ECAI 2023
K. Gal et al. (Eds.)

doi:10.3233/FAIA230332

Incomplete Bipolar Argumentation Frameworks

Bettina Fazzinga[a;*], Sergio Flesca[b] and Filippo Furfaro[b]

[a]DICES - University of Calabria
[b]DIMES - University of Calabria

Abstract. We introduce *Incomplete Bipolar Argumentation Frameworks* (iBAFs), the extension of Dung's *Abstract Argumentation Frameworks* (AAFs) allowing the simultaneous presence of *supports* (borrowed from BAFs - *Bipolar AAFs*) and of *uncertain* elements of the argumentation graph (borrowed from iAAFs - *incomplete AAFs*). We investigate the computational complexity of verification problem (under the possible perspective) and the acceptance problem, by studying its sensitivity to the semantics of supports and the semantics of extensions. On the one hand, we show that adding supports on top of incompleteness does not affect the complexity of the acceptance. On the other hand, surprisingly, we show that the joint use of bipolarity and incompleteness has a deep impact on the complexity of the verification: for the semantics under which the verification over AAFs is polynomial-time solvable, although moving from AAFs to BAFs or to iAAFs does not change the complexity, the complexity of the verification over iBAFs may increase up to NP-complete.

1 Introduction

In the recent years, owing to the usefulness exhibited by Dung's *Abstract Argumentation Framework* (AAF [18]) in various contexts (ranging from legal disputes [40] to process mining [25]), many extensions of AAFs have been proposed to enhance their modeling capability. Two directions have attracted particular attention: encoding *support relationships* between arguments, and taking into account possible *uncertainty* involving arguments and attacks. The efforts in the former direction have led to *Bipolar AAFs* (BAFs), a family of variants of AAFs differing from one another in the semantics of supports. Two well-established semantics of supports (to which those in [38, 39, 27] are added) are the *abstract semantics* [12], where a support encodes a positive interaction between arguments (semantically opposite to the meaning of attack), and the *deductive semantics* [9], where supports encode a "deductive" correlation: "a *supports* b" means that the acceptance of a implies the acceptance of b. The introduction of supports following these semantics has called for extending the reasoning paradigm defined over AAFs, since combining supports and attacks yields "implicit" attacks, called *supported* and *supermediated* attacks, as shown in Example 1.

Example 1 *The BAF in Figure 1(a) has "explicit" attacks (solid-line arrows) and supports (double-line arrows), as well as "implicit" attacks (not shown in the figure). In fact, whatever the semantics of supports (abstract or deductive), the facts that a supports b and that b attacks c imply the so called "*supported attack*" (a, c). Moreover, under the deductive semantics, since the acceptance of a supporting argument implies the acceptance of the supported argument, also the so called "*supermediated attack*" (e, a) should be considered, implied by the facts that b is attacked by e and supported by a.*

The second direction of research has led to several variants of AAFs, where the uncertainty affecting arguments and attacks is modeled *qualitatively*, by allowing the uncertain aspects of the dispute to be specified, or *quantitatively*, by also allowing the specification of the extent of this uncertainty. *Incomplete AAFs* (iAAFs) [8] are prominent representatives of qualitative approaches. Basically, an iAAF is an AAF where some arguments and/or attacks are marked as *uncertain*, as their occurrence in the dispute is not guaranteed. To address this form of incompleteness, the reasoning paradigm of AAFs has been refined to consider the multiple scenarios for the argumentation graph implied by the uncertainty, as shown in Example 2.

Example 2 *Figure 1(b) is an iAAF where the argument a and the attack (c, d) are uncertain. This iAAF compactly represents the four AAFs in Figure 1(c), called "completions", representing the alternative combinations of presence/absence for a and (c, d). The existence of alternative scenarios has been taken into account in the literature by revising the notion of "extension" into "i^*-extension", under two perspectives: a set is a "possible" (resp., "necessary") i^*-extension if it is an extension (in the classical sense) in at least one (resp., every) completion of the iAAF. Thus, under the complete semantics, $\{a, c\}$ is a possible (but not necessary) i^*-extension (as it is an extension only in the rightmost completion), while $\{b, d\}$ is a necessary (thus, also possible) i^*-extension.*

In the literature, there are several works investigating suitable adaptions of the traditional extension-based reasoning paradigm to BAFs, iAAFs [8] and their variants. A fundamental result [14, 13] is that reasoning over BAFs, in terms of verifying extensions, is the same as over AAFs, under the so called "Dungean" semantics (i.e. the admissible, stable, grounded, complete, preferred semantics), independently from the semantics of supports (abstract or deductive) and the coherence condition underlying the semantics – in fact, in the presence of supports, it may make sense to consider more general coherence conditions than the conflict-freeness. This means that

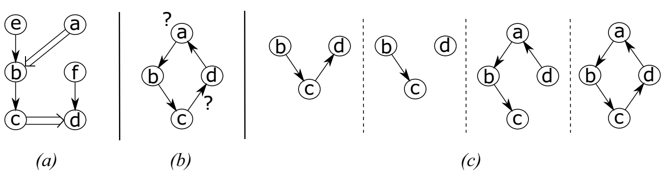

Figure 1: A BAF (a), an iAAF (b) and its completions (c)

* Corresponding Author. Email: bettina.fazzinga@unical.it

<table>
<tr><th rowspan="2"></th><th colspan="3">admissible</th><th colspan="3">stable</th><th colspan="3">complete</th><th colspan="3">grounded</th><th colspan="3">preferred</th></tr>
<tr><th>d-ad</th><th>s-ad</th><th>c-ad</th><th>d-st</th><th>s-st</th><th>c-st</th><th>d-co</th><th>s-co</th><th>c-co</th><th>d-gr</th><th>s-gr</th><th>c-gr</th><th>d-pr</th><th>s-pr</th><th>c-pr</th></tr>
<tr><td>AAF, BAF</td><td colspan="3">P</td><td colspan="3">P</td><td colspan="3">P</td><td colspan="3">P</td><td colspan="3">coNP-compl</td></tr>
<tr><td>iAAF</td><td colspan="3">P</td><td colspan="3">P</td><td colspan="3">P</td><td colspan="3">P</td><td colspan="3">Σ_2^p-compl</td></tr>
<tr><td>d-iBAF</td><td colspan="2">NP-compl</td><td>P</td><td colspan="3">P</td><td colspan="3">NP-compl</td><td colspan="3">NP-compl</td><td colspan="3">Σ_2^p-compl</td></tr>
<tr><td>s-iBAF</td><td>P</td><td>NP-compl</td><td>P</td><td colspan="3">P</td><td colspan="2">NP-compl</td><td>in NP</td><td colspan="2">NP-compl</td><td>in NP</td><td colspan="3">Σ_2^p-compl</td></tr>
</table>

Table 1: Complexity of verifying extensions (for iAAFs and iBAFs, the possible perspective is taken into account). The prefixes `d-`, `s-` and `c-` denote the coherence conditions considered for BAFs in the literature (specifically, `d-` consists in the conflict-freeness condition).

the verification problem over BAFs is coNP-complete under the preferred semantics and in P under every other Dungean semantics. The same holds for iAAFs, with the only exception of the preferred semantics under the possible perspective [20]: in this particular case, the verification problem's complexity moves from coNP-complete to Σ_2^p-complete, but is the same as over AAFs in the other cases: it is in coNP under the preferred semantics in the necessary perspective, and in P under the other Dungean semantics, whatever the perspective. In other words, using BAFs to specify supports or iAAFs to specify uncertain arguments/attacks does not make the reasoning computationally more complex than over AAFs.

Pushed by the popularity of BAFs and iAAFs, we introduce *Incomplete Bipolar AAFs* (iBAFs), that augment AAFs with (abstract or deductive) supports and the possibility of specifying which arguments, attacks, and supports are uncertain. The need for introducing uncertain supports was thoroughly discussed in [37], where an empirical study investigating how different people perceive the relationships between arguments was presented and showed that, analogously to the case of attacks, it often happens that, for a pair of arguments a, b, some people may see a support from a to b, while others may not. This clearly calls for resorting to the notion of uncertain support when merging different subjective views into a single AAF. As a real-life example motivating uncertain supports under the deductive perspective (the case of abstract supports is covered by several examples in [37]), consider the arguments below, claimed in a trial against a company PharmaX producing a vaccine:

a: "*As drug preparations must be sterile, mercury was used by PharmaX as an excipient of the vaccine*";
b: "*As the vaccine contains harmful ingredients, it should not have been used on humans.*"

Now: 1) some jurors believing that mercury is harmful may see a deductive support (a, b), as they believe that if a is accepted then b is accepted too; 2) other jurors believing that mercury is dangerous may not see this support, since they adopt a more "scientific" reasoning: they think that even if a is accepted, b may not be accepted, since b does not take into account that toxicity may be dosage-dependent; 3) jurors believing that mercury is safe probably see no correlation between a and b, and so on.

Starting from this, we thoroughly address the computational complexity of the reasoning over iBAFs, and consider both *the verification problem* (for which we investigate the variant IBVER under the possible perspective) and the *acceptance problem*. Interestingly, we show that, while adding bipolarity on top of incompleteness does not affect the complexity of the acceptance problem (as it can be easily shown to coincide with the case of iAAFs, where supports cannot be specified), the simultaneous presence of bipolarity and incompleteness makes things much more intricate for the verification than the case where these aspects do not occur or do not co-exist, as the complexity of IBVER depends on the combination $\langle$*semantics of supports*, *semantics of extensions*$\rangle$:
1) under the stable and the preferred semantics, IBVER is in P and Σ_2^p-complete, respectively, meaning that there is no additional source of complexity compared with BAFs and iAAFs;
2) under the other Dungean semantics, IBVER may become intractable (NP-complete), depending on the semantics of supports and on the coherence condition underlying the semantics of extension, meaning that in these cases new sources of complexity can arise from the joint use of supports and incompleteness.

Our results on the verification problem are summarized in Table 1, looking into which it is possible to appreciate that in some cases IBVER under the (variants of) the stable semantics is in P while it is NP-complete under (variants of) the admissible semantics. This is rather unusual in the context of frameworks generalizing AAFs, as typically the verification problems under the admissible and stable semantics lie in the same complexity class. Overall, the complexity analysis, besides being relevant *per se*, provides further contributions: 1) the hardness proofs give insights into conditions that can be expressed in iBAFs and that could not be expressed in BAFs or in iAAFs; 2) the algorithms used in the proofs of polynomial-time solvability provide practical solutions to the verification problem.

2 Preliminaries

We assume that the reader is familiar with *Abstract Argumentation Frameworks* (AAFs) and review *Bipolar Abstract Argumentation Frameworks* (BAFs). We do not provide specific preliminaries for *incomplete AAFs* (iAAFs) since what said in the introduction about them provides a sufficient background.

Definition 1 [BAF] *A bipolar abstract argumentation framework (BAF) is a tuple $\mathcal{F} = \langle \mathcal{A}, \mathcal{R}_a, \mathcal{R}_s \rangle$, where $\mathcal{A}$ is a set of arguments, while $\mathcal{R}_a \subseteq \mathcal{A} \times \mathcal{A}$ and $\mathcal{R}_s \subseteq \mathcal{A} \times \mathcal{A}$ are relations whose elements are called* attacks *and* supports*, respectively.*

If $\mathcal{R}_a$ (resp., $\mathcal{R}_s$) contains (a, b), we say "a attacks b" or "$a \to b$" (resp., "a supports b" or "$a \Rightarrow b$"). Moreover, for every (a, b) in the closure of $\mathcal{R}_s$, we say that "a transitively supports b" or $a \Rightarrow^+ b$. Since in the following we will introduce forms of attack other than those encoded by $\mathcal{R}_a$, those in $\mathcal{R}_a$ will be referred to as "*direct attacks*". *In what follows, all the definitions and notions regarding BAFs will be reviewed by assuming that a BAF $\mathcal{F} = \langle \mathcal{A}, \mathcal{R}_\mathbf{a}, \mathcal{R}_\mathbf{s} \rangle$ is given.*

In the first proposal of BAF [12], supports were given an *abstract semantics*, that is the opposite of the traditional semantics of attack in AAFs. According to this semantics. the combination of supports and direct attacks implies further attacks, called *supported attacks*.

Definition 2 [Supported attack] *Let $a, b \in \mathcal{A}$. There is a* supported attack *from a to b (written $a \to^s b$) iff there is $c \in \mathcal{A}$ such that $a \Rightarrow^+ c \wedge c \to b$.*

Among the other semantics for supports in the literature, we consider the well-established *deductive semantics* [9]: "*a supports b*" means that if a is accepted, then b is accepted. As observed in [13], adopting this semantics calls for considering not only direct and supported attacks, but also *supermediated* attacks.

Definition 3 [supermediated attack] *Let* $a, b \in \mathcal{A}$. *There is a* supermediated attack *from* a *to* b *(written as* $a \to^{m} b$*) iff there is an argument* $c \in \mathcal{A}$ *such that* $b \Rightarrow^{+} c \wedge (a \to^{s} c \vee a \to c)$.

Example 3 *The BAF in Figure 1*(a) *has the "direct" attacks* (e, b), (b, c), (f, d), *the supports* (a, b), (c, d), *the supported attack* (a, c), *and the supermediated attacks* (e, a), (f, c).

We partition BAFs into s-BAFs and d-BAFs: s-BAFs (resp., d-BAFs) adopt the abstract (resp., deductive) semantics, so only direct and supported attacks (resp., direct, supported and supermediated) are considered. In BAFs of any type, $a \to^{*} b$ will denote the existence of a "generic attack" from a to b: that is, over s-BAFs, $a \to^{*} b \equiv a \to b \vee a \to^{s} b$, while, over d-BAFs, $a \to^{*} b \equiv a \to b \vee a \to^{s} b \vee a \to^{m} b$.

The classical notion of defense becomes: given $a, b \in \mathcal{A}$ such that $a \to^{*} b$, the set S *defends* b *against* a if $\exists s \in S$ s.t. $s \to^{*} a$. In turn, $a \in \mathcal{A}$ is *acceptable* w.r.t. $S \subset \mathcal{A}$ if $\forall b \in \mathcal{A}$ such that $b \to^{*} a$, S defends a against b.

2.1 Semantics of extensions

Three different requirements for coherence were defined: "conflict-freeness", "safety", and "support closedness".

Definition 4 [Conflict-freeness, safety, support closedness] *A set of arguments* S *is* conflict-free *iff* $\nexists a \in S$ *such that* $S \to^{*} a$, *is* safe *iff* $\nexists a \in \mathcal{A}$ *such that* $S \to^{*} a \wedge (S \Rightarrow^{+} a \vee a \in S)$, *and is* closed for $\mathcal{R}_s$ *iff* $\nexists a \in \mathcal{A} \setminus S$ *such that* $S \Rightarrow a$.

When moving from AAFs to BAFs, the classical semantics of extensions can be adapted by imposing different combinations of coherence requirements. In particular, we consider the following classical semantics (denoted as "Dungean semantics"): admissible (`ad`), stable (`st`), complete (`co`), grounded (`gr`), preferred (`pr`), and start with reviewing the adaptions of the admissible one. A set $S \subseteq \mathcal{A}$ that defends all of its arguments is 1) a *d-admissible* (Dung-admissible) or `d-ad` extension iff S is conflict-free; 2) an *s-admissible* (Safe-admissible) or `s-ad` extension iff S is safe; 3) a *c-admissible* (Closure-admissible) or `c-ad` extension iff S is conflict-free and closed for $\mathcal{R}_s$. Given this, S is an extension of type:

- *d-complete* or `d-co` (resp., `s-co`, `c-co`) iff S is `d-ad` (resp., `s-ad`, `c-ad`) and contains every acceptable argument;
- *d-grounded* or `d-gr` (resp., `s-gr`, `c-gr`) iff S is a minimal (w.r.t. $\subseteq$) `d-co` (resp., `s-co`, `c-co`) extension;
- a *d-preferred* or `d-pr` (resp., `s-pr`, `c-pr`) iff S is a maximal (w.r.t. $\subseteq$) `d-co` (resp., `s-co`, `c-co`) extension.

Analogously, a set $S \subseteq \mathcal{A}$ is a *d-stable* or `d-st` (resp., `s-st`, `c-st`) extension iff $\forall a \in \mathcal{A} \setminus S$ $S \to^{*} a$ and S is conflict-free (resp., is safe; is conflict free and closed for $\mathcal{R}_s$).

Example 4 *In the BAF in Figure 1*(a), $\{a, b, f\}$ *is conflict-free and safe, but is not a* `d-ad` *extension for s-BAF or d-BAF (since* b *is not defended). Furthermore, for the s-BAF case,* $\{a, f\}$ *is a* `s-ad`, `s-gr` *and* `s-pr` *extension,* $\{a, e, f\}$ *is a* `st`, `d-ad`, `d-gr`, *and* `d-pr` *extension,* $\{e, f\}$ *is a* `c-pr` *and* `c-gr` *extension. For the d-BAF case* $\{e, f\}$ *is the unique* `st` *extension, that is also* `c-pr` *and* `c-gr`.

Remark 1 *Safety implies conflict-freeness, and these two requirements are equivalent in d-BAFs (but not in s-BAFs). So, in d-BAFs, the d- and s- variants of every semantics coincide. Furthermore, in d-BAFs, the three variants of the semantics* `st` *coincide, as the conflict-freeness and the requirement that an extension* S *attacks every argument outside* S *implies that* S *can support no argument outside itself.*

We denote as BVER$^{\sigma}(\mathcal{F}, S)$ the fundamental problem of verifying if the set S of arguments is an extension over the BAF $\mathcal{F}$ under the semantics σ. It is well known that the computational complexity of BVER under any `x`-variant (with `x` $\in \{$`s`, `d`, `c`$\}$) of a classical Dungean semantics σ is the same as the verification problem VER for AAFs under σ, as observed in [21]. The reason is that checking if S is an extension of a BAF $\mathcal{F}$ under `x`-σ can be done by first checking if S is a σ-extension of the AAF obtained from $\mathcal{F}$ by materializing all the implicit attacks and removing the supports, and then checking over $\mathcal{F}$ the coherence requirements of `x`.

3 Incomplete BAFs (iBAFs)

We extend *incomplete AAFs* (iAAFs) [4, 8, 7, 19, 20, 23] to *incomplete BAFs* (iBAFs).

Definition 5 (iBAF) *An* incomplete Bipolar Abstract Argumentation Framework *is a tuple* $\langle \mathcal{A}, \mathcal{A}^{?}, \mathcal{R}_a, \mathcal{R}_a^{?}, \mathcal{R}_s, \mathcal{R}_s^{?} \rangle$, *where* $\mathcal{A}$, $\mathcal{A}^{?}$ *are disjoint sets of arguments,* $\mathcal{R}_a$ *and* $\mathcal{R}_a^{?}$ *disjoint sets of attacks between arguments in* $\mathcal{A} \cup \mathcal{A}^{?}$, *and* $\mathcal{R}_s$ *and* $\mathcal{R}_s^{?}$ *disjoint sets of supports between arguments in* $\mathcal{A} \cup \mathcal{A}^{?}$.

The arguments in $\mathcal{A}$ are said to be *certain* (they are definitely known to exist), while those in $\mathcal{A}^{?}$ *uncertain* (it is not known for sure if they occur in the argumentation or not). Analogously, the attacks in $\mathcal{R}_a$ and the supports in $\mathcal{R}_s$ are said to be *certain* (they are known to occur, if both the incident arguments exist), while those in $\mathcal{R}_a^{?}$ and $\mathcal{R}_s^{?}$ *uncertain* (they may not occur, even if both the incident arguments exist). An iBAF represents alternative scenarios, called *completions*, corresponding to the different combinations of occurrence/non-occurrence of uncertain arguments/attacks/supports.

Definition 6 (Completion) *A completion for an iBAF* $\mathcal{IF} = \langle \mathcal{A}, \mathcal{A}^{?}, \mathcal{R}_a, \mathcal{R}_a^{?}, \mathcal{R}_s, \mathcal{R}_s^{?} \rangle$ *is a BAF* $\mathcal{F} = \langle \mathcal{A}', \mathcal{R}_a', \mathcal{R}_s' \rangle$ *where* $\mathcal{A} \subseteq \mathcal{A}' \subseteq (\mathcal{A} \cup \mathcal{A}^{?})$ *and* $\mathcal{R}_a \cap (\mathcal{A}' \times \mathcal{A}') \subseteq \mathcal{R}_a' \subseteq (\mathcal{R}_a \cup \mathcal{R}_a^{?}) \cap (\mathcal{A}' \times \mathcal{A}')$ *and* $\mathcal{R}_s \cap (\mathcal{A}' \times \mathcal{A}') \subseteq \mathcal{R}_s' \subseteq (\mathcal{R}_s \cup \mathcal{R}_s^{?}) \cap (\mathcal{A}' \times \mathcal{A}')$.

As happens for iAAFs, the notion of extension in AAFs can be adapted to iBAFs under a possible and a necessary perspective, where the condition imposed by the semantics is required to be true in *at least one* and *every* completion, respectively. This yields the definition of i*-extension below.

Definition 7 (i*-extension) *Given an iBAF* $\mathcal{IF}$ *and a semantics* σ, *a set* S *is a possible (resp., necessary) i*-extension for* $\mathcal{IF}$ *(under* σ*) if, for at least one (resp., for every) completion* F *of* $\mathcal{IF}$, S *is an extension of* F *under* σ.

In this paper, we focus on **possible i*-extensions**, and address IBVER$^{\sigma}(\mathcal{IF}, S)$, that is the adaption of the verification problem to iBAFs and asks if S is a possible i*-extension of $\mathcal{IF}$ under σ.

Example 5 *Consider the iBAF* $\mathcal{IF}$ *obtained from the BAF in Figure 1 imposing that the argument c, the attack* (b, c) *and the support* (c, d) *are uncertain. It is easy to see that, under both the abstract and deductive semantics of supports,* $\{a, c, e\}$ *is a possible i*-extension of* $\mathcal{IF}$ *under* `d-ad` *(as it is a* `d-ad` *extension in the completion where* (b, c) *and* (c, d) *are not present), but it is not a possible i*-extension under* `s-ad` *or* `c-ad` *(as, in every completion,* $e \to b$ *and* $a \Rightarrow b$*).*

Given an iBAF $\mathcal{IF} = \langle \mathcal{A}, \mathcal{A}^{?}, \mathcal{R}_a, \mathcal{R}_a^{?}, \mathcal{R}_s, \mathcal{R}_s^{?} \rangle$, every symbol $\diamond \in \{\to, \to^{*}, \Rightarrow, \Rightarrow^{+}\}$ introduced for BAFs will be used for $\mathcal{IF}$ with the semantics: $a \diamond b$ over $\mathcal{IF}$ means that $a \diamond b$ over the BAF $\langle \mathcal{A} \cup \mathcal{A}^?, \mathcal{R}_a \cup \mathcal{R}_a^?, \mathcal{R}_s \cup \mathcal{R}_s^? \rangle$. Moreover, we write $a \twoheadrightarrow^* b$ to say that $a \rightarrow^* b$ in the BAF $\mathcal{F}' = \langle \mathcal{A}', \mathcal{R}_a', \mathcal{R}_s' \rangle$, where $\mathcal{A}' = \mathcal{A} \cup \{a, b\}$, $\mathcal{R}_a' = \mathcal{R}_a \cap (\mathcal{A}' \times \mathcal{A}')$, $\mathcal{R}_s' = \mathcal{R}_s \cap (\mathcal{A}' \times \mathcal{A}')$, i.e. there is a generic attack from a to b even if every uncertain attack, support, and argument (different from both a and b) is removed from $\mathcal{IF}$. For instance, consider the d-iBAF obtained from the BAF in Figure 1(b) by adding the support (a, c). Then, $b \twoheadrightarrow^* a$, but not $a \twoheadrightarrow^* d$ (since if the uncertain attack (c, d) is removed, the supported attack (a, d) cannot be triggered).

Attacks and supports will be also considered from sets of arguments to arguments, and from arguments to sets of arguments: for any $\diamond \in \{\rightarrow, \stackrel{s}{\rightarrow}, \stackrel{m}{\rightarrow}, \rightarrow^*, \Rightarrow, \Rightarrow^+, \twoheadrightarrow^*\}$, $S \diamond a$ (resp., $a \diamond S$) means that there is $b \in S$ such that $b \diamond a$ (resp., $a \diamond b$). Moreover, we write $\mathcal{IF} \setminus X$, where $X \subseteq \mathcal{A}^?$, to denote the iBAF obtained from $\mathcal{IF}$ by removing all the arguments in X and the attacks and supports involving an argument in X.

4 Computational Complexity of Verifying i*-extensions

We here provide a thorough analysis of the complexity of reasoning over iBAFs (in terms of solving IBVER), where the sensitivity to the semantics of supports and the semantics of extensions is studied. Theorem 1 states a general upper bound for the complexity of IBVER.

Theorem 1 *For both s- and d- iBAFS.* $\text{IBVER}^\sigma(\mathcal{IF}, S)$ *is in NP under every semantics* $\texttt{x}$-σ*, with* $\texttt{x} \in \{\texttt{s}, \texttt{d}, \texttt{c}\}$ *and* $\sigma \in \{\texttt{ad}, \texttt{st}, \texttt{co}, \texttt{gr}\}$*, and in* Σ_2^p *under every semantics* $\texttt{x}$-$\texttt{pr}$*, with* $\texttt{x} \in \{\texttt{s}, \texttt{d}, \texttt{c}\}$.

Proof. $\text{IBVER}^\sigma(\mathcal{IF}, S)$ can be solved by guessing a completion $\mathcal{F}$ of $\mathcal{IF}$ and then solving $\text{BVER}^\sigma(\mathcal{F}, S)$. Thus, IBVER is in NP (resp., Σ_2^p) when BVER is in P (resp., coNP). □

Under the variants of the preferred semantics, it is easy to see that the upper bound of Theorem 1 is tight, so IBVER is Σ_2^p-complete. In fact, the Σ_2^p-hardness holds for iAAFs [8], that are iBAFs where supports are not used. We now focus on the variants of the other semantics, and show when the upper bound is tight or not. We consider s-and d- iBAFs separately, in order to highlight the role of the supports' semantics in the complexity.

4.1 *Reasoning over s-iBAFs*

Theorem 2 states under which semantics IBVER over s-iBAFs is NP-complete. This result (and all the NP-hardness results in this paper) is proved by showing reductions from the problem 3-SAT(ϕ) of deciding the satisfiability of a 3-CNF formula ϕ. To avoid repetitions, in all these reductions we assume that the formula ϕ has the form $\phi = C_1 \wedge \cdots \wedge C_m$, where each clause C_j is a disjunction of literals $l_j^1 \vee l_j^2 \vee l_j^3$, and each literal is a variable x_i or its negation $\neg x_i$ from the set of variables $X = \{x_1, \ldots, x_n\}$.

Theorem 2 *Over s-iBAFs,* IBVER^σ *is NP-complete under* $\sigma \in \{\texttt{s-ad}, \texttt{d-co}, \texttt{s-co}, \texttt{d-gr}, \texttt{s-gr}\}$.

Proof. Reducing 3-SAT(ϕ) to $\text{IBVER}^\sigma(\mathcal{IF}, S)$ under $\sigma = \texttt{s-ad}$. Let $\mathcal{IF} = \langle \mathcal{A}, \mathcal{A}^?, \mathcal{R}_a, \mathcal{R}_a^?, \mathcal{R}_s, \mathcal{R}_s^? \rangle$ be the iBAF where:
– $\mathcal{A}$ contains the arguments: 1) ϕ; 2) $\overline{C_j}$, for each $j \in [1..m]$; 3) x_i, $\neg x_i$, $\hat{x}_i$, $\neg \hat{x}_i$, for each $i \in [1..n]$;
– $\mathcal{R}_a$ contains: 1) $(\overline{C_j}, \phi)$ for each $j \in [1..m]$, 2) $(\hat{x}_i, \overline{C_j})$ (resp., $(\neg \hat{x}_i, \overline{C_j})$) whenever the literal x_i (resp., $\neg x_i$) occurs in C_j, 3) $(\neg \hat{x}_i, \hat{x}_i)$, for each $i \in [1..n]$;
– $\mathcal{R}_s^?$ contains $(x_i, \hat{x}_i)$, $(\neg x_i, \neg \hat{x}_i)$ for each $i \in [1..n]$;
– $\mathcal{A}^? = \mathcal{R}_a^? = \mathcal{R}_s = \emptyset$.
Figure 2 shows $\mathcal{IF}$ for $\phi = C_1 \wedge C_2 = (x_1 \vee x_2 \vee x_3) \wedge (\neg x_1 \vee \neg x_2 \vee \neg x_3)$. We prove the correctness of the reduction by proving "*ϕ is satisfiable*" $\Leftrightarrow$ "*$S = \{\phi, x_1, \neg x_1, \ldots, x_n, \neg x_n\}$ is an* $\texttt{s-ad}$ *extension of* $\mathcal{IF}$".
($\Rightarrow$): Let t be a satisfying truth assignment over X, and $\mathcal{F}$ the completion of $\mathcal{IF}$ containing, for each $i \in [1..n]$, the support $(x_i, \hat{x}_i)$ (resp., $(\neg x_i, \neg \hat{x}_i)$) iff $t(x_i) = \text{TRUE}$ (resp., $t(x_i) = \text{FALSE}$). As t makes ϕ true, each $\overline{C_j}$ is the target of at least one supported attack from some x_i or $\neg x_i$, hence ϕ is defended by S against the attacks from $\overline{C_1}, \ldots, \overline{C_m}$, thus S is an $\texttt{s-ad}$ extension of $\mathcal{F}$.
($\Leftarrow$): Let $\mathcal{F}$ be a completion of $\mathcal{IF}$ such that S is an $\texttt{s-ad}$ extension of $\mathcal{F}$. Consider the following relation t between X and $\{\text{TRUE}, \text{FALSE}\}$: $t = \{(x_i, \text{TRUE}) | (x_i, \hat{x}_i)$ is a support in $\mathcal{F}\} \cup \{(x_i, \text{FALSE}) | (\neg x_i, \neg \hat{x}_i)$ is a support in $\mathcal{F}\} \cup \{(x_i, \text{FALSE}) |$ neither $(x_i, \hat{x}_i)$ nor $(\neg x_i, \neg \hat{x}_i)$ are supports in $\mathcal{F}\}$. Since S is an $\texttt{s-ad}$ extension of $\mathcal{F}$, there is no $\hat{x}_i$ such that $x_i \Rightarrow \hat{x}_i$ and $\neg x_i \rightarrow^* \hat{x}_i$. Thus, there is no $i \in [1..n]$ such that $\mathcal{F}$ contains both the supports $(x_i, \hat{x}_i)$ and $(\neg x_i, \neg \hat{x}_i)$ (the latter implies $\neg x_i \stackrel{s}{\rightarrow} \hat{x}_i$). Hence, t is a truth assignment over X, as it assigns exactly one truth value to every x_i. In $\mathcal{F}$, since S is an $\texttt{s-ad}$ extension, ϕ is defended by S against the attacks from $\overline{C_1}, \ldots, \overline{C_m}$, thus every $\overline{C_j}$ must be the target of at least one supported attack from some x_i or $\neg x_i$. In turn, this means that, for every clause C_j, t makes at least one literal in C_j true, meaning that t makes ϕ satisfied.

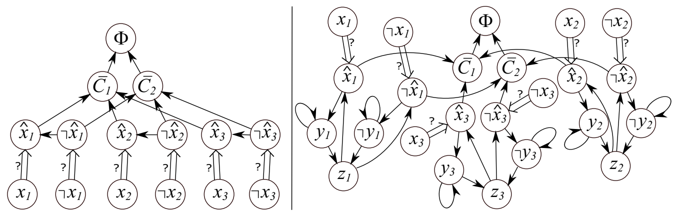

(a) Case $\texttt{s-ad}$ (b) Case $\texttt{s-co}$

Figure 2: $\phi = C_1 \wedge C_2 = (x_1 \vee x_2 \vee x_3) \wedge (\neg x_1 \vee \neg x_2 \vee \neg x_3)$ (supports are depicted as dashed lines)

Reducing 3-SAT(ϕ) to $\text{IBVER}^\sigma(\mathcal{IF}, S)$ under $\sigma \in \{\texttt{d-co}, \texttt{s-co}, \texttt{d-gr}, \texttt{s-gr}\}$. We focus on $\sigma = \texttt{s-co}$, as a similar reasoning over the same construction works in the other cases. Let $\mathcal{IF}$ be the iBAF $\langle \mathcal{A}, \mathcal{A}^?, \mathcal{R}_a, \mathcal{R}_a^?, \mathcal{R}_s, \mathcal{R}_s^? \rangle$ where:
– $\mathcal{A}$ contains the arguments: 1) ϕ; 2) $\overline{C_j}$, for each $j \in [1..m]$; 3) $x_i, \neg x_i, \hat{x}_i, \neg \hat{x}_i, y_i, \neg y_i, z_i$ for each $i \in [1..n]$;
– $\mathcal{R}_a$ contains: 1) $(\overline{C_j}, \phi)$ for each $j \in [1..m]$, 2) $(\hat{x}_i, \overline{C_j})$ (resp., $(\neg \hat{x}_i, \overline{C_j})$) whenever the literal x_i (resp., $\neg x_i$) occurs in C_j, 3) $(z_i, \hat{x}_i)$, $(z_i, \neg \hat{x}_i)$, (y_i, y_i), $(\neg y_i, \neg y_i)$, $(\hat{x}_i, y_i)$, $(\neg \hat{x}_i, \neg y_i)$, (y_i, z_i), $(\neg y_i, z_i)$, for each $i \in [1..n]$;
– $\mathcal{R}_s^?$ contains $(x_i, \hat{x}_i)$, $(\neg x_i, \neg \hat{x}_i)$ for each $i \in [1..n]$;
– $\mathcal{A}^? = \mathcal{R}_a^? = \mathcal{R}_s = \emptyset$.
We prove the correctness of the reduction by proving the equivalence: "ϕ is satisfiable" $\Leftrightarrow$ "$S = \{\phi, x_1, \neg x_1, \ldots, x_n, \neg x_n\}$ is a $\texttt{s-co}$ extension of $\mathcal{IF}$".
($\Rightarrow$): Let t be a satisfying truth assignment over $x_1, \ldots, x_n$, and $\mathcal{F}$ the completion of $\mathcal{IF}$ containing, for each $i \in [1..n]$ the support $(x_i, \hat{x}_i)$ (resp., $(\neg x_i, \neg \hat{x}_i)$) iff $t(x_i) = \text{TRUE}$ (resp., $t(x_i) = \text{FALSE}$). It is easy to see that S is an $\texttt{s-ad}$ extension of $\mathcal{F}$, since it is conflict-free, it attacks no argument supported by

itself (that is, no $\hat{x}_i$ or $\neg\hat{x}_i$) and defends its arguments, since the fact that t makes ϕ true implies that each $\overline{C_j}$ is the target of at least one supported attack from some x_i or $\neg x_i$, thus ϕ is defended by S against the attacks from $\overline{C_1}, \dots, \overline{C_m}$. Moreover, there is no argument outside S acceptable w.r.t. S, since: 1) every $\overline{C_j}$ is attacked by S; 2) every $\hat{x}_i$ and $\neg\hat{x}_i$ is not defended by S against the attack from z_i; 3) every $\hat{y}_i$ and every $\neg\hat{y}_i$ attacks itself; 4) every z_i is defended by S against only one of the attacks from y_i and $\neg y_i$, since $S \rightarrow^{s} y_i$ iff $t(x_i) = \text{TRUE}$, and $S \rightarrow^{s} \neg y_i$ iff $t(x_i) = \text{FALSE}$. Hence, S is an s-co extension of $\mathcal{F}$.

($\Leftarrow$): Let $\mathcal{F}$ be a completion of $\mathcal{IF}$ such that S is an s-co extension of $\mathcal{F}$. Consider the following relation t between X and $\{\text{TRUE}, \text{FALSE}\}$: $t = \{(x_i, \text{TRUE}) | (x_i, \hat{x}_i) \text{is a support in } \mathcal{F}\} \cup \{(x_i, \text{FALSE}) | (\neg x_i, \neg\hat{x}_i)$ is a support in $\mathcal{F}\} \cup \{(x_i, \text{FALSE})|$ neither $(x_i, \hat{x}_i)$ nor $(\neg x_i, \neg\hat{x}_i)$ are supports in $\mathcal{F}\}$.

Since S is an s-co extension of $\mathcal{F}$, no argument z_i is acceptable w.r.t. S, meaning that for every $i \in [1..n]$ at least one of the supports $(x_i, \hat{x}_i)$, $(\neg x_i, \neg\hat{x}_i)$ is not in $\mathcal{F}$ (otherwise, $\mathcal{F}$ would contain both the supported attacks (x_i, y_i), $(\neg x_i, \neg y_i)$, that defend z_i against the attacks from y_i and $\neg y_i$). Hence, t is a truth assignment over X, as it assigns exactly one truth value to every x_i. Since S is an s-co extension, ϕ is defended by S against the attacks from $\overline{C_1}, \dots, \overline{C_m}$, thus every $\overline{C_j}$ must be the target of at least one supported attack from some x_i or $\neg x_i$. In turn, this means that t makes true at least one literal in every C_j, meaning that t makes ϕ satisfied. □

We now address the semantics not covered above, and show that under them IBVER is in P. Preliminarily, we introduce some properties that may be satisfied or not by a given s-iBAF, and a lemma, that will help the reasoning in the main theorem's proof.

Definition 8 *Properties $p1, p2, p3$ are satisfied by an s-iBAF $\mathcal{IF} = \langle \mathcal{A}, \mathcal{A}^?, \mathcal{R}_a, \mathcal{R}_a^?, \mathcal{R}_s, \mathcal{R}_s^? \rangle$ and a set $S \subseteq \mathcal{A} \cup \mathcal{A}^?$ iff:*
– $p1$: $\mathcal{R}_a^? \cap \left(\left(\mathcal{A} \cup \mathcal{A}^? \right) \times S \right) = \emptyset$, i.e. there are no uncertain attacks towards S;
– $p2$: $\nexists a \in \mathcal{A}^? \setminus S$ s.t. $a \twoheadrightarrow^* S$, i.e. there are no uncertain arguments that will (generically) attack S even if every other uncertain argument/attack/support is removed;
– $p3$: $\nexists (a, y) \in \mathcal{R}_s^?$ s.t. $y \in \mathcal{A} \wedge y \twoheadrightarrow^* S$, i.e. there are no uncertain supports towards certain arguments that will (generically) attack S even if every other uncertain argument/attack/support is removed.

Lemma 1 *Let $\mathcal{IF} = \langle \mathcal{A}, \mathcal{A}^?, \mathcal{R}_a, \mathcal{R}_a^?, \mathcal{R}_s, \mathcal{R}_s^? \rangle$ be an s-iBAF and $S \subseteq \mathcal{A} \cup \mathcal{A}^?$. Then:*
I) under $\sigma \in \{$d-ad, c-ad, d-st, s-st, c-st$\}$, S is a possible i*-extension of $\mathcal{IF}$ iff it is a possible i*-extension of $\mathcal{IF}' = \langle \mathcal{A}, \mathcal{A}^?, \mathcal{R}_a, \mathcal{R}_a^? \setminus \left(\left(\mathcal{A} \cup \mathcal{A}^? \right) \times S \right), \mathcal{R}_s, \mathcal{R}_s^? \rangle$;
II) under $\sigma \in \{$d-ad, d-st$\}$, if $p1$ is satisfied, then, $\forall a \in \mathcal{A}^? \setminus S$ s.t. $a \twoheadrightarrow^* S$, S is a possible i*-extension of $\mathcal{IF}$ iff S is a possible i*-extension of $\mathcal{IF}' = \mathcal{IF} \setminus \{a\}$;
III) under $\sigma \in \{$d-ad, d-st$\}$, if $p1$, $p2$ are satisfied, then $\forall (a, y) \in \mathcal{R}_s^?$ s.t. $y \in A \wedge y \twoheadrightarrow^* S$, S is a possible i*-extension of $\mathcal{IF}$ iff S is a possible i*-extension of $\mathcal{IF}' = \mathcal{IF} \setminus \{(a, y)\}$.

The theorem below states under which semantics IBVER is in P. Its proof (for the case $\sigma =$ d-ad) relies on the correctness of Algorithm 1, that first enforces properties $p1, p2, p3$ by removing the uncertain attacks/arguments/supports making them violated, and then returns as answer the result of the verification performed over the completion containing all the uncertain arguments/attacks/supports that were not removed. The other cases can be solved via minor changes to the strategy of Algorithm 1.

As for c-co and c-gr, the general NP upper bound of Theorem 1 holds and a tighter characterization is left to future work.

Theorem 3 *Over s-iBAFs, IBVER$^\sigma$ is in P under $\sigma \in \{$d-ad, c-ad, d-st, s-st, c-st$\}$.*

Proof. ($\sigma =$ d-ad, the other cases are analogous). We first prove that Algorithm 1 is correct (i.e. it returns true iff S is a possible i*-extension for $\mathcal{IF}$ under $\sigma =$d-ad). Lemma 1 implies that, under $\sigma =$ d-ad, S is a possible i*-extension for $\mathcal{IF}$ iff S is a possible i*-extension for the iBAF $\mathcal{IF}'$ constructed by Lines 1- 5. Thus, to prove the correctness of Algorithm 1, it suffices to prove the equivalence *EQ*: "*S is a* d-ad *extension of $\mathcal{F}'$*"$\Leftrightarrow$ "*S is a possible i*-extension for $\mathcal{IF}'$ under* d-ad", where $\mathcal{F}'$ is the BAF built at Line 6 and $\mathcal{IF}'$ the iBAF buillt by Lines 1- 5. The $\Rightarrow$ direction is straightforward. We prove the $\Leftarrow$ direction reasoning by contradiction. Assume that S is a possible i*-extension for $\mathcal{IF}'$ under d-ad and S is not a d-ad extension of $\mathcal{F}'$. Let $\mathcal{F}$ be a completion of $\mathcal{IF}'$ such that S is a d-ad extension of $\mathcal{F}$. Since S is not a d-ad extension of $\mathcal{F}'$, (at least) one of the following holds:

i) $\exists a, b \in S$ such that $a \rightarrow^* b$,
ii) $\exists a \in S$ and $\exists b \in A' \setminus S$ such that $b \rightarrow^* a$ and S does not defend a against b.

We show that *i*) and *ii*) yield contradictions. If *i*) holds, $a \rightarrow^* b$ also over $\mathcal{F}$, as all the supports/attacks and arguments that imply the presence of the attack (a, b) over $\mathcal{F}'$ are certain in $\mathcal{IF}'$, as implied by $p1$, $p2$, $p3$. Hence, S is not a d-ad extension of $\mathcal{F}$, as it is not conflict free, which is a contradiction. If *ii*) holds, properties $p1, p2$ and $p3$ imply that $b \rightarrow^* a$ over $\mathcal{F}$. Hence, since S is a d-ad extension of $\mathcal{F}$, $S \rightarrow b$ over $\mathcal{F}$, and this, in turn, implies that $S \rightarrow b$ over $\mathcal{F}'$, as every uncertain attack in $\mathcal{IF}'$ from S towards arguments not in S is present in $\mathcal{F}'$. Therefore, S defends a against b, which is a contradiction. This completes the proof of *EQ* and of the correctness of Algorithm 1. Then, the statement follows from the fact that that Algorithm 1 runs in polynomial time. □

4.2 Reasoning over d-iBAFs

Theorem 4 *Over d-iBAFs, IBVER$^\sigma$ is NP-complete under $\sigma \in \{$d-ad, s-ad, d-co, s-co, c-co, d-gr, s-gr, c-gr$\}$.*

Proof. Reducing 3-SAT(ϕ) to IBVER$^\sigma(\mathcal{IF}, S)$ under $\sigma \in \{$ d-ad, s-ad $\}$. Since the two semantics d-ad, s-ad coincide in d-iBAFs (see Remark 1), we consider $\sigma =$ s-ad only. Let

Algorithm 1 Solving IBVER$^\sigma(\mathcal{IF}, S)$ under $\sigma =$ d-ad

Input: $\mathcal{IF} = \langle \mathcal{A}, \mathcal{A}^?, \mathcal{R}_a, \mathcal{R}_a^?, \mathcal{R}_s, \mathcal{R}_s^? \rangle$; $S \subseteq A \cup A^?$
Output: TRUE iff S is a possible i*-extension for $\mathcal{IF}$ under d-ad;

Initialize $\mathcal{IF}'$ as the maximal "sub-iBAF" of $\mathcal{IF}$ that satisfies $p1$
1: $\mathcal{IF}' = \langle \mathcal{A}', \mathcal{A}'^?, \mathcal{R}_a', \mathcal{R}_a'^?, \mathcal{R}_s', \mathcal{R}_s'^? \rangle$, where $\mathcal{A}' = \mathcal{A}$; $\mathcal{A}'^? = \mathcal{A}^?$; $\mathcal{R}_a' = \mathcal{R}_a$; $\mathcal{R}_s' = \mathcal{R}_s$; $\mathcal{R}_s'^? = \mathcal{R}_s^?$; $\mathcal{R}_a'^? = \mathcal{R}_a^? \setminus \left(\left(A \cup A^? \right) \times S \right)$
Enforce $p2$
2: **while** $\exists a \in \mathcal{A}'^? \setminus S$ s.t. $\exists s \in S$ where $a \twoheadrightarrow^* s$ in $\mathcal{IF}'$ **do**
3: $\quad \mathcal{IF}' = \mathcal{IF}' \setminus \{a\}$
Enforce $p3$
4: **while** $\exists (a, y) \in \mathcal{R}_s'^?$ s.t. $y \in \mathcal{A}' \wedge y \twoheadrightarrow^* S$ in $\mathcal{IF}'$ **do**
5: $\quad \mathcal{R}_s'^? = \mathcal{R}_s'^? \setminus \{(a, y)\}$
6: $\mathcal{F}' = \langle \mathcal{A}' \cup \mathcal{A}'^?, \mathcal{R}_a' \cup \mathcal{R}_a'^?, \mathcal{R}_s' \cup \mathcal{R}_s'^? \rangle$
7: **return** TRUE if S is a d-ad extension of $\mathcal{F}'$; FALSE otherwise

Algorithm 2 **Verifying c-ad extensions over d-iBAFs**

Input: $\mathcal{IF} = \langle \mathcal{A}, \mathcal{A}^?, \mathcal{R}_a, \mathcal{R}_a^?, \mathcal{R}_s, \mathcal{R}_s^? \rangle$; $S \subseteq \mathcal{A} \cup \mathcal{A}^?$;
Output: TRUE iff S is a possible i*-extension for $\mathcal{IF}$ under c-ad;
1: $\mathcal{F}' = \langle \mathcal{A}', \mathcal{R}_a', \mathcal{R}_s' \rangle$: where
$\quad \mathcal{A}' = \mathcal{A} \cup \mathcal{A}^? \setminus \{a \in \mathcal{A}^? \mid \exists x \in S \text{ s.t. } (x, a) \in \mathcal{R}_s\}$
$\quad \mathcal{R}_a' = \mathcal{R}_a \cup (\mathcal{R}_a^? \setminus (S \times S))$
$\quad \mathcal{R}_s' = \mathcal{R}_s \cup (\mathcal{R}_s^? \setminus (S \times (\mathcal{A}' \setminus S)))$
2: **if** S is not conflict free and closed for $\mathcal{R}_s'$ in $\mathcal{F}'$ **then**
3: $\quad$ **return** FALSE
4: $\mathcal{R}_a' = \mathcal{R}_a' \setminus \{(x, s) \in \mathcal{R}_a^? \mid s \in S \wedge \text{ not } S \rightarrow^* y\}$
5: **while** $\exists (y, s) \in \mathcal{R}_a'$ s.t. $s \in S$ and not $S \rightarrow^* y$ **do**
6: $\quad$ **if** $y \in \mathcal{A}^?$ **then** $\mathcal{F}' = \mathcal{F}' \setminus \{y\}$ **else return** FALSE
7: **return** TRUE

$\mathcal{IF} = \langle \mathcal{A}, \mathcal{A}^?, \mathcal{R}_a, \mathcal{R}_a^?, \mathcal{R}_s, \mathcal{R}_s^? \rangle$ be the iBAF where:
– $\mathcal{A}^? = \emptyset$, while $\mathcal{A}$ contains the arguments: 1) ϕ; 2) $\overline{C_j}$, for each $j \in [1..m]$; 3) $x_i, \neg x_i, \hat{x}_i$, for each $i \in [1..n]$;
– $\mathcal{R}_a$ contains: 1) $(\overline{C_j}, \phi)$ for each $j \in [1..m]$, 2) $(\hat{x}_i, \overline{C_j})$ whenever the literal $\neg x_i$ occurs in C_j;
– $\mathcal{R}_s$ contains $(\overline{C_j}, \hat{x}_i)$ whenever the literal x_i occurs in C_j;
– $\mathcal{R}_a^?$ contains $(x_i, \hat{x}_i)$, for each $i \in [1..n]$;
– $\mathcal{R}_s^?$ contains $(\neg x_i, \hat{x}_i)$, for each $i \in [1..n]$.
Similarly to the previous NP-hardness results, the correctness of the reduction is due to the equivalence: "*ϕ is satisfiable*" $\Leftrightarrow$ "*$S = \{\phi, x_1, \neg x_1, \ldots, x_n, \neg x_n\}$ is a* s-ad *extension of $\mathcal{IF}$*". The proof for $\sigma \in \{$d-co, s-co, c-co, d-gr, s-gr, c-gr$\}$ uses an analogous reasoning. ☐

As for the other semantics, Theorem 5 below states that IBVER is in P for the semantics not considered in Theorem 4 (under which IBVER is NP-complete). For the case $\sigma =$ c-ad, the core of the tractability result of Theorem 5 is Algorithm 2. This algorithm first computes the completion $\mathcal{F}$ containing all the certain and uncertain arguments/attacks/supports except: 1) the uncertain attacks between arguments in S, 2) the uncertain arguments that are supported (via certain supports) by S, and 3) the uncertain supports from S towards arguments outside S. After checking in $\mathcal{F}$ if S is conflict free and closed for supports (otherwise, it returns FALSE), it considers every undefended attack (a, s) towards S and tries to remove it from $\mathcal{F}$: it accomplishes this in the two cases where (a, s) was uncertain in $\mathcal{IF}$ (in this case, (a, s) is directly removed from $\mathcal{F}$) or a was uncertain in $\mathcal{IF}$ (in this case, (a, s) is removed as the side effect of removing a). Finally, it returns TRUE if and only if the so obtained $\mathcal{F}$ contains no more undefended attacks towards S. Basically, the proof of Theorem 5 consists in showing that this strategy is correct: as $\mathcal{F}$ is obviously a completion for which S is a c-ad extension, the hard part of the proof consists in showing that, if the algorithm returns FALSE, there is no completion (different from $\mathcal{F}$) for which S is a c-ad extension.

Theorem 5 *Over d-iBAFs,* IBVER$^\sigma$ *is in P under* $\sigma \in \{$c-ad, d-st, s-st, c-st$\}$.

Proof. Case $\sigma =$ c-ad (the other cases can be proved with similar reasoning). We first prove that Algorithm 2 is correct (it returns TRUE iff S is a possible c-ad i*-extension of $\mathcal{IF}$). The only if direction is straightforward as Algorithm 2 returns TRUE iff S is a c-ad extension of $\mathcal{F}'$, which is a completion of $\mathcal{IF}$. We prove the if direction reasoning by contradiction. Assume that S is a possible c-ad extension of $\mathcal{IF}$, but Algorithm 2 returns FALSE. If FALSE is returned at line 3, at least one of the following cases holds:

$c1: \ \exists (x, y) \in \mathcal{R}_a \cap (S \times S)$;

$c2: \ \exists (x, y) \in \mathcal{R}_s \cap (S \times (\mathcal{A} \setminus S))$.

As in both cases there is no completion of $\mathcal{IF}$ where S is conflict free and closed for supports, S is not a possible c-ad extension of $\mathcal{IF}$, which is a contradiction. This means that FALSE is returned at line 6. We use the following claim.

Claim 1 *Assume that Algorithm 2 returns* FALSE *at some iteration of the* **while** *loop (lines 5-6), and let j be the iteration at which* FALSE *is returned. Let $i < j$ and $\mathcal{F}^i = \langle \mathcal{A}^i, \mathcal{R}_a^i, \mathcal{R}_s^i \rangle$ be the completion generated at the end of the i-th step of the while loop. For any completion $\mathcal{F}'' = \langle \mathcal{A}'', \mathcal{R}_a'', \mathcal{R}_s'' \rangle$ of $\mathcal{IF}$, if $\mathcal{A}'' \setminus \mathcal{A}^i \neq \emptyset$ or $\mathcal{R}_a'' \setminus \mathcal{R}_a^i \neq \emptyset$ or $\mathcal{R}_s'' \setminus \mathcal{R}_s^i \neq \emptyset$, then S is not a* c-ad *extension of $\mathcal{F}''$.*

Let n be the **while** step before the one at which Algorithm 2 returns FALSE. Now, for any completion $\mathcal{F}''$ of $\mathcal{IF}$ such that $\mathcal{A}'' \setminus \mathcal{A}^n \neq \emptyset$ or $\mathcal{R}_a'' \setminus \mathcal{R}_a^n \neq \emptyset$ or $\mathcal{R}_s'' \setminus \mathcal{R}_s^n \neq \emptyset$, Claim 1 implies that S is not a c-ad extension of $\mathcal{F}''$. Moreover, for any completion $\mathcal{F}''$ of $\mathcal{IF}$ such that $\mathcal{A}'' \setminus \mathcal{A}^n = \emptyset$ and $\mathcal{R}_a'' \setminus \mathcal{R}_a^n = \emptyset$ and $\mathcal{R}_s'' \setminus \mathcal{R}_s^n = \emptyset$, it holds that $\exists (y, x) \in \mathcal{R}_a''$ s.t. $x \in S$ and not $S \rightarrow^* y$ and $(y, x) \notin \mathcal{R}_a^?$ and $y \notin \mathcal{A}^?$, and then S is not a c-ad extension of $\mathcal{F}''$ either, as y attacks S but is not attacked by S in $\mathcal{F}''$. Thus, every completion $\mathcal{F}''$ of $\mathcal{IF}$ is such that S is not a c-ad extension of $\mathcal{F}''$, which is a contradiction. This completes the proof of the correctness of Algorithm 2. The fact that Algorithm 2 runs in polynomial time concludes the proof for $\sigma =$ c-ad. The proof for $\sigma \in \{$d-st, s-st, c-st$\}$ is analogous. ☐

5 The Acceptance Problem

We start from the classical definitions of *credulous acceptance* (CA) and *skeptical acceptance* (SA) over AAFs:
"CA$^\sigma(a, F)$ *(resp.,* SA$^\sigma(a, F)$*) is the problem of deciding if the argument a is in at least one (resp., every) extension of the AAF F*"

Then, we consider the adaption of the acceptance problem introduced in [7] for iAAFs (where supports are not considered). That is, we consider the four variants of the acceptance problem PCA,PSA,NCA,NSA, where the credulous (C) and skeptical (S) perspectives implied by the presence of multiple extensions are combined with the possible (P) and necessary (N) perspectives implied by the presence of multiple completions:
"PXA$^\sigma(a, IF)$ *(resp.,* NXA$^\sigma(a, IF)$*), where $X \in \{$*C,S*$\}$, is the problem of deciding if in at least one (resp., every) completion F of IF the answer of* XA(a, F, σ) *is yes.*"

Now, we naturally extend the acceptance problem to the case of iBAFs by using the same definition above, with the only difference that now IF is an iBAF.

In order to provide a complexity characterization of the acceptance problems over iBAFs, we start from what is known about the same problems over iAAFs. A complete picture of the complexity of PCA, NCA, PSA, NSA over iAAFs is given by Table 2, taken from [7]. From Table 2, it turns out that, except for the trivial case of PSA and NSA under σ =ad, under the other Dungean semantics PCA, NCA, PSA, NSA over iAAFs are complete for complexity classes above P in the polynomial hierarchy. An immediate consequence is that, except for the above mentioned trivial cases, PCA, NCA, PSA, NSA over (s- and d-) iBAFs are hard for the same classes as over iAAFs, since iAAFs are iBAFs without supports. Moreover, the guess-and-check strategies used in [7] to prove the memberships work also in the case of s- and d- iBAFs, after a minor change to the check phase. In fact, in the case of iBAFs, the guessed completion is a BAF (while in the case of iAAFs, the guessed completion is an AAF). Hence, the check phase over iBAFs simply consists in performing the same

Semantics σ	PCA	NCA	PSA	NSA
ad	NP-c	Π_2^p-c	trivial	trivial
st	NP-c	Π_2^p-c	Σ_2^p-c	coNP-c
gr	NP-c	coNP-c	NP-c	coNP-c
co	NP-c	Π_2^p-c	NP-c	coNP-c
pr	NP-c	Π_2^p-c	Σ_3^p-c	Π_2^p-c

Table 2: Computational complexity of the (variants of) the acceptance problem over iAAFs and iBAFs. As for iBAFs, the results hold for s-iBAFs and d-iBAFs and for all the coherence conditions (d-, s-, c-).

check phase done over iAAFs after two preliminary steps: 1) translating the guessed completion/BAF into an equivalent AAF where the implicit attacks are materialized; 2) checking the coherence condition (conflict-freeness, safety, or support closedness). Since 1) and 2) can be done in polynomial time, the check phase over iBAFs has the same complexity as over iAAFs, thus also the memberships reported in the table above hold for iBAFs.

What said above, along with the fact that, under σ = x-ad (with x $\in$ {d, s, c}), it is straightforward to see that PSA and NSA over iBAFs are trivial (as over iAAFs), proves the following statement.

Theorem 6 *The computational complexities of* PCA, NSA, PSA *and* NSA *over d-iBAFs and s-iBAFs and for all the coherence conditions (d-, s-, c-) are reported in Table 2.*

6 Related Work

BAFs (along with the abstract semantics of supports) were first introduced in [12], and then revisited in [9, 35, 36], where the *deductive*, *necessary*, *evidential* semantics (that are comprehensively reviewed in [13]) were introduced, respectively. We observe that the results in this paper hold also under the necessary semantics, that is dual to the deductive one (as necessary supports are deductive supports in the opposite direction). More recently, alternative supports' semantics have been proposed in [38, 39], where supporting arguments are capable of defending the supported arguments. Analogously, the semantics in [27] pursues the idea of enforcing some monotonicity, in the sense that supports are prevented from decreasing the acceptance degree of arguments (from skeptically to credulously accepted, or even rejected). In fact, the various semantics in the literature catch different intuitions on the meaning of supports, so there can hardly be consensus on which is the most natural one. In this context, this paper has focused on two traditional and well-established semantics of supports, and proposes a line of research that is worth investigating under the other semantics in the literature.

Further related works are those where correlations similar to supports have been investigated, such as the *subarguments* in [33] (which are closely related with necessary supports [15]), the dependencies in [10, 43], the pro arguments in [29], as well as the acceptance conditions of *Abstract Dialectical Frameworks* [10], that can express different forms of supports.

As for iAAFs, they were introduced in [8], and their semantics based on i*-extensions in [20]. Previously, *Partial Argumentation Frameworks* (PAFs) had been introduced in [11] to encode incompleteness affecting attacks. The possibility of specifying correlations over iAAFs between uncertain arguments/attacks was studied in [19, 23, 31], where correlations are expressed in terms of constraints restricting the set of completions to be considered in the reasoning (thus, in spirit, these constraints are different from constraints and preferences studied in the deterministic setting [5, 2, 1]). A recent survey discussing how iAAFs are related to the forms of incompleteness encoded in other variants of AAFs (such as *Control Argumentation Frameworks* [34]) can be found in [32]. The reasoning over AAFs in the presence of incompleteness is also related to revising AAFs to enforce the existence of an extension [6], or to make a set an extension [16], even when information on the agents who claimed the arguments is available [26]. In this regard, this paper suggests that the enforcement problem over BAFs is worth investigating, since it may raise issues that are not present over AAFs (as the enforcement over BAFs would require the insertion/removal of supports to be considered as a new primitive, that comprises the insertion/removal of groups of attacks).

iBAFs are also related to *probabilistic BAFs* (prBAFs) [21], that merge BAFs with *prAAFs* (i.e. *probabilistic AAFs* adopting the constellations approach [3, 17, 22, 24, 28, 30]): basically, iBAFs can be viewed as prBAFs where all the uncertain terms of the dispute are independent one from another, and no measure of this uncertainty is given. However, the results in this work are not subsumed by those in [21], where the problem P-EXT of computing the probability of extensions (that is the probabilistic counterpart of the verification) has been shown to be highly intractable under independence (i.e. FP$^{\#P}$-complete) for every combination of semantics of supports and extensions. Indeed, the results obtained for prBAFs neither subsume those in this paper nor make them less surprising: on the one hand, the role of assembling supports and probabilities in the high complexity of P-EXT over prBAFs is blurred by the fact that P-EXT is already FP$^{\#P}$-hard under most of Dungean semantics over prAAFs (without supports). On the other hand, the FP$^{\#P}$-hardness of P-EXT gives no hint for a tight complexity characterization of IBVER, as the source of complexity related to the counting mechanism underlying probability evaluation is absent in the verification problem (in fact, the literature contains several functional problems complete for "hard" complexity classes, whose decision counterparts are in P).

Finally, the framework in [42] is related to iBAFs since it mixes probabilities, supports (in the form of subarguments [41]) and a form of incompleteness: it allows labelings where some arguments are marked as "OFF" (thus, excluded from the reasoning), but the semantics of "OFF" differs from iBAFs' "non-occurrence", since the "OFF" label is propagated to the supporting arguments (and this does not happen with non-occurrence).

7 Conclusions and Future Work

We have augmented Abstract Argumentation Frameworks with the possibility of simultaneously specifying supports between arguments (according to the traditional abstract and deductive supports' semantics) and the presence of uncertain elements of the argumentation graph. We have studied the computational complexity of fundamental reasoning problems, and obtained a surprising result regarding the verification problem under the possible semantics: we have shown that, although bipolarity and incompleteness do not affect the complexity of this problem under the Dungean semantics (except for the preferred semantics) if considered separately, they may have a deep impact on the tractability/intractability of this problem when jointly used. Future work will focus on: 1) the two cases (under the variants c-co and c-gr of the complete and grounded semantics) for which the complexity characterization in this paper was not tight, and 2) extending the study to the necessary perspective of the verification problem, and 3) extending iBAFs with the possibility of specifying correlations between uncertain arguments/attacks/supports, as done in [19, 23] for arguments and attacks only.

References

[1] G. Alfano, S. Greco, F. Parisi, and I. Trubitsyna, 'Abstract argumentation framework with conditional preferences', in *Proc. of AAAI*, pp. 6218–6227, (2023).

[2] G. Alfano, S. Greco, F. Parisi, and I. Trubitsyna. Preferences and constraints in abstract argumentation. *Preprint of IJCAI*, 2023.

[3] Gianvincenzo Alfano, Marco Calautti, Sergio Greco, Francesco Parisi, and Irina Trubitsyna, 'Explainable acceptance in probabilistic and incomplete abstract argumentation frameworks', *Artificial Intelligence*, **323**, 103967, (2023).

[4] Gianvincenzo Alfano, Sergio Greco, Francesco Parisi, and Irina Trubitsyna, 'Incomplete argumentation frameworks: Properties and complexity', in *Proc. of AAAI*, pp. 5451–5460, (2022).

[5] Gianvincenzo Alfano, Sergio Greco, Francesco Parisi, and Irina Trubitsyna, 'On preferences and priority rules in abstract argumentation', in *Proc. of IJCAI*, pp. 2517–2524, (2022).

[6] Ringo Baumann and Markus Ulbricht, 'If nothing is accepted - repairing argumentation frameworks', *J. Artif. Intell. Res.*, **66**, 1099–1145, (2019).

[7] Dorothea Baumeister, Matti Järvisalo, Daniel Neugebauer, Andreas Niskanen, and Jörg Rothe, 'Acceptance in incomplete argumentation frameworks', *Artif. Intell.*, **295**, 103470, (2021).

[8] Dorothea Baumeister, Daniel Neugebauer, Jörg Rothe, and Hilmar Schadrack, 'Verification in incomplete argumentation frameworks', *Artif. Intell.*, **264**, 1–26, (2018).

[9] Guido Boella, Dov M. Gabbay, Leendert W. N. van der Torre, and Serena Villata, 'Support in abstract argumentation', in *Proc. of COMMA 2010*, pp. 111–122, (2010).

[10] Gerhard Brewka and Stefan Woltran, 'Abstract dialectical frameworks', in *Proc. of KR*, pp. 102–111, (2010).

[11] Claudette Cayrol, Caroline Devred, and Marie-Christine Lagasquie-Schiex, 'Handling ignorance in argumentation: Semantics of partial argumentation frameworks', in *Proc. Conf. Symbolic and Quantitative Approaches to Reasoning with Uncertainty (ECSQARU), Hammamet, Tunisia*, pp. 259–270, (2007).

[12] Claudette Cayrol and Marie-Christine Lagasquie-Schiex, 'On the acceptability of arguments in bipolar argumentation frameworks', in *Proc. of ECSQARU*, pp. 378–389, (2005).

[13] Claudette Cayrol and Marie-Christine Lagasquie-Schiex, 'Bipolarity in argumentation graphs: Towards a better understanding', *Int. J. Approx. Reasoning*, **54**(7), 876–899, (2013).

[14] Claudette Cayrol and Marie-Cristine Lagasquie-Schiex, 'Bipolar abstract argumentation systems', in *Argum. in Artif. Intell.*, 65–84, (2009).

[15] Andrea Cohen, Sebastian Gottifredi, Alejandro Javier García, and Guillermo Ricardo Simari, 'A survey of different approaches to support in argumentation systems', *Know. Eng. Review*, **29**(5), 513–550, (2014).

[16] Sylvie Coste-Marquis, Sébastien Konieczny, Jean-Guy Mailly, and Pierre Marquis, 'Extension enforcement in abstract argumentation as an optimization problem', in *Proc. Int. Conf. on Artificial Intelligence (IJCAI)*, pp. 2876—-2882, (2015).

[17] Dragan Doder and Stefan Woltran, 'Probabilistic argumentation frameworks - A logical approach', in *Proc. of SUM*, pp. 134–147, (2014).

[18] Phan Minh Dung, 'On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games', *Artif. Intell.*, **77**(2), 321–358, (1995).

[19] Bettina Fazzinga, Sergio Flesca, and Filippo Furfaro, 'Reasoning over argument-incomplete aafs in the presence of correlations', in *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI 2021, Virtual Event / Montreal, Canada, 19-27 August 2021*, ed., Zhi-Hua Zhou, pp. 189–195.

[20] Bettina Fazzinga, Sergio Flesca, and Filippo Furfaro, 'Revisiting the notion of extension over incomplete abstract argumentation frameworks', in *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI 2020*, ed., Christian Bessiere, pp. 1712–1718.

[21] Bettina Fazzinga, Sergio Flesca, and Filippo Furfaro, 'Probabilistic bipolar abstract argumentation frameworks: complexity results', in *Proc. Int. Joint Conf. on Artificial Intelligence (IJCAI)*, pp. 1803–1809, (2018).

[22] Bettina Fazzinga, Sergio Flesca, and Filippo Furfaro, 'Complexity of fundamental problems in probabilistic abstract argumentation: Beyond independence', *Artif. Intell.*, **268**, 1–29, (2019).

[23] Bettina Fazzinga, Sergio Flesca, and Filippo Furfaro, 'Reasoning over attack-incomplete aafs in the presence of correlations', in *Proceedings of the 18th International Conference on Principles of Knowledge Representation and Reasoning, KR 2021, Online event, November 3-12, 2021*, eds., Meghyn Bienvenu, Gerhard Lakemeyer, and Esra Erdem, pp. 301–311, (2021).

[24] Bettina Fazzinga, Sergio Flesca, and Filippo Furfaro, 'Abstract argumentation frameworks with marginal probabilities', in *Proc. 31st Int. Joint Conference on Artificial Intelligence (IJCAI), Vienna, Austria, 23-29 July 2022*, pp. 2613–2619, (2022).

[25] Bettina Fazzinga, Sergio Flesca, and Filippo Furfaro, 'Taking into account "who said what" in abstract argumentation: Complexity results', *Artif. Intell.*, **318**, 103885, (2023).

[26] Bettina Fazzinga, Sergio Flesca, Filippo Furfaro, and Luigi Pontieri, 'Process mining meets argumentation: Explainable interpretations of low-level event logs via abstract argumentation', *Inf. Syst.*, **107**, 101987, (2022).

[27] Anis Gargouri, Sébastien Konieczny, Pierre Marquis, and Srdjan Vesic, 'On a notion of monotonic support for bipolar argumentation frameworks', in *Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS)*, pp. 546–554, (2021).

[28] Anthony Hunter, 'Probabilistic qualification of attack in abstract argumentation', *Int. J. Approx. Reasoning*, **55**(2), 607–638, (2014).

[29] Benjamin Irwin, Antonio Rago, and Francesca Toni, 'Forecasting argumentation frameworks', in *Proc. Int. Conf. on Principles of Knowledge Representation and Reasoning, (KR)*, (2022).

[30] Hengfei Li, Nir Oren, and Timothy J. Norman, 'Probabilistic argumentation frameworks', in *Proc.of TAFA*, (2011).

[31] Jean-Guy Mailly, 'Constrained incomplete argumentation frameworks', in *Proc. Eur. Conf. on Symbolic and Quantitative Approaches to Reasoning with Uncertainty (ECSQARU)*, pp. 103–116, (2021).

[32] Jean-Guy Mailly, 'Yes, no, maybe, I don't know: Complexity and application of abstract argumentation with incomplete knowledge', *Argument Comput.*, **13**(3), 291–324, (2022).

[33] Diego C. Martínez, Alejandro Javier García, and Guillermo Ricardo Simari, 'On acceptability in abstract argumentation frameworks with an extended defeat relation', in *Proc. of COMMA*, pp. 273–278, (2006).

[34] Andreas Niskanen, Daniel Neugebauer, and Matti Järvisalo, 'Controllability of control argumentation frameworks', in *Proc. Int. Joint Conference on Artificial Intelligence (IJCAI)*, ed., Christian Bessiere, pp. 1855–1861, (2020).

[35] Farid Nouioua and Vincent Risch, 'Bipolar argumentation frameworks with specialized supports', in *Proc. of ICTAI*, pp. 215–218, (2010).

[36] Nir Oren and Timothy J. Norman, 'Semantics for evidence-based argumentation', in *Proc. of COMMA*, pp. 276–284, (2008).

[37] Sylwia Polberg and Anthony Hunter, 'Empirical evaluation of abstract argumentation: Supporting the need for bipolar and probabilistic approaches', *Int. J. Approx. Reason.*, **93**, 487–543, (2018).

[38] Nico Potyka, 'Bipolar abstract argumentation with dual attacks and supports', in *Proc. Int. Conf. on Principles of Knowledge Representation and Reasoning (KR)*, pp. 677–686, (2020).

[39] Nico Potyka, 'Generalizing complete semantics to bipolar argumentation frameworks', in *Proc. Int. Conf. on Symbolic and Quantitative Approaches to Reasoning with Uncertainty (ECSQARU)*, pp. 130–143, (2021).

[40] Henry Prakken and Giovanni Sartor, 'Law and logic: A review from an argumentation perspective', *Artificial Intelligence*, **227**, 214–245, (2015).

[41] Régis Riveret, 'On searching explanatory argumentation graphs', *J. Appl. Non Class. Logics*, **30**(2), 123–192, (2020).

[42] Régis Riveret and Nir Oren, 'On probabilistic argumentation and subargument-completeness', *J. Log. Comput.*, **32**(1), 115–128, (2022).

[43] Bart Verheij, 'The toulmin argument model in artificial intelligence', in *Argum. in Artificial Intelligence*, 219–238, (2009).