Verification in Incomplete Argumentation Frameworks

Dorothea Baumeister, Daniel Neugebauer, Jörg Rothe, and Hilmar Schadrack

Abstract

We tackle the problem of expressing incomplete knowledge in abstract argumentation frameworks originally introduced by Dung [15]. In applications, incomplete argumentation frameworks may arise as intermediate states in an elicitation process, or when merging different beliefs about an argumentation framework's state, or in cases where complete information cannot be obtained. We consider two specific models of incomplete argumentation frameworks, one focusing on attack incompleteness [4] and the other on argument incompleteness [5], and we also provide a general model of incomplete argumentation framework that subsumes both specific models. In these models, we study the computational complexity of variants of the verification problem with respect to common semantics of argumentation frameworks.

1 Introduction

Abstract argumentation frameworks are a simple, yet powerful tool for non-monotonic reasoning that were originally introduced by Dung [15]. In this model, individual arguments are considered to be abstract entities, disregarding their internal structure and focusing only on the attack relation between them. Various semantics defined by Dung and others allow to investigate the acceptability status of sets of arguments based on the attack relation. However, abstract argumentation frameworks are suitable to describe an argumentation's state only in an optimal situation—they require that *all* relevant arguments are included and that there is no uncertainty regarding the attacks between them. If these conditions are not met, the existing methods for semantic analysis cannot be applied.

This paper merges and extends our previous work on attack-incomplete [4] and argument-incomplete argumentation frameworks [5] where we have introduced and studied two specific models for describing incomplete information in argumentation frameworks, the latter allowing to express uncertainty about the set of arguments and the former about the attack relation between arguments. These formalizations of uncertainty capture various real-world phenomena like intermediate states of an evolving argumentation, partial-information settings (and, in particular, permanently unavailable information), and the task of merging individual views on an argumentation.

After describing the two specific models of attack- and argument-incomplete argumentation frameworks, we combine them to allow simultaneous uncertainty about attacks and arguments. Our objective in each model is to analyze how the computational complexity of certain variants of the verification problem (to be formally defined in Section 2) is affected by introducing uncertainty. However, this problem is defined in Dung's original model of argumentation framework, so we first need to adjust it to our extended models. A natural way to adapt a decision problem in the face of incomplete knowledge is to ask whether the answer is *possibly* (respectively, *necessarily*) "yes"—i.e., given all possible completions of the incomplete state, to ask whether *at least one* such completion (respectively, whether *all* these completions) are yes-instances of the original problem. This approach has already been taken in various areas of computational social choice, namely voting [21, 33, 10, 1, 2], fair division [7], algorithmic game theory [22], and judgment aggregation [3], but is new to argumentation theory: this paper's predecessors [4, 5] were the first to define and study possible and necessary verification for certain semantics in incomplete argumentation frameworks.

In related work, incomplete knowledge about the attack relation was first introduced by Coste-Marquis et al. [11] and analyzed with respect to argument acceptability by Cayrol et al. [8]. Opposed to us, however, they develop new semantics for attack-incomplete argumentation frameworks and thus put a lot of focus on the incomplete framework itself, rather than on its completions. Other work on incomplete knowledge in abstract argumentation includes *probabilistic argumentation frameworks* (see, for example, the work of Li al. [23], Rienstra [28], Fazzinga et al. [17, 18], Hunter [20], and Doder and Woltran [14]), where arguments and/or attacks have an associated probability as a quantified notion of uncertainty.

A related concept to incomplete knowledge is that of dynamic change. Cayrol et al. [9] study belief revision, which allows the addition or deletion of one single argument or several arguments, together with a respective change in the attack relation. Their work focuses on how this can alter the set of extensions of the given argumentation framework. Liao et al. [24] investigate the complexity of computing the status of an argument (i.e., whether it is accepted, rejected, or undecided) upon changing the arguments and attacks. Coste-Marquis et al. [12] study how belief revision postulates can be applied to argumentation systems. Boella et al. [6] address the question of which arguments or attacks can be removed without changing the set of extensions. Another dynamic setting is that of merging or aggregating different argumentation frameworks. Coste-Marquis et al. [11] study incomplete argumentation frameworks as a possible result of merging individual views. Tohmé et al. [31] discuss criteria for methods that aggregate several attack relations into a single attack relation (without uncertainty).

This paper is structured as follows. In Section 2, we provide the formal model of standard argumentation framework. Sections 3.1 and 3.2 introduce, respectively, the attack-incomplete and argument-incomplete model extensions, which are then combined into a universal incompleteness model in Section 3.3. In Section 4, we summarize our results and point out some interesting tasks that could be tackled in future work.

2 Preliminaries

In this section, we give formalizations of the basic notions of abstract argumentation. While we adapt some notation from the book chapter by Dunne and Wooldridge [16], the underlying concepts are due to Dung [15].

Definition 1. An *argumentation framework* AF is a pair $\langle \mathscr{A}, \mathscr{R} \rangle$, where $\mathscr{A}$ is a finite set of *arguments*, and $\mathscr{R} \subseteq \mathscr{A} \times \mathscr{A}$ is a binary relation. We say that *a attacks b* if $(a,b) \in \mathscr{R}$.

We will use the common representation of argumentation frameworks by graphs: Every argumentation framework $AF = \langle \mathscr{A}, \mathscr{R} \rangle$ can be seen as a directed graph $G_{AF} = (V,E)$ by identifying arguments with vertices and attacks with directed edges, i.e., $V = \mathscr{A}$ and $E = \mathscr{R}$.

Example 2. Figure 1 displays the graph representation of the argumentation framework $AF = \langle \mathscr{A}, \mathscr{R} \rangle$ with $\mathscr{A} = \{a,b,c\}$ and $\mathscr{R} = \{(a,b),(c,a)\}$. It will be used—and extended along the way—as a running example throughout the paper.

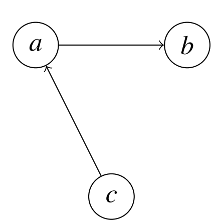

Figure 1: A simple argumentation framework

In the literature, many *semantics* have been defined which allow to evaluate the acceptability status of sets of arguments. We use the semantics introduced by Dung in his seminal paper [15]:

Definition 3. Let $AF = \langle \mathscr{A}, \mathscr{R} \rangle$ be an argumentation framework. A set $S \subseteq \mathscr{A}$ is

- *conflict-free* if there are no $a, b \in S$ such that $(a,b) \in \mathscr{R}$,
- *admissible* if S is conflict-free and every $a \in S$ is acceptable with respect to S, where an argument $a \in \mathscr{A}$ is *acceptable with respect to* $S \subseteq \mathscr{A}$ if for each $b \in \mathscr{A}$ with $(b,a) \in \mathscr{R}$ there is a $c \in S$ such that $(c,b) \in \mathscr{R}$,
- *preferred* if S is a maximal (with respect to set inclusion) admissible set,
- *stable* if S is conflict-free and for every $b \in \mathscr{A} \setminus S$ there is an $a \in S$ with $(a,b) \in \mathscr{R}$,
- *complete* if S is admissible and contains all $a \in \mathscr{A}$ that are acceptable with respect to S, and
- *grounded* if S is the least (with respect to set inclusion) fixed point of the characteristic function of AF, where the *characteristic function* $F_{AF} : 2^{\mathscr{A}} \to 2^{\mathscr{A}}$ of AF is defined by $F_{AF}(S) = \{a \in \mathscr{A} \mid a \text{ is acceptable with respect to } S\}$.

The characteristic function always has a least fixed point, since it is monotonic with respect to set inclusion, so the existence of the (unique) grounded set is guaranteed. The complete sets of an argumentation framework can be characterized as the fixed points of F_{AF}—in particular, the grounded set is complete. Dung [15] also proved several other correlations between his semantics. In particular, he showed that every admissible set is a subset of a preferred set, and that there always is at least one preferred set (which may be the empty set). Also, every stable set is preferred, and every preferred set is complete. It is easy to find examples that a preferred or grounded set does not have to be stable, and it is easy to show that each of the above defined semantics entails conflict-freeness and admissibility.

Dung uses the notion of *extensions* of an argumentation framework as a term for those subsets that fulfill the criteria of a given semantics. For example, a set of arguments is called a *preferred extension of the argumentation framework* if it is a preferred set of the given argumentation framework. Dung considers conflict-freeness and admissibility to be basic requirements rather than semantics, and therefore did not call conflict-free or admissible sets "extensions"—for convenience, however, we will do so sometimes.

We also need some of the basic notions from complexity theory. We assume the reader to be familiar with the complexity classes P, NP, and coNP, as well as hardness, completeness, polynomial-time-reducibility, $\leq_{\mathrm{m}}^{\mathrm{p}}$, and (oracle) Turing machines. DP is a complexity class introduced by Papadimitriou and Yannakakis [27] as the class of differences of any two NP problems; it also is the second level of the boolean hierarchy over NP. Problems that are solvable by a nondeterministic oracle Turing machine with access to an NP oracle belong to $\Sigma_2^p = \mathrm{NP}^{\mathrm{NP}}$; this class constitutes, together with $\Pi_2^p = \mathrm{coNP}^{\mathrm{NP}}$, the second level of the polynomial hierarchy, and was introduced by Meyer and Stockmeyer [25, 30]. It is known that $\mathrm{P} \subseteq \mathrm{NP} \subseteq \mathrm{DP} \subseteq \Sigma_2^p$, but it is still unknown whether any of these inclusions is strict. For further details, see, e.g., [26, 29].

Dunne and Wooldridge [16] investigated several decision problems defined for argumentation frameworks, many of which are hard to decide, as they are complete for NP, coNP, DP, or even Π_2^p. Here, we will focus on the verification problem, which is coNP-complete for the preferred semantics [13], but can be decided in polynomial time for all other semantics defined above, which follows immediately from the work of Dung [15].

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

In our notation, the boldfaced letter **s** is a placeholder for any of the six semantics defined earlier. For better readability, we will sometimes shorten their names and write CF for *conflict-freeness*, AD for *admissibility*, PR for *preferredness*, ST for *stability*, CP for *completeness*, and GR for *groundedness*.

3 Incomplete Argumentation Frameworks

In this section, we introduce three different notions of incompleteness for argumentation frameworks. We start with attack incompleteness in Section 3.1, followed by argument incompleteness in Section 3.2. In Section 3.3, both approaches are combined to provide a general model of incompleteness in argumentation frameworks.

3.1 Attack Incompleteness

The first notion of incompleteness we consider concerns the attack relation between arguments [4]. While Dung's original model only allows to express whether an attack (a,b) exists ($(a,b) \in \mathscr{R}$) or doesn't exist ($(a,b) \notin \mathscr{R}$), the extended model also allows to explicitly express lack of information about an attack. *Attack-incomplete argumentation frameworks* were originally proposed by Coste-Marquis et al. [11]—we employ their model, but use a slightly modified notation.

Definition 4. 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 $\mathscr{R}^?$ contains all possible additional attacks not (yet) known to exist. The set of attacks that are known to never exist is denoted by $\mathscr{R}^- = (\mathscr{A} \times \mathscr{A}) \setminus (\mathscr{R} \cup \mathscr{R}^?)$.

Example 5. Extending the argumentation framework from Example 2 by three possible attacks, $\mathscr{R}^? = \{(a,a),(b,a),(b,c)\}$, yields the attack-incomplete argumentation framework $\langle \mathscr{A}, \mathscr{R}, \mathscr{R}^? \rangle$ the graph representation of which is given in Figure 2c. This incomplete framework might be the product of several *individual (subjective)* views that share a common set of arguments but may have different attacks. Figures 2a and 2b show two such individual argumentation frameworks, which are merged into the attack-incomplete argumentation framework of Figure 2c by including those attacks that exist in all individual views as definite attacks ($\mathscr{R}$), and including attacks that exist in some but not all individual views as possible attacks ($\mathscr{R}^?$).

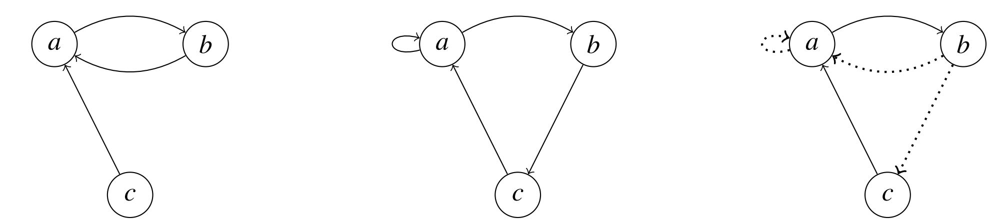

(a) Agent 1's individual view (b) Agent 2's individual view (c) Merging the views of 1 and 2

Figure 2: Attack incompleteness from merging two individual argumentation frameworks

In an attack-incomplete argumentation framework $\langle \mathscr{A}, \mathscr{R}, \mathscr{R}^? \rangle$, for each possible but as yet unknown attack in $\mathscr{R}^?$, when deciding whether or not the attack will be included, one obtains a standard argumentation framework that can be seen as a *completion* of $\langle \mathscr{A}, \mathscr{R}, \mathscr{R}^? \rangle$.

Definition 6. Let $AtIAF = \langle \mathscr{A}, \mathscr{R}, \mathscr{R}^? \rangle$ 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* $AtIAF$.

The number of possible completions for a given attack-incomplete argumentation framework is clearly $2^{|\mathscr{R}^?|}$. For $\mathscr{R}^? = \emptyset$, there is no uncertainty and only one completion exists, which coincides with the attack-incomplete framework itself. In general, however, the number of completions may be exponential in relation to the instance's size.

In an attack-incomplete argumentation framework $AtIAF$, we say that a property defined for standard argumentation frameworks (e.g., a semantics) holds *possibly* if there exists a completion $AtIAF^*$ of $AtIAF$ for which the property holds, and a property holds *necessarily* if it holds for all completions of $AtIAF$. Accordingly, we can define two variants of the verification problem in the attack-incomplete case for each given semantics **s**:

s-ATT-INC-POSSIBLE-VERIFICATION (**s**-ATTINCPV)	
Given:	An attack-incomplete argumentation framework $AtIAF = \langle \mathscr{A}, \mathscr{R}, \mathscr{R}^? \rangle$ and a set $S \subseteq \mathscr{A}$.
Question:	Is there a completion AF^* of $AtIAF$ such that S is an **s** extension of AF^*?

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

Both problems are potentially harder than standard verification, since they add an existential (respectively, universal) quantifier over a potentially exponential space of solutions. However, in most cases we were able to prove that the complexity in fact does not increase.

We now give our proof for the complexity of **s**-ATTINCPV for $\mathbf{s} \in \{\text{CF}, \text{AD}, \text{ST}\}$, which uses a single critical completion to reduce the problem to standard VERIFICATION. The other proofs are deferred to the appendix. Our results from [4] for both problems are stated in Table 1 in Section 4.

Definition 7. Let $AtIAF = \langle \mathscr{A}, \mathscr{R}, \mathscr{R}^? \rangle$ be an attack-incomplete argumentation framework and let $S \subseteq \mathscr{A}$. The *optimistic completion of AtIAF for S* is $AtIAF_S^{\text{opt}} = \langle \mathscr{A}, \mathscr{R}_S^{\text{opt}} \rangle$ with $\mathscr{R}_S^{\text{opt}} = \mathscr{R} \cup \{(a,b) \in \mathscr{R}^? \mid b \notin S\}$.

Example 8. Consider again the attack-incomplete argumentation framework from Figure 2c. Its optimistic completion for the set $S = \{b,c\}$ is given in Figure 3b. Arguments in S are highlighted by a boldfaced circle, and the possible attacks added to the set of attacks in the optimistic completion are displayed as boldfaced arcs.

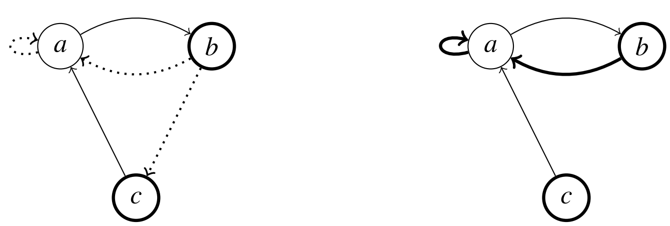

(a) An attack-incomplete argumentation framework

(b) The optimistic completion for the set $S = \{b,c\}$ of arguments

Figure 3: Optimistic completion of an attack-incomplete argumentation framework.

Lemma 9. *Let $AtIAF = \langle \mathscr{A}, \mathscr{R}, \mathscr{R}^? \rangle$ be an attack-incomplete argumentation framework, let $S \subseteq \mathscr{A}$, and let $AtIAF_S^{\text{opt}}$ be the optimistic completion of AtIAF for S.*

1. *S is possibly conflict-free in AtIAF if and only if S is a conflict-free extension of $AtIAF_S^{\text{opt}}$.*
2. *$a \in S$ is possibly acceptable with respect to S in AtIAF if and only if a is acceptable with respect to S in $AtIAF_S^{\text{opt}}$.*
3. *S is possibly admissible in AtIAF if and only if S is an admissible extension of $AtIAF_S^{\text{opt}}$.*
4. *S is possibly stable in AtIAF if and only if S is a stable extension of $AtIAF_S^{\text{opt}}$.*

Proof. The converse is trivial in all cases: If S fulfills a given criterion in $AtIAF_S^{\text{opt}}$, this immediately yields that S possibly fulfills the criterion in $AtIAF$. 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 $AtIAF_S^{\text{opt}}$, 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 $AtIAF$. Thus S is not a possibly conflict-free set in $AtIAF$.
2. If there is some $a \in S$ that is not acceptable with respect to S in $AtIAF_S^{\text{opt}}$, 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 $AtIAF$. Since a cannot be acceptable with respect to S in any completion of $AtIAF$, a is not possibly acceptable with respect to S in $AtIAF$.
3. Assume that S is not an admissible extension in $AtIAF_S^{\text{opt}}$, i.e., S is not conflict-free in $AtIAF_S^{\text{opt}}$ or there is some $a \in S$ that is not acceptable with respect to S in $AtIAF_S^{\text{opt}}$. In either case, the previous results imply that S is not conflict-free in any completion of $AtIAF$ or a is not acceptable with respect to S in any completion of $AtIAF$. Thus S is not a possibly admissible extension in $AtIAF$.
4. If a set S of arguments is not stable in $AtIAF_S^{\text{opt}}$, S is necessarily not conflict-free in $AtIAF$ or there is an $a \in \mathscr{A} \smallsetminus S$ that is not attacked by S in $AtIAF_S^{\text{opt}}$, and therefore—by construction of $AtIAF_S^{\text{opt}}$—a cannot be attacked by S in any completion of $AtIAF$. In both cases, there is no completion of $AtIAF$ for which S is stable, so S is not a possibly stable extension of $AtIAF$.

This completes the proof. ❑

Theorem 10. *For* $\mathbf{s} \in \{\text{CF}, \text{AD}, \text{ST}\}$, **s**-ATTINCPV *is in* P.

Proof. The optimistic completion can obviously be constructed in polynomial time. As already mentioned, the problem **s**-VERIFICATION can be solved in polynomial time for a given completion. Lemma 9 then provides that the answer to **s**-ATTINCPV is the same as that to **s**-VERIFICATION for the optimistic completion. ❑

3.2 Argument Incompleteness

In our second model, which we proposed previously [5], we allow uncertainty about the set of arguments. While the total set of arguments that may take part in the argumentation is known (and finite), there is uncertainty for some of these arguments as to whether or not they actually exist in the argumentation—they may not be constructible given a certain knowledge base, they may not be applicable, or they may simply not be brought forward by any agent. Note that this notion of possible non-existence is different from that of (in)acceptability.

Definition 11. An *argument-incomplete argumentation framework* is a triple $\langle \mathscr{A}, \mathscr{A}^?, \mathscr{R} \rangle$, where $\mathscr{A}$ and $\mathscr{A}^?$ are disjoint sets of arguments and $\mathscr{R}$ is a subset of $(\mathscr{A} \cup \mathscr{A}^?) \times (\mathscr{A} \cup \mathscr{A}^?)$. $\mathscr{A}$ is the set of arguments that are known to definitely exist, while $\mathscr{A}^?$ contains all possible additional arguments that are not (yet) known to exist.

Note that, in this model, there is no uncertainty regarding the attack relation—even though attacks may be indirectly excluded by excluding an incident argument. As an example, consider a discussion where each agent has a private set of arguments that they can bring forward, but they may also choose to not introduce some of the arguments that they know of—maybe for strategic purposes.

However, for the “outcome” of the argumentation, only those arguments that were explicitly stated by some agent are considered. Such a situation could be modeled using an argument-incomplete argumentation framework.

Example 12. Extending the argumentation framework from Example 2 by two possible arguments $\mathscr{A}^? = \{d,e\}$ together with an extension of the attack relation, by including the attacks (d,b), (d,c), (b,d), and (e,c), yields the argument-incomplete argumentation framework $\langle \mathscr{A}, \mathscr{A}^?, \mathscr{R} \rangle$ the graph representation of which is given in Figure 4c. As already discussed in Example 5, this incomplete framework might result from merging several individual views, which agree on all attacks over those arguments that are known to all agents but may have different argument sets. Figures 4a and 4b show two such individual argumentation frameworks, which are then merged into the argument-incomplete argumentation framework of Figure 4c by including all arguments that are known in every agent's argumentation framework as definite arguments ($\mathscr{A}$), and including arguments that exist in some but not in all agents' argumentation frameworks as possible arguments ($\mathscr{A}^?$). Note that there is no choice of whether or not we include attacks: An attack is included if and only if there is at least one agent who has this attack in her individual argumentation framework.

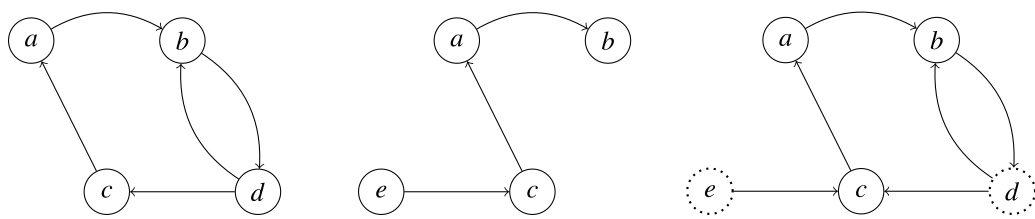

(a) Agent 1's individual view (b) Agent 2's individual view (c) Merging the views of 1 and 2

Figure 4: Argument incompleteness from merging two individual argumentation frameworks

Also for argument-incomplete argumentation frameworks, we can define completions quite similar to those of Definition 6:

Definition 13. Let $ArIAF = \langle \mathscr{A}, \mathscr{A}^?, \mathscr{R} \rangle$ be an argument-incomplete argumentation framework. For a set $\mathscr{A}^*$ of arguments with $\mathscr{A} \subseteq \mathscr{A}^* \subseteq \mathscr{A} \cup \mathscr{A}^?$, define the *restriction of* $\mathscr{R}$ *to* $\mathscr{A}^*$ by

$$\mathscr{R}|_{\mathscr{A}^*} = \{(a,b) \in \mathscr{R} \mid a,b \in \mathscr{A}^*\}.$$

Then, an argumentation framework $AF^* = \langle \mathscr{A}^*, \mathscr{R}|_{\mathscr{A}^*} \rangle$ is called a *completion of* $ArIAF$.

Obviously, the total number of possible completions is again exponential—this time in the number of possible new arguments, i.e., there can be up to $2^{|\mathscr{A}^?|}$ possible completions.

Let us now define the two variants of the verification problem in argument-incomplete argumentation frameworks for each given semantics **s**:

	s-ARG-INC-POSSIBLE-VERIFICATION (**s**-ARGINCPV)
Given:	An argument-incomplete argumentation framework $ArIAF = \langle \mathscr{A}, \mathscr{A}^?, \mathscr{R} \rangle$ and a set $S \subseteq \mathscr{A} \cup \mathscr{A}^?$.
Question:	Is there a completion $AF^* = \langle \mathscr{A}^*, \mathscr{R}|_{\mathscr{A}^*} \rangle$ of $ArIAF$ such that $S|_{\mathscr{A}^*} = S \cap \mathscr{A}^*$ is an **s** extension of AF^*?

	s-ARG-INC-NECESSARY-VERIFICATION (**s**-ARGINCNV)
Given:	An argument-incomplete argumentation framework $ArIAF = \langle \mathscr{A}, \mathscr{A}^?, \mathscr{R} \rangle$ and a set $S \subseteq \mathscr{A} \cup \mathscr{A}^?$.
Question:	For all completions $AF^* = \langle \mathscr{A}^*, \mathscr{R}|_{\mathscr{A}^*} \rangle$ of $ArIAF$, is $S|_{\mathscr{A}^*} = S \cap \mathscr{A}^*$ an **s** extension of AF^*?

We now present two selected proofs that showcase how the complexity of both problems can be different even for the same semantics. Theorem 14 is taken from our previous work [5], while Theorem 15 presents a new result. The remaining proofs for the argument-incomplete model are deferred to the appendix. Again, our complexity results—including those from [5] —are stated in Table 1 in Section 4.

Theorem 14. AD-ARGINCPV *is* NP-*complete.*

Proof. NP-membership follows from the fact that, given a completion, it can be verified in polynomial time whether a set S is admissible in that completion. To show NP-hardness, we reduce from the following NP-complete problem (see the book by Garey and Johnson [19]):

EXACT-COVER-BY-3-SETS (X3C)	
Given:	A set $B = \{b_1, \dots, b_{3k}\}$ and a family $\mathscr{S}$ of subsets of B, with $\|S_j\| = 3$ for all $S_j \in \mathscr{S}$.
Question:	Does there exist a subfamily $\mathscr{S}' \subseteq \mathscr{S}$ of size k that exactly covers B, i.e., $\bigcup_{S_j \in \mathscr{S}'} S_j = B$?

Given an instance $(B, \mathscr{S}) = (\{b_1, \dots, b_{3k}\}, \{S_1, \dots, S_m\})$ of X3C, we construct an instance $(\langle \mathscr{A}, \mathscr{A}^?, \mathscr{R} \rangle, S)$ of AD-ARGINCPV as follows:[1]

$$
\begin{aligned}
\mathscr{A} &= \{x\} \cup B, \\
\mathscr{A}^? &= \mathscr{S}, \\
\mathscr{R} &= \{(b_i, x) \mid b_i \in B\} \cup \\
&\quad \{(S_j, b_{j_1}), (S_j, b_{j_2}), (S_j, b_{j_3}) \mid S_j = \{b_{j_1}, b_{j_2}, b_{j_3}\} \in \mathscr{S}\} \cup \\
&\quad \{(S_i, S_j), (S_j, S_i) \mid S_i, S_j \in \mathscr{S} \text{ and } S_i \cap S_j \neq \emptyset\}, \\
S &= \{x\} \cup \mathscr{S}.
\end{aligned}
$$

In particular, $\mathscr{A} \cup \mathscr{A}^?$ contains one argument b_i for every element $b_i \in B$, $1 \leq i \leq 3k$, one argument S_j for every set S_j in $\mathscr{S}$, $1 \leq j \leq m$, and one additional argument x. All arguments corresponding to elements of B attack x, and each argument S_j attacks the three arguments corresponding to those elements of B that belong to S_j in $\mathscr{S}$. Additionally, there are attacks between S_i and S_j if the corresponding sets in $\mathscr{S}$ are not disjoint. Finally, $\mathscr{A}$ and S act as opponents: x belongs to both, but the arguments corresponding to elements in B belong to $\mathscr{A}$ only, whereas the arguments corresponding to the sets in $\mathscr{S}$ belong to S only. See Figure 5 for two examples of this construction: Figure 5a shows a yes-instance of AD-ARGINCPV created from a yes-instance of X3C, and Figure 5b shows a no-instance of AD-ARGINCPV created from a no-instance of X3C.

We claim that $(B, \mathscr{S}) \in$ X3C if and only if $(\langle \mathscr{A}, \mathscr{A}^?, \mathscr{R} \rangle, S) \in$ AD-ARGINCPV.

($\Longrightarrow$) Clearly, if $(B, \mathscr{S})$ is a yes-instance of X3C, we can add exactly those arguments S_i to $\mathscr{A}$ that correspond to an exact cover of B. Let $\mathscr{A}^*$ be the argument set of this completion. In $\mathscr{A}^*$, every b_i, $1 \leq i \leq 3k$, is attacked by exactly one argument S_j, $1 \leq j \leq m$, as of the exact cover. Hence, $x \in S|_{\mathscr{A}^*}$ is defended against every attack. Additionally, the arguments S_j in $\mathscr{A}^*$ have no attacks between them, because the corresponding sets are pairwise disjoint, which implies that no new attacks on the elements of $S|_{\mathscr{A}^*}$ are introduced. But this means that $S|_{\mathscr{A}^*}$ is admissible in $\langle \mathscr{A}^*, \mathscr{R}|_{\mathscr{A}^*} \rangle$.

($\Longleftarrow$) If there is a completion with the argument set $\mathscr{A}^*$, this completion must defend x against every b_i, $1 \leq i \leq 3k$. This means that there must exist a cover of the elements of B by the sets of $\mathscr{S}$. But because the arguments S_j attack each other whenever they are not disjoint, this cover must be exact; otherwise, the set $S|_{\mathscr{A}^*}$ would not be conflict-free. Hence, there exists an exact cover of B. ❑

[1] We slightly abuse notation and use the same identifiers for both instances; it will always be clear from the context, though, which instance an element belongs to.

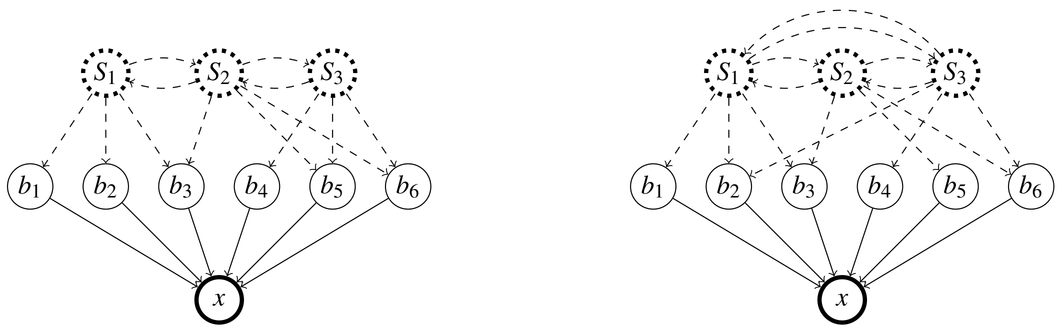

(a) $\mathscr{S} = \{\{b_1,b_2,b_3\},\{b_3,b_5,b_6\},\{b_4,b_5,b_6\}\}$. $(B,\mathscr{S})$ is a yes-instance of X3C that yields a yes-instance of AD-ARGINCPV.

(b) $\mathscr{S} = \{\{b_1,b_2,b_3\},\{b_3,b_5,b_6\},\{b_2,b_4,b_6\}\}$. $(B,\mathscr{S})$ is a no-instance of X3C that yields a no-instance of AD-ARGINCPV.

Figure 5: Two examples of the reduction from X3C to AD-ARGINCPV. Both X3C instances have $B = \{b_1,\dots,b_6\}$. $\mathscr{A}$ contains the solid arguments, the dotted arguments belong to $\mathscr{A}^?$, and the thick arguments are part of S.

Theorem 15. AD-ARGINCNV *and* ST-ARGINCNV *are in* P.

Proof. Let $I = (\langle \mathscr{A}, \mathscr{A}^?, \mathscr{R}\rangle, S)$ be an instance for AD-ARGINCNV. If S is not necessarily conflict-free in $\langle \mathscr{A}, \mathscr{A}^?, \mathscr{R}\rangle$, it is not necessarily admissible in $\langle \mathscr{A}, \mathscr{A}^?, \mathscr{R}\rangle$, either. Since CF-ARGINCNV is in P, this can be checked in polynomial time. In the following, we may assume that S is necessarily conflict-free.

Let $\mathscr{A}_0 = \mathscr{A} \cup (\mathscr{A}^? \setminus S)$ and $C_0 = \langle \mathscr{A}_0, \mathscr{R}|_{\mathscr{A}_0}\rangle$, and for each argument $a \in \mathscr{A}^? \cap S$, let $\mathscr{A}_a = \mathscr{A}_0 \cup \{a\}$ and $C_a = \langle \mathscr{A}_a, \mathscr{R}|_{\mathscr{A}_a}\rangle$. If, for some $x \in \{0\} \cup (\mathscr{A}^? \cap S)$, $S|_{\mathscr{A}_x}$ is not admissible in the completion C_x, we clearly have $I \notin$ AD-ARGINCNV. Since the number of these completions is bounded by the number of arguments (plus one), this can again be verified in polynomial time. We may now assume that, in each completion C_x, $S|_{\mathscr{A}_x}$ is admissible.

Note that each of these completions includes *all* possible attacks against the respective set $S|_{\mathscr{A}_x}$, because the completions include all possibly harmful arguments (members of $\mathscr{A}_0$) and because there cannot be any attacks among members of S. This yields that $S|_{\mathscr{A}_0}$ defends all attacks against its elements in *any* completion, and, for all $a \in \mathscr{A}^? \cap S$, $S|_{\mathscr{A}_a}$ defends all attacks against a in *any* completion. Finally, since in any completion $C^* = \langle \mathscr{A}^*, \mathscr{R}|_{\mathscr{A}^*}\rangle$, it holds that $S|_{\mathscr{A}^*} \subseteq \bigcup_x S|_{\mathscr{A}_x}$, we can conclude that each element of $S|_{\mathscr{A}^*}$ is acceptable with respect to $S|_{\mathscr{A}^*}$ in C^*, so S is necessarily admissible in $\langle \mathscr{A}, \mathscr{A}^?, \mathscr{R}\rangle$ and $I \in$ AD-ARGINCNV.

For ST-ARGINCNV, the same construction as above can be used. We can again conclude that $I \notin$ ST-ARGINCNV in all cases where we had $I \notin$ AD-ARGINCNV, since each stable set needs to be admissible. In addition, it is easy to see that, in order for S to be necessarily stable, the set $S|_{\mathscr{A}_0}$ in the completion C_0 as defined above needs to attack all arguments in $\mathscr{A}_0 \setminus S$. However, since $\mathscr{A}_0 \setminus S = \mathscr{A} \setminus S$ ($\mathscr{A}_0$ contains all arguments that are not in S) and further $S|_{\mathscr{A}_0}$ is a subset of $S|_{\mathscr{A}^*}$ for any completion with argument set $\mathscr{A}^*$, this already yields that $S|_{\mathscr{A}^*}$ necessarily attacks all arguments outside of $S|_{\mathscr{A}^*}$ in any completion, and we have $I \in$ ST-ARGINCNV. ❑

3.3 General Incompleteness

We now combine the two given models by allowing incomplete knowledge about both the attack relation and the set of arguments at the same time.

Definition 16. An *incomplete argumentation framework* is a quadruple $\langle \mathscr{A}, \mathscr{A}^?, \mathscr{R}, \mathscr{R}^?\rangle$, where $\mathscr{A}$ and $\mathscr{A}^?$ are disjoint sets of arguments and $\mathscr{R}$ and $\mathscr{R}^?$ are disjoint subsets of $(\mathscr{A} \cup \mathscr{A}^?) \times (\mathscr{A} \cup \mathscr{A}^?)$.

$\mathscr{A}$ (respectively, $\mathscr{R}$) is the set of arguments (respectively, the set of attacks) that are known to definitely exist, while $\mathscr{A}^?$ (respectively, $\mathscr{R}^?$) contains all possible additional arguments (respectively, all possible additional attacks) not (yet) known to exist.

Again, an incomplete argumentation framework can be the result of merging a number of individual argumentation frameworks. Recall that in Section 3.1 we only allowed those argumentation frameworks to be merged that share a common set of arguments, i.e., we could aggregate only those argumentation frameworks $AF_1 = \langle \mathscr{A}_1, \mathscr{R}_1 \rangle, \dots, AF_n = \langle \mathscr{A}_n, \mathscr{R}_n \rangle$ for which $\mathscr{A}_i = \mathscr{A}_j$ holds for any $i, j \in \{1, \dots, n\}$. And in Section 3.2 we restricted ourselves to those argumentation frameworks that agree on all attacks between common arguments. Formally, this can be expressed by requiring $\mathscr{R}_i|_{\mathscr{A}_i \cap \mathscr{A}_j} = \mathscr{R}_j|_{\mathscr{A}_i \cap \mathscr{A}_j}$ for all $i, j \in \{1, \dots, n\}$.

In this section, however, we do not restrict the input anymore. Hence, we need to specify how we can merge argumentation frameworks that were not mergeable before, namely those over possibly different sets of arguments regarding attack incompleteness, and those over possibly different attack relations in the case of argument incompleteness.

Definition 17. The merging operation for n individual argumentation frameworks $AF_1, \dots, AF_n$ is defined to be the following incomplete argumentation framework $\langle \mathscr{A}, \mathscr{A}^?, \mathscr{R}, \mathscr{R}^? \rangle$: $\mathscr{A}$ consists of all arguments that belong to *all* $AF \in \{AF_1, \dots, AF_n\}$. $\mathscr{A}^?$ consists of all arguments that belong to *at least one* (but not to all) $AF \in \{AF_1, \dots, AF_n\}$. $\mathscr{R}$ consists of all attacks (a, b) that belong to *all* $AF \in \{AF_1, \dots, AF_n\}$ containing both a and b. $\mathscr{R}^?$ consists of all attacks (a, b) that belong to *at least one* (but not to all) $AF \in \{AF_1, \dots, AF_n\}$ that contain both a and b.

(a) Another agent's individual view

(b) Merging the argumentation frameworks of Figures 2a, 2b, 4a, 4b, and 6a

Figure 6: Incompleteness from merging five individual argumentation frameworks

Example 18. Extending the argumentation framework from Example 2 the same way we did in Examples 5 and 12, we obtain the incomplete argumentation framework $\langle \mathscr{A}, \mathscr{A}^?, \mathscr{R}, \mathscr{R}^? \rangle$ the graph representation of which is given in Figure 6b. This incomplete argumentation framework is the result of merging the individual argumentation frameworks from Figures 2a, 2b, 4a, 4b, and 6a according to Definition 17.

The given merge operation is a strict generalisation of those in Sections 3.1 and 3.2. If we restrict the input of the merging operation the same way we restricted the input in Section 3.1 (that is, requiring $\mathscr{A}_i = \mathscr{A}_j$ for all $i, j \in \{1, \dots, n\}$), we have $\mathscr{A}^? = \emptyset$ and the same merging operation as defined there. On the other hand, if we restrict the input the same way we did in Section 3.2 (that is, requiring $\mathscr{R}_i|_{\mathscr{A}_i \cap \mathscr{A}_j} = \mathscr{R}_j|_{\mathscr{A}_i \cap \mathscr{A}_j}$ for all $i, j \in \{1, \dots, n\}$), we have $\mathscr{R}^? = \emptyset$ and the same merging operation as defined there.

The merging operation we defined above regarding the argument sets can be seen as a global merging: If an argument is contained in all input argumentation frameworks, put it into $\mathscr{A}$, otherwise into $\mathscr{A}^?$. In contrast, the merging operation regarding the attack relation is a local merging: If an attack (a, b) is contained in all those inputs that actually have an opinion over both a and b, put it into $\mathscr{R}$, otherwise into $\mathscr{R}^?$. This conforms to the way in which *consensual expansion* as defined by Coste-Marquis et al. [11] handles the merging of attacks.

In the general model of incomplete argumentation framework, a notion of *completion* can now be defined as follows.

Definition 19. Let $IAF = \langle \mathscr{A}, \mathscr{A}^?, \mathscr{R}, \mathscr{R}^? \rangle$ be a given incomplete argumentation framework. An argumentation framework $AF^* = \langle \mathscr{A}^*, \mathscr{R}^* \rangle$ with $\mathscr{A} \subseteq \mathscr{A}^* \subseteq \mathscr{A} \cup \mathscr{A}^?$ and $\mathscr{R}|_{\mathscr{A}^*} \subseteq \mathscr{R}^* \subseteq \left(\mathscr{R} \cup \mathscr{R}^?\right)|_{\mathscr{A}^*}$ is called a *completion of IAF*.

Finally, for each given semantics **s**, the variants of the verification problem adapted to incomplete argumentation frameworks are defined analogously to those in the attack-incomplete and argument-incomplete setting.

s-INC-POSSIBLE-VERIFICATION (**s**-INCPV)	
Given:	An incomplete argumentation framework $IAF = \langle \mathscr{A}, \mathscr{A}^?, \mathscr{R}, \mathscr{R}^? \rangle$ and a set $S \subseteq \mathscr{A} \cup \mathscr{A}^?$.
Question:	Is there a completion $AF^* = \langle \mathscr{A}^*, \mathscr{R}^* \rangle$ of IAF such that $S|_{\mathscr{A}^*} = S \cap \mathscr{A}^*$ is an **s** extension of AF^*?

s-INC-NECESSARY-VERIFICATION (**s**-INCNV)	
Given:	An incomplete argumentation framework $IAF = \langle \mathscr{A}, \mathscr{A}^?, \mathscr{R}, \mathscr{R}^? \rangle$ and a set $S \subseteq \mathscr{A} \cup \mathscr{A}^?$.
Question:	For all completions $AF^* = \langle \mathscr{A}^*, \mathscr{R}^* \rangle$ of IAF, is $S|_{\mathscr{A}^*} = S \cap \mathscr{A}^*$ an **s** extension of AF^*?

Let us start with a collection of straightforward results, the proofs of which are basically the same as the ones of the individual models (see [4] and [5]). These results are possible due to the fact that the space of possible solutions is increased by a factor of, respectively, $2^{|\mathscr{R}^?|}$ and $2^{|\mathscr{A}^?|}$ in the attack-incomplete and argument-incomplete models. In the combined model, the total number of binary choices is $|\mathscr{R}^?| + |\mathscr{A}^?|$, so this factor is at most $2^{|\mathscr{R}^?|+|\mathscr{A}^?|}$, where the exponent is still polynomial in the input size.[2]

Lemma 20. *1.* CF-INCPV *and* CF-INCNV *both are in* P.

2. For $\mathbf{s} \in \{\text{AD}, \text{ST}, \text{CP}, \text{GR}\}$, **s**-INCPV *is in* NP.

3. For $\mathbf{s} \in \{\text{AD}, \text{ST}, \text{CP}, \text{GR}\}$, **s**-INCNV *is in* coNP.

4. PR-INCPV *is in* Σ_2^p.

5. PR-INCNV *is in* coNP.

Proof. Given an incomplete argumentation framework $IAF = \langle \mathscr{A}, \mathscr{A}^?, \mathscr{R}, \mathscr{R}^? \rangle$ and a set $S \subseteq \mathscr{A} \cup \mathscr{A}^?$ of arguments, S is possibly conflict-free in IAF if and only if $S|_{\mathscr{A}}$ is conflict-free in the *minimal completion* $\langle \mathscr{A}, \mathscr{R}|_{\mathscr{A}} \rangle$ of $ArIAF$, which discards all additional arguments and attacks. Similary, S is necessarily conflict-free in IAF if and only if S is conflict-free in the *maximal completion* $\langle \mathscr{A} \cup \mathscr{A}^?, \mathscr{R} \cup \mathscr{R}^? \rangle$ of $ArIAF$, which includes all additional arguments and attacks. Since both the minimal and maximal completion can clearly be constructed in polynomial time, we have P-membership for both problems.

The remaining results of this lemma follow directly from the quantifier representations of the given problems: In the possible case, we start with an existential quantifier, and in the necessary case

[2] Since excluding possible arguments may implicitly also exclude possible attacks, it may be that not all $2^{|\mathscr{R}^?|+|\mathscr{A}^?|}$ combinations are feasible.

with an universal quantifier. For $\mathbf{s} \in \{\text{AD}, \text{ST}, \text{CP}, \text{GR}\}$ it is checkable in polynomial time whether the given subset is an **s** extension, which provides the results of Items 2 and 3. The standard verification problem for the preferred semantics belongs to coNP, hence it can be written as a universal quantifier followed by a statement checkable in polynomial time. Therefore, we have two alternating quantifiers in the case of PR-INCPV (Item 4), and two collapsing universal quantifier in the case of PR-INCNV (Item 5). This completes the proof. ❑

It is easy to see that all lower bounds obtained for both individual incomplete models carry over to the combined model. For example, consider an instance $IAF = \langle \mathscr{A}, \mathscr{A}^?, \mathscr{R}, \mathscr{R}^? \rangle$ where the set $\mathscr{A}^?$ of unknown arguments is empty, then *IAF* is equivalent to the attack-incomplete argumentation framework $\langle \mathscr{A}, \mathscr{R}, \mathscr{R}^? \rangle$ and both **s**-INCPV and **s**-INCNV collapse to **s**-ATTINCPV and **s**-ATTINCNV, respectively. An analogous comment applies to $\mathscr{R}^? = \emptyset$ and the argument-incomplete model. We conclude those results in the following corollary.

Corollary 21. *1. For* $\mathbf{s} \in \{\text{AD}, \text{ST}, \text{CP}, \text{GR}\}$, **s**-INCPV *is* NP-*hard.*

2. PR-INCPV *is* DP-*hard.*

3. PR-INCNV *is* coNP-*hard.*

4 Conclusion and Open Questions

We introduced a general model of incompleteness in argumentation frameworks, subsuming two previous models, one focusing on attack incompleteness and the other on argument incompleteness. We then have studied the computational complexity of variants of the verification problem with respect to common semantics of argumentation frameworks.

Table 1: Overview of complexity results in the standard model [15, 13], the attack-incomplete model [4], the argument-incomplete model ([5] and Theorem 15), and the combined model (Lemma 20 and Corollary 21).

s	VER	ATTINCPV	ATTINCNV	ARGINCPV	ARGINCNV	INCPV	INCNV
CF	in P	in P	in P	in P	in P	in P	in P
AD	in P	in P	in P	NP-c.	in P	NP-c.	in coNP
ST	in P	in P	in P	NP-c.	in P	NP-c.	in coNP
CP	in P	in P	in P	NP-c.	in coNP	NP-c.	in coNP
GR	in P	in P	in P	NP-c.	in coNP	NP-c.	in coNP
PR	coNP-c.	coNP-h., in Σ_2^p	coNP-c.	DP-h., in Σ_2^p	coNP-c.	DP-h., in Σ_2^p	coNP-c.

Table 1 gives an overview of the complexity results for the verification problem in the standard model and in the three incompleteness models considered in this paper. A task for future work is to determine the exact complexity in those cases where we haven't found tight bounds yet, e.g., for PR-ATTINCPV. Also, we would like to analyze the complexity of possible and necessary variants of other decision problems than verification.

Acknowledgments

This paper merges and extends two earlier ADT-2015 papers, and we thank the 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] D. Baumeister, M. Roos, and J. Rothe. Computational complexity of two variants of the possible winner problem. In *Proceedings of the 10th International Conference on Autonomous Agents and Multiagent Systems*, pages 853–860. IFAAMAS, May 2011.

[2] D. Baumeister, M. Roos, J. Rothe, L. Schend, and L. Xia. The possible winner problem with uncertain weights. In *Proceedings of the 20th European Conference on Artificial Intelligence*, pages 133–138. IOS Press, August 2012.

[3] D. Baumeister, G. Erdélyi, O. Erdélyi, and J. Rothe. Complexity of manipulation and bribery in judgment aggregation for uniform premise-based quota rules. *Mathematical Social Sciences*, 76:19–30, 2015.

[4] D. Baumeister, D. Neugebauer, and J. Rothe. Verification in attack-incomplete argumentation frameworks. In *Proceedings of the 4th International Conference on Algorithmic Decision Theory*, pages 341–358. Springer-Verlag *Lecture Notes in Artificial Intelligence*, September 2015.

[5] D. Baumeister, J. Rothe, and H. Schadrack. Verification in argument-incomplete argumentation frameworks. In *Proceedings of the 4th International Conference on Algorithmic Decision Theory*, pages 359–376. Springer-Verlag *Lecture Notes in Artificial Intelligence*, September 2015.

[6] G. Boella, S. Kaci, and L. van der Torre. Dynamics in argumentation with single extensions: Abstraction principles and the grounded extension. In *Proceedings of the 10th European Conference on Symbolic and Quantitative Approaches to Reasoning with Uncertainty*, pages 107–118. Springer-Verlag *Lecture Notes in Artificial Intelligence #5590*, July 2009.

[7] S. Bouveret, U. Endriss, and J. Lang. Fair division under ordinal preferences: Computing envy-free allocations of indivisible goods. In *Proceedings of the 19th European Conference on Artificial Intelligence*, pages 387–392. IOS Press, August 2010.

[8] C. Cayrol, C. Devred, and M. Lagasquie-Schiex. Handling ignorance in argumentation: Semantics of partial argumentation frameworks. In *Proceedings of the 9th European Conference on Symbolic and Quantitative Approaches to Reasoning with Uncertainty*, pages 259–270. Springer-Verlag *Lecture Notes in Artificial Intelligence #4724*, November 2007.

[9] C. Cayrol, F. de Saint-Cyr, and M. Lagasquie-Schiex. Change in abstract argumentation frameworks: Adding an argument. *Journal of Artificial Intelligence Research*, 38:49–84, 2010.

[10] Y. Chevaleyre, J. Lang, N. Maudet, J. Monnot, and L. Xia. New candidates welcome! Possible winners with respect to the addition of new candidates. *Mathematical Social Sciences*, 64(1): 74–88, 2012.

[11] S. Coste-Marquis, C. Devred, S. Konieczny, M. Lagasquie-Schiex, and P. Marquis. On the merging of Dung's argumentation systems. *Artificial Intelligence*, 171(10):730–753, 2007.

[12] S. Coste-Marquis, S. Konieczny, J. Mailly, and P. Marquis. On the revision of argumentation systems: Minimal change of arguments statuses. In *Proceedings of the 14th International Conference on Principles of Knowledge Representation and Reasoning*. AAAI Press, July 2014.

[13] Y. Dimopoulos and A. Torres. Graph theoretical structures in logic programs and default theories. *Theoretical Computer Science*, 170(1):209–244, 1996.

[14] D. Doder and S. Woltran. Probabilistic argumentation frameworks – a logical approach. In *Proceedings of the 8th International Conference on Scalable Uncertainty Management*, pages 134–147. Springer-Verlag *Lecture Notes in Artificial Intelligence #8720*, September 2014.

[15] P. Dung. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and *n*-person games. *Artificial Intelligence*, 77(2):321–357, 1995.

[16] P. Dunne and M. Wooldridge. Complexity of abstract argumentation. In I. Rahwan and G. Simari, editors, *Argumentation in Artificial Intelligence*, chapter 5, pages 85–104. Springer, 2009.

[17] B. Fazzinga, S. Flesca, and F. Parisi. On the complexity of probabilistic abstract argumentation. In *Proceedings of the 23rd International Joint Conference on Artificial Intelligence*, pages 898–904. AAAI Press/IJCAI, August 2013.

[18] B. Fazzinga, S. Flesca, and F. Parisi. Efficiently estimating the probability of extensions in abstract argumentation. In *Proceedings of the 7th International Conference on Scalable Uncertainty Management*, pages 106–119. Springer-Verlag *Lecture Notes in Artificial Intelligence #8078*, 2013.

[19] M. Garey and D. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman and Company, 1979.

[20] A. Hunter. Probabilistic qualification of attack in abstract argumentation. *International Journal of Approximate Reasoning*, 55(2):607–638, 2014.

[21] K. Konczak and J. Lang. Voting procedures with incomplete preferences. In *Proceedings of the Multidisciplinary IJCAI-05 Workshop on Advances in Preference Handling*, pages 124–129, July/August 2005.

[22] J. Lang, A. Rey, J. Rothe, H. Schadrack, and L. Schend. Representing and solving hedonic games with ordinal preferences and thresholds. In *Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems*, pages 1229–1237. IFAAMAS, May 2015.

[23] H. Li, N. Oren, and T. Norman. Probabilistic argumentation frameworks. In *Proceedings of the 1st International Workshop on Theory and Applications of Formal Argumentation*, pages 1–16. Springer-Verlag *Lecture Notes in Artificial Intelligence #7132*, July 2011.

[24] B. Liao, L. Jin, and R. Koons. Dynamics of argumentation systems: A division-based method. *Artificial Intelligence*, 175(11):1790–1814, 2011.

[25] A. Meyer and L. Stockmeyer. The equivalence problem for regular expressions with squaring requires exponential space. In *Proceedings of the 13th IEEE Symposium on Switching and Automata Theory*, pages 125–129. IEEE Computer Society Press, 1972.

[26] C. Papadimitriou. *Computational Complexity*. Addison-Wesley, 2nd edition, 1995. Reprinted with corrections.

[27] C. Papadimitriou and M. Yannakakis. The complexity of facets (and some facets of complexity). *Journal of Computer and System Sciences*, 28(2):244–259, 1984.

[28] T. Rienstra. Towards a probabilistic Dung-style argumentation system. In *Proceedings of the First International Conference on Agreement Technologies*, pages 138–152, October 2012.

[29] J. Rothe. *Complexity Theory and Cryptology. An Introduction to Cryptocomplexity*. EATCS Texts in Theoretical Computer Science. Springer-Verlag, 2005.

[30] L. Stockmeyer. The polynomial-time hierarchy. *Theoretical Computer Science*, 3(1):1–22, 1976.

[31] F. Tohmé, G. Bodanza, and G. Simari. Aggregation of attack relations: A social-choice theoretical analysis of defeasibility criteria. In *Foundations of Information and Knowledge Systems*, pages 8–23. Springer, 2008.

[32] K. Wagner. More complicated questions about maxima and minima, and some closures of NP. *Theoretical Computer Science*, 51(1–2):53–80, 1987.

[33] L. Xia and V. Conitzer. Determining possible and necessary winners given partial orders. *Journal of Artificial Intelligence Research*, 41:25–67, 2011.

Dorothea Baumeister, Daniel Neugebauer, Jörg Rothe, and Hilmar Schadrack
Institut für Informatik
Heinrich-Heine-Universität Düsseldorf
Düsseldorf, Germany
Email: {baumeister,neugebauer,rothe,schadrack}@cs.uni-duesseldorf.de

A Appendix

A.1 Proofs for the verification problems in attack-incomplete argumentation frameworks

Definition 22. Let $AtIAF = \langle \mathscr{A}, \mathscr{R}, \mathscr{R}^? \rangle$ be an attack-incomplete argumentation framework and let $S \subseteq \mathscr{A}$. The *pessimistic completion of AtIAF for S* is $AtIAF_S^{\text{pes}} = \langle \mathscr{A}, \mathscr{R}_S^{\text{pes}} \rangle$ with $\mathscr{R}_S^{\text{pes}} = \mathscr{R} \cup \{(a,b) \in \mathscr{R}^? \mid b \in S\}$.

Lemma 23. *Let $AtIAF = \langle \mathscr{A}, \mathscr{R}, \mathscr{R}^? \rangle$ be an attack-incomplete argumentation framework, $S \subseteq \mathscr{A}$, and let $AtIAF_S^{\text{pes}}$ be the pessimistic completion of AtIAF for S.*

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

Proof. Here, the left-to-right implications are trivial: If S necessarily fulfills a criterion in $AtIAF$, it must fulfill it in particular in the pessimistic completion. We prove the other direction of the implications individually:

1. If S is conflict-free in $AtIAF_S^{\text{pes}}$, then all interior attacks among elements of S are in $\mathscr{R}^-$, because if such an attack were in $\mathscr{R}$, S would not be conflict-free in any completion of $AtIAF$, and if such an attack was in $\mathscr{R}^?$, it would have been included in $\mathscr{R}_S^{\text{pes}}$, which contradicts our assumption that S is conflict-free in $AtIAF_S^{\text{pes}}$. Since all interior attacks among elements of S are in $\mathscr{R}^-$, S is necessarily conflict-free in $AtIAF$.
2. If each $a \in S$ is acceptable with respect to S in $AtIAF_S^{\text{pes}}$, then S defends each of these arguments against all their attackers. By construction, $\mathscr{R}_S^{\text{pes}}$ contains all possible attacks from $\mathscr{R}^?$ that attack elements of S, and no possible attacks that can defend S. Therefore, all attacks in $\mathscr{R}_S^{\text{pes}}$ that defend elements of S against possible or definite attacks are already in $\mathscr{R}$, otherwise they could not be in $\mathscr{R}_S^{\text{pes}}$, and are therefore in $\mathscr{R}^*$ for any completion $AtIAF^*$. This implies that each $a \in S$ is necessarily acceptable with respect to S in $AtIAF$.
3. Assume that S is an admissible extension of $AtIAF_S^{\text{pes}}$, i.e., S is conflict-free in $AtIAF_S^{\text{pes}}$ and each $a \in S$ is acceptable with respect to S in $AtIAF_S^{\text{pes}}$. The previous results then imply that S is necessarily conflict-free in $AtIAF$ and each $a \in S$ is necessarily acceptable with respect to S in $AtIAF$, which immediately yields that S is necessarily admissible in $AtIAF$.
4. Assume that S is a stable extension of $AtIAF_S^{\text{pes}}$, i.e., S is conflict-free in $AtIAF_S^{\text{pes}}$ and S attacks each element $b \notin S$ in $AtIAF_S^{\text{pes}}$. Again, this implies that S is necessarily conflict-free in $AtIAF$. Further, since $\mathscr{R}_S^{\text{pes}}$ only contains attacks by S that are already in $\mathscr{R}$, S necessarily attacks each element $b \notin S$ in $AtIAF$. Combined, we have that S is necessarily stable in $AtIAF$.

This completes the proof. ❑

Theorem 24. *For* $\mathbf{s} \in \{\text{CF}, \text{AD}, \text{ST}\}$*,* **s**-ATTINCNV *is in* P.

Proof. The pessimistic completion can obviously be constructed in polynomial time. As already mentioned, the problem **s**-VERIFICATION can be solved in polynomial time for a given completion. Lemma 23 then provides that the answer to **s**-ATTINCNV is the same as that to **s**-VERIFICATION for the pessimistic completion. ❑

Definition 25. Let $AtIAF = \langle \mathscr{A}, \mathscr{R}, \mathscr{R}^? \rangle$ be an attack-incomplete argumentation framework and $S \subseteq \mathscr{A}$. The *fixed completion* $AtIAF_S^{\text{fix}}$ of $AtIAF$ is the completion that is obtained by the following algorithm. The algorithm defines a finite sequence $(AtIAF_i)_{i \geq 0}$ of attack-incomplete argumentation frameworks, with the fixed completion being the minimal completion of the sequence's last element.

1. Include definite attacks: Let $AtIAF_0 = AtIAF$.
2. Include external conflicts: Let $AtIAF_1 = \langle \mathscr{A}, \mathscr{R}_1, \mathscr{R}_1^? \rangle$ with $\mathscr{R}_1 = \mathscr{R} \cup \{(a,b) \in \mathscr{R}^? \mid a \notin S$ and $b \notin S\}$ and $\mathscr{R}_1^? = \mathscr{R}^? \setminus \mathscr{R}_1$.
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 $AtIAF_2 = \langle \mathscr{A}, \mathscr{R}_2, \mathscr{R}_2^? \rangle$ with $\mathscr{R}_2 = \mathscr{R}_1 \cup \{(a,b) \in \mathscr{R}_1^? \mid a \in S$ and $b \in T\}$ and $\mathscr{R}_2^? = \mathscr{R}_1^? \setminus \mathscr{R}_2$.
4. Avoid arguments outside of S to be acceptable with respect to S: For the current i (initially, $i = 2$), let $AtIAF_i^{\min}$ be the minimal completion of $AtIAF_i$ and $T_i = F_{AtIAF_i^{\min}}(S) \setminus S$ (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 $AtIAF_{i+1} = \langle \mathscr{A}, \mathscr{R}_{i+1}, \mathscr{R}_{i+1}^? \rangle$ with$\mathscr{R}_{i+1} = \mathscr{R}_i \cup \{(a,b) \in \mathscr{R}_i^? \mid a \in S$ and $b \in T_i\}$ and $\mathscr{R}_{i+1}^? = \mathscr{R}_i^? \setminus \mathscr{R}_{i+1}$, and set $i \leftarrow i+1$.
5. Repeat Step 4 until no more attacks are added.
6. The fixed completion of $AtIAF$ is $AtIAF_S^{\text{fix}} = \langle \mathscr{A}, \mathscr{R}_S^{\text{fix}} \rangle$ with $\mathscr{R}_S^{\text{fix}} = \mathscr{R}_i$.

Definition 26. Let $AtIAF = \langle \mathscr{A}, \mathscr{R}, \mathscr{R}^? \rangle$ be an attack-incomplete argumentation framework and $S \subseteq \mathscr{A}$. The *unfixed completion* $AtIAF_S^{\text{unf}}$ of $AtIAF$ is the completion that is obtained by the following algorithm. The algorithm defines a finite sequence $(AtIAF_i)_{i \geq 0}$ of attack-incomplete argumentation frameworks, with the unfixed completion being the minimal completion of the sequence's last element.

1. Include definite attacks: Let $AtIAF_0 = AtIAF$.
2. Include attacks against S: Let $AtIAF_1 = \langle \mathscr{A}, \mathscr{R}_1, \mathscr{R}_1^? \rangle$ with $\mathscr{R}_1 = \mathscr{R} \cup \{(a,b) \in \mathscr{R}^? \mid b \in S\}$ and $\mathscr{R}_1^? = \mathscr{R}^? \setminus \mathscr{R}_1$.
3. Exclude external conflicts: Let $AtIAF_2 = \langle \mathscr{A}, \mathscr{R}_2, \mathscr{R}_2^? \rangle$ with $\mathscr{R}_2 = \mathscr{R}_1$ and $\mathscr{R}_2^? = \mathscr{R}_1^? \setminus \{(a,b) \in \mathscr{R}_1^? \mid a \notin S$ 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 $AtIAF_3 = \langle \mathscr{A}, \mathscr{R}_3, \mathscr{R}_3^? \rangle$ with $\mathscr{R}_3 = \mathscr{R}_2$ and $\mathscr{R}_3^? = \mathscr{R}_2^? \setminus \{(a,b) \in \mathscr{R}_2^? \mid a \in S$ 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 $AtIAF_{i,S'}^{\text{opt}}$ be the optimistic completion of $AtIAF_i$ for S' and let $AtIAF_i^{\min}$ be the minimal completion of $AtIAF_i$.
   (b) If t_j is acceptable with respect to S in $AtIAF_{i,S'}^{\text{opt}}$, but not acceptable with respect to S in $AtIAF_i^{\min}$, let $AtIAF_{i+1} = \langle \mathscr{A}, \mathscr{R}_{i+1}, \mathscr{R}_{i+1}^? \rangle$ with $\mathscr{R}_{i+1} = \mathscr{R}_i \cup \{(a,b) \in \mathscr{R}_i^? \mid a \in S$ and $(b,t_j) \in \mathscr{R}_i\}$ and $\mathscr{R}_{i+1}^? = \mathscr{R}_i^? \setminus \mathscr{R}_{i+1}$, 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 $AtIAF$ is $AtIAF_S^{\text{unf}} = \langle \mathscr{A}, \mathscr{R}_S^{\text{unf}} \rangle$ with $\mathscr{R}_S^{\text{unf}} = \mathscr{R}_i$.

Lemma 27. *For an attack-incomplete argumentation framework $AtIAF = \langle \mathscr{A}, \mathscr{R}, \mathscr{R}^? \rangle$ and a set $S \subseteq \mathscr{A}$ of arguments, the fixed completion $AtIAF_S^{\text{fix}}$ and the unfixed completion $AtIAF_S^{\text{unf}}$ 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 $AtIAF$, 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. ❑

Lemma 28. *Let $AtIAF = \langle \mathscr{A}, \mathscr{R}, \mathscr{R}^? \rangle$ be an attack-incomplete argumentation framework, $S \subseteq \mathscr{A}$, and let $AtIAF_S^{\text{fix}}$ be the fixed completion of AtIAF for S.*

(1) S is a possibly complete extension of AtIAF if and only if S is a complete extension of $AtIAF_S^{\text{fix}}$.

(2) S is a possibly grounded extension of AtIAF if and only if S is the grounded extension of $AtIAF_S^{\text{fix}}$.

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

(1) Assume that S is not a complete extension of $AtIAF_S^{\text{fix}}$, i.e., S is not a fixed point of $F_{AtIAF_S^{\text{fix}}}$. We will show that this implies that S is not possibly complete in $AtIAF$. Since S is not a fixed point of $F_{AtIAF_S^{\text{fix}}}$, there is an argument $b \notin S$ which is acceptable with respect to S in $AtIAF_S^{\text{fix}}$. We prove that, then, there must be some $c \notin S$ for which all attackers of c are attacked by S in any completion $AtIAF^*$ ($c = b$ may or may not be the case) by individually covering all cases in which attacks are added to $\mathscr{R}_S^{\text{fix}}$:

All attacks from $\mathscr{R}^?$ between arguments outside of S, which are added to $\mathscr{R}_S^{\text{fix}}$ 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 $AtIAF_S^{\text{fix}}$ and not in $AtIAF^*$, but this happens only if S attacks an attacker (or several attackers) of b in $AtIAF_S^{\text{fix}}$ that would otherwise be acceptable with respect to S, and which therefore must be acceptable with respect to S in $AtIAF^*$. In either case, if an argument outside of S is acceptable with respect to S in $AtIAF_S^{\text{fix}}$, then some argument outside of S must be acceptable with respect to S in each completion $AtIAF^*$ of $AtIAF$ in which S is admissible. Therefore, if S is not a complete extension of $AtIAF_S^{\text{fix}}$, it is not a complete extension of any completion $AtIAF^*$ of $AtIAF$, and therefore not a possibly complete extension of $AtIAF$.

(2) Let $AtIAF^*$ be an arbitrary completion of $AtIAF$ and assume that S is its grounded extension. We prove that, then, S is also the grounded extension of $AtIAF_S^{\text{fix}}$. Let $A_i = F^i_{AtIAF^*}(\emptyset)$ and $B_i = F^i_{AtIAF_S^{\text{fix}}}(\emptyset)$, where F^i is the i-fold composition of the respective characteristic function F. Since S is grounded in $AtIAF^*$, it is complete in $AtIAF_S^{\text{fix}}$ 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 $AtIAF_S^{\text{fix}}$.

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 $AtIAF^*$ (respectively, in $AtIAF_S^{\text{fix}}$). 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 $AtIAF^*$ are unattacked in $AtIAF_S^{\text{fix}}$, 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 $AtIAF^*$, but not acceptable with respect to B_k in $AtIAF_S^{\text{fix}}$. We know that—since $A_{k+1} \subseteq S$—no possible attacks against A_{k+1} (and in particular, against a) are included in $AtIAF_S^{\text{fix}}$ and all possible defending attacks by arguments in A_{k+1} against arguments outside of S are included in $AtIAF_S^{\text{fix}}$. Further, no element of S attacks a in $AtIAF_S^{\text{fix}}$, since $a \in S$ and S is complete in $AtIAF_S^{\text{fix}}$. Therefore, a is acceptable with respect to A_k in $AtIAF_S^{\text{fix}}$; otherwise it could not be acceptable with respect to A_k in $AtIAF^*$. Now, the only way for a to not be acceptable with respect to B_k in $AtIAF_S^{\text{fix}}$ 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 $AtIAF^*$, since a is acceptable with respect to A_k in $AtIAF^*$. This implies that $b \notin S$, since S is conflict-free in $AtIAF^*$. 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 $AtIAF_S^{\text{fix}}$. Therefore, a must be acceptable with respect to B_k in $AtIAF_S^{\text{fix}}$, 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 $B_i \not\subseteq S$ for some $i \geq 0$. Then it also holds that $G_S^{\text{fix}} \not\subseteq S$ for the grounded extension G_S^{fix} of $AtIAF_S^{\text{fix}}$. 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 $AtIAF_S^{\text{fix}}$, since the grounded extension G_S^{fix} of $AtIAF_S^{\text{fix}}$ 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. ❑

Lemma 29. *Let $AtIAF = \langle \mathscr{A}, \mathscr{R}, \mathscr{R}^? \rangle$ be an attack-incomplete argumentation framework, $S \subseteq \mathscr{A}$, and let $AtIAF_S^{\text{unf}}$ be the unfixed completion of AtIAF for S.*

(1) S is a necessarily complete extension of AtIAF if and only if S is a complete extension of $AtIAF_S^{\text{unf}}$.

(2) S is a necessarily grounded extension of AtIAF if and only if S is the grounded extension of $AtIAF_S^{\text{unf}}$.

Proof. Again, the left-to-right implication is trivial in both cases. We prove the other direction of the implications individually. First, if S is not necessarily admissible in $AtIAF$, S is not admissible either (and therefore, neither complete nor grounded) in $AtIAF_S^{\text{unf}}$, because $AtIAF_S^{\text{unf}}$ includes all possible attacks against arguments in S and excludes all defending attacks by arguments in S. We may therefore assume that S is necessarily admissible in $AtIAF$.

(1) Assume that S is not necessarily complete in $AtIAF$. We prove that S is not complete in $AtIAF_S^{\text{unf}}$: Since S is necessarily admissible but not necessarily complete in $AtIAF$, there is a completion $AtIAF^*$ of $AtIAF$ in which there exists some $b' \in \mathscr{A} \setminus S$ that is acceptable with respect to S in $AtIAF^*$. Obviously, this means that b' is possibly acceptable with respect to S in $AtIAF$. We will prove that, after each step of the algorithm, if there is some $b \in \mathscr{A} \setminus S$ that is acceptable with respect to S in $AtIAF_i$, then there also is some $c \in \mathscr{A} \setminus S$ that is acceptable with respect to S in $AtIAF_{i+1}$ ($c = b$ may or may not be the case).

- After Step 1, b' is possibly acceptable with respect to S in $AtIAF_0$, since $AtIAF_0 = AtIAF$.
- After Step 2, b' is possibly acceptable with respect to S in $AtIAF_1$, because including attacks against S has no influence on whether S possibly attacks all attackers of b'.
- After Step 3, b' is possibly acceptable with respect to S in $AtIAF_2$, because excluding attacks between arguments in $\mathscr{A} \setminus S$ can only make it more likely for S to attack all attackers of b'.
- Step 4 has no effect on instances where S is necessarily admissible, because there are no possible defending attacks by S against $\mathscr{A} \setminus S$ that could be excluded, since in such an instance S necessarily defends itself against all possible attacks.
- The only way for an argument $b \in \mathscr{A} \setminus S$ to no longer be possibly acceptable with respect to S in $AtIAF_{i+1}$ after an iteration of Step 5 is if an attack by some $a \in S$ against b is included. If

this is the case, the defended argument t_j is possibly acceptable with respect to S in $AtIAF_{i+1}$. Either way, the previously possibly acceptable argument b or the new argument t_j is possibly acceptable with respect to S in $AtIAF_{i+1}$.

After Step 4, the only attacks that are not yet definite are attacks by arguments in S against arguments in $\mathscr{A} \setminus S$. Therefore, the only way for the condition in Step 5b to be met—i.e., t_j is possibly, but not currently accepted by S—is if there is an attack $(a,b) \in \mathscr{R}_i^?$ with $a \in S$ and $(b,t_j) \in \mathscr{R}_i$, which proves that $AtIAF_{i+1} \neq AtIAF_i$. So, when the algorithm terminates in Step 7, we know that there is an argument $b \in \mathscr{A} \setminus S$ that is possibly acceptable with respect to S in $AtIAF_i$ (as proven earlier) and that is also acceptable with respect to S in $AtIAF_i$'s minimal completion, because otherwise the condition in Step 5b would have been met. Since the unfixed completion is $AtIAF_i$'s minimal completion, this establishes that there is an argument in $\mathscr{A} \setminus S$ that is acceptable with respect to S in $AtIAF_S^{\text{unf}}$, which implies that S is not complete in $AtIAF_S^{\text{unf}}$, and concludes the proof of the first item.

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

First, we prove that $B_i \subseteq A_i$ 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 $AtIAF^*$ (respectively, in $AtIAF_S^{\text{unf}}$). We know that $B_1 \subseteq S$. Since the unfixed completion includes all possible attacks against elements of S, all $a \in S$ that are unattacked in $AtIAF_S^{\text{unf}}$ are necessarily unattacked, and therefore unattacked in $AtIAF^*$, too, which proves $B_1 \subseteq A_1$. If we now have $B_k \subseteq A_k$ for some $k \geq 1$, this implies $B_{k+1} \subseteq A_{k+1}$: Assume that this is not true, i.e., that $B_k \subseteq A_k$, but there is an argument $b \in B_{k+1}$ with $b \notin A_{k+1}$. Then, b is acceptable with respect to B_k in $AtIAF_S^{\text{unf}}$, but not acceptable with respect to A_k in $AtIAF^*$. Recall that all possible attacks against B_{k+1} (and in particular, against b) are included in $AtIAF_S^{\text{unf}}$ and no possible defending attacks by arguments in B_{k+1} against arguments outside of S are included in $AtIAF_S^{\text{unf}}$. Therefore, since b is acceptable with respect to $B_k \subseteq S$ in $AtIAF_S^{\text{unf}}$, it is necessarily acceptable with respect to B_k and, in particular, acceptable with respect to B_k in $AtIAF^*$. Now, the only way for b to not be acceptable with respect to A_k in $AtIAF^*$ is if there were some $a \in A_k \setminus B_k$ that possibly attacks b. Then there would have to be a defending attack by an argument $d \in B_k$ against a in $AtIAF_S^{\text{unf}}$, since b is acceptable with respect to B_k in $AtIAF_S^{\text{unf}}$. This implies that $a \notin S$, since S is conflict-free in $AtIAF_S^{\text{unf}}$. Finally, since (d,a) is a necessary attack, it holds in particular that $(d,a) \in \mathscr{R}^*$, which contradicts that A_k is admissible in $AtIAF^*$. Therefore, b must be acceptable with respect to A_k in $AtIAF^*$, which proves that $B_{k+1} \subseteq A_{k+1}$.

Now we prove that $A_i \subseteq S$ for all $i \geq 0$: Assume that $A_i \not\subseteq S$ for some $i \geq 0$. Then it also holds that $G^* \not\subseteq S$ for the grounded extension G^* of $AtIAF^*$. It further holds that $S \subset G^*$, since there exists a $j \geq 0$ such that $S \subseteq A_i$ for all $i \geq j$, as established before. However, this contradicts the fact that S is complete in $AtIAF^*$, since the grounded extension G^* of $AtIAF^*$ is its least complete extension with respect to set inclusion [15] and cannot be a strict subset of the complete extension S. This completes the proof. ❑

Theorem 30. *For* $\mathbf{s} \in \{\text{CP}, \text{GR}\}$*, both* **s**-ATTINCPV *and* **s**-ATTINCNV *are in* P.

Proof. Lemma 27 provides polynomial-time constructability for the fixed and unfixed completion. Given a completion, **s**-VERIFICATION can be solved in polynomial time, and Lemmas 28 and 29 imply that the answer to, respectively, **s**-ATTINCPV and **s**-ATTINCNV is the same as that to **s**-VERIFICATION for the respective completion. ❑

Theorem 31. *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 $AtIAF = \langle \mathscr{A}, \mathscr{R}, \mathscr{R}^? \rangle$ and a set $S \subseteq \mathscr{A}$, there is a completion $AtIAF^* = \langle \mathscr{A}, \mathscr{R}^* \rangle$ such that S is preferred in $AtIAF^*$. To check whether S is preferred in $AtIAF^*$, one has to check whether for all sets $S' \subseteq \mathscr{A}$ 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 $AtIAF^*$ of the given attack-incomplete $AtIAF$ 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 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 [13]. 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 $(\langle \mathscr{A}, \mathscr{R}, \emptyset \rangle, S)$. The only completion of $\langle \mathscr{A}, \mathscr{R}, \emptyset \rangle$ 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 $(\langle \mathscr{A}, \mathscr{R}, \emptyset \rangle, S) \in$ PR-ATTINCPV, which in turn is equivalent to $(\langle \mathscr{A}, \mathscr{R}, \emptyset \rangle, S) \in$ PR-ATTINCNV. ❑

A.2 Proofs for the verification problems in argument-incomplete argumentation frameworks

Proposition 32. PR-ARGINCPV *is* coNP-*hard and* PR-ARGINCNV *is* coNP-*complete.*

Proof. We show coNP-hardness by a reduction from the coNP-complete problem PR-VERIFICATION. Let $(\langle \mathscr{A}, \mathscr{R} \rangle, S)$ be a given instance of PR-VERIFICATION, and construct from it $(\langle \mathscr{A}, \emptyset, \mathscr{R} \rangle, S)$, considered as an instance of both PR-ARGINCPV and PR-ARGINCNV. In the argument-incomplete argumentation framework, there are no arguments that can possibly join the discussion. Hence, the only completion in both cases is the argumentation framework $\langle \mathscr{A}, \mathscr{R} \rangle$. Now, it is easy to see that

$$
\begin{aligned}
& (\langle \mathscr{A}, \mathscr{R} \rangle, S) \in \text{PR-VERIFICATION} \\
\iff \quad & (\langle \mathscr{A}, \emptyset, \mathscr{R} \rangle, S) \in \text{PR-ARGINCPV} \\
\iff \quad & (\langle \mathscr{A}, \emptyset, \mathscr{R} \rangle, S) \in \text{PR-ARGINCNV}.
\end{aligned}
$$

This completes the proof. ❑

Theorem 33. *For* $\mathbf{s} \in \{\text{ST}, \text{CP}, \text{GR}\}$*,* **s**-ARGINCPV *is* NP-*complete, and* PR-ARGINCPV *is* NP-*hard.*

Proof. Membership of the three former problems in NP is clear. It remains to show hardness for all four problems. We do this by showing that the reduction used in Theorem 14 also works for those four problems. To this end, we will prove that

$$
\begin{aligned}
& (\langle \mathscr{A}, \mathscr{A}^?, \mathscr{R} \rangle, S) \in \text{AD-ARGINCPV} \\
\iff \quad & (\langle \mathscr{A}, \mathscr{A}^?, \mathscr{R} \rangle, S) \in \text{ST-ARGINCPV} \\
\iff \quad & (\langle \mathscr{A}, \mathscr{A}^?, \mathscr{R} \rangle, S) \in \text{PR-ARGINCPV} \\
\iff \quad & (\langle \mathscr{A}, \mathscr{A}^?, \mathscr{R} \rangle, S) \in \text{GR-ARGINCPV} \\
\iff \quad & (\langle \mathscr{A}, \mathscr{A}^?, \mathscr{R} \rangle, S) \in \text{CP-ARGINCPV}
\end{aligned}
$$

holds for the instance $(\langle \mathscr{A}, \mathscr{A}^?, \mathscr{R}\rangle, S)$ constructed in the proof of Theorem 14.

$(\langle \mathscr{A}, \mathscr{A}^?, \mathscr{R}\rangle, S) \in$ AD-ARGINCPV implies $(\langle \mathscr{A}, \mathscr{A}^?, \mathscr{R}\rangle, S) \in$ ST-ARGINCPV: If $S|_{\mathscr{A}^*}$ is admissible for a completion $\langle \mathscr{A}^*, \mathscr{R}|_{\mathscr{A}^*}\rangle$, it in particular is conflict-free. We know from the reduction that $\langle \mathscr{A}^*, \mathscr{R}|_{\mathscr{A}^*}\rangle$ only contains arguments S_j that do not attack each other, and all these arguments belong to $S|_{\mathscr{A}^*}$. Hence, the only arguments outside of $S|_{\mathscr{A}^*}$ are the b_i's. But all of them are attacked, as explained in the proof of Theorem 14. Therefore, $S|_{\mathscr{A}^*}$ is a stable extension of $\langle \mathscr{A}^*, \mathscr{R}|_{\mathscr{A}^*}\rangle$.

$(\langle \mathscr{A}, \mathscr{A}^?, \mathscr{R}\rangle, S) \in$ PR-ARGINCPV implies $(\langle \mathscr{A}, \mathscr{A}^?, \mathscr{R}\rangle, S) \in$ GR-ARGINCPV: If $S|_{\mathscr{A}^*}$ is preferred for a completion $\langle \mathscr{A}^*, \mathscr{R}|_{\mathscr{A}^*}\rangle$, it is admissible, and thus the only arguments that are not attacked by any other argument are those S_j that correspond to an exact cover. This means for the characteristic function of this completion $\langle \mathscr{A}^*, \mathscr{R}|_{\mathscr{A}^*}\rangle$ that the output of the first step is the set that contains exactly those S_j. In the second step, we add argument x, because all those S_j defend x against all attacks from the arguments b_i. No new arguments are added in step three. Therefore, this set is the grounded extension of the argumentation framework $\langle \mathscr{A}^*, \mathscr{R}|_{\mathscr{A}^*}\rangle$. But this set is exactly the set $S|_{\mathscr{A}^*}$. Hence, $S|_{\mathscr{A}^*}$ is the grounded extension of $\langle \mathscr{A}^*, \mathscr{R}|_{\mathscr{A}^*}\rangle$.

It is easy to see the three remaining implications needed to prove these five statements equivalent: Every stable set is preferred, every grounded set is complete, and every complete set is admissible. This completes the proof. ❑

We now strengthen the NP-hardness lower bound for PR-ARGINCPV given in Theorem 33 to DP-hardness. The following lemma due to Wagner [32] gives a sufficient condition for proving hardness for DP.

Lemma 34 (Wagner [32]). *Let A be some* NP*-hard problem, and let B be any set. If there exists a polynomial-time computable function f such that, for any two instances z_1 and z_2 of A for which $z_2 \in A$ implies $z_1 \in A$, we have*

$$(z_1 \in A \text{ and } z_2 \notin A) \iff f(z_1, z_2) \in B,$$

then B is DP*-hard.*

Theorem 35. PR-ARGINCPV *is* DP*-hard.*

Proof. We will use Wagner's lemma to show DP-hardness: Let PR-ARGINCPV be the set B from Wagner's lemma, and let X3C be the NP-complete problem A in that lemma. Let z_1 and z_2 be two instances of X3C such that $z_2 \in$ X3C implies $z_1 \in$ X3C. We construct an instance $(\langle \mathscr{A}, \mathscr{A}^?, \mathscr{R}\rangle, S)$ of PR-ARGINCPV as follows:

- Construct an instance $(\langle \mathscr{A}_1, \mathscr{A}_1^?, \mathscr{R}_1\rangle, S_1)$ from the X3C instance z_1 exactly as in the proof of Theorem 14.
- The construction of an instance $(\langle \mathscr{A}_2, \mathscr{A}_2^?, \mathscr{R}_2\rangle, S_2)$ from the X3C instance z_2, however, is obtained as the composition of two reductions: Since PR-VERIFICATION is coNP-complete and X3C is NP-complete, there exists a reduction f such that $z_2 \notin$ X3C if and only if $f(z_2) \in$ PR-VERIFICATION. Now, letting g be the reduction from Proposition 32, we have $z_2 \notin$ X3C if and only if $g(f(z_2)) \in$ PR-ARGINCPV.
- Given two instances of PR-ARGINCPV, $(\langle \mathscr{A}_1, \mathscr{A}_1^?, \mathscr{R}_1\rangle, S_1)$ and $(\langle \mathscr{A}_2, \mathscr{A}_2^?, \mathscr{R}_2\rangle, S_2)$, let $(\langle \mathscr{A}, \mathscr{A}^?, \mathscr{R}\rangle, S) = (\langle \mathscr{A}_1 \cup \mathscr{A}_2, \mathscr{A}_1^? \cup \mathscr{A}_2^?, \mathscr{R}_1 \cup \mathscr{R}_2\rangle, S_1 \cup S_2)$ if $(\mathscr{A}_1 \cup \mathscr{A}_1^?) \cap (\mathscr{A}_2 \cup \mathscr{A}_2^?) = \emptyset$ (otherwise, simply rename the elements in one instance). Hence, this new instance consists of two disconnected argument-incomplete argumentation frameworks.

This completes the reduction. We claim that ($z_1 \in$ X3C and $z_2 \notin$ X3C) if and only if $(\langle \mathscr{A}, \mathscr{A}^?, \mathscr{R}\rangle, S) \in$ PR-ARGINCPV.

$(\Longrightarrow)$ If $z_1 \in$ X3C and $z_2 \notin$ X3C, then $(\langle \mathscr{A}_1, \mathscr{A}_1^?, \mathscr{R}_1\rangle, S_1)$ and $(\langle \mathscr{A}_2, \mathscr{A}_2^?, \mathscr{R}_2\rangle, S_2)$ both are yes-instances of PR-ARGINCPV. Thus we must have a completion for the first and a completion for the second argument-incomplete argumentation framework such that S_1 restricted to the arguments in this first completion and S_2 restricted to the arguments in the second completion are preferred in their respective completion. But then, using the same completions for each part of $\langle \mathscr{A}, \mathscr{A}^?, \mathscr{R} \rangle$, we have that S restricted to those arguments must be preferred in this argumentation framework. This is true because no new attacks are introduced in $\langle \mathscr{A}, \mathscr{A}^?, \mathscr{R} \rangle$ and, therefore, neither are any new conflicts added nor do the elements of S have to be defended by any other arguments than before. Hence, $(\langle \mathscr{A}, \mathscr{A}^?, \mathscr{R} \rangle, S)$ is a yes-instance of PR-ARGINCPV.

($\Longleftarrow$) Conversely, assume that $(\langle \mathscr{A}, \mathscr{A}^?, \mathscr{R} \rangle, S)$ is a yes-instance of PR-ARGINCPV, and assume further that $(\langle \mathscr{A}_i, \mathscr{A}_i^?, \mathscr{R}_i \rangle, S_i)$ is a no-instance of PR-ARGINCPV for some $i \in \{1,2\}$. Then there is no completion $\langle \mathscr{A}_i^*, \mathscr{R}_i|_{\mathscr{A}_i^*} \rangle$ of $\langle \mathscr{A}_i, \mathscr{A}_i^?, \mathscr{R}_i \rangle$ such that $S_i|_{\mathscr{A}_i^*}$ is preferred in it. That means that for every completion $\langle \mathscr{A}_i^*, \mathscr{R}_i|_{\mathscr{A}_i^*} \rangle$, $S_i|_{\mathscr{A}_i^*}$ either is not conflict-free, or is not admissible, or that there exists a superset of $S_i|_{\mathscr{A}_i^*}$ in $\langle \mathscr{A}_i^*, \mathscr{R}_i|_{\mathscr{A}_i^*} \rangle$ that is admissible. We consider these cases separately:

1. If $S_i|_{\mathscr{A}_i^*}$ is not conflict-free in $\langle \mathscr{A}_i^*, \mathscr{R}_i|_{\mathscr{A}_i^*} \rangle$, this conflict also exists in $S|_{\mathscr{A}^*}$ for any completion $\langle \mathscr{A}^*, \mathscr{R}|_{\mathscr{A}^*} \rangle$ of $\langle \mathscr{A}, \mathscr{A}^?, \mathscr{R} \rangle$ with $\mathscr{A}^* \cap (\mathscr{A}_i \cup \mathscr{A}_i^?) = \mathscr{A}_i^*$.

2. If $S_i|_{\mathscr{A}_i^*}$ is not admissible in $\langle \mathscr{A}_i^*, \mathscr{R}_i|_{\mathscr{A}_i^*} \rangle$, there must be an undefended attack. However, by the same argument as above, this attack is still undefended in any completion $\langle \mathscr{A}^*, \mathscr{R}|_{\mathscr{A}^*} \rangle$ of $\langle \mathscr{A}, \mathscr{A}^?, \mathscr{R} \rangle$ with $\mathscr{A}^* \cap (\mathscr{A}_i \cup \mathscr{A}_i^?) = \mathscr{A}_i^*$.

3. If there is a superset of $S_i|_{\mathscr{A}_i^*}$ preventing it from being preferred in $\langle \mathscr{A}_i^*, \mathscr{R}_i|_{\mathscr{A}_i^*} \rangle$, this superset translates into a superset of $S|_{\mathscr{A}^*}$ for any completion $\langle \mathscr{A}^*, \mathscr{R}|_{\mathscr{A}^*} \rangle$ of $\langle \mathscr{A}, \mathscr{A}^?, \mathscr{R} \rangle$ with $\mathscr{A}^* \cap (\mathscr{A}_i \cup \mathscr{A}_i^?) = \mathscr{A}_i^*$, thus also preventing $S|_{\mathscr{A}^*}$ from being preferred in $\langle \mathscr{A}^*, \mathscr{R}|_{\mathscr{A}^*} \rangle$.

Hence, none of these cases can happen, because $(\langle \mathscr{A}, \mathscr{A}^?, \mathscr{R} \rangle, S)$ is a yes-instance of PR-ARGINCPV. But this means that $S_i|_{\mathscr{A}_i^*}$ is a preferred extension of a completion $\langle \mathscr{A}_i^*, \mathscr{R}_i|_{\mathscr{A}_i^*} \rangle$ of $\langle \mathscr{A}_i, \mathscr{A}_i^?, \mathscr{R}_i \rangle$, a contradiction. ❑