr{R}_S^{\text{opt}}$ and there is no attack from an element of S against b in $\mathscr{R}_S^{\text{opt}}$. By construction, $\mathscr{R}_S^{\text{opt}}$ does not contain any possible attacks (members of $\mathscr{R}^?$) that attack elements of S, and it contains all possible attacks that can defend S. Therefore, all attacks in $\mathscr{R}_S^{\text{opt}}$ against elements of S are already in $\mathscr{R}^+$, so the undefended attack from b against a is in every completion of AF. Since a cannot be acceptable with respect to S in any completion of AF, a is not possibly acceptable with respect to S in AF.
3. Assume that S is not an admissible extension in AF_S^{opt}, i.e., S is not conflict-free in AF_S^{opt} or there is some $a \in S$ that is not acceptable with respect to S in AF_S^{opt}. In either case, the previous results imply that S is not conflict-free in any completion of AF or a is not acceptable with respect to S in any completion of AF. Thus S is not a possibly admissible extension in AF.
4. If a set S of arguments is not stable in AF_S^{opt}, S is necessarily not conflict-free in AF or there is an $a \in \mathscr{A} \setminus S$ that is not attacked by S in AF_S^{opt}, and therefore—by construction of AF_S^{opt}—a cannot be attacked by S in any completion of AF. In both cases, there is no completion of AF for which S is stable, so S is not a possibly stable extension of AF.

This completes the proof. □

An analogous result holds for the pessimistic completion and the same properties holding necessarily. The proof of Lemma 2 is omitted due to space constraints.

Lemma 2. *Let* $AF = \langle \mathscr{A}, \mathscr{R}^+, \mathscr{R}^- \rangle$ *be an attack-incomplete argumentation framework,* $S \subseteq \mathscr{A}$*, and let* AF_S^{pes} *be the pessimistic completion of* AF *for* S*.*

1. *S is necessarily conflict-free in AF if and only if S is a conflict-free extension of AF_S^{pes}.*
2. *$a \in S$ is necessarily acceptable with respect to S in AF if and only if a is acceptable with respect to S in AF_S^{pes}.*
3. *S is necessarily admissible in AF if and only if S is an admissible extension of AF_S^{pes}.*
4. *S is necessarily stable in AF if and only if S is a stable extension of AF_S^{pes}.*

Note that in the second part of Lemmas 1 and 2 it is required that $a \in S$; the properties do not hold for the general case where $a \in \mathscr{A}$.

Finally, we can conclude that **s**-AttIncPV and **s**-AttIncNV are in P for conflict-freeness, admissibility, and stability.

Theorem 1. *For* $\mathbf{s} \in \{\text{CF}, \text{AD}, \text{ST}\}$*, both* **s**-AttIncPV *and* **s**-AttIncNV *are in* P*.*

Proof. The optimistic and pessimistic completions can obviously be constructed in polynomial time. As already mentioned, the problem **s**-Verification can be solved in polynomial time for a given completion. Lemmas 1 and 2 then provide that the answer to, respectively, **s**-AttIncPV and **s**-AttIncNV is the same as that to **s**-Verification for the respective completion. □

4.2 Verifying Groundedness and Completeness

Recall that, for a given argumentation framework AF, the set of complete extensions is the set of fixed points of the characteristic function F_{AF}, and the (unique) grounded extension is the fixed point of the characteristic function F_{AF} when starting with the empty set. For the complete and the grounded semantics, a critical completion of an attack-incomplete argumentation framework for a given set S of arguments can be constructed by choosing attacks in a way that makes it most likely (respectively, most unlikely) for S to be a fixed point of F_{AF}. We call a completion in which S is most likely to be a fixed point of F_{AF} a "fixed completion," and a completion in which it is most unlikely to be a fixed point an "unfixed completion."

Definition 3. *Let* $AF = \langle \mathscr{A}, \mathscr{R}^+, \mathscr{R}^- \rangle$ *be an attack-incomplete argumentation framework and* $S \subseteq \mathscr{A}$*. The* fixed completion AF_S^{fix} *of* AF *is the completion that is obtained by the following algorithm. The algorithm defines a finite sequence* $(AF_i)_{i \geq 0}$ *of attack-incomplete argumentation frameworks, with the fixed completion being the minimal completion of the sequence's last element.*

1. *Include definite attacks: Let* $AF_0 = AF$.
2. *Include external conflicts: Let* $AF_1 = \langle \mathscr{A}, \mathscr{R}_1^+, \mathscr{R}^- \rangle$ *with* $\mathscr{R}_1^+ = \mathscr{R}^+ \cup \{(a,b) \in \mathscr{R}^? \mid a \notin S \text{ and } b \notin S\}$.
3. *Include defending attacks: Let* $T = \{t \in \mathscr{A} \setminus S \mid \exists s \in S : (t,s) \in \mathscr{R}_1^+\}$ *(i.e., each argument in* T *necessarily attacks* S*) and let* $AF_2 = \langle \mathscr{A}, \mathscr{R}_2^+, \mathscr{R}^- \rangle$ *with* $\mathscr{R}_2^+ = \mathscr{R}_1^+ \cup \{(a,b) \in \mathscr{R}_1^? \mid a \in S \text{ and } b \in T\}$.
4. *Avoid arguments outside of* S *to be acceptable with respect to* S*: For the current* i *(initially,* $i = 2$*), let* $AF_i^{\min}$ *be the minimal completion of* AF_i *and* $T_i = F_{AF_i^{\min}}(S) \setminus S$ *(i.e.,* T_i *is the set of arguments that are not in* S*, but that are acceptable with respect to* S *in the current minimal completion). Let* $AF_{i+1} = \langle \mathscr{A}, \mathscr{R}_{i+1}^+, \mathscr{R}^- \rangle$ *with* $\mathscr{R}_{i+1}^+ = \mathscr{R}_i^+ \cup \{(a,b) \in \mathscr{R}_i^? \mid a \in S \text{ and } b \in T_i\}$*, and set* $i \leftarrow i + 1$.
5. *Repeat Step 4 until no more attacks are added.*
6. *The fixed completion of* AF *is* $AF_S^{\text{fix}} = \langle \mathscr{A}, \mathscr{R}_S^{\text{fix}} \rangle$ *with* $\mathscr{R}_S^{\text{fix}} = \mathscr{R}_i^+$.

Definition 4. *Let* $AF = \langle \mathscr{A}, \mathscr{R}^+, \mathscr{R}^- \rangle$ *be an attack-incomplete argumentation framework and* $S \subseteq \mathscr{A}$*. The* unfixed completion AF_S^{unf} *of* AF *is the completion that is obtained by the following algorithm. The algorithm defines a finite sequence* $(AF_i)_{i \geq 0}$ *of attack-incomplete argumentation frameworks, with the unfixed completion being the minimal completion of the sequence's last element.*

1. *Include definite attacks: Let* $AF_0 = AF$.
2. *Include attacks against* S*: Let* $AF_1 = \langle \mathscr{A}, \mathscr{R}_1^+, \mathscr{R}_1^- \rangle$ *with* $\mathscr{R}_1^+ = \mathscr{R}^+ \cup \{(a,b) \in \mathscr{R}^? \mid b \in S\}$ *and* $\mathscr{R}_1^- = \mathscr{R}^-$.
3. *Exclude external conflicts: Let* $AF_2 = \langle \mathscr{A}, \mathscr{R}_2^+, \mathscr{R}_2^- \rangle$ *with* $\mathscr{R}_2^+ = \mathscr{R}_1^+$ *and* $\mathscr{R}_2^- = \mathscr{R}_1^- \cup \{(a,b) \in \mathscr{R}_1^? \mid a \notin S \text{ and } b \notin S\}$.
4. *Exclude defending attacks: Let* $T = \{t \in \mathscr{A} \setminus S \mid \exists s \in S : (t,s) \in \mathscr{R}_2^+\}$ *(i.e., each argument in* T *necessarily attacks* S*) and let* $AF_3 = \langle \mathscr{A}, \mathscr{R}_3^+, \mathscr{R}_3^- \rangle$ *with* $\mathscr{R}_3^+ = \mathscr{R}_2^+$ *and* $\mathscr{R}_3^- = \mathscr{R}_2^- \cup \{(a,b) \in \mathscr{R}_2^? \mid a \in S \text{ and } b \in T\}$.
5. *Try to let arguments outside of* S *be acceptable with respect to* S*: Let* $T = \mathscr{A} \setminus S = \{t_1, \ldots, t_k\}$*. For the current* i *(initially,* $i = 3$*) and for each* $t_j \in T$*, do:*
   (a) For $S' = S \cup \{t_j\}$*, let* $AF_{i,S'}^{\text{opt}}$ *be the optimistic completion of* AF_i *for* S' *and let* $AF_i^{\min}$ *be the minimal completion of* AF_i.
   (b) If t_j *is acceptable with respect to* S *in* $AF_{i,S'}^{\text{opt}}$*, but not acceptable with respect to* S *in* $AF_i^{\min}$*, let* $AF_{i+1} = \langle \mathscr{A}, \mathscr{R}_{i+1}^+, \mathscr{R}_{i+1}^- \rangle$ *with* $\mathscr{R}_{i+1}^+ = \mathscr{R}_i^+ \cup \{(a,b) \in \mathscr{R}_i^? \mid a \in S \text{ and } (b, t_j) \in \mathscr{R}_i^+\}$ *and* $\mathscr{R}_{i+1}^- = \mathscr{R}_i^-$*, and set* $i \leftarrow i+1$*. (To accept an argument* t_j *that is not currently accepted by* S *but possibly accepted by* S*, include all possible attacks by* S *against* t_j*'s attackers.)*
6. *Repeat Step 5 until no more attacks are added.*
7. *The unfixed completion of* AF *is* $AF_S^{\text{unf}} = \langle \mathscr{A}, \mathscr{R}_S^{\text{unf}} \rangle$ *with* $\mathscr{R}_S^{\text{unf}} = \mathscr{R}_i^+$.

Lemma 3. *For an attack-incomplete argumentation framework* $AF = \langle \mathscr{A}, \mathscr{R}^+, \mathscr{R}^- \rangle$ *and a set* $S \subseteq \mathscr{A}$ *of arguments, the fixed completion* AF_S^{fix} *and the unfixed completion* AF_S^{unf} *can be constructed in polynomial time.*

Proof. All individual steps in both constructions can obviously be carried out in time polynomial in the number of arguments. It remains to prove that the loops in, respectively, Step 4 and Step 5 run at most a polynomial number of times. For the fixed completion, in each execution of a loop there is either (at least) one possible attack that is added to $\mathscr{R}^{+}_{i+1}$, or no action is taken in which case the loop terminates. Therefore, the number of times a loop is executed is bounded by the number of possible attacks in the attack-incomplete argumentation framework AF, which is at most n^2, where n is the number of arguments. For the unfixed completion, the only difference is the sub-loop in Step 5, which however has a predefined number of iterations that is bounded by the number n of arguments. Therefore, the total number of loop iterations in the construction of the unfixed completion is bounded by n^3. This completes the proof. □

Now we prove that the fixed completion indeed is a critical completion for the complete and the grounded semantics.

Lemma 4. *Let $AF = \langle \mathscr{A}, \mathscr{R}^+, \mathscr{R}^- \rangle$ be an attack-incomplete argumentation framework, $S \subseteq \mathscr{A}$, and let AF^{fix}_S be the fixed completion of AF for S. (1) S is a possibly complete extension of AF if and only if S is a complete extension of AF^{fix}_S. (2) S is a possibly grounded extension of AF if and only if S is the grounded extension of AF^{fix}_S.*

Proof. Again, the converse is trivial in both cases. Further, if S is not an admissible extension in AF^{fix}_S, then S is not admissible in any completion of AF, due to the same arguments that we used for the optimistic completion and, therefore, neither possibly complete nor possibly grounded in AF. So, we may assume that S is admissible in AF^{fix}_S.

(1) Assume that S is not a complete extension of AF^{fix}_S, i.e., S is not a fixed point of $F_{AF^{\text{fix}}_S}$. We will show that this implies that S is not possibly complete in AF. Since S is not a fixed point of $F_{AF^{\text{fix}}_S}$, there is an argument $b \notin S$ which is acceptable with respect to S in AF^{fix}_S.

We prove that, then, there must be some $c \notin S$ for which all attackers of c are attacked by S in AF^* ($c = b$ may or may not be the case) by individually covering all cases in which attacks are added to $\mathscr{R}^{\text{fix}}_S$:

All attacks from $\mathscr{R}^?$ between arguments outside of S, which are added to $\mathscr{R}^{\text{fix}}_S$ in Step 2, cannot make an argument $b \notin S$ acceptable with respect to S: If S did not attack all attackers of an argument before, it cannot do so after *more* attackers are added.

All attacks that are added in Step 3 are crucial for S to be admissible, and must therefore also be included in $\mathscr{R}^*$. In a case where multiple arguments in S attack a single attacker of S, it would be sufficient to include one of these defending attacks, but including all of them does not make a difference, since the criterion of being acceptable with respect to S does not distinguish between different elements of S.

All attacks that are added in Step 4 are attacks by S against arguments that are currently acceptable with respect to S. Since all possible attacks among arguments outside of S were already included in Step 2, the only way to destroy acceptability of these arguments is by S directly attacking them. Therefore, none of the attacks added in Step 4 can be omitted without making the respective argument acceptable with respect to S (again, it is not necessary to distinguish between multiple attacks by different arguments in S against the same argument). It is possible for a given $b \notin S$ to be acceptable with respect to S in AF_S^{fix} and not in AF^*, but this happens only if S attacks an attacker (or several attackers) of b in AF_S^{fix} that would otherwise be acceptable with respect to S, and which therefore must be acceptable with respect to S in AF^*. In either case, if an argument outside of S is acceptable with respect to S in AF_S^{fix}, then some argument outside of S must be acceptable with respect to S in each completion AF^* of AF in which S is admissible. Therefore, if S is not a complete extension of AF_S^{fix}, it is not a complete extension of any completion AF^* of AF, and therefore not a possibly complete extension of AF.

(2) Let AF^* be an arbitrary completion of AF and assume that S is its grounded extension. We prove that, then, S is also the grounded extension of AF_S^{fix}. Let $A_i = F^i_{AF^*}(\emptyset)$ and $B_i = F^i_{AF_S^{\text{fix}}}(\emptyset)$, where F^i is the i-fold composition of the respective characteristic function F. Since S is grounded in AF^*, it is complete in AF_S^{fix} due to our previous result, and it holds that $A_i \subseteq S$ for all $i \geq 0$ and there exists a $j \geq 0$ such that for all $i \geq j$, it holds that $A_i = S$. We will prove that $A_i \subseteq B_i \subseteq S$ for all $i \geq 0$. Combined, these statements show that there exists some j such that $B_i = S$ for all $i \geq j$, which is equivalent to S being the grounded extension of AF_S^{fix}.

First, we prove that $A_i \subseteq B_i$ for all $i \geq 0$. For $i = 0$, we have $A_i = B_i = \emptyset$. For $i = 1$, A_i (respectively, B_i) is the set of all unattacked arguments in AF^* (respectively, in AF_S^{fix}). We know that $A_1 \subseteq S$. Since the fixed completion does not include any possible attacks against elements of S, all $a \in S$ that are unattacked in AF^* are unattacked in AF_S^{fix}, too, which proves $A_1 \subseteq B_1$. If we now have $A_k \subseteq B_k$ for some $k \geq 1$, this implies $A_{k+1} \subseteq B_{k+1}$: Assume that this were not true, i.e., that $A_k \subseteq B_k$, but there is an argument $a \in A_{k+1}$ with $a \notin B_{k+1}$. Then, a is acceptable with respect to A_k in AF^*, but not acceptable with respect to B_k in AF_S^{fix}. We know that—since $A_{k+1} \subseteq S$—no possible attacks against A_{k+1} (and in particular, against a) are included in AF_S^{fix} and all possible defending attacks by arguments in A_{k+1} against arguments outside of S are included in AF_S^{fix}. Further, no element of S attacks a in AF_S^{fix}, since $a \in S$ and S is complete in AF_S^{fix}. Therefore, a is acceptable with respect to A_k in AF_S^{fix}; otherwise it could not be acceptable with respect to A_k in AF^*. Now, the only way for a to not be acceptable with respect to B_k in AF_S^{fix} is if there were some $b \in B_k \setminus A_k$ that necessarily attacks a. Then there would have to be a defending attack by an argument $d \in A_k$ against b in AF^*, since a is acceptable with respect to A_k in AF^*. This implies that $b \notin S$, since S is conflict-free in AF^*. Finally, since (d, b) is a possible (or even a necessary) defending attack by an element of S against $b \notin S$, $(d, b) \in \mathscr{R}_S^{\text{fix}}$ holds by construction of the fixed completion, which contradicts that B_k is admissible in AF_S^{fix}. Therefore, a must be acceptable with respect to B_k in AF_S^{fix}, which proves that $A_{k+1} \subseteq B_{k+1}$.

Now we prove that $B_i \subseteq S$ for all $i \geq 0$: Assume that $B_i \nsubseteq S$ for some $i \geq 0$. Then it also holds that $G_S^{\text{fix}} \nsubseteq S$ for the grounded extension G_S^{fix} of AF_S^{fix}. It further holds that $S \subset G_S^{\text{fix}}$, since there exists a $j \geq 0$ such that $S \subseteq B_i$ for all $i \geq j$, as established before. However, this contradicts the fact that S is complete in AF_S^{fix}, since the grounded extension G_S^{fix} of AF_S^{fix} is its least complete extension with respect to set inclusion and the complete set S cannot be a strict subset of G_S^{fix}. This completes the proof. □

Analogously, the unfixed completion is a critical completion for the complete and the grounded semantics. Due to limitation of space, the proof of Lemma 5 is omitted.

Lemma 5. *Let $AF = \langle \mathscr{A}, \mathscr{R}^+, \mathscr{R}^- \rangle$ be an attack-incomplete argumentation framework, $S \subseteq \mathscr{A}$, and let AF_S^{unf} be the unfixed completion of AF for S. (1) S is a necessarily complete extension of AF if and only if S is a complete extension of AF_S^{unf}. (2) S is a necessarily grounded extension of AF if and only if S is the grounded extension of AF_S^{unf}.*

Finally, our results allow us to establish that **s**-AttIncPV and **s**-AttIncNV are in P for the complete and grounded semantics.

Theorem 2. *For* $\mathbf{s} \in \{\text{CP}, \text{GR}\}$*, both* **s**-AttIncPV *and* **s**-AttIncNV *are in* P.

Proof. Lemma 3 provides polynomial-time constructability for the fixed and unfixed completion. Given a completion, **s**-Verification can be solved in polynomial time, and Lemmas 4 and 5 imply that the answer to, respectively, **s**-AttIncPV and **s**-AttIncNV is the same as that to **s**-Verification for the respective completion. □

4.3 Verifying Preferredness

As mentioned above, the Verification problem for the preferred semantics is coNP-complete. For pr-AttIncPV and pr-AttIncNV, we have the following results.

Theorem 3. *The problem* pr-AttIncPV *is in* Σ_2^p *and* coNP*-hard, and* pr-AttIncNV *is* coNP*-complete.*

Proof. In pr-AttIncPV one has to check whether, given an attack-incomplete argumentation framework $AF = \langle \mathscr{A}, \mathscr{R}^+, \mathscr{R}^- \rangle$ and a set $S \subseteq \mathcal{A}$, there is a completion $AF^* = \langle \mathcal{A}, \mathcal{R}^* \rangle$ such that S is preferred in AF^*. To check whether S is preferred in AF^*, one has to check whether for all sets $S' \subseteq \mathcal{A}$ with $S \subset S'$ it holds that S is an admissible extension and S' is not an admissible extension. Thus this problem is in Σ_2^p.

To see that pr-AttIncNV is in coNP, consider the complementary problem. Here one has to check whether there is a completion AF^* of the given attack-incomplete AF such that the given set S is *not* preferred. To see this, it is enough to check whether there is a strict superset of S that is admissible or whether S

Table 2. Overview of complexity results both in the standard model (**s**-VERIFICATION) and in the attack-incomplete model of this paper (**s**-ATTINCPV and **s**-ATTINCNV)

s	VERIFICATION	ATTINCPV		ATTINCNV	
CF	in P	in P	[10]	in P	[10]
AD	in P	in P	(Theorem 1)	in P	[10]
ST	in P	in P	(Theorem 1)	in P	(Theorem 1)
CP	in P	in P	(Theorem 2)	in P	(Theorem 2)
GR	in P	in P	(Theorem 2)	in P	(Theorem 2)
PR	coNP-complete	coNP-hard, in Σ_2^p	(Theorem 3)	coNP-complete	(Theorem 3)

itself is not admissible. Since admissibility can be checked in polynomial time, the complement of PR-ATTINCNV is in NP and hence PR-ATTINCNV is in coNP.

On the other hand, coNP-hardness for both problems follows by a direct reduction from the original PR-VERIFICATION problem, which is coNP-complete [15]. For a given instance $(\langle \mathscr{A}, \mathscr{R} \rangle, S)$ of PR-VERIFICATION, the constructed instance of both pr-ATTINCPV and pr-ATTINCNV is $(\langle \mathscr{A}, \mathscr{R}, (\mathscr{A} \times \mathscr{A}) \setminus \mathscr{R} \rangle, S)$. The only completion of $\langle \mathscr{A}, \mathscr{R}, (\mathscr{A} \times \mathscr{A}) \setminus \mathscr{R} \rangle$ is $\langle \mathscr{A}, \mathscr{R} \rangle$. Now, it is easy to see that $(\langle \mathscr{A}, \mathscr{R} \rangle, S) \in$ PR-VERIFICATION if and only if $(\langle \mathscr{A}, \mathscr{R}, (\mathscr{A} \times \mathscr{A}) \setminus \mathscr{R} \rangle, S) \in$ *pr*-ATTINCPV, which in turn is equivalent to $(\langle \mathscr{A}, \mathscr{R}, (\mathscr{A} \times \mathscr{A}) \setminus \mathscr{R} \rangle, S) \in$ *pr*-ATTINCNV. □

5 Conclusions and Open Questions

We have investigated argumentation frameworks in a setting where we don't have full knowledge of the attacks. We adapted the **s**-VERIFICATION decision problems with respect to notions of possibility and necessity to fit the model of Cayrol et al. [10], and we analyzed their complexity for the fundamental semantics admissibility, stability, completeness, groundedness, and preferredness. This may be useful to predict those sets of arguments that will be "good" solutions once all attacks are known eventually.

Table 2 summarizes our results, and also gives the previously known results for argumentation frameworks without uncertainty that are due to Dimopoulos and Torres [15], Dung [17], and Dunne and Wooldridge [18], as well as the results for incomplete argumentation frameworks provided by Cayrol et al. [10]. We have shown positive results (characterizations) for all considered semantics except preferredness, for which the exact complexity in the case of possible verification remains open. As a task for future work, we propose to generalize other decision problems like **s**-CREDULOUS-ACCEPTANCE, **s**-SKEPTICAL-ACCEPTANCE, **s**-EXISTENCE, and **s**-NONEMPTINESS to fit the model of attack-incompleteness and analyze their complexity. Additionally, one could have a closer look at other semantics like semi-stable, ideal, or prudent semantics (see [18] for the definition of these decision problems and semantics).

Acknowledgments. We thank the anonymous reviewers for their helpful comments. This work was supported in part by an NRW grant for gender-sensitive universities and the project "Online Participation," both funded by the NRW Ministry for Innovation, Science, and Research.

References

1. Baumeister, D., Rothe, J.: Preference aggregation by voting (chapter 4). In: Rothe, J. (ed.) Economics and Computation: An Introduction to Algorithmic Game Theory, Computational Social Choice, and Fair Division, pp. 197–325. Springer, Heidelberg (2015). doi:10.1007/978-3-662-47904-9_4
2. Baumeister, D., Roos, M., Rothe, J.: Computational complexity of two variants of the possible winner problem. In: Proceedings of AAMAS 2011, pp. 853–860. IFAAMAS (2011)
3. Baumeister, D., Roos, M., Rothe, J., Schend, L., Xia, L.: The possible winner problem with uncertain weights. In: Proceedings of ECAI 2012, pp. 133–138. IOS Press (2012)
4. Baumeister, D., Erdélyi, G., Erdélyi, O., Rothe, J.: Complexity of manipulation and bribery in judgment aggregation for uniform premise-based quota rules. Math. Soc. Sci. **76**, 19–30 (2015)
5. Baumeister, D., Rothe, J., Schadrack, H.: Verification in argument-incomplete argumentation frameworks. In: Walsh, T. (ed.) ADT 2015. LNCS (LNAI), vol. 9346, pp. xx–yy. Springer, Heidelberg (2015)
6. Boella, G., Kaci, S., van der Torre, L.: Dynamics in with single extensions: abstraction principles and the grounded extension. In: Sossai, C., Chemello, G. (eds.) ECSQARU 2009. LNCS, vol. 5590, pp. 107–118. Springer, Heidelberg (2009)
7. Boutilier, C., Rosenschein, J.: Incomplete information and communication in voting (chapter 10). In: Brandt, F., Conitzer, V., Endriss, U., Lang, J., Procaccia, A. (eds.) Handbook of Computational Social Choice. Cambridge University Press, Cambridge (2015)
8. Bouveret, S., Endriss, U., Lang, J.: Fair division under ordinal preferences: computing envy-free allocations of indivisible goods. In: Proceedings of ECAI 2010, pp. 387–392. IOS Press (2010)
9. Brandt, F., Conitzer, V., Endriss, U.: Computational social choice. In: Weiß, G. (ed.) Multiagent Systems, 2nd edn, pp. 213–283. MIT Press, Cambridge (2013)
10. Cayrol, C., Devred, C., Lagasquie-Schiex, M.C.: Handling ignorance in argumentation: semantics of partial argumentation frameworks. In: Mellouli, K. (ed.) ECSQARU 2007. LNCS (LNAI), vol. 4724, pp. 259–270. Springer, Heidelberg (2007)
11. Cayrol, C., de Saint-Cyr, F., Lagasquie-Schiex, M.: Revision of an argumentation system. In: Proceedings of KR 2008, pp. 124–134. AAAI Press (2008)
12. Cayrol, C., de Saint-Cyr, F., Lagasquie-Schiex, M.: Change in abstract argumentation frameworks: adding an argument. J. Artif. Intell. Res. **38**, 49–84 (2010)
13. Chevaleyre, Y., Lang, J., Maudet, N., Monnot, J., Xia, L.: New candidates welcome! Possible winners with respect to the addition of new candidates. Math. Soc. Sci. **64**(1), 74–88 (2012)
14. Coste-Marquis, S., Konieczny, S., Mailly, J., Marquis, P.: On the revision of argumentation systems: minimal change of arguments statuses. In: Proceedings of KR 2014. AAAI Press (2014)

15. Dimopoulos, Y., Torres, A.: Graph theoretical structures in logic programs and default theories. Theoret. Comput. Sci. **170**(1), 209–244 (1996)
16. Doder, D., Woltran, S.: Probabilistic argumentation frameworks – a logical approach. In: Straccia, U., Calì, A. (eds.) SUM 2014. LNCS, vol. 8720, pp. 134–147. Springer, Heidelberg (2014)
17. Dung, P.: On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. Artif. Intell. **77**(2), 321–357 (1995)
18. Dunne, P., Wooldridge, M.: Complexity of abstract argumentation (chapter 5). In: Rahwan, I., Simari, G. (eds.) Argumentation in Artificial Intelligence, pp. 85–104. Springer, New York (2009)
19. Fazzinga, B., Flesca, S., Parisi, F.: On the complexity of probabilistic abstract argumentation. In: Proceedings of IJCAI 2013, pp. 898–904. AAAI Press/IJCAI (2013)
20. Fazzinga, B., Flesca, S., Parisi, F.: Efficiently estimating the probability of extensions in abstract argumentation. In: Liu, W., Subrahmanian, V.S., Wijsen, J. (eds.) SUM 2013. LNCS, vol. 8078, pp. 106–119. Springer, Heidelberg (2013)
21. Hemaspaandra, E., Hemaspaandra, L., Rothe, J.: The complexity of online manipulation of sequential elections. J. Comput. Syst. Sci. **80**(4), 697–710 (2014)
22. Hunter, A.: Probabilistic qualification of attack in abstract argumentation. Int. J. Approximate Reasoning **55**(2), 607–638 (2014)
23. Konczak, K., Lang, J.: Voting procedures with incomplete preferences. In: Proceedings of Multidisciplinary IJCAI 2005 Workshop on Advances in Preference Handling, pp. 124–129 (2005)
24. Lang, J., Rey, A., Rothe, J., Schadrack, H., Schend, L.: Representing and solving hedonic games with ordinal preferences and thresholds. In: Proceedings of AAMAS 2015. IFAAMAS (2015)
25. Li, H., Oren, N., Norman, T.J.: Probabilistic argumentation frameworks. In: Modgil, S., Oren, N., Toni, F. (eds.) TAFA 2011. LNCS, vol. 7132, pp. 1–16. Springer, Heidelberg (2012)
26. Liao, B., Jin, L., Koons, R.: Dynamics of argumentation systems: a division-based method. Artif. Intell. **175**(11), 1790–1814 (2011)
27. Meyer, A., Stockmeyer, L.: The equivalence problem for regular expressions with squaring requires exponential space. In: Proceedings of FOCS 1972, pp. 125–129. IEEE Press (1972)
28. Papadimitriou, C.: Computational Complexity, 2nd edn. Addison-Wesley, New York (1995). Reprinted with corrections
29. Parkes, D., Procaccia, A.: Dynamic social choice with evolving preferences. In: Proceedings of AAAI 2013, pp. 767–773. AAAI Press (2013)
30. Rahwan, I., Simari, G. (eds.): Argumentation in Artificial Intelligence. Springer, New York (2009)
31. Rienstra, T.: Towards a probabilistic Dung-style argumentation system. In: Proceedings of AT 2012, pp. 138–152 (2012)
32. Rothe, J.: Complexity Theory and Cryptology: An Introduction to Cryptocomplexity. EATCS Texts in Theoretical Computer Science. Springer, Heidelberg (2005)
33. Stockmeyer, L.: The polynomial-time hierarchy. Theoret. Comput. Sci. **3**(1), 1–22 (1976)
34. Xia, L., Conitzer, V.: Determining possible and necessary winners given partial orders. J. Artif. Intell. Res. **41**, 25–67 (2011)