Consiglio Nazionale delle Ricerche
Istituto di Calcolo e Reti ad Alte Prestazioni

Complexity of fundamental problems in probabilistic abstract argumentation: beyond independence

Bettina Fazzinga, Sergio Flesca, Filippo Furfaro

RT-ICAR-CS-18-01 **Gennaio 2018**

Consiglio Nazionale delle Ricerche, Istituto di Calcolo e Reti ad Alte Prestazioni (ICAR)
– Sede di Cosenza, Via P. Bucci 7-11C, 87036 Rende, Italy, URL: www.icar.cnr.it
– Sezione di Napoli, Via P. Castellino 111, 80131 Napoli, URL: www.icar.cnr.it
– Sezione di Palermo, Via Ugo La Malfa, 153, 90146 Palermo, URL: www.icar.cnr.it

Complexity of fundamental problems in probabilistic abstract argumentation: beyond independence

Bettina Fazzinga[a], Sergio Flesca[b], Filippo Furfaro[b]

[a]*ICAR - CNR - ITALY*
[b]*DIMES - Università della Calabria - ITALY*

Abstract

The complexity of the probabilistic counterparts of the classical verification and acceptance problems is investigated over *probabilistic Abstract Argumentation Frameworks* (prAAFs). Most of the popular semantics of extensions (*admissible*, *stable*, *preferred*, *complete*, *grounded*, *ideal-set*, *ideal* and *semi-stable*) are considered, and the sensitivity of the complexity to several aspects (such as the semantics of the extension, the representation paradigm for encoding the prAAF, and the types of correlations between arguments/defeats) is studied. The complexity of the problems is shown to range from FP to $FP^{\#P}$-complete, with $FP^{||NP}$-complete cases, depending on the semantics of the extensions and the imposed correlations.

Keywords: Probabilistic Abstract Argumentation, Complexity

1. Introduction

Abstract Argumentation Frameworks (AAFs)

In the last decade, several argumentation frameworks have been proposed, with the aim of suitably modeling disputes between two or more parties. Typically, argumentation frameworks model both the possibility of parties to present arguments supporting their theses, and the possibility that some arguments rebut other arguments. Although argumentation is strongly related to philosophy and law, it has gained remarkable interest in AI as a reasoning model for representing dialogues, making decisions, and handling inconsistency and uncertainty [1, 2, 3].

A powerful yet simple argumentation framework is that proposed in the seminal paper [4], called *abstract argumentation framework* (AAF). An AAF is a representation of a dispute in terms of an *argumentation graph* $\langle A, D\rangle$, where A

is the set of nodes (each called *argument*) and D is the set of edges (each called *defeat* or, equivalently, *attack*). Basically, an argument is an abstract entity that may attack and/or be attacked by other arguments, and an attack expresses the fact that an argument rebuts/weakens another argument.

Example 1. *(inspired from Example* 1 *in [5]) The defense attorney of Mary and Marc wants to reason about the possible outcome of the trial of the robbery case involving his clients. The arguments of the case are the following, where Ann is a potential witness:*

- a: *"Mary says she was at the park when the robbery took place, and therefore denies being involved in the robbery";*
- b: *"Marc says he was at home when the robbery took place, and therefore denies being involved in the robbery";*
- c: *"Ann says that she is certain that he saw Mary outside the bank just before the robbery took place, and she also thinks that possibly she saw Marc there too".*

The arguments a *and* b *support the innocence of the defendants, and* c *means that a potential witness instills doubts about the innocence of both Mary and Marc.*

This scenario can be modeled by the AAF $\mathcal{A}$*, whose set of arguments is* $\{a, b, c\}$*, and whose defeat relation consists of the defeats* $\delta_{ac} = (a, c)$*,* $\delta_{ca} = (c, a)$*,* $\delta_{bc} = (b, c)$ *and* $\delta_{cb} = (c, b)$*, meaning that arguments* a *and* b *are both attacked by* c *and they both counter-attack* c*.* $\square$

Several semantics for AAFs, such as *admissible*, *stable*, *preferred*, *complete*, *grounded*, *ideal-set*, *ideal* and *semi-stable* have been proposed [4, 6, 7, 8] to identify "reasonable" sets of arguments, called *extensions*. Basically, each of these semantics corresponds to some properties that "certify" whether a set of arguments can be profitably used to support a point of view in a discussion. For instance, under the *admissible* semantics, a set S of arguments is an extension if S is "*conflict-free*" (i.e., there is no defeat between arguments in S) and is "robust" against the other arguments (i.e., every argument outside S attacking an argument in S is counterattacked by an argument in S). This means that who uses the set of arguments S in a discussion does not contradict her/himself, and can rebut to the arguments possibly presented by the other parties. The other semantics correspond to other ways of determining whether a set of arguments would be a "good point" in a dispute, and will be described in the core of the paper.

Uncertainty in argumentation: Probabilistic AAFs

As a matter of fact, in the real world, arguments and defeats are often *uncertain*. For instance, consider an argument a (or a defeat δ) encoding an interpretation or a translation of the description of a fact reported in a reference text. Then, a (or δ) may be uncertain in the sense that the original paragraph may have interpretations other than that encoded by a (or δ). In the legal scenario, this happens for arguments encoding different interpretations of the law.

Another form of uncertainty for an argument a corresponds to the fact that it is not guaranteed that a will be actually presented in the dispute (for instance, a witness may not decide to show up) or accepted by the referee coordinating the dispute (for instance, a judge may decide to rule out arguments presented by the parties). Analogously, a defeat between two arguments in a legal dispute can be uncertain: for instance, a jury can be instructed to consider an argument rebutted by another depending on the way of interpreting legal precedents.

Thus, several proposals have been made to model uncertainty in AAFs, by considering weights, preferences, or probabilities associated with arguments and/or defeats. One of the most popular approaches based on probability theory for modeling the uncertainty is the so called *constellations approach* [9, 10, 11, 12, 13, 5, 14, 15, 16]: the dispute is represented by means of a *Probabilistic Argumentation Framework* (prAAF), that consists in a set of alternative scenarios, each represented by a (deterministic) AAF (or, equivalently, an argumentation graph) associated with a probability. The various works in the literature investigating prAAFs can differ in the assumption on how the probability distribution function (pdf) over the scenarios is specified. For instance, in [9], the pdf is defined "extensively", by enumerating all the possible scenarios and, for each of them, the value of its probability. This form of prAAF will be denoted as EX (shorthand for "*extensive*"), and is "maximally expressive": it imposes no restriction on the probability assignments, thus allowing any kind of correlations between arguments/defeats to be expressed. On the other hand, in [14], the restriction that arguments and defeats are independent is assumed, and this is exploited to simplify the way probabilities are specified: the pdf is not explicitly specified, as it is implied by the marginal probabilities associated with arguments and defeats. This form of prAAF will be denoted as IND (shorthand for "*independent*"), and is obviously less expressive than EX (since it allows no correlation to be expressed) but much more compact.

Reasoning over argumentation frameworks: from the deterministic to the probabilistic scenario

In the deterministic setting, two classical problems supporting the reasoning over AAFs are:

- $\textsc{Ext}^{sem}(S)$: the verification problem of deciding whether a set of arguments S is an extension according to the semantics *sem*;
- $\textsc{Acc}^{sem}(a)$: the problem of deciding whether the argument a is *acceptable*, i.e., it belongs to at least one extension under the semantics *sem*.[1]

Basically, the relevance of these problems is that solving $\textsc{Ext}^{sem}(S)$ supports the decision on whether presenting a set of arguments in a dispute is a reasonable strategy, while solving $\textsc{Acc}^{sem}(a)$ focuses this analysis on single arguments.

In the probabilistic setting, there are multiple scenarios to be taken into account, and a set S (resp., an argument a) can be an extension (resp., acceptable) in some scenarios, but not in others. Thus, the natural probabilistic counterparts P-$\textsc{Ext}^{sem}(S)$ and P-$\textsc{Acc}^{sem}(a)$ of the above-mentioned problems $\textsc{Ext}^{sem}(S)$ and $\textsc{Acc}^{sem}(a)$ consist in evaluating the overall probability that S is an extension and a is acceptable, respectively, where "overall" means summing the probabilities of the scenarios where the property is verified. Intuitively, solving instances of P-$\textsc{Ext}^{sem}(S)$ and P-$\textsc{Acc}^{sem}(a)$ over sets of arguments and single arguments, respectively, can support the exploration of the search space of the possible strategies than can be adopted in the dispute, with the aim of composing a set of arguments providing good chances of success.

In the core of this paper, we will provide a thorough complexity characterization of P-$\textsc{Ext}^{sem}(S)$ and P-$\textsc{Acc}^{sem}(a)$, as better explained in what follows.

Contribution

In the literature, the complexity of $\textsc{Ext}^{sem}(S)$ an $\textsc{Acc}^{sem}(a)$ has been extensively studied, and a summary of the results is provided in Table 1 and Table 2. Much less is known about the complexity of the counterparts of the same problems in the probabilistic setting. Up to our knowledge, the only results in the literature are those in [16] and [17] referring to the complexity of P-$\textsc{Ext}^{sem}(S)$

[1] The acceptance problem can be stated also under a sceptical semantics, where the answer is *yes* iff a belongs to *all* the extensions. In Section 8 we will elaborate on how easily our results can be extended to cover the sceptical semantics.

and P-$\text{Acc}^{sem}(a)$ over prAAFs of form IND, while no tight characterization has been provided over prAAFs of form EX.

In this regard, our first contribution is the complexity characterization of P-$\text{Ext}^{sem}(S)$ and P-$\text{Acc}^{sem}(a)$ over the form EX: we show that, depending on the semantics of the extension, P-$\text{Ext}^{sem}(S)$ and P-$\text{Acc}^{sem}(a)$ are either in FP (i.e., they can be solved in polynomial-time) or $FP^{||NP}$-complete or in $FP^{||\Sigma_2^p}$ (where $FP^{||\mathcal{C}}$ is the class of problems solvable by a deterministic polynomial-time Turing machine extended with parallel invocations of an oracle for the class $\mathcal{C}$).

Starting from this, in order to give an insight on the source of complexity of the considered problems, we investigate the sensitivity of the complexity to the type of correlations (between arguments and/or defeats) encoded in the prAAF. As a matter of fact, to perform this analysis, both the paradigms EX and IND are inadequate. On the one hand, IND does not allow correlations to be expressed. On the other hand, in EX correlations can be encoded, but implicitly: for instance, one can impose that two arguments must co-exist by assigning 0 probability to possible scenarios where only of them is present, but the reverse process of inferring the correlations from the probabilities is hard to accomplish. Hence, we introduce a new prAAF (called GEN) where the constructs used for defining the pdf over the possible scenarios allow for explicitly specifying correlations (such as mutual exclusion and co-occurrence) between arguments and defeats. This framework is based on the well-known paradigm of *world-set descriptors* (*wsd*s) and *ws-sets*, that was shown to be a complete and succinct formalism for specifying pdfs over possible worlds in [18, 19] and has been used profitably in the context of probabilistic databases. Interestingly, we exploit the fact that different syntactic restrictions on wsds correspond to allowing different forms of correlations to be expressed, and provide a complexity analysis of P-$\text{Ext}^{sem}(S)$ and P-$\text{Acc}^{sem}(a)$ for each syntactic class, thus showing the sensitivity of their complexity to the presence of different forms of correlations between arguments/defeats.

As a collateral contribution, it is worth noting that GEN and its subclasses not only serve the purpose of allowing a thorough complexity characterization of P-$\text{Ext}^{sem}(S)$ and P-$\text{Acc}^{sem}(a)$, but also can be viewed as prAAFs of their own validity. We will discuss how their nice combination of expressiveness, compactness and aptitude to allow an easy specification of different correlations makes them valid paradigms for modeling uncertainty in argumentation in practical scenarios.

2. PRELIMINARIES

2.1. Abstract Argumentation Frameworks (AAFs)

An *abstract argumentation framework* [4] (*AAF*) is a pair $\langle A, D\rangle$, where A is a finite set, whose elements are called *arguments*, and $D \subseteq A \times A$ is a binary relation over A, whose elements are called *defeats* (or *attacks*). An argument is an abstract entity whose role is determined by its relationships with other arguments.

Given arguments $a, b \in A$, we say that a *defeats* b iff there is $(a, b) \in D$. Similarly, a set $S \subseteq A$ *defeats* an argument $b \in A$ iff there is $a \in S$ such that a *defeats* b; and argument a *defeats* S iff there is $b \in S$ such that a *defeats* b. Given a set $S \subseteq A$ of arguments, we define S^+ as the set of arguments that are defeated by S, that is, $S^+ = \{a \in A \text{ s.t. } S \text{ defeats } a\}$.

A set $S \subseteq A$ of arguments is said to be *conflict-free* if there are no $a, b \in S$ such that a *defeats* b. An argument a is said to be *acceptable w.r.t. a set of arguments* $S \subseteq A$ iff $\forall b \in A$ such that b *defeats* a, there is $c \in S$ such that c *defeats* b.

An AAF can be naturally viewed as a graph, called *argumentation graph*, whose nodes are the arguments in A and whose edges are the defeats in D. Thus, we will often use the terms "AAF" and "argumentation graph" as synonyms.

2.2. Semantics of the AAFs and fundamental problems: $\textsc{Ext}^{sem}(S)$ *and* $\textsc{Acc}^{sem}(a)$

Several semantics for AAFs have been proposed to identify "reasonable" sets of arguments, called *extensions*. We consider the following well-known semantics: *admissible* (`ad`), *stable* (`st`), *complete* (`co`), *grounded* (`gr`), *preferred* (`pr`) [4], *ideal-set* (`ids`), *ideal* (`ide`) [6], and *semi-stable* (`sst`) [8].
A set $S \subseteq A$ is said to be:

- an *admissible extension* iff S is conflict-free and all its arguments are acceptable w.r.t. S;
- a *stable extension* iff S is conflict-free and S defeats each argument in $A \setminus S$;
- a *complete extension* iff S is admissible and S contains all the arguments that are acceptable w.r.t. S;
- a *grounded extension* iff S is a minimal (w.r.t. $\subseteq$) complete set of arguments;
- a *semi-stable extension* iff S is a complete extension where $S \cup S^+$ is maximal (w.r.t. $\subseteq$);

- a *preferred extension* iff S is a maximal (w.r.t. $\subseteq$) complete set of arguments;
- an *ideal-set extension* iff S is admissible and S is contained in every preferred set of arguments;
- an *ideal extension* iff S is a maximal (w.r.t. $\subseteq$) ideal-set extension.

In the following, we denote as SEM the set $\{\texttt{ad}, \texttt{st}, \texttt{co}, \texttt{gr}, \texttt{pr}, \texttt{ids}, \texttt{ide}, \texttt{sst}\}$ consisting of the above-listed semantics.

Example 2. *Consider the AAF ⟨A, D⟩ of Example 1, where the set A of arguments is $\{a, b, c\}$, and the set D of defeats is $\{\delta_{ac} = (a, c), \delta_{ca} = (c, a), \delta_{bc} = (b, c), \delta_{cb} = (c, b)\}$. As $S = \{a, b\}$ is conflict-free and both a and b are acceptable w.r.t. S, it is the case that S is admissible. It is easy to see that sets $\emptyset$, $S_1 = \{c\}$, $S_2 = \{a\}$ and $S_3 = \{b\}$ are admissible extensions, while set $S_4 = \{a, c\}$ is not admissible since it is not conflict-free. Both S_2 and S_3 are not complete, since they do not contain all the acceptable arguments: b (resp., a) is acceptable w.r.t S_2 (S_3). Sets $\emptyset$, S and S_1 are complete extensions. Moreover, S and S_1 are also stable, semi-stable and preferred extensions. $\emptyset$ is the unique grounded extension, and it is also both an ideal-set and an ideal extension.* □

Given an AAF $\alpha = \langle A, D\rangle$, a set $S \subseteq A$ of arguments, and a semantics *sem* $\in SEM$, we define the function *ext*(α, sem, S) that returns *true* if S is an extension according to sem, *false* otherwise.

The fundamental problem of verifying whether a set S of arguments is an extension over a given AAF according to a semantics $sem \in SEM$ will be denoted as $\textsc{Ext}^{sem}(S)$. Basically, solving an instance of $\textsc{Ext}^{sem}(S)$ means checking whether a set of arguments is a reasonable strategy in the dispute, where the meaning of "reasonable" is encoded in the chosen semantics. Focusing on a single argument, rather than on a set of arguments, is the rationale behind the acceptance problem $\textsc{Acc}^{sem}(a)$, that is the problem of verifying whether the argument a belongs to at least one extension, under the specified semantics.

2.3. *Probabilistic AAFs (prAAFs):* EX *and* IND.

A well-established way of modeling uncertainty in abstract argumentation is the so-called *constellations approach*, that consists in considering alternative possible scenarios, and assigning a probability to each of them. Basically, each scenario is an AAF (formally called "*possible AAF*") consisting in a subset of the arguments and defeats that can occur in the dispute, and can be viewed as a hypothesis on which arguments will be actually presented in the dispute, and on which "rebuttal" relationships between them will occur. Thus, a *probabilistic AAF* (prAAF) $\mathcal{F}$ is triple $\langle A, D, P\rangle$ where A and D are sets of arguments and defeats, respectively, and P is a probability distribution function (pdf) over the set $\{\langle A', D'\rangle \mid A' \subseteq A \wedge D' \subseteq (A' \times A') \cap D\}$ consisting of the possible AAFs having subsets of A and D as arguments and defeats, respectively.

In the literature, different proposals of prAAFs can be found. The main difference is in the way the pdf P over the possible AAFs is encoded, and we denote as EX and IND the most popular representation paradigms for P.

In particular, EX is a form of prAAF (that is at the basis of the frameworks in [9, 10, 12]) where the pdf over the possible AAFs is specified "extensively", by indicating one by one the scenarios with non-zero probability, and the probability of each of them. That is, a prAAF $\mathcal{F}$ of form EX is a tuple $\langle A, D, \vec{\alpha}, \vec{P}\rangle$, where A and D are sets of arguments and defeats, while $\vec{\alpha}$ is the sequence $\vec{\alpha} = \alpha_1, \dots, \alpha_m$ of the possible AAFs that are assigned non-zero probability and $\vec{P} = P(\alpha_1), \dots, P(\alpha_m)$ are their probabilities. The size of a prAAF $\langle A, D, \vec{\alpha}, \vec{P}\rangle$ of form EX is thus $O\Big((|A| + |D|) \cdot (|\vec{\alpha}| + |\vec{P}|)\Big)$.

Example 3 (An example of prAAF of form EX). *Consider the sets of arguments $A = \{a, b, c\}$ and of defeats $D = \{\delta_{ac}, \delta_{ca}, \delta_{bc}, \delta_{cb}\}$ introduced in Example 1, and assume that the lawyer thinks that only the following 4 scenarios are possible:*
S_1*: "Ann will not testify";*
S_2*: "Ann will testify, and the jury will deem that her argument c undermines those of Mary and Marc (arguments a, b), and vice versa";*
S_3*: "Ann will testify, and the jury will deem that her argument c undermines Mary's and Marc's arguments a, b, while, owing to the bad reputations of Mary and Marc, a and b will be not perceived as strong enough to undermine argument c";*
S_4*: "Ann will testify, and the jury will deem that her argument c undermines Mary's argument a but not Marc's argument b, since Ann was uncertain about Marc's presence. In the other direction, a and b will be not perceived as strong enough to undermine c".*

Each scenario S_i is encoded by the AAF α_i in the following list:
$\alpha_1 = \langle\{a, b\}, \emptyset\rangle$, $\alpha_2 = \langle\{a, b, c\}, \{\delta_{ac}, \delta_{ca}, \delta_{bc}, \delta_{cb}\}\rangle$,
$\alpha_3 = \langle\{a, b, c\}, \{\delta_{ca}, \delta_{cb}\}\rangle$, $\alpha_4 = \langle\{a, b, c\}, \{\delta_{ca}\}\rangle$.

Basically, the form of prAAF EX *allows the lawyer to define, one by one, which* *scenarios are possible, and then to assign a probability to the AAF corresponding to each scenario, on the basis of her/his perception of how likely the scenario is. For instance, the pdf set by the lawyer could be such that:* $P(\alpha_1) = 0.1$ *and* $P(\alpha_2) = P(\alpha_3) = P(\alpha_4) = (1 - P(\alpha_1))/3 = 0.3$*, meaning that the lawyer thinks that there is* 10% *probability that Ann will not manage to testify (owing to her ill-health), and that, in the case she testifies, the other three scenarios are equi-probable.*

This uncertainty is modeled by the following prAAF of form EX*:* $\mathcal{F} = \langle A, D, \vec{\alpha}, \vec{P} \rangle$*, where:* $A = \{a, b, c\}$*,* $D = \{\delta_{ca}, \delta_{cb}, \delta_{ac}, \delta_{bc}\}$*,* $\vec{\alpha} = [\alpha_1, \ldots, \alpha_4]$ *and* $\vec{P} = [0.1, 0.3, 0.3, 0.3]$. □

We now focus on the form of prAAF denoted as IND [11, 14, 15, 16]. Here, the possible AAFs and the pdf over them are implicitly defined, as they are implied by assigning marginal probabilities to arguments and defeats, and assuming independence between them. That is, a prAAF of type IND is a tuple $\langle A, D, P_A, P_D \rangle$, where $A = \{a_1, \ldots, a_m\}$ and $D = \{\delta_1, \ldots, \delta_n\}$ are the sets of arguments and defeats, and $P_A = \{P(a_1), \ldots, P(a_m)\}$, $P_D = \{P(\delta_1), \ldots, P(\delta_n)\}$, are the marginal probabilities of arguments and defeats. The pdf P over the possible scenarios that is implied by the independence assumption and the marginal probabilities P_A, P_D is as follows. For each possible AAF $\alpha_i = \langle A_i, D_i \rangle$, with $A_i \subseteq A$, and $D_i \subseteq (A_i \times A_i) \cap D$, the probability $P(\alpha_i)$ is:

$$P(\alpha_i) = \prod_{a \in A_i} P(a) \times \prod_{a \in A \setminus A_i} \Big(1 - P(a)\Big) \times \prod_{\delta \in D_i} P(\delta) \times \prod_{\delta \in \overline{D}(\alpha_i) \setminus D_i} \Big(1 - P(\delta)\Big), \quad (1)$$

where $\overline{D}(\alpha_i)$ is the set of all the defeats in D between arguments in A_i, that is $\overline{D}(\alpha_i) = D \cap (A_i \times A_i)$.

The size of a prAAF of type IND is $O(|A| + |D| + |P_A| + |P_D|)$.

Example 4 (An example of prAAF of form IND). *Consider the sets of arguments* $A = \{a, b, c\}$ *and of defeats* $D = \{\delta_{ac}, \delta_{ca}, \delta_{bc}, \delta_{cb}\}$ *introduced in Example 1. In the case that the lawyer thinks that the occurrences of any argument does not influence the presence of other arguments (and the same for defeats), independence can be assumed between the terms of the dispute. Hence, the lawyer can focus on the probability of occurrence of each single argument and defeat separately from the others. For instance, the lawyer may set* $P(c) = 0.9$ *(meaning that there is* 10% *probability that Ann will not manage to testify) and* $P(a) = P(b) = 1$

(meaning that Mary and Marc will certainly testify). Moreover, she/he could set $P(\delta_{ca}) = 1$ *(meaning that she/he is certain that the jury will consider Ann's argument as a solid rebuttal of Mary's argument). Analogously, she/he could set* $P(\delta_{cb}) = 0.8$ *and* $P(\delta_{ac}) = P(\delta_{bc}) = 0.4$.

Given this, since the arguments are considered independent, the possible scenarios modeled by IND *are all the AAFs* $\langle A_i, D_i\rangle$ *where* A_i *is a subset of* A *and* D_i *a subset of the defeats in* D *between the arguments in* A_i*. Specifically, there are* 9 *possible AAFs, where* 4 *out of* 9 *are equal to* $\alpha_1, \ldots \alpha_4$ *of Example 3, and the others are* $\alpha_5 = \langle\{a, b, c\}, \{\delta_{ac}, \delta_{ca}\}\rangle$, $\alpha_6 = \langle\{a, b, c\}, \{\delta_{ca}, \delta_{bc}\}\rangle$, $\alpha_7 = \langle\{a, b, c\}, \{\delta_{ca}, \delta_{cb}, \delta_{ac}\}\rangle$, $\alpha_8 = \langle\{a, b, c\}, \{\delta_{ca}, \delta_{cb}, \delta_{bc}\}\rangle$, $\alpha_9 = \langle\{a, b, c\}, \{\delta_{ca}, \delta_{ac}, \delta_{bc}\}\rangle$*. Moreover, the probability assigned to each AAF* $\langle A_i, D_i\rangle$ *is the product of the probabilities (resp., the complements of the probabilities) of the arguments in* A_i *(resp., in* A *but not in* A_i*) and the probabilities (resp., the complements of the probabilities) of the defeats in* D_i *(resp., of the defeats between arguments in* A_i *that are in* D *but not in* D_i*). For instance,* $P(\alpha_1) = P(a) \times P(b) \times (1 - P(c)) = 0.1$ *and* $P(\alpha_3) = P(a) \times P(b) \times P(c) \times P(\delta_{ca}) \times P(\delta_{cb}) \times (1 - P(\delta_{ac})) \times (1 - P(\delta_{bc})) = 0.26$. □

In what follows, given a prAAF $\mathcal{F} = \langle A, D, P\rangle$ of any kind (thus, independently from the way P is encoded), we denote as $\mathcal{F}.\vec{\alpha} = \alpha_1, \ldots, \alpha_m$ the possible AAFs that are assigned non-zero probability by P, and as $\mathcal{F}.\vec{P} = P(\alpha_1), \ldots, P(\alpha_m)$ their probabilities. For instance, if $\mathcal{F}$ is encoded as a prAAF of form EX, then $\mathcal{F}.\vec{\alpha}$ and $\mathcal{F}.\vec{P}$ are exactly the terms $\vec{\alpha}$ and $\vec{P}$ in the tuple $\langle A, D, \vec{\alpha}, \vec{P}\rangle$ encoding $\mathcal{F}$. Analogously, if $\mathcal{F}$ is of form IND, then $\mathcal{F}.\vec{\alpha}$ is the set of all the possible AAFs definable over A and D, and $\mathcal{F}.\vec{P}$ is the pdf defined in Equation 1.

It is worth noting that IND and EX can be viewed as extreme and opposite ways of allowing probabilities to be specified. On the one hand, IND is *compact*, as it requires probabilities to be specified only for each argument and defeat, whose number n_{AD} is typically less than the number n_S of scenarios (observe that n_S is $O(2^{n_{AD}})$). However, IND assumes independence between every pair of arguments and/or defeats, thus it cannot be used when one wants to specify correlations. On the other hand, EX is *complete*, as it allows probabilities to be directly specified on each possible scenario, but is *verbose*, as more and more probabilities must be specified as the number of scenarios to be considered increases. In the following, we will introduce a new prAAF, called GEN, where the pdf over the possible scenarios is encoded using the paradigm of *world-set descriptors*. This paradigm is somehow intermediate between the forms EX and IND: it has the same expressiveness as the form EX (since it allows any pdf to be encoded, thus allowing any correlation to be expressed), but is more compact, as its encoding does not require the explicit enumeration of the scenarios with non-zero probability.

2.4. Reasoning in the probabilistic setting: P-EXT$^{sem}(S)$ *and* P-ACC$^{sem}(a)$

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

Definition 1 (P-EXT$^{sem}(S)$ and $P^{sem}(S)$). *Given a prAAF $\mathcal{F}$, a set S of arguments, and a semantics sem $\in SEM$,* P-EXT$^{sem}(S)$ *is the problem of computing the probability $P^{sem}(S)$ that S is an extension according to sem, that is:*

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

Example 5. *Considering Example 3, the probability $P^{ad}(\{a,b\})$ that set $\{a,b\}$ (resp., $\{c\}$) is an admissible extension is given by $P(\alpha_1) + P(\alpha_2) = 0.4$. Analogously, $P^{ad}(\{c\}) = P(\alpha_2) + P(\alpha_3) + P(\alpha_4) = 0.9$. Considering Example 4, instead, $P^{ad}(\{a,b\})$ is equal to $P(\alpha_1) + P(\alpha_4) + P(\alpha_5) + P(\alpha_6) + P(\alpha_7) + P(\alpha_8) + P(\alpha_9) = 0.676$, and $P^{ad}(\{c\})$ is equal to $P(\alpha_2) + P(\alpha_3) + P(\alpha_4) + P(\alpha_5) + P(\alpha_7) + P(\alpha_8) = 0.829$.* $\square$

Analogously, when moving from the deterministic to the probabilistic setting, the acceptance problem ACC$^{sem}(a)$ becomes the problem of evaluating the probability that the acceptance of a is verified, by taking into account all the alternative scenarios.

Definition 2 (P-ACC$^{sem}(a)$ and $P^{sem}_{acc}(a)$). *Given a prAAF $\mathcal{F}$, an argument a, and a semantics sem $\in SEM$,* P-ACC$^{sem}(a)$ *is the problem of computing the probability $P^{sem}_{acc}(a)$ that a belongs to at least one extension according to sem, that is:*

$$P^{sem}_{acc}(a) = \sum_{\alpha \,\in\, \mathcal{F}.\vec{\alpha} \,\wedge\, \exists S | (ext(\alpha, sem, S) \,\wedge\, a \in S)} \mathcal{F}.P(\alpha) \qquad (3)$$

In the following, given a form of prAAF F, we will denote as P-EXT$^{sem}_{\mathrm{F}}(S)$ and P-ACC$^{sem}_{\mathrm{F}}(a)$ the same problems P-EXT$^{sem}(S)$ and P-ACC$^{sem}(a)$ restricted to the case that the input prAAF is of form F, respectively.

2.5. Complexity classes

Compared with other papers providing results on the complexity of the decisional problems $\textsc{Ext}^{sem}(S)$ and $\textsc{Acc}^{sem}(a)$, here we have to resort also to *functional* complexity classes (as done in [16]), that are more suitable for characterizing the complexity of P-$\textsc{Ext}^{sem}(S)$ and P-$\textsc{Acc}^{sem}(a)$, since they are intrinsically research problems. Hence, we here briefly recall the meaning of both the decisional and functional complexity classes that are used in the rest of the paper.

P (resp. NP) is the class of the decision problems that can be solved by a deterministic (resp. non-deterministic) Turing machine in polynomial time w.r.t. the size of the input of the problem. $coNP$ is the class of the decision problems whose complement is in NP. It is conjectured that $P \subset NP$, $P \subset coNP$, and $NP \neq coNP$. The *polynomial hierarchy* is the following sequence of classes: let $\Delta_0^P = \Sigma_0^P = \Pi_0^P = P$, for all $i \geq 0$, $\Delta_{i+1}^P = P^{\Sigma_i^P}$, $\Sigma_{i+1}^P = NP^{\Sigma_i^P}$, $\Pi_{i+1}^P = coNP^{\Sigma_i^P}$, where X^Y denotes the class of the problems that can be solved by an algorithm in class X that calls an oracle in class Y. Thus, in the first level ($i = 1$) we have $\Delta_1^P = P$, $\Sigma_1^P = NP$, and $\Pi_1^P = coNP$, and in the second level ($i = 2$) we have $\Delta_2^P = P^{NP}$, $\Sigma_2^P = NP^{NP}$, and $\Pi_2^P = coNP^{NP}$. At each level $i > 0$, the three classes are related by the same inclusions that holds for P, NP, and $coNP$. Moreover, each class at each level includes all classes at previous levels. The (cumulative) polynomial hierarchy is the class $PH = \bigcup_{i\geq 0} \Sigma_i^P$. The complexity class Θ_2^p is a subclass of Δ_2^P and is also known as $P^{||NP}$, that is the class of decision problems that can be solved by a polynomial-time Turing machine asking queries in *parallel* to an NP oracle (here, parallel means that the invocations of the oracle can be done in a non-adaptive way, in the sense that what asked at any invocation does not depend from the result of any other invocation). It has been shown that $P^{||NP}$ coincides with $P^{NP[\log n]}$ (in this case, at most $O(\log n)$ adaptive queries can be asked by the oracle machine).

FP is the class of the function problems that can be solved by a deterministic Turing machine in polynomial time (w.r.t. the size of the input of the problem). In this paper, we in particular deal with the classes $FP^{\#P}$ and $FP^{||NP}$. $FP^{\#P}$ is the class of functions computable by a polynomial-time Turing machine with a $\#P$ oracle. $\#P$ is the complexity class of the functions f such that f counts the number of accepting paths of a nondeterministic polynomial-time Turing machine [20]. Analogously, $FP^{||NP}$ is the class of functions computable by a polynomial-time Turing machine with access to an NP oracle, whose calls are non-adaptive (as explained above, this is equivalent to saying that the oracle invocations take place in parallel).

We will exploit the following results regarding the relationships between different functional complexity classes: (i) for each complexity class $\#\mathcal{C} \in \#PH$, it holds that $FP^{\#P} = FP^{\#\mathcal{C}}$, since $\#PH \subseteq FP^{\#P[1]}$ under polynomial time 1-Turing reductions [21], and $FP^{FP^{\#P[1]}} \subseteq FP^{\#P}$; (ii) a function is $FP^{\#P}$-hard iff it is $\#P$-hard [2], and thus to prove that a problem is $FP^{\#P}$-hard it suffices to reduce a $\#P$-hard problem to it.

3. COMPLEXITY OF P-EXT$^{sem}_{\text{EX}}(S)$ and P-ACC$^{sem}_{\text{EX}}(a)$

We here provide the first contribution of our paper. Given that the complexities of P-EXT$^{sem}(S)$ and P-ACC$^{sem}(a)$ over prAAFs of form IND have been characterized in [16, 17], we here complete the picture of the complexity of these fundamental problems by characterizing their complexities over prAAFs of form EX. In particular, we will prove that P-EXT$^{sem}_{\text{EX}}(S)$ is polynomial-time solvable for any semantics sem whose deterministic counterpart EXT$^{sem}(S)$ is polynomial-time decidable, and that it is $FP^{||NP}$-complete for all the other semantics in SEM (whose deterministic counterpart EXT$^{sem}(S)$ is either NP-complete or Θ_2^p-complete – see Table 1).

Theorem 1. P-EXT$^{sem}_{\text{EX}}(S)$ *is in* FP *for* $sem \in \{\mathtt{ad}, \mathtt{st}, \mathtt{gr}, \mathtt{co}\}$.

Proof. The fact that computing Equation 2 can be done in polynomial time straightforwardly follows from the facts that: $i.$ for every $sem \in \{\mathtt{ad}, \mathtt{st}, \mathtt{gr}, \mathtt{co}\}$, deciding whether S is an extension in a deterministic AAF (i.e., solving EXT$^{sem}(S)$) is in PTIME; $ii.$ the number of possible AAFs over which this check must be performed is linear in the input (we recall that the size of EX is proportional to the number of possible AAFs). □

Theorem 2. P-EXT$^{sem}_{\text{EX}}(S)$ *is* $FP^{||NP}$*-complete for* $sem \in \{\mathtt{pr}, \mathtt{ids}, \mathtt{ide}, \mathtt{sst}\}$.

---

[2] Obviously a problem is $\#P$-hard if it is $FP^{\#P}$-hard. As for the vice versa, it can be shown that any problem $A \in FP^{\#P}$ can be reduced to a $\#P$-hard problem B, under polynomial time 1-Turing reductions reasoning as follows. $A \in FP^{\#P}$ implies that there is a polynomial-time algorithm M_A that solves A using exactly one call to a $\#P$ oracle C (as $FP^{\#P} \subseteq FP^{\#P[1]}$). Moreover, since B is $\#P$-hard, there is a polynomial-time algorithm M_C solving the oracle C (used in M_A) that uses exactly one call to a B oracle. Therefore, by just combining M_A and M_C we obtain a polynomial time algorithm which solves A by using exactly one call to a B oracle, and this implies that A can be reduced to B under polynomial-time 1-Turing reductions.

sem	$\textsc{Ext}^{sem}(S)$	P-$\textsc{Ext}^{sem}(S)$			
		IND	EX	GEN, BOOL, MON, IND-A	IND-D
admissible	*P*	*FP*	*FP*	$FP^{\#P}$-c	*FP*
stable	*P*	*FP*	*FP*	$FP^{\#P}$-c	*FP*
complete	*P*	$FP^{\#P}$-c	*FP*	$FP^{\#P}$-c	$FP^{\#P}$-c
grounded	*P*	$FP^{\#P}$-c	*FP*	$FP^{\#P}$-c	$FP^{\#P}$-c
semi-stable	*coNP*-c	$FP^{\#P}$-c	$FP^{\|\|NP}$-c	$FP^{\#P}$-c	$FP^{\#P}$-c
preferred	*coNP*-c	$FP^{\#P}$-c	$FP^{\|\|NP}$-c	$FP^{\#P}$-c	$FP^{\#P}$-c
ideal-set	*coNP*-c	$FP^{\#P}$-c	$FP^{\|\|NP}$-c	$FP^{\#P}$-c	$FP^{\#P}$-c
ideal	in Θ_2^p, *coNP*-h	$FP^{\#P}$-c	$FP^{\|\|NP}$-c	$FP^{\#P}$-c	$FP^{\#P}$-c

Table 1: Complexity of $\textsc{Ext}^{sem}(S)$ and P-$\textsc{Ext}^{sem}(S)$ for different forms of prAAFs (the results for $\textsc{Ext}^{sem}(S)$ are from the literature)

<table>
<tr><th rowspan="2">sem</th><th rowspan="2">$\textsc{Acc}^{sem}(a)$</th><th colspan="3">P-$\textsc{Acc}^{sem}(a)$</th></tr>
<tr><th>IND</th><th>EX</th><th>GEN, BOOL, IND-D MON, IND-A</th></tr>
<tr><td>admissible</td><td>NP-c</td><td rowspan="8">$FP^{\#P}$-c</td><td>$FP^{||NP}$-c</td><td rowspan="8">$FP^{\#P}$-c</td></tr>
<tr><td>stable</td><td>NP-c</td><td>$FP^{||NP}$-c</td></tr>
<tr><td>complete</td><td>NP-c</td><td>$FP^{||NP}$-c</td></tr>
<tr><td>grounded</td><td>P</td><td>FP</td></tr>
<tr><td>semi-stable</td><td>Σ_p^2-c</td><td>in $FP^{||\Sigma_p^2}$, $FP^{||NP}$-h</td></tr>
<tr><td>preferred</td><td>NP-c</td><td>$FP^{||NP}$-c</td></tr>
<tr><td>ideal-set</td><td>in Θ_2^p, coNP-h</td><td>$FP^{||NP}$-c</td></tr>
<tr><td>ideal</td><td>in Θ_2^p, coNP-h</td><td>$FP^{||NP}$-c</td></tr>
</table>

Table 2: Complexity of $\textsc{Acc}^{sem}(a)$ and P-$\textsc{Acc}^{sem}(a)$ for different forms of prAAFs (the results for $\textsc{Ext}^{sem}(S)$ are from the literature)

Proof.
Membership in $FP^{||NP}$. For the semantics `sst`, `pr` and `ids` (for which $\textsc{Ext}^{sem}(S)$ is coNP-complete), the membership in $FP^{||NP}$ follows from the fact that P-$\textsc{Ext}^{sem}_{\textsc{ex}}(S)$ can be solved by performing as many parallel invocations to NP oracles (each solving an instance of $\textsc{Ext}^{sem}(S)$ over a possible AAF with non-zero probability) as the number of possible AAFs encoded in the prAAF. For the semantics `ide`, the membership to $FP^{||NP}$ still holds, since a polynomial time Turing machine with parallel invocations to Θ_2^p oracles can be easily converted into a polynomial time Turing machine with parallel invocations to NP oracles (this follows from the fact that $\Theta_2^p = P^{||NP}$).

Hardness for $FP^{||NP}$. We will show a reduction from the $FP^{||NP}$-hard problem $sup(\phi)$, that is the problem of computing the *supremum* of the satisfying assignments for a 3CNF Boolean formula $\phi(x_1, \ldots, x_n)$. We recall that the supremum $sup(\phi)$ is the assignment where, for each $i \in [1..n]$, the variable x_i is assigned with *true* iff there exists a satisfying assignment of ϕ wherein $x_i =$ *true*.

For the sake of presentation, we first show a reduction for $sem \in \{\texttt{ids}, \texttt{ide}\}$, then we discuss the case of preferred and semi-stable semantics.

Let $\phi = C_1 \wedge C_2 \wedge \ldots \wedge C_k$ be the 3-CNF boolean formula in the instance of supremum, over the set $X = \{x_1, \ldots, x_n\}$ of variables. We build n formulas $\phi_1, \ldots, \phi_n$ from ϕ, where each ϕ_i has the form $\phi_i = C_{i,1} \wedge C_{i,2} \wedge \ldots \wedge C_{i,k_i}$ and is obtained from ϕ by assigning $x_i =$ *true*. Herein, $k_i \leq k$, since clauses of ϕ containing x_i evaluate to true, thus are removed from ϕ_i. Moreover, when constructing ϕ_i, any occurrence of $\neg x_i$ is deleted, thus any clause $C_{i,j}$ of ϕ_i can have either $|C_{i,j}| = 2$ or $|C_{i,j}| = 3$ literals (w.l.o.g. we assume that each original clause of ϕ contains literals of distinct variables). We denote the subset of X containing the variables still occurring in ϕ_i as $X(\phi_i)$, and the cardinality of $X(\phi_i)$ as n_i. For instance, starting from $\phi = (\neg x_1 \vee x_2 \vee x_3) \wedge (\neg x_1 \vee \neg x_2 \vee x_3) \wedge (x_1 \vee x_4 \vee x_5)$, where $X = \{x_1, x_2, x_3, x_4, x_5\}$ and $k = 3$, after assigning $x_1 =$ *true* we obtain $\phi_1 = (x_2 \vee x_3) \wedge (\neg x_2 \vee x_3)$, where $X(\phi_1) = \{x_2, x_3\}$ and $k_1 = 2$.

Then, we define the prAAF $\mathcal{F}_\phi = \langle A, D, \vec{\alpha}, \vec{P} \rangle$ of form EX whose components are as follows:

- $A = A_0 \cup \ldots \cup A_n$, where A_0 consists of an argument s, whereas each other A_i contains (i) an argument $C_{i,j}$ for each clause $C_{i,j}$ appearing in ϕ_i; (ii) two arguments x_l and $\neg x_l$ for each variable $x_l \in X(\phi_i)$; (iii) two arguments ψ_i, ϕ_i; (iv) an argument s.

- $D = D_0 \cup \ldots \cup D_n$, where D_0 contains only the defeat $\delta_0 = (s, s)$, whereas each other D_i contains (i) two defeats $\delta_{\psi_i} = (\psi_i, \psi_i)$ and $\delta_{\phi_i} = (\phi_i, \psi_i)$; (ii) two defeats $\delta_1^i = (s, \phi_i)$ and $\delta_2^i = (\phi_i, s)$; (iii) for each clause $C_{i,j}$ of ϕ_i, a defeat $\delta_{i,j} = (C_{i,j}, \phi_i)$; (iv) for each clause $C_{i,j}$ and each literal $l_{i,j}^u$ occurring in $C_{i,j}$ (with $u \in [1..|C_{i,j}|]$), a defeat $\delta_{i,j}^u = (l_{i,j}^u, C_{i,j})$, where $l_{i,j}^u$ is an argument of the form x_l or $\neg x_l$, with $x_l \in X(\phi_i)$; (v) two defeats $\delta_1^{x_l} = (x_l, \neg x_l)$ and $\delta_2^{x_l} = (\neg x_l, x_l)$ for each variable $x_l \in X(\phi_i)$; (vi) two defeats $\delta_{l,1}^{\psi_i} = (\psi_i, x_l)$ and $\delta_{l,2}^{\psi_i} = (\psi_i, \neg x_l)$ for each variable $x_l \in X(\phi_i)$.

- $\vec{\alpha} = \alpha_0, \alpha_1, \ldots, \alpha_n$, where $\alpha_i = \langle A_i, D_i \rangle$, for each $i \in [1..n]$;

- $\vec{P} = P_0, P_1, \ldots, P_n$, where $P_0 = 1/2^n$ and, for each $i \in [1..n]$, $P_i = 2^{i-1}/2^n$.

We consider the bijection β from $\phi_1, \ldots, \phi_n$ to the AAFs $\langle A_1, D_1 \rangle \ldots, \langle A_n, D_n \rangle$ such that, for every ϕ_i, $\beta(\phi_i) = \langle A_i, D_i \rangle$.

The graphical representation of the AAF $\langle A_1, D_1 \rangle = \beta(\phi_1)$, where $\phi_1 = (x_2 \vee x_3) \wedge (\neg x_2 \vee x_3)$, is reported in Figure 1.

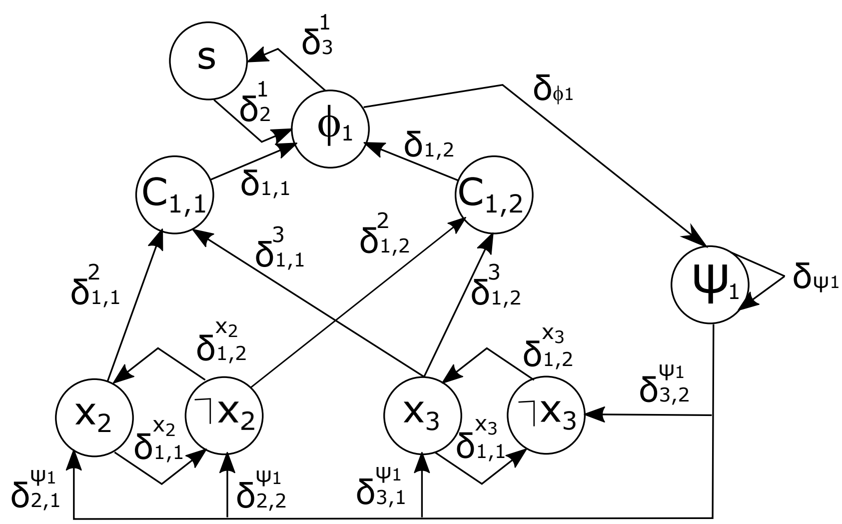

Figure 1: Graphical representation of the AAF $\langle A_1, D_1 \rangle = \beta(\phi_1)$, where $\phi_1 = (x_2 \vee x_3) \wedge (\neg x_2 \vee x_3)$ is obtained from the 3-CNF formula $\phi = (\neg x_1 \vee x_2 \vee x_3) \wedge (\neg x_1 \vee \neg x_2 \vee x_3) \wedge (x_1 \vee x_4 \vee x_5)$ by assigning $x_1 = \mathit{true}$.

It is easy to see that $\{s\}$ is not an ideal-set extension (and, thus, it is not an ideal extension) in the AAF $\langle A_i, D_i \rangle$ iff there exists a truth assignment t for

$x_1, \ldots, x_{i-1}, x_{i+1}, \ldots, x_n$ making ϕ_i evaluate to true. In fact, if such a t exists, it is possible to find a conflict-free set L of n_i arguments that defend ϕ_i from every attack from $C_{i,1}, \ldots, C_{i,k_i}$. Specifically, each argument in L is either a literal of the form x_m or $\neg x_m$ (with $x_m \in X(\phi_i)$), depending on whether $t(x_m)$= *true* or $t(x_m)$= *false*, respectively. Since ϕ_i attacks ψ_i, the set L is defended from the attacks from ψ_i, thus the set $L \cup \{\phi_i\}$ is an admissible extension in $\langle A_i, D_i \rangle$, and it is also a preferred extension, since no other argument is acceptable in it. Then, although $\{s\}$ is an admissible extension (and it is also a preferred extension), it is not an ideal-set extension or an ideal extension, as it is not contained in $L \cup \{\phi\}$. On the contrary, it is easy to see that, if there is no truth assignment t for $x_1, \ldots, x_{i-1}, x_{i+1}, \ldots, x_n$ making ϕ_i evaluate to true, there is no conflict-free set of arguments defending ϕ_i from the attacks from $C_{i,1}, \ldots, C_{i,k_i}$. In this case, it is easy to check that $\{s\}$ is the unique admissible extension, implying that it is also a preferred extension, an ideal-set extension and an ideal extension.

It is important to note that $\{s\}$ is not an ideal-set extension in $\langle A_0, D_0 \rangle$, since it is not even admissible due to $\delta_0 = (s, s)$.

Given this, a reduction to our problem from the $FP^{||NP}$-hard problem $sup(\phi)$ can be defined as follows. First, we construct an instance of P-$\text{EXT}_{\text{EX}}^{sem}(S)$, with $sem \in \{\texttt{ide}, \texttt{ids}\}$, by defining a prAAF $\mathcal{F}_\phi$ according to the above-described construction and choosing $S = \{s\}$. Then, we take the answer of P-$\text{EXT}_{\text{EX}}^{sem}(S)$, that is the probability $P^{sem}(S)$. Since, by construction, the probabilities of the possible AAFs compose a series of powers of 2 multiplied by 2^{-n}, from the binary encoding of $2^n \cdot P^{sem}(S)$ it is straightforward to extract the set $I = \{i | \{s\}$ is not an extension over $\langle A_i, D_i \rangle\}$. For what shown above, $I \setminus \{0\}$ coincides with the set of variables x_i in X such that there is a truth assignment for X, where x_i is assigned *true*, that makes ϕ satisfied. Hence, I is returned as the answer of $sup(\phi)$.

It is straightforward to see that the above described reduction is computable in polynomial time, since the PrAF $\mathcal{F}_\phi$ defined in the construction can be computed in polynomial time from the original Boolean $3CNF$ formula ϕ, and computing the encoding of $2^n \cdot P^{sem}(S)$ can be done in linear time w.r.t. the number of variables in X.

For $sem \in \{\texttt{pr}, \texttt{sst}\}$, a reduction from $sup(\phi)$ can be defined by constructing a prAAF $\mathcal{F}'_\phi = \langle A'_0 \cup \ldots \cup A'_n, D'_0 \cup \ldots \cup D'_n, \vec{\alpha'}, \vec{P'} \rangle$ with the following relationship with the $\mathcal{F}_\phi$ used above: (i) $A'_i = A_i$, for each $i \in [0..n]$; (ii) $D'_0 = D_0$ and $D'_i = D_i \setminus \{\delta_1^i = (s, \phi_i), \delta_2^i = (\phi_i, s)\}$; (iii) $\alpha_i = \langle A'_i, D'_i \rangle$, for each $i \in [0..n]$; (iv) $\vec{P'} = \vec{P}$.

For any $sem \in \{\texttt{pr}, \texttt{sst}\}$, it is easy to see that, for each $i \in [1..n]$, the set $S = \{s\}$ is an extension over $\langle A'_i, D'_i \rangle$ iff there is no truth assignment t for $x_1, \ldots x_{i-1}, x_{i+1}, \ldots, x_n$ making ϕ_i evaluate to true. Indeed, if no truth assignment t exists, $S = \{s\}$ is the unique preferred extension in $\langle A'_i, D'_i \rangle$, and is also a semi-stable extension, since it is the unique complete extension. Otherwise, if such a t exists, analogously to what seen before, it is possible to find a conflict-free set L of arguments defending ϕ_i from the attacks from $C_{i,1}, \ldots, C_{i,k_i}$, where each argument in L is either of the form x_m or $\neg x_m$, depending on whether $t(x_m) = \textit{true}$ or not. Thus, $S = \{s\}$ is not a preferred extension anymore, as $L \cup \{\phi_i, s\}$ is now a preferred extension. In the case that t exists, we have also that although $S = \{s\}$ is a complete extension, it is not a semi-stable extension since $S \cup S^+$ is not maximal: in fact, $(S' \cup S'^+) \supset (S \cup S^+)$ is a semi-stable extension, where $S' = L \cup \{\phi_i, s\}$. □

The first four columns of Table 1 provide a synopsis of the complexity of $\textsc{Ext}^{sem}(S)$ and P-$\textsc{Ext}^{sem}(S)$ for the different semantics in SEM (the results for P-$\textsc{Ext}^{sem}_{\textsc{ind}}(S)$ are taken from [16], and those for $\textsc{Ext}^{sem}(S)$ from [22, 23, 24, 25, 26, 27]). For now, the reader is asked to disregard the last two columns of Table 1, since they refer to the complexity of P-$\textsc{Ext}^{sem}(S)$ for different forms of prAAFs, that will be discussed in the next sections of this paper. Interestingly, the leftmost part of the table shows that solving P-$\textsc{Ext}^{sem}_{\textsc{ex}}(S)$ is never harder than P-$\textsc{Ext}^{sem}_{\textsc{ind}}(S)$. At a first glance, this is quite surprising, since EX allows correlations to be expressed while IND does not (since it assumes independence between arguments and defeats). However, this result must be read starting from the fact that IND is more succinct than EX. In fact, given a prAAF $\mathcal{F}^{\textsc{ex}}$ of form EX and a prAAF $\mathcal{F}^{\textsc{ind}}$ of form IND, such that $\mathcal{F}^{\textsc{ex}}$ encodes the same pdf over the possible AAFs as that implicitly encoded by $\mathcal{F}^{\textsc{ind}}$, the size of $\mathcal{F}^{\textsc{ex}}$ is exponentially larger than $\mathcal{F}^{\textsc{ind}}$.

We now focus our attention on P-$\textsc{Acc}^{sem}_{\textsc{ex}}(a)$. Its complexity, under the different semantics, is stated in the following three theorems.

Theorem 3. P-$\textsc{Acc}^{sem}_{\textsc{ex}}(a)$ *is in* FP *for* $sem = \texttt{gr}$.

Proof. Analogously to Theorem 1, the statement straightforwardly follows from the fact that the deterministic counterpart $\textsc{Acc}^{sem}(a)$ is polynomial time decidable for the form EX under the grounded semantics and that the number of possible AAFs and the size of their probabilities are linear w.r.t. the size of the input. □

Theorem 4. P-$\textsc{Acc}^{sem}_{\textsc{ex}}(a)$ *is in* $FP^{||NP}$ *for* $sem \in \{\texttt{ad}, \texttt{st}, \texttt{co}, \texttt{pr}, \texttt{ids}, \texttt{ide}\}$, *and in* $FP^{||\Sigma_p^2}$ *for* $sem = \texttt{sst}$.

Proof. The membership can be shown by reasoning analogously to the proof of Theorem 2. Thus, the membership in $FP^{||NP}$ holds for any semantics under which the deterministic counterpart $\textsc{Acc}^{sem}(a)$ over the form $\textsc{ex}$ is either NP-complete, or coNP-complete, or Θ_p^2-complete, that is for each $sem \in \{\mathtt{ad}, \mathtt{st}, \mathtt{co}, \mathtt{pr}, \mathtt{ids}, \mathtt{ids}\}$. For $sem = \mathtt{sst}$ the membership in $FP^{||\Sigma_p^2}$ derives from the fact that under this semantics $\textsc{Acc}^{sem}(a)$ is Σ_p^2-complete. □

Theorem 5. P-$\textsc{Acc}_{\text{EX}}^{sem}(a)$ *is* $FP^{||NP}$*-hard for* $sem \in \{\mathtt{ad}, \mathtt{st}, \mathtt{co}, \mathtt{pr}, \mathtt{ids}, \mathtt{ide}, \mathtt{sst}\}$

As done in the proof of Theorem 2, we show a reduction from *supremum*. Consider the instance of P-$\textsc{Acc}_{\text{EX}}^{sem}(a)$ consisting in the pair prAAF/argument $\langle \mathcal{F}, w\rangle$, where $\mathcal{F}$ is the prAAF of form $\textsc{ex}$ obtained by taking that used in the proof of Theorem 2 (see also Fig. 1) and augmenting the set of arguments A with the new argument w and the set of defeats D with the defeat $\delta_{sw} = (s, w)$. Moreover, each possible AAF $\langle A_i, D_i\rangle$ is modified so that the new argument w is put into A_i, and the defeat $\delta_{sw} = (s, w)$ into D_i. Reasoning as in the proof of Theorem 2, it can be seen that under any $sem \in \{\mathtt{ad}, \mathtt{st}, \mathtt{co}, \mathtt{pr}\}$, the satisfiability of ϕ_i implies the existence of a set L such that $S = \{w, \phi_i\} \cup L$ is an extension over $\langle A_i, D_i\rangle$. Specifically, L can be constructed by taking a truth assignment t making ϕ_i satisfied and then putting x_j into L iff $t(x_j)$ =*true*, and $\neg x_j$ otherwise. Now, it is easy to see that S is first of all conflict-free, and then that it is admissible (since all its arguments are defended from defeats coming from external arguments), and semi-stable (since all the arguments in A_i are either in S or attacked by S) and stable (since all the arguments outside S are attacked by some argument in S). Since S is stable, we also have that it is preferred and complete.

On the other hand, if ϕ_i is not satisfiable, there can be no extension S' over $\langle A_i, D_i\rangle$ that contains w. In fact, under $sem = \mathtt{ad}$, S' should contain also ϕ_i (which is the only argument defending w from the attack from s), and, in turn, this implies that S' should contain arguments (corresponding to the literals of the form x or $\neg x$, with $x \in X(\phi_i)$) defending ϕ_i from the attacks from arguments of the form $C_{i,j}$. This would contradict the unsatisfiability of ϕ_i, since S' would still correspond to a truth assignment satisfying ϕ_i. The same reasoning holds under the semantics $\mathtt{co}$, $\mathtt{sst}$ and $\mathtt{pr}$, since they require that S' is admissible, and for $sem = \mathtt{st}$, as if S' were stable, it would be also admissible.

Under any semantics $sem \in \{\mathtt{ad}, \mathtt{st}, \mathtt{co}, \mathtt{pr}\}$, what said so far means that, for each $i \in [1..n]$, the formula ϕ_i is satisfiable iff there exists an extension S over $\langle A_i, D_i\rangle$ with $w \in S$.

Hence, the answer to the instance of *supremum* can be eventually reconstructed from the binary encoding of the probability returned by a solver of the instance $\langle \mathcal{F}, w \rangle$ of P-$\text{ACC}^{sem}_{\text{EX}}(a)$, analogously to what done in the proof of Theorem 2.

For the semantics $\texttt{ids}$ and $\texttt{ide}$, the reduction can be shown to an instance $\langle \mathcal{F}, s \rangle$ of P-$\text{ACC}^{sem}_{\text{EX}}(a)$ over the same prAAF used in the proof of Theorem 2. In fact, in the proof of Theorem 2, we have already proved that, for each $i \in [1..n]$:

$i.$ if ϕ_i is satisfiable, then there is at least one preferred extension S over $\langle A_i, D_i \rangle$ such that $\phi_i \in S$;

$ii.$ if ϕ_i is not satisfiable, there is a unique preferred extension, consisting in $\{s\}$.

Now, $i.$ implies i': if ϕ_i is satisfiable, then there is at least one preferred extension that does not contain s (since s and ϕ are in conflict), hence any ideal-set or ideal extension does not contain s. Moreover, $ii.$ implies ii' : if ϕ_i is not satisfiable, then $\{s\}$ is the unique non-empty ideal-set extension and the unique ideal extension.

Hence, i' and ii' imply that ϕ_i is not satisfiable iff s belongs to an extension, under $sem \in \{\texttt{ids}, \texttt{ide}\}$. Therefore, an instance of *supremum* can be solved by determining, from the binary encoding of the probability returned by the solver of the so-constructed instance of P-$\text{ACC}^{sem}_{\text{EX}}(a)$, the indexes of the possible AAFs $\langle A_i, D_i \rangle$ such that s belongs to no extension over $\langle A_i, D_i \rangle$. □

4. A GENERAL PROBABILISTIC ABSTRACT ARGUMENTATION FRAMEWORK

In this section we introduce a new form of prAAF (called GEN) based on a paradigm for defining the pdf over the possible AAFs different from those of EX and IND. This paradigm is that of *world-set descriptors* and *world-set sets*, and will be explained in a short time. The first motivation for introducing GEN is that this paradigm allows correlations to be easily specified into and detected from the encoding of the pdf. This amenity makes GEN suitable for our aim of giving an insight on the sources of complexity of P-$\text{EXT}^{sem}(S)$ and P-$\text{ACC}^{sem}(a)$, as it will allow us to investigate the sensitivity of the complexity to different forms of correlations between arguments/defeats. This kind of analysis could not be done by working on EX and IND, since IND does not model the presence of correlations, while in EX correlations can be encoded, but inferring the correlations from the encoding of the pdf over the possible AAFs is a hard task, that would impede our sensitivity analysis. On the other hand, our new form prAAF can be also viewed as a distinguished contribution: we will see that its nice combination of expressiveness, compactness and aptitude to allow an easy specification of different correlations makes GEN a valid probabilistic framework for modeling uncertainty in argumentation in practical scenarios. As a matter of fact, the paradigm of *world-set descriptors* and *world-set sets* has been popular for decades in the context of probabilistic databases as an effective mechanism for associating probabilities to the data [18, 19].

The rest of this section is organized as follows. First, in Section 4.1, we recall the paradigm of *world-set descriptors* and *world-set sets*. Then, in Section 4.2, we introduce the prAAF GEN, that exploits this paradigm to specify pdfs over possible AAFs. Finally, in Section 4.4, we introduce several sub-classes of GEN, generated by different syntactic restrictions over the world-set descriptors in GEN. We will see that these restrictions correspond to different limitations on the type of correlations that can be imposed between arguments and defeats. Starting from this, in Section 5, we will be able to study how the complexity of P-EXT$^{sem}(S)$ and P-ACC$^{sem}(a)$ is affected by moving from one sub-class to another, thus also giving an insight on the sensitiveness of the complexity on the presence of different forms of correlations between arguments and defeats.

4.1. The probabilistic model of world-set descriptors and world-set sets

In this section, we recall the notion of *world-set descriptor (wsd)* and *world-set set*, that were shown in [18, 19] to be a succinct paradigm for specifying any probability distribution function over a finite set of scenarios (in particular, this paradigm was shown to be exponentially more succinct than the well-known paradigms of *world-set decompositions* [28] and ULDBs [29]).

Formally, each scenario is called *possible world*, and the set of possible worlds representing all the possible scenarios is modeled as follows. First, a finite set V of variables is given, where $\forall x \in V$, the domain of x is denoted as Dom_x and is finite. Then, each possible world is encoded as a total *valuation* (i.e., value assignment) of these variables. For instance, over the set of binary variables $V = \{x, y\}$, the following four possible worlds are defined: $w_1 = \{x \to 0, y \to 0\}$; $w_2 = \{x \to 0, y \to 1\}$; $w_3 = \{x \to 1, y \to 0\}$; $w_4 = \{x \to 1, y \to 1\}$.

In order to assign probabilities to the possible worlds, the variables are interpreted as independent random variables. This means assigning, for each $x \in V$ and $i \in Dom_x$, a probability $P(\{x \mapsto i\})$ to the assignment $x \mapsto i$, so that the probabilities of all the assignments for x sum up to one. In turn, the probability of a possible world w is the product of the probabilities of the assignments defining w. For instance, the probability of the possible world w_1 above is

$P(\{x \to 0\}) \cdot P(\{y \to 0\})$.

In [19], the set of variables, their domains, and probability distributions are represented in a world-table W consisting of all the possible triples (x, i, p), where x is a variable, $i \in Dom_x$, and p is the associated probability $P(\{x \mapsto i\})$.

The assignment-based mechanism for encoding the possible worlds can be naturally used to compactly represent sets of possible worlds, via *world-set descriptors* (*wsd*s) and *world-set sets* (ws-sets). Regarding the former, a *wsd* over V is a valuation of a subset of the variables in V. Thus, if a wsd d is *total* (i.e., it contains an assignment for every $x \in V$), then it identifies a unique possible world. Otherwise, it describes the set of the possible worlds identified by total functions f that can be obtained by extension of d, i.e., for all x on which d is defined, $d(x) = f(x)$. For instance, in the above-discussed case that $V = \{x, y\}$ and x, y are binary variables, the wsd $d = \{y \mapsto 0\}$ describes the set of possible worlds $\{w_1, w_3\}$, as both the total assignments identifying w_1 and w_3 agree on assigning 0 to y. If $d = \emptyset$, then d denotes the set of all possible worlds. The set of possible worlds identified by d is denoted as $\omega(d)$. Because of the independence of the variables, the aggregate probability of the worlds in $\omega(d)$ is $P(d) = \prod_{\{x \mapsto i\} \subseteq d} P(\{x \mapsto i\})$.

As regards *ws-sets*, a ws-set S is a set of ws-descriptors and represents the set of possible worlds resulting from the union of the sets of possible worlds represented by the ws-descriptors in S. The semantics of ws-sets is defined by using the (herewith overloaded) function ω extended to ws-sets, $\omega(S) = \cup_{d \in S} \omega(d)$. For instance, in the example running in this section, the ws-set $S = \{\{x \mapsto 0\}, \{x \mapsto 1, y \mapsto 1\}\}$ describes the set of possible worlds $\omega(S) = \{w_1, w_2, w_4\}$. Finally, the probability $P(S)$ of a ws-set S is the sum of the probabilities of the possible worlds in $\omega(S)$.

In the following, given a world table W, we will denote the variables occurring in W as $Var(W)$, and the set of wsds and ws-sets over $Var(W)$ as $wsd(W)$ and $WS(W)$, respectively.

4.2. *The prAAF* GEN *based on wsds and ws-sets*

In this section, we define the prAAF GEN, that exploits wsds and ws-sets for encoding the pdf over the possible AAFs. We show that this prAAF subsumes both EX and IND, and we also define some sub-classes of GEN of practical interest, that allow different forms of correlations to be imposed and that will be used in the following section to give a more insightful characterization of the complexity of P-EXT$^{sem}(S)$.

Definition 3 (prAAFs of form GEN). *A prAAF of form* GEN *is a tuple* $\mathcal{F} = \langle A, D, W, \lambda \rangle$*, where* A *is a set of arguments,* D *a set of defeats,* W *a world table, and* $\lambda : A \cup D \to WS(W)$ *is a function assigning every argument and defeat with a ws-set over* W*.*

We now formally explain the semantics of a prAAF $\mathcal{F}$ of type GEN (with "semantics", we mean the set $\mathcal{F}.\vec{\alpha}$ of possible AAFs encoded by $\mathcal{F}$, along with the probabilities $\mathcal{F}.\vec{P}$ assigned to these possible AAFs). We first introduce the notion of *support*: we say that a possible world $w \in \omega(W)$ *supports* the possible AAF $\alpha = \langle A', D' \rangle$ (denoted as $w \models \alpha$) if every argument/defeat $\sigma \in A' \cup D'$ is such that $w \in \omega(\lambda(\sigma))$, and there are no argument $a \in A \setminus A'$ such that $w \in \omega(\lambda(a))$ and no defeat $\delta \in (A' \times A') \setminus D'$ such that $w \in \omega(\lambda(\delta))$. In other words, the fact that w supports α means that:
1) the arguments of α are all and only the arguments a in A such that w is one of the possible worlds identified by $\lambda(a)$; and
2) the defeats of α are all and only the defeats δ that can be defined over pairs of arguments in α and that are such that w is one of the possible worlds identified by $\lambda(\delta)$.

Given this, the semantics of $\mathcal{F}$ is as follows. As regards $\mathcal{F}.\vec{\alpha}$, it contains every possible AAF $\alpha_i = \langle A_i, D_i \rangle$ such that:
i) $A_i \subseteq A$ and $D_i \subseteq D \cap (A_i \times A_i)$, and
ii) there is $w \in \omega(W)$ that supports α_i.
As regards $\mathcal{F}.\vec{P}$, every possible AAF $\alpha = \langle A_i, D_i \rangle \in \mathcal{F}.\vec{\alpha}$ is associated with the probability $P(\alpha_i) = \sum_{w \in \omega(W) \wedge w \models \alpha_i} P(w)$, that is the sum of the probabilities of the possible worlds in $\omega(W)$ supporting α_i.

The following example shows how the so-defined semantics of prAAFs of form GEN can be exploited to define prAAFs suitably modeling real-life situations.

Example 6. *Consider again the scenario of Example 1 and suppose that the lawyer considers only the following scenarios as possible:* $\alpha_1 = \langle \{a, b\}, \emptyset \rangle$, $\alpha_2 = \langle \{a, b, c\}, \emptyset \rangle$*, and* $\alpha_3 = \langle \{a, b, c\}, \{\delta_{ca}, \delta_{cb}\} \rangle$*, with probability* $p(\alpha_1) = 0.2$, $p(\alpha_2) = 0.4$ *and* $p(\alpha_3) = 0.4$*. To model this situation, we define a* $\mathcal{F} = \langle A, D, W, \lambda \rangle$ *such that* $A = \{a, b, c\}$*,* $D = \{\delta_{ca}, \delta_{cb}\}$*,* W *consists of the triples* $\langle x, true, 1 \rangle$*,* $\langle y, true, 0.8 \rangle$*,* $\langle y, false, 0.2 \rangle$*,* $\langle z, true, 0.5 \rangle$*, and* $\langle z, false, 0.5 \rangle$*, and* λ *is such that* $\lambda(a) = \lambda(b) = \{\{x \mapsto true\}\}$*,* $\lambda(c) = \{\{y \mapsto true\}\}$*,* $\lambda(\delta_{ca}) = \lambda(\delta_{cb}) = \{\{z \mapsto true\}\}$*. This way, the set of possible worlds is* $\omega(W) = \{w_1, w_2, w_3, w_4\}$*, where* $w_1 = \{x \mapsto true, y \mapsto true, z \mapsto true\}$*,*

$w_2 = \{x \mapsto true, y \mapsto true, z \mapsto false\}$*,* $w_3 = \{x \mapsto true, y \mapsto false, z \mapsto true\}$ *and* $w_4 = \{x \mapsto true, y \mapsto false, z \mapsto false\}$*, with probabilities* $p(w_1) = 1 \times 0.8 \times 0.5 = 0.4$*,* $p(w_2) = 0.4$*,* $p(w_3) = 0.1$ *and* $p(w_4) = 0.1$*. The set of possible AAFs is* $\mathcal{F}.\vec{\alpha} = \{\alpha_1, \alpha_2, \alpha_3\}$*, where* α_1*,* α_2*, and* α_3 *are defined above. It is easy to see that* α_1 *is supported by* w_3 *and* w_4*,* α_2 *is supported by* w_2*, and* α_3 *is supported by* w_1*. Hence,* $p(\alpha_1) = p(w_3) + p(w_4) = 0.2$*,* $p(\alpha_2) = p(w_2) = 0.4$ *and* $p(\alpha_3) = p(w_1) = 0.4$*.*

4.3. Expressiveness and succinctness of GEN

In this section we analyze the expressiveness and succinctness of GEN in comparison with the "traditional" forms EX and IND. The following theorem states that the form GEN can encode any prAAF $\mathcal{F}$ of form IND or EX within the same space as $\mathcal{F}$. The proof is constructive, and its steps define an algorithm that can be used for translating any prAAF of form EX or IND into one of form GEN.

Proposition 1. *For any prAAF* $\mathcal{F}$ *of form* EX *or* IND*, there is a prAAF* $\mathcal{F}^{\text{GEN}}$ *of form* GEN *such that* $\mathcal{F}^{\text{GEN}}.\vec{\alpha} = \mathcal{F}.\vec{\alpha}$ *and* $\mathcal{F}^{\text{GEN}}.\vec{P} = \mathcal{F}.\vec{P}$ *and* $\text{size}(\mathcal{F}^{\text{GEN}}) = O(\text{size}(\mathcal{F}))$.

Proof. We first prove that the form GEN can encode any prAAF of form EX within the same space. Given a prAAF $\mathcal{F}^{\text{EX}} = \langle A, D, \vec{\alpha}, \vec{P} \rangle$ of form EX, an equivalent prAAF $\mathcal{F}^{\text{GEN}} = \langle A, D, W, \lambda \rangle$ of form GEN can be defined as follows. We take a variable x over the domain $Dom_x = \{1, \dots, k\}$, where k is the cardinality of $\mathcal{F}^{\text{EX}}.\vec{\alpha}$. Then, for each $i \in [1..k]$, we put into W the row $(x, i, \mathcal{F}^{\text{EX}}.\vec{P}[i])$. Finally, for each argument or defeat σ, we define: $\lambda(\sigma) = \{\ \{x \mapsto i\} \mid \sigma$ is in the possible AAF $\alpha_i \in \mathcal{F}^{\text{EX}}.\vec{\alpha}\}$. It is easy to see that, since each value of x corresponds to a possible AAF encoded by $\mathcal{F}^{\text{EX}}$, the possible AAFs encoded in $\mathcal{F}^{\text{GEN}}$ are the same as $\mathcal{F}^{\text{EX}}$, and the probabilities assigned to these AAFs by the two prAAFs are the same. Moreover, it is easy to check that the size of $\mathcal{F}^{\text{GEN}}$ (that is, the sum of the size of W and of the ws-sets associated to each argument/defeat) is $O(\textit{size}(\mathcal{F}^{\text{EX}}))$.

We now show the encoding of a prAAF $\mathcal{F}^{\text{IND}} = \langle A, D, P_A, P_D \rangle$ of form IND into an equivalent prAAF $\mathcal{F}^{\text{GEN}} = \langle A, D, W, \lambda \rangle$ of form GEN. For each argument/defeat $\sigma \in A \cup D$, we first generate a boolean variable x_σ and put into W the row $(x_\sigma, \textit{true}, P(\sigma))$, where $P(\sigma)$ is the probability assigned to σ in $P_A \cup P_D$. Then, we define $\lambda(\sigma) = \{\{x_\sigma \mapsto \textit{true}\}\}$. The equivalence of the semantics of $\mathcal{F}^{\text{IND}}$ and the so-obtained $\mathcal{F}^{\text{GEN}}$ is straightforward, as well as the fact that

$size(\mathcal{F}^{\text{GEN}}) = O(size(\mathcal{F}))$. □

From the proposition above, it straightforwardly follows that GEN is *strictly more expressive* than IND (where no correlations can be specified) and *has the same expressiveness* as EX. In fact, EX can represent any pdf over the possible AAFs, like GEN. The point is that, while for any prAAF of form EX there is an equivalent one of form GEN using the "same" space (as stated in Proposition 1), the vice versa does not hold: given a prAAF of form GEN, encoding it in the form EX can require exponential space. An example of this compactness property characterizing GEN can be easily constructed from what stated in Proposition 1: starting from a prAAF $\mathcal{F}^{\text{IND}}$ of form IND, the equivalent prAAF $\mathcal{F}^{\text{GEN}}$ of form GEN has the same size as $\mathcal{F}^{\text{IND}}$, while it is straightforward to see that the size of the equivalent prAAF $\mathcal{F}^{\text{EX}}$ of form EX has exponential size w.r.t. $\mathcal{F}^{\text{IND}}$ and, thus, $\mathcal{F}^{\text{GEN}}$ (since the number of possible worlds is exponential in the number of arguments/defeats). The amenity of combining completeness and compactness is one of the strong points making GEN a suitable choice for modeling uncertainty in argumentation. We will elaborate more on this aspect in Section 6.

4.4. Restricted forms of GEN

We now introduce some syntactic restrictions over GEN, i.e., over the the ws-descriptors and the ws-sets that can be associated with the arguments/defeats. These restrictions will be useful in the context of our complexity analysis, since they correspond to allowing different forms of correlations between arguments/defeats. Thus, studying if and how the complexity varies when imposing a restriction gives an insight on the sources of complexity intrinsic in the different forms of correlations. When introducing each restriction, we will also discuss its relevance as a "practical" form of prAAF, by pointing out whether it corresponds to a form of prAAF already proposed in the literature, or whether it exhibits specific amenities that make it a candidate for being profitably used as the core of an argumentation system.

***Boolean prAAF* (BOOL).** The first restriction is that of allowing variables in the world table to take values from the boolean domain only. Specifically, a prAAF $\mathcal{F} = \langle A, D, W, \lambda \rangle$ of form GEN is said to be *boolean* (or, equivalently, of form BOOL) if, for each $x \in Var(W)$, Dom_x is the boolean domain. Intuitively, boolean prAAFs allow the occurrences of arguments and defeats within a dispute to be defined in terms of boolean formulas over a set of "elementary" independent probabilistic events. In fact, any "boolean" ws-set $\{wsd_1, \ldots, wsd_n\}$ encodes the DNF formula $c(wsd_1) \vee \ldots \vee c(wsd_n)$, where, in turn, for each $wsd_i = \{x_1 \mapsto$

true$, \ldots, x_k \mapsto$ *true*$, x_{k+1} \mapsto$ *false*$, \ldots, x_m \mapsto$ *false*$\}$, the term $c(wsd_i)$ is the conjunction $x_1 \wedge \ldots \wedge x_k \wedge \neg x_{k+1} \wedge \ldots \wedge \neg x_m$. This way of defining arguments and defeats is at the basis of one of the prAAF models analyzed in [16]. Observe that the prAAF defined in Example 6 is of form BOOL.

For the sake of simplicity, in the descriptors in boolean prAAFs, we use the literals x and $\neg x$ as shorthands for $x \mapsto$ *true* and $x \mapsto$ *false*, respectively. Thus, for instance, the ws-set $\{\ \{x \mapsto \textit{true}, y \mapsto \textit{false}\}, \{y \mapsto \textit{true}\}\ \}$ is written as $\{\ \{x, \neg y\}, \{y\}\ \}$.

***Monadic prAAF* (MON).** The second restriction is that of requiring that a prAAF is boolean and λ associates every argument/defeat with a ws-set consisting of a unique ws-descriptor containing only one variable assignment. Thus, a prAAF $\mathcal{F} = \langle A, D, W, \lambda \rangle$ is said to be *monadic* (or, equivalently, of form MON) iff it is boolean and, for each $\sigma \in A \cup D$, $\lambda(\sigma) = \{\{x \mapsto v\}\}$, where $x \in Var(W)$ and $v \in Dom_x$. Monadic prAAFs allow us to express co-existence of arguments/defeats and mutual exclusiveness between pairs of arguments and defeats in terms of *xor constraints*. Specifically, an xor constraint between two arguments/defeats σ_1 and σ_2 states that any possible AAF contains either σ_1 or σ_2, but not both. Basically, the co-existence of a set of arguments/defeats $\{\sigma_1, \ldots, \sigma_k\}$ is imposed by assigning the same ws-set to each of them, i.e., $\lambda(\sigma_1) = \ldots = \lambda(\sigma_k)$, where $\lambda(\sigma_1)$ is of the form $\{\ \{x\}\ \}$ or $\{\ \{\neg x\}\ \}$. On the other hand, an xor constraint over a pair σ_1, σ_2 of arguments/defeats is imposed by using the negation of the literal describing σ_1 as descriptor for σ_2 (for instance, if $\lambda(\sigma_1) = \{\ \{x\}\ \}$, then $\lambda(\sigma_2) = \{\ \{\neg x\}\ \}$). Observe that the prAAF defined in Example 6 is monadic (therein, the co-existence of the defeats δ_{ca} and δ_{cb} is imposed).

In Section 6, we will see that the characteristics of MON make it well-suited for being used in practice to support the specification of probabilities and correlations in a user-friendly way (even without the need to explicitly define wsds and ws-sets).

***Monadic prAAF with independent defeats* (IND-D).** This restriction over monadic prAAFs imposes that a variable occurring in the ws-set of any defeat cannot be used to describe some other arguments/defeats. This means that every defeat is described by a variable that is independent from those used to describe the other terms of the dispute. Formally, a monadic prAAF $\mathcal{F} = \langle A, D, W, \lambda \rangle$ is said to be *monadic with independent defeats* (or, equivalently, of form IND-D) iff, for each $\delta \in D$, given that $\lambda(\delta) = \{\{x \mapsto v\}\}$, there is no argument/defeat $\sigma \in A \cup D$ such that $\lambda(\sigma) = \{\{x \mapsto v'\}\}$, with $v, v' \in \{\textit{true, false}\}$.

On the one hand, prAAFs of form IND-D still allow us to impose the co-existence of arguments and xor constraints over pairs of arguments. On the other hand, in this form of prAAF, defeats are modeled as conditionally independent from one another, given the occurrence of the arguments over which they are defined. Remarkably, this case has been considered in some well-known frameworks in the literature, such as the framework addressed in [12], where defeats (but not the arguments) are assumed to be certain (obviously, "certainty" is a particular case of independence).

***Monadic prAAF with independent arguments* (IND-A).** Finally, we consider the restriction dual to the last one, where the independence refers to the variables describing the arguments instead of the defeats. Thus, a monadic prAAF $\mathcal{F} = \langle A, D, W, \lambda \rangle$ is said to be *monadic with independent arguments* (or, equivalently, of form IND-A) iff, for each $a \in A$, given that $\lambda(a) = \{\{x \mapsto v\}\}$, there is no argument/defeat $\sigma \in A \cup D$ such that $\lambda(\sigma) = \{\{x \mapsto v'\}\}$, with $v, v' \in \{$*true, false*$\}$.

It is easy to see that prAAFs of type IND-A allow us to impose the co-existence of defeats and xor constraints over pairs of defeats, while modeling the occurrences of different arguments within the dispute as independent events. Interestingly, the case that arguments are independent while defeats can be correlated is at the basis of the study in [5], where a framework for probabilistically modeling attacks (while arguments are certain) has been introduced.

5. COMPLEXITY RESULTS FOR P-EXT$^{sem}(S)$ AND P-ACC$^{sem}(a)$ OVER GEN AND ITS SUBCLASSES

We here provide the complexity characterization of P-EXT$^{sem}(S)$ and P-ACC$^{sem}(a)$ in the case that GEN or one of its subclasses is used as form of prAAF. For every $F \in \{\text{GEN}, \text{BOOL}, \text{MON}, \text{IND-A}, \text{IND-D}\}$, we will use the symbol F as subscript of P-EXT$^{sem}(S)$ and P-ACC$^{sem}(a)$ to denote the restrictions of these problems over prAAFs of form F. The results are summarized in Table 1 and Table 2. Given that these subclasses of GEN correspond to different restrictions on the type of correlations between arguments/defeats, studying what changes when moving from a subclass to another will allow us to give an insight on the sources of complexity of these problems (a conclusive discussion will be provided the end of this section, in Section 5.4).

5.1. Upper bound

We start by showing an upper bound for the complexity of P-$\text{EXT}^{sem}_{\text{GEN}}(S)$ and P-$\text{ACC}^{sem}_{\text{GEN}}(a)$. Specifically, we will show that these problems are in $FP^{\#P}$. Obviously, this upper bound holds for all the subclasses of GEN. In the next section, we will see that $FP^{\#P}$ is a tight bound (i.e., it is also a lower bound) for P-$\text{ACC}^{sem}(a)$ independently from the chosen semantics and the subclass of GEN, while for P-$\text{EXT}^{sem}(S)$ things are more intricate, since for some combinations subclass/semantics it is possible to obtain a polynomial bound.

Theorem 6. *For any $sem \in SEM$,* P-$\text{EXT}^{sem}_{\text{GEN}}(S)$ *and* P-$\text{ACC}^{sem}_{\text{GEN}}(a)$ *are in* $FP^{\#P}$.

Proof. We provide the detailed proof for P-$\text{EXT}^{sem}_{\text{GEN}}(S)$ in the case of the admissible semantics, and then show that the same result holds for the same problem under the other semantics, and for P-$\text{ACC}^{sem}_{\text{GEN}}(a)$ under any semantics.

In order to prove the membership in $FP^{\#P}$ for P-$\text{EXT}^{sem}_{\text{GEN}}(S)$ under $sem =$ ad we show that $P^{\text{ad}}(S)$ can be computed by the following polynomial time algorithm $\mathcal{A}$ with access to a $\#P$ oracle.

Observe that $P^{\text{ad}}(S)$ is a rational number whose denominator d is equal to $\prod_{x \in Var(W)} \text{LCD}\left(\{p_x | \langle x, i, p_x\rangle \in W\}\right)$, where LCD denotes the least common denominator of a set of rational numbers. Hence, $\mathcal{A}$ first computes d in polynomial time w.r.t. the size of $\mathcal{F}$ and then it calls a $\#P$ oracle for computing the numerator n of $P^{\text{ad}}(S)$. The oracle counts the number of accepting paths of a nondeterministic polynomial-time Turing machine M such that:

(i) M nondeterministically guesses a possible word w in $\omega(W)$. Recall that w uniquely determines a possible AAF α for $\mathcal{F}$ such that $w \models \alpha$.

(ii) At each leaf of the computation of the previous step, let w be the guessed possible world in $\omega(W)$, $P(w)$ its probability, and α the AAF supported by w. The computation tree is then split again $d \cdot P(w)$ times to reflect the probability of the possible world w (herein, $P(w)$ is a rational number whose denominator is d, and it can be computed in polynomial time w.r.t. the size of $\mathcal{F}$).

(iii) Finally, M checks in polynomial time if S is an admissible set of arguments in the AAF α.

It is easy to see that the number of accepting paths of M is $d \cdot \sum_{\substack{w \in \omega(\mathcal{F}) \wedge \alpha \in \alpha(\mathcal{F}) \wedge \\ w \models \alpha \wedge ext(\alpha, ad, S)}} P(w)$, that is the numerator n of $P^{\mathtt{ad}}(S)$. Finally, Algorithm $\mathcal{A}$ returns both n and d. This completes the $FP^{\#P}$-membership proof for P-$\textsc{Ext}^{sem}_{\text{GEN}}(S)$ under $sem = \mathtt{ad}$.

The membership of P-$\textsc{Ext}^{sem}_{\text{GEN}}(S)$ in $FP^{\#P}$ for $\mathtt{sem} \in \{\mathtt{st}, \mathtt{co}, \mathtt{gr}, \mathtt{sst}, \mathtt{pr}, \mathtt{ids}, \mathtt{ide}\}$ can be obtained by modifying algorithm $\mathcal{A}$ and the machine M described above so that at step (iii) of the computation of M a $\mathcal{C}$ oracle is invoked to decide an instance of the deterministic problem $\textsc{Ext}^{sem}(S)$, where $\mathcal{C}$ is the complexity class of $\textsc{Ext}^{sem}(S)$. Then, the statement follows from the fact that $FP^{\#P} = FP^{\#\mathcal{C}}$ for any complexity class in the polynomial hierarchy (see Section 2.5).

The very same reasoning proves the membership of P-$\textsc{Acc}^{sem}_{\text{GEN}}(a)$ in $FP^{\#P}$, independently from the semantics: it suffices to replace the oracle invoked at (iii) with an oracle over the class of complexity of the problem $\textsc{Acc}^{sem}(a)$. □

5.2. Lower bounds: hard cases

We first focus on P-$\textsc{Ext}^{sem}(S)$ and show that the upper bound $FP^{\#P}$ is also a lower bound (independently from the semantics) even in the case that the restriction IND-A (boolean, monadic, with independent arguments) of GEN is considered.

Theorem 7. *For any $sem \in SEM$,* P-$\textsc{Ext}^{sem}_{\text{IND-A}}(S)$ *is $FP^{\#P}$-hard.*

Proof. The $FP^{\#P}$-hardness for $sem \in \{\mathtt{co}, \mathtt{gr}, \mathtt{pr}, \mathtt{ids}, \mathtt{ide}, \mathtt{sst}\}$ is implied by the fact that P-$\textsc{Ext}^{sem}_{\text{IND}}(S)$ is $FP^{\#P}$-complete [16] and IND can be seen as a further restriction of IND-A. As regards the case that $sem \in \{\mathtt{ad}, \mathtt{st}\}$, we will prove that P-$\textsc{Ext}^{sem}_{\text{IND-A}}(S)$ is $FP^{\#P}$-hard by showing a reduction to our problem from the $\#P$-hard problem $\#P2CNF$, that is, the problem of counting the number of satisfying assignments of a CNF formula where each clause consists of exactly 2 positive literals.

Let $\phi = C_1 \wedge C_2 \wedge \ldots C_k$ be a $P2CNF$, where $X = \{x_1, \ldots, x_n\}$ is the set of its propositional variables. We define the prAAF $\mathcal{F}_\phi = \langle A, D, W, \lambda \rangle$ of form IND-A where:

- The set A consists of: (i) three arguments A_j with $j \in [1..3]$; and (ii) an argument C_i for each clause C_i appearing in ϕ;
- The relation D contains, for each clause C_i (with $i \in [1..k]$), a defeat $\delta^1_{C_i} = (A_1, C_i)$, a defeat $\delta^2_{C_i} = (A_2, C_i)$, and a defeat $\delta^3_{C_i} = (C_i, A_3)$.

- The world table W contains a triple $\langle x, true, 1\rangle$, and, for each $i \in [1..n]$, the two triples $\langle x_i, true, \frac{1}{2}\rangle$ and $\langle x_i, false, \frac{1}{2}\rangle$;
- Function λ is defined as follows: $i)$ for each $a \in A$, $\lambda(a) = \{x \mapsto true\}$; and $ii)$ for each $i \in [1..k]$, $\lambda(\delta^3_{C_i}) = \{x \mapsto true\}$; and $iii)$ for each $i \in [1..k]$, $\lambda(\delta^1_{C_i}) = ws_{i1} = \{x_j \mapsto true\}$, and $\lambda(\delta^2_{C_i}) = ws_{i2} = \{x_h \mapsto true\}$, where $C_i = x_j \vee x_h$.

The graphical representation of the prAAF corresponding to a formula ϕ is reported in Figure 2.

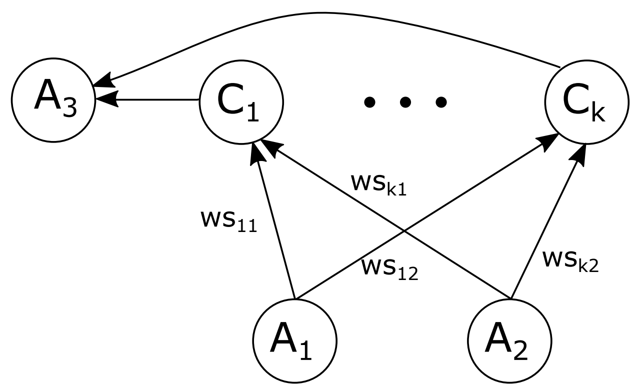

Figure 2: Graphical representation of the prAAF corresponding to a formula ϕ

We consider the bijection β from the truth assignments for the propositional variables $x_1, \ldots, x_n$ and the possible worlds $w \in \omega(W)$ such that, for each truth assignment t for the propositional variables $x_1, \ldots, x_n$, $\beta(t) = w$ is the possible world defined as follows: $i)$ x is assigned $true$; and $ii)$ x_i is assigned $t(x_i)$ for each $i \in [1..n]$.

It is easy to see that ϕ evaluates to true under t iff $\{A_1, A_2, A_3\}$ is an admissible extension in the possible AAFs α for $\mathcal{F}_\phi$ such that $\beta(t) \models \alpha$. Indeed, $\{A_1, A_2, A_3\}$ is an admissible extension in a possible AAF $\alpha = \langle A', D'\rangle$ iff for each argument C_i with $i \in [1..k]$ at least one of the defeats $\delta^1_{C_i}$ or $\delta^2_{C_i}$ belongs to D'. This implies that $\{A_1, A_2, A_3\}$ is an admissible extension in all and only the possible AAFs α such that there exists a truth assignment t for $x_1, \ldots, x_n$ which makes ϕ evaluate to true and such that $\beta(t) \models \alpha$.

Moreover, it is easy to see that for each possible AAF α for $\mathcal{F}_\phi$, $\{A_1, A_2, A_3\}$ is an admissible extension iff $\{A_1, A_2, A_3\}$ is a stable extension. Hence, from now on we continue the proof considering the admissible semantics only, as the same reasoning apply to the case of the stable semantics.

Given this, a reduction to our problem from the $\#P$-hard problem $\#P2CNF$ can be defined as follows. Given an instance ϕ of $\#P2CNF$, we construct an instance of P-$\text{EXT}^{sem}_{\text{IND-A}}(S)$ by defining a prAAF $\mathcal{F}_\phi$ of the form IND-A as in the above-described construction and choosing $S = \{A_1, A_2, A_3\}$. Next we return as number of satisfying assignments the value $2^n \cdot P^{\mathtt{ad}}(S)$, where n is the number of propositional variables of ϕ. It is straightforward to see that the above described reduction is computable in polynomial time, since the prAAF $\mathcal{F}_\phi$ defined in the construction can be computed in polynomial time from a positive $2CNF$ formula ϕ and the formula $2^n \cdot P^{\mathtt{ad}}(S))$ can be computed in polynomial time given that the value $P^{\mathtt{ad}}(S)$ is the value returned by the oracle.

We now prove that the previously defined reduction is sound. Since $S = \{A_1, A_2, A_3\}$ is an admissible extension in all and only the possible AAFs α such that there exists a truth assignment t for $x_1, \ldots, x_n$ which makes ϕ evaluate to true and such that $\beta(t) \models \alpha$, in order to show that the reduction is sound, it suffices to prove that the number $\#\phi$ of satisfying assignments of ϕ is equal to $2^n \cdot P^{\mathtt{ad}}(S))$.

We now prove that $\#\phi = 2^n \cdot P^{\mathtt{ad}}(S))$. First, observe that the probability of every possible world $w \in \omega(W)$ is equal to $\frac{1}{2^n}$. Next, note that $P^{\mathtt{ad}}(S)$ is the sum of the probabilities of the possible worlds $w \in \omega(W)$ such that there is a possible AAF $\alpha \in \alpha(\mathcal{F})$ with $w \models \alpha$ in which S is an admissible extension, that is, the sum of the probabilities of possible worlds w such that $\tau = \beta^{-1}(w)$ is a truth assignment for the variables of ϕ making it satisfied. Hence, $P^{\mathtt{ad}}(S) = \frac{1}{2^n} \cdot \#\phi$, which implies that $\#\phi = 2^n \cdot P^{\mathtt{ad}}(S))$.

Hence, the above described reduction from the $\#P$-hard problem $\#P2CNF$ to P-$\text{EXT}^{sem}_{\text{IND-A}}(S)$, with *sem* $\in \{\mathtt{ad}, \mathtt{st}\}$, is a Cook reduction, and this suffices to prove that our problem is $FP^{\#P}$-hard, since a problem is $FP^{\#P}$-hard iff it is $\#P$-hard (see Section 2.5 for a discussion of this property). $\square$

The results stated above, along with the hierarchy of the restrictions at the basis of BOOL, MON, and IND-A, straightforwardly entail the following result.

Corollary 1. *For any* $sem \in SEM$, P-$\text{EXT}^{sem}_{\text{GEN}}(S)$, P-$\text{EXT}^{sem}_{\text{BOOL}}(S)$, P-$\text{EXT}^{sem}_{\text{MON}}(S)$, *and* P-$\text{EXT}^{sem}_{\text{IND-A}}(S)$ *are* $FP^{\#P}$*-complete.*

The results stated so far give a complete picture of the complexity of P-$\text{EXT}^{sem}(S)$ for GEN and all its restrictions but IND-D. We now focus on IND-D and report a first result. In particular, from the fact shown in [16] that P-$\text{EXT}^{sem}_{\text{IND}}(S)$ is $FP^{\#P}$-complete for every $sem \in \{\mathtt{co}, \mathtt{gr}, \mathtt{pr}, \mathtt{ids}, \mathtt{ide}, \mathtt{sst}\}$ (see also Table 1), we have that P-$\text{EXT}^{sem}_{\text{IND-D}}(S)$ under these semantics is still $FP^{\#P}$-complete, since IND can be seen as a further restriction of IND-D, where both arguments and defeats are (conditionally) independent from one another.

Fact 1. *For any* $sem \in \{\mathtt{co}, \mathtt{gr}, \mathtt{pr}, \mathtt{ids}, \mathtt{ide}, \mathtt{sst}\}$*,* P-EXT$^{sem}_{\text{IND-D}}(S)$ *is* $FP^{\#P}$*-complete.*

Only the case of P-EXT$^{sem}_{\text{IND-D}}(S)$ with $sem \in \{\mathtt{ad}, \mathtt{st}\}$ is not included in the characterization provided so far: it will be addressed in the following section, where we will prove that is tractable.

We conclude this section concerning the discussion of lower bounds symptomatic of intractability by turning our attention to P-ACC$^{sem}(a)$. In this case, we have that the $FP^{\#P}$ upper bound stated in the previous section is tight for every semantics and for every subclass of GEN. This can be viewed as a corollary of Theorem 6, since it straightforwardly derives from it and from the fact that P-ACC$^{sem}(a)$ is already $FP^{\#P}$-hard over IND [17].

Corollary 2. *For every* $sem \in SEM$ *and* $F \in \{$GEN, BOOL, MON, IND-A, IND-D $\}$*,* P-ACC$^{sem}_{F}(a)$ *is* $FP^{\#P}$*-complete.*

5.3. Upper bounds: tractable cases

In this section, we show that P-EXT$^{sem}(S)$ can be solved in polynomial time in the case that the prAAF is of form IND-D and *sem* $\in \{\mathtt{ad}, \mathtt{st}\}$. We exploit the same idea used in [30] and in [16], that is that of associating sets of arguments with propositional formulas such that the set is an extension under a certain semantics iff the set is a model of the formula.

In brief, the reason behind the tractability of this case is that the fact that S is admissible (resp., stable) can be expressed as a probabilistic event $E_{\mathtt{ad}}(S)$ (resp., $E_{\mathtt{st}}(S)$) whose probability $P(E_{\mathtt{ad}}(S))$ (resp., $P(E_{\mathtt{st}}(S))$) is equal to $P^{\mathtt{ad}}(S)$ (resp., $P^{\mathtt{st}}(S)$) and can be computed in polynomial time.

Before providing our results, we introduce some notation. Let σ be an argument or a defeat. We denote as $Lit(\sigma)$ the literal appearing in the ws-set of σ, that is, $Lit(\sigma) = x$ (resp., $Lit(\sigma) = \neg x$) iff $\lambda(\sigma) = \{\{x\}\}$ (resp., $\lambda(\sigma) = \{\{\neg x\}\}$) – we recall that we are considering prAAFs of form IND-D, where the "monadic" restriction holds.

We now model the fact that a set S of arguments is an admissible (resp., stable) extension as a "complex" event $E_{\mathtt{ad}}(S)$ (resp., $E_{\mathtt{st}}(S)$), constructed from the "basic" events $Lit(a)$ and $Lit(\delta)$ (where a and δ are arguments and defeats of the prAAF).

Definition 4 **($E_{\mathtt{ad}}(S)$).** *Given a prAAF $\mathcal{F} = \langle A, D, W, \lambda\rangle$ of form* IND-D *and a set $S \subseteq A$ of arguments, the event that S is an admissible extension is $E_{\mathtt{ad}}(S) = e_1(S) \wedge e_2(S) \wedge e_3(S)$ where:*

- $e_1(S) = \bigwedge_{a \in S} Lit(a)$
- $e_2(S) = \bigwedge_{\substack{\delta = (a,b) \in D \\ \wedge a \in S \wedge b \in S}} \neg Lit(\delta)$
- $e_3(S) = \bigwedge_{d \in A \setminus S} (e_{31}(S,d) \vee e_{32}(S,d) \vee e_{33}(S,d))$ *where:*
  - $e_{31}(S,d) = \neg Lit(d)$
  - $e_{32}(S,d) = Lit(d) \;\wedge \bigwedge_{\substack{\delta = (d,b) \in D \\ \wedge b \in S}} \neg Lit(\delta)$
  - $e_{33}(S,d) = Lit(d) \;\wedge \bigvee_{\substack{\delta = (d,b) \in D \\ \wedge b \in S}} Lit(\delta) \;\wedge \bigvee_{\substack{\delta = (a,d) \in D \\ \wedge a \in S}} Lit(\delta)$

$\square$

The rationale of the expression in Definition 4 is as follows. The event $E_{\mathtt{ad}}(S)$, encoding the fact that S is an admissible extension, occurs iff $e_1(S)$, $e_2(S)$ and $e_3(S)$ simultaneously occur, where:

(i) $e_1(S)$ is the event that all of the arguments in S occur;

(ii) $e_2(S)$ is the event that no defeat (a, b) (with $a, b \in S$) occurs, meaning that S is conflict-free;

(iii) $e_3(S)$ is the event that for all the arguments d in $A \setminus S$, exactly one of the following events occurs:

- $e_{31}(S,d)$: d does not occur, or
- $e_{32}(S,d)$: d occurs and no defeat (d,b) such that $b \in S$ occurs.
- $e_{33}(S,d)$: d occurs, there is at least one argument $b \in S$ such that (d,b) occurs, and there is at least one argument $a \in S$ such that (a,d) occurs.

Given this, it is easy to see that, given a prAAF $\mathcal{F} = \langle A, D, W, \lambda \rangle$ of form IND-D, for each possible world $w \in \omega(W)$, $E_{\texttt{ad}}(S)$ is true w.r.t. w iff S is an admissible extension in the possible AAF α supported by w. Hence, since:

- $P(E_{\texttt{ad}}(S))$ is given by the sum of the probabilities of the possible worlds $w \in \omega(W)$ such that $E_{\texttt{ad}}(S)$ is true [3], and
- $P^{\texttt{ad}}(S)$ is given by the sum of the probabilities of the possible AAFs α such that S is an admissible extension in α, and $P(\alpha)$ is given by the sum of the probabilities of the possible worlds in $\omega(W)$ supporting α,

then it holds that $P(E_{\texttt{ad}}(S)) = P^{\texttt{ad}}(S)$.

Similar to the case of the admissible semantics, the fact that a set S of arguments is a stable extension is expressed by the probabilistic event $E_{\texttt{st}}(S) = e_1(S) \wedge e_2(S) \wedge e'_3(S)$, where:

- $e_1(S)$ and $e_2(S)$ are the events introduced in Definition 4, and
- $e'_3(S)$ is the event that for all the arguments d in $A \setminus S$, either d does not occur, or d occurs and it is defeated by S.

The following definition provides the formalization of event $E_{\texttt{st}}(S)$.

Definition 5 ($E_{\texttt{st}}(S)$)**.** *Given a prAAF $\mathcal{F} = \langle A, D, W, \lambda \rangle$ and a set $S \subseteq A$, the event that S is a stable extension is defined as:*

$$E_{\texttt{st}}(S) = e_1(S) \wedge e_2(S) \wedge e'_3(S)$$

where $e_1(S)$ and $e_2(S)$ are the events introduced in Definition 4, and

$$e'_3(S) = \bigwedge_{d \in A \setminus S} e_{31}(S, d) \vee e_{32}(S, d)$$

where $e_{31}(S, d) = \neg Lit(d)$ and $e_{32}(S, d) = Lit(d) \wedge \bigvee_{\substack{\delta = (a, d) \in D \\ \wedge a \in S}} Lit(\delta)$. □

Reasoning similarly to the case of the admissible semantics, it is easy to see that $P(E_{\texttt{st}}(S)) = P^{\texttt{st}}(S)$.

[3] We say that $E_{\texttt{ad}}(S)$ is true w.r.t. a possible world w iff the formula defining $E_{\texttt{ad}}(S)$ is true in the case that every variable x appearing in $E_{\texttt{ad}}(S)$ is replaced with the value assigned to x in w.

Theorem 8. *For any $sem \in \{\mathtt{ad}, \mathtt{st}\}$,* P-EXT$^{sem}_{\text{IND-D}}(S)$ *is in FP.*

Proof. We first prove the tractability for P-EXT$^{sem}_{\text{IND-D}}(S)$ under $sem = \mathtt{ad}$ by showing that $P(E_{\mathtt{ad}}(S))$ can be computed in FP.

We show that one of the two following cases holds:

i) there is no possible world supporting a possible AAF $\alpha_i = \langle A_i, D_i \rangle$ where $S \subseteq A_i$, or

ii) it is possible to rewrite $E_{\mathtt{ad}}(S)$ into a boolean expression $\textsc{rew}(E_{\mathtt{ad}}(S))$ equivalent to $E_{\mathtt{ad}}(S)$ having the following form:

$$\begin{array}{c} x_1 \wedge \ldots \wedge x_n \wedge \neg x_{n+1} \wedge \ldots \wedge \neg x_{n+m} \wedge \\ (E_1 \wedge \ldots \wedge E_k) \wedge \\ \big((x_{n+m+1} \wedge E_{k+1}) \vee \big(\neg x_{n+m+1} \wedge E'_{k+1}\big)\big) \wedge \ldots \\ \ldots \wedge \big((x_{n+m+l} \wedge E_{k+l}) \vee \big(\neg x_{n+m+l} \wedge E'_{k+l}\big)\big) \end{array} \tag{4}$$

where:

– for each $i, j \in [1..n+m+l]$, with $i \neq j$, we have $x_i \neq x_j$;
– for each $i \in [1..k+h+l]$, E_i (resp. E'_i) is a conjunction of boolean formulas, i.e., $E_i = E_{i1} \wedge \ldots \wedge E_{ih}$ (resp., $E'_i = E'_{i1} \wedge \ldots \wedge E'_{ih'}$), where every E_{ij} (resp. E'_{ij}) is a boolean formula of the form $E^*_{ij} \vee \neg E^*_{ij} \wedge E^{\#}_{ij}$. Herein, each E^*_{ij} and each $E^{\#}_{ij}$ are boolean formulas of the form $E^*_{ij} = \bigwedge_{i'=1}^{r} y_{i'} \wedge \bigwedge_{i'=r+1}^{r+r'} \neg y_{i'}$ and $E^{\#}_{ij} = \bigvee_{j'=1}^{s} z_{j'} \vee \bigvee_{j'=s+1}^{s+s'} \neg z_{j'}$, where each variable $y_{i'}$ (with $i' \in [1..r+r']$) and each variable $z_{j'}$ (with $j' \in [1..s+s']$) are fresh variables having no other occurrences in the whole formula $\textsc{rew}(E_{\mathtt{ad}}(S))$.

As regards case i), it is easy to see that there is no possible world supporting a possible AAF $\alpha_i = \langle A_i, D_i \rangle$ where $S \subseteq A_i$ if and only if there exist $a, b \in S$ such that $Lit(a) = x_i$ and $Lit(b) = \neg x_i$. In this case $P(E_{\mathtt{ad}}(S)) = 0$.

As regards case ii), in what follows we show how to obtain a boolean expression $\textsc{rew}(E_{\mathtt{ad}}(S))$ equivalent to $E_{\mathtt{ad}}(S)$ of the form described in Equation (4) by rewriting the conjunction $e_1(S) \wedge e_2(S) \wedge e_3(S)$ of Definition 4, that is $\textsc{rew}(E_{\mathtt{ad}}(S)) = \textsc{rew}(e_1(S)) \wedge \textsc{rew}(e_2(S)) \wedge \textsc{rew}(e_3(S))$.

First, $e_1(S)$ is rewritten as $\textsc{rew}(e_1(S)) = x_1 \wedge \ldots \wedge x_{n'} \wedge \neg x_{n'+1} \wedge \ldots \wedge \neg x_{n'+m'}$, where $x_i \neq x_j$, for each i, j in $[1..n'+m']$ by removing duplicate literals in the formula $e_1(S) = \bigwedge_{a \in S} Lit(a)$.

Moreover, it is easy to see that $e_2(S)$ can be rewritten in a expression $\text{REW}(e_2(S)) = x_{n'+m'+1} \wedge \ldots \wedge x_{n'+m'+h'} \wedge \neg x_{n'+m'+h'+1} \wedge \ldots \wedge \neg x_{n'+m'+h'+k'}$ where for each $i, j \in [n'+m'+1..n'+m'+h'+k']$ it holds that $x_i \neq x_j$ if $i \neq j$ since

$$e_2(S) = \bigwedge_{\substack{\delta = (a,b) \in D \\ \wedge a \in S \wedge b \in S}} \neg Lit(\delta)$$

and every $Lit(\delta)$ where $\delta = (a, b) \in D$ and $a \in S \wedge b \in S$ mentions a different variable. Moreover, there is no variable mentioned in $\text{REW}(e_2(S))$ that is mentioned in $\text{REW}(e_1(S))$. Therefore $\text{REW}(e_1(S)) \wedge \text{REW}(e_2(S))$ is an expression of the form $x_1 \wedge \ldots \wedge x_n \wedge \neg x_{n+1} \wedge \ldots \wedge \neg x_{n+m}$ where for each $i, j \in [1..n+m]$ $x_i \neq x_j$ if $i \neq j$.

Furthermore, we define the rewriting $\text{REW}(e_3(S))$ as follows. First, for each d, observe that $e_3(S, d)$ is defined as follows (see Definition 4):

$$e_3(S,d) = e_{31}(S,d) \vee e_{32}(S,d) \vee e_{33}(S,d) = \neg Lit(d) \vee (Lit(d) \wedge (E^* \vee (\neg E^* \wedge E^{\#})),$$

where $E^* = \bigwedge_{\substack{\delta=(d,b)\in D \\ \wedge b \in S}} \neg Lit(\delta)$ and $E^{\#} = \bigvee_{\substack{\delta = (a,d) \in D \\ \wedge a \in S}} Lit(\delta)$. Morevoer, since a variable mentioned in a literal associated ot a defeat is not mentioned in the literal of any other argument/defeat then it holds that each variable x mentioned in either E^* and $E^{\#}$ satisfies the following conditions:

- x is mentioned only once in E^* and $E^{\#}$, and
- x does not appear elsewhere in $\text{REW}(E_{\mathtt{ad}}(S))$.

Moreover, $e_3(S, d)$ is further rewritten as follows:

- (i) if $Lit(d) = x$ then $\text{REW}(e_3(S, d)) = (x \wedge E) \vee (\neg x \wedge E')$ where $E = true$ and $E' = E^* \vee \neg E^* \wedge E^{\#}$,
- otherwise (that is $Lit(d) = \neg x$), $\text{REW}(e_3(S, d)) = (x \wedge E) \vee (\neg x \wedge E')$ where $E = E^* \vee \neg E^* \wedge E^{\#}$ and $E' = true$.

Finally, we set $\text{REW}(e_3(S))$ equal to $\bigwedge_{d \in A \setminus S} \text{REW}(e_3(S, d))$.

Next, $\text{REW}(e_3(S))$ is further simplified grouping together conjunctions having the same literal as first argument. That is, denoting as $X(A \setminus S)$ the set of x such that $d \in A \setminus S$ exists with $Lit(d) = x$ or $Lit(d) = \neg x$, $e_3(S)$ is rewritten so that it contains, for each $x \in X(A \setminus S)$, a unique term of the form $(x \wedge (E_1 \wedge E_2 \wedge$

$\ldots \wedge E_k)) \vee (\neg x \wedge (E'_1 \wedge E'_2 \wedge \ldots E'_k))$, where k is the number of arguments d in $A \setminus S$ such that $Lit(d) = x$ or $Lit(d) = \neg x$.

Moreover, if some x in $X(A \setminus S)$ is such that x or $\neg x$ are mentioned in $\text{REW}(e_1(S)) \wedge \text{REW}(e_2(S))$, then $\text{REW}(e_3(S))$ is further simplified as follows:

- if x is mentioned in $\text{REW}(e_1(S)) \wedge \text{REW}(e_2(S))$, the formula $(x \wedge (E_1 \wedge E_2 \wedge \ldots \wedge E_k)) \vee (\neg x \wedge (E'_1 \wedge E'_2 \wedge \ldots E'_k))$ is replaced by $(E_1 \wedge E_2 \wedge \ldots \wedge E_k)$, since $\text{REW}(e_1(S)) \wedge \text{REW}(e_2(S)) \wedge (\neg x \wedge (E'_1 \wedge E'_2 \wedge \ldots E'_k))$ is not satisfiable and the literal x appears in $\text{REW}(e_1(S)) \wedge \text{REW}(e_2(S))$,

- otherwise, if $\neg x$ is mentioned in $\text{REW}(e_1(S)) \wedge \text{REW}(e_2(S))$, the formula $(x \wedge (E_1 \wedge E_2 \wedge \ldots \wedge E_k)) \vee (\neg x \wedge (E'_1 \wedge E'_2 \wedge \ldots E'_k))$ is replaced by $(E'_1 \wedge E'_2 \wedge \ldots E'_k))$, since $\text{REW}(e_1(S)) \wedge \text{REW}(e_2(S)) \wedge (x \wedge (E_1 \wedge E_2 \wedge \ldots \wedge E_k))$ is not satisfiable and the literal $\neg x$ appears in $\text{REW}(e_1(S)) \wedge \text{REW}(e_2(S))$.

From the definition of $\text{REW}(E_{\mathtt{ad}}(S))$ it straightforwardly follows that $\text{REW}(E_{\mathtt{ad}}(S))$ is a boolean expression of the form (4) that is equivalent to $E_{\mathtt{ad}}(S)$, whose size is polynomially bounded by the size of the prAAF and is constructed from $E_{\mathtt{ad}}(S)$ in time polynomial w.r.t. the size of the prAAF (as every rewriting operation applied to $E_{\mathtt{ad}}(S)$ can be done in polynomial time w.r.t. the size of the prAAF).

Now, we conclude the proof by showing that the probability of a boolean expression of the form (4) can be evaluated in polynomial time w.r.t. the size of the formula.

Specifically, it is easy to see that the probability of a formula of the form (4), that is

$$P\left(\begin{array}{c} x_1 \wedge \ldots \wedge x_n \wedge \neg x_{n+1} \wedge \ldots \wedge \neg x_{n+m} \wedge (E_1 \wedge \ldots \wedge E_k)) \wedge \\ \big((x_{n+m+1} \wedge E_{k+1}) \vee \big(\neg x_{n+m+1} \wedge E'_{k+1}\big)\big) \wedge \ldots \\ \ldots \wedge \big((x_{n+m+l} \wedge E_{k+l}) \vee \big(\neg x_{n+m+l} \wedge E'_{k+l}\big)\big) \end{array}\right)$$

corresponds to the following expression

$$\begin{array}{c} P(x_1) \times \ldots \times P(x_n) \times (1 - P(x_{n+1})) \times \ldots \\ \times (1 - P(x_{n+m})) \times (P(E_1) \times \ldots \times P(E_k))) \times \\ \Big((P(x_{n+m+1}) \times P(E_{k+1})) + \Big((1 - P(x_{n+m+1})) \times P(E'_{k+1})\Big)\Big) \times \\ \ldots \times \Big((P(x_{n+m+l}) \times P(E_{k+l})) + \Big((1 - P(x_{n+m+l})) \times P(E'_{k+l})\Big)\Big) \end{array} \tag{5}$$

where, for each $i \in [1..n+m+l]$, $P(x_i)$ is the probability assigned to $x_i \mapsto true$ in the world table W and, for each $j \in [1..k+l]$, $P(E_j)$ (resp. $P(E'_j)$, is defined as follows.

Recall that $E_j = E_{j1} \wedge \ldots \wedge E_{jh}$ (resp. $E'_j = E'_{j1} \wedge \ldots \wedge E'_{jh}$), where for each $l \in [1..h]$ it holds that $E_{jl} = E^*_{jl} \vee \neg E^*_{jl} \wedge E^\#_{jl}$ (resp. $E'_{jl} = E'^*_{jl} \vee \neg E'^*_{jl} \wedge E'^\#_{jl}$). Hence, $P(E_j)$ is equal to $P(E_{j1}) \times \ldots \times P(E_{jh})$ (as every E_{jl} mentions different variables) and $P(E_{ij})$ is equal to $P(E^*_{jl}) + (1 - P(E^*_{jl})) \times P(E^\#_{jl})$ (as E^*_{jl} and $\neg E^*_{jl}$ are mutually exclusive events and $E^\#_{jl}$ mentions only variables not appearing in E^*_{jl}).

Finally, we define $P(E^*_{jl})$ and $P(E^\#_{jl})$ as follows. Recall that

$$E^*_{jl} = \bigwedge_{\substack{\delta=(d,b)\in D\\ \wedge b \in S}} \neg Lit(\delta) = \bigwedge_{i=1}^{r} y_i \wedge \bigwedge_{j=r+1}^{r+r'} \neg y_i$$

and

$$E^\#_{jl} = \bigvee_{\substack{\delta=(a,d)\in D\\ \wedge a \in S}} Lit(\delta) = \bigvee_{i=1}^{s} z_i \vee \bigvee_{j=s+1}^{s+s'} \neg z_i,$$

where the variables mentioned in E^*_{jl} and $E^\#_{jl}$ are all distinct and do not appear elsewhere in $\text{REW}(E_{\mathtt{ad}}(S))$. Hence it is easy to see that,

$$P(E^*_{jl}) = \prod_{i=1}^{r} P(y_i) \times \prod_{j=r+1}^{r+r'} (1 - P(y_i))$$

and

$$P(E^\#_{jl}) = 1 - \prod_{i=1}^{s} (1 - P(z_i)) \times \prod_{j=s+1}^{s+s'} P(z_i)$$

Hence, by substituting the above defined values of $P(E^*_{jl})$ and $P(E^\#_{jl})$ into Equation (5) we obtain an expression that can be computed in polynomial time w.r.t. the size of the prAAF. $\square$

5.4. Discussion of the results

The results obtained so far about the complexity of P-EXT$^{sem}(S)$ and P-ACC$^{sem}(a)$ are summarized in Table 1 and Table 2. Analogously to what observed in Section 3, it is worth noting that, except for the tractable cases, solving P-EXT$^{sem}(S)$ and P-ACC$^{sem}(a)$ over GEN and its subclasses is more complex than over EX (as $FP^{\#P}$ contains $FP^{\mathcal{C}}$, for any $\mathcal{C}$ in the polynomial hierarchy).

This result must be read by keeping in mind that GEN is exponentially more succinct than EX, thus, intuitively, the cost of solving a problem over EX may benefit from a "discount" consisting in a "level of exponentiality", compared with solving the same problem encoded over GEN.

Another interesting point is that the presence of different forms of correlations has an impact on the complexity of P-$\text{EXT}^{sem}(S)$ under the admissible and stable semantics. In fact, in these cases, P-$\text{EXT}^{sem}(S)$ can be polynomially solved when defeats are independent, and correlations (in terms of co-existence and xor constraints) are expressed between arguments (see column IND-D in Table 1). Interestingly, switching things, by making arguments independent and defeats correlated, makes the complexity explode. Observe this behavior is not implied by the fact that the deterministic version $\text{EXT}^{sem}(S)$ is polynomially decidable under the admissible and stable semantics. In fact, P-$\text{EXT}^{sem}(S)$ over IND-D is $FP^{\#P}$-complete under the complete and the grounded semantics, under which $\text{EXT}^{sem}(S)$ is polynomially decidable.

As for P-$\text{ACC}^{sem}(a)$, the sensitivity analysis to the sub-classese of GEN does not allow us to find tractability islands related to the types of correlations imposed over the terms of the dispute, since P-$\text{ACC}^{sem}(a)$ is already $FP^{\#P}$-complete under every semantics over IND. However, the result that P-$\text{ACC}^{sem}(a)$ over GEN is still in $FP^{\#P}$ allows us to conclude that imposing correlations does not further increase the complexity of the problem.

It is worth noting that the succinctness property of the paradigm of ws-descriptors and ws-sets makes our complexity characterization over the form GEN describe the intrinsic complexity of P-$\text{EXT}^{sem}(S)$ and P-$\text{ACC}^{sem}(a)$ (since the results refer to a compact representation of the input). Interestingly, the validity on the same results shown over GEN can be easily proven to hold when using other popular paradigms for representing pdfs and correlations, such as Bayesian Networks. In this case, for instance, the results of $FP^{\#P}$-hardness and the FP-membership for different combinations $\langle$semantics/types of correlations$\rangle$ over Bayesian Networks can be obtained from those proved for GEN by observing that encoding xor and co-existence relationships in a Bayesian Network requires the same space as in GEN.

In brief, what reported in Table 1 and Table 2 gives a complete picture of how the complexity is sensitive to the semantics, the way of encoding the pdf, and the types of correlations between arguments/defeats. In the future work, we will devote our attention to searching for islands of tractability not related to syntactic restrictions, but to structural properties of a suitable graph-based representation of the prAAF.

6. Using GEN in practice: ongoing research

As highlighted in Section 4.2, the form GEN exhibits a nice trade-off between expressiveness and compactness, as its paradigm for encoding the pdf over the possible AAFs allows any pdf to be represented, while being exponentially more succinct than EX. In particular, the form GEN is likely to be a good choice at least when the number of scenarios (possible AAFs) is not small (so that the enumeration of the possible AAFs required by the form EX is impractical), and the independence between all the arguments/defeats cannot be assumed (so that IND cannot be used).

Moreover, the form GEN can be used in a very user-friendly way in the case that the only correlations are XOR and co-existence constraints: as it will be clearer in what follows, this case avoids the need to explicitly specify the ws-sets of each term of the dispute, as the ws-sets can be automatically obtained if only the marginal probabilities of arguments/defeats and their correlations are specified. In this regard, we are currently developing an argumentation system providing a visual interface that allows the user to specify arguments, defeats, probabilities and correlations as components of a particular argumentation graph, that is eventually translated into a prAAF of form MON (as discussed in Section4.4, this subclass of GEN suffices to express the considered forms of correlations). In particular, the graph has the following characteristics (see Figure 3):

- the nodes represent arguments, and can be grouped into macro-nodes in order to impose a co-existence constraint over sets of them (see nodes a_1, a_2, a_3 in Figure 3);
- the edges represent defeats, and are pairwise connected with a dotted line marked with "AND" when a coexistence constraint holds over them (see the defeats involving a_1, a_2, and a_3 in Figure 3);
- pairs of arguments and pairs of defeats such that an XOR constraint holds between them are connected with a dotted line marked with "XOR" (see the defeats between a_5 and a_6 in Figure 3);
- each node/defeat is associated with a marginal probability. In particular, for nodes belonging to the same macro-node, the probability is specified only once over the macro-node; moreover, for terms linked by dotted lines marked with AND or XOR, the probabilities are due to be equal or complimentary, respectively.

The semantics of the graph is the prAAF obtained by assuming independence wherever not explicitly specified through co-existence and XOR constraints. For

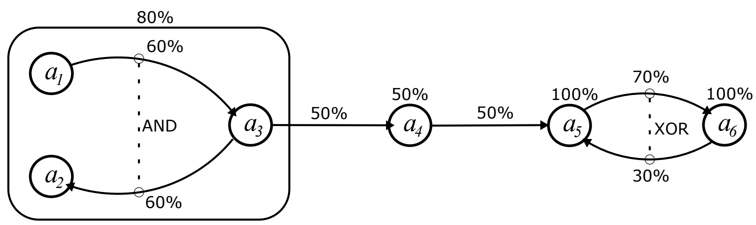

Figure 3: FIX

instance, the fact that the node a_4 in the graph of Figure 3 is included in no macro-node and connected to no argument via XOR links means that argument a_4 is independent from all the other arguments (that is, its occurrence in the actual dispute is not influenced by the presence of other arguments). Analogously, the macro-node containing a_1, a_2, a_3 in the graph in Figure 3 means that each of these three arguments has 80% probability to occur in the dispute, and that they are not independent: if one of them occurs, all the others do. Intuitively, the semantics of this form of graph in term of a prAAF of form MON is straightforward: the co-existence constraints between arguments can be expressed by associating the arguments in a macro-node with the same descriptor, while XOR constraints by associating the two arguments with alternative values of the same variable. The interested reader can find a complete "interpretation" of the graph in Figure 3 in Example 7 below.

Remarkably, this way of specifying correlations and probabilities resembles the way of assigning probabilities that is used in IND (which is very user-friendly, since the user may focus on each sub-group of the dispute separately), while allowing common types of correlations to be specified, thus overcoming a limit in expressiveness of IND.

Example 7. *The translation of the argumentation graph in Figure 3 into a prAAF $\mathcal{F} = \langle A, D, W, \lambda \rangle$ is as follows. Each node a_i is put in A and a binary variable x_i is created. Next, the variables corresponding to arguments in a same macro-node and/or connected via XOR links are merged into a unique variable where the subscripts are concatenated. In this case, we create variables x_{123}, x_4, x_5, x_6. Then, we put into W the triples reported in the table on left-hand side below, and partially define λ as specified in the table on the right-hand side below:*

Var	*value*	p
x_{123}	0	0.2
x_{123}	1	0.8
x_4	0	0.5
x_4	1	0.5
x_5	1	1
x_6	1	1

Term	$\lambda($*Term*$)$
a_1	$\{\{x_{123} \to 1\}\}$
a_2	$\{\{x_{123} \to 1\}\}$
a_3	$\{\{x_{123} \to 1\}\}$
a_4	$\{\{x_4 \to 1\}\}$
a_5	$\{\{x_5 \to 1\}\}$
a_6	$\{\{x_6 \to 1\}\}$

The same rationale is used to populate D and define the variables corresponding to the defeats. In this case, denoting as δ_{ij} the defeat from a_i to a_j, we define the variables $y_{13,32}$, y_{34}, y_{45}, $y_{56,65}$. Then, we complete W and the definition of λ with the following entries:

Var	*value*	p
$y_{13,32}$	0	0.4
$y_{13,32}$	1	0.6
y_{34}	0	0.5
y_{34}	1	0.5
y_{45}	0	0.5
y_{45}	1	0.5
$y_{56,65}$	0	0.7
$y_{56,65}$	1	0.3

Term	$\lambda(Term)$
δ_{13}	$\{\{y_{13,32} \to 1\}\}$
δ_{32}	$\{\{y_{13,32} \to 1\}\}$
δ_{34}	$\{\{y_{34} \to 1\}\}$
δ_{45}	$\{\{y_{45} \to 1\}\}$
δ_{56}	$\{\{y_{56,65} \to 0\}\}$
δ_{65}	$\{\{y_{56,65} \to 1\}\}$

It is easy to see how the correlations expressed in the graph have been encoded. For instance, the fact that λ associates δ_{56} with the ws-set $\{\{y_{56,65} \to 0\}\}$ and δ_{65} with the "complimentary" ws-set $\{\{y_{56,65} \to 1\}\}$ translates the XOR constraints over these two defeats. Similarly, the fact that λ associates δ_{13} and δ_{32} with the same ws-set $\{\{y_{13,32} \to 1\}\}$ implies that these defeats will be in the same possible AAFs.

To summarize, the reader can check that the prAAF of form MON *reported above represents the set of 26 possible AAFs in the following table (each row describes the composition of a possible AAF), whose enumeration is clearly harder than drawing the graph.*

α	a_1	a_2	a_3	a_4	a_5	a_6	δ_{13}	δ_{32}	δ_{34}	δ_{45}	δ_{56}	δ_{65}
α_1	1	1	1	0	1	1	1	1	0	0	1	0
α_2	1	1	1	0	1	1	0	0	0	0	1	0
α_3	0	0	0	0	1	1	0	0	0	0	1	0
α_4	1	1	1	0	1	1	1	1	0	0	0	1
α_5	1	1	1	0	1	1	0	0	0	0	0	1
α_6	0	0	0	0	1	1	0	0	0	0	0	1
α_7	1	1	1	1	1	1	1	1	0	0	1	0
α_8	1	1	1	1	1	1	0	0	0	0	1	0
α_9	0	0	0	1	1	1	0	0	0	0	1	0
α_{10}	1	1	1	1	1	1	1	1	0	0	0	1
α_{11}	1	1	1	1	1	1	0	0	0	0	0	1
α_{12}	0	0	0	1	1	1	0	0	0	0	0	1
α_{13}	1	1	1	1	1	1	1	1	1	0	1	0
α_{14}	1	1	1	1	1	1	0	0	1	0	1	0
α_{15}	1	1	1	1	1	1	1	1	1	0	0	1
α_{16}	1	1	1	1	1	1	0	0	1	0	0	1
$\cdots$	$\cdots$	$\cdots$	$\cdots$	$\cdots$	$\cdots$	$\cdots$	$\cdots$	$\cdots$	$\cdots$	$\cdots$	$\cdots$	$\cdots$
α_{26}	1	1	1	1	1	1	0	0	1	1	0	1

The probability of each possible AAF can be easily understood by looking at the graph. For instance, the probability of α_7 is the product of the following factors (in brackets we report the event corresponding to the probability): 0.8 *(a_1, a_2, a_3 occur)* $\times$ 0.5 *(a_4 occurs)* $\times$ 1 *(a_5 occurs)* $\times$ 1 *(a_6 occurs)* $\times$ 0.6 *(δ_{12} and δ_{32}, given that a_1, a_2, a_3 occur)* $\times$ 0.5 *(δ_{34} does not occur, given that a_3 and a_4 occur)* $\times$ 0.5 *(δ_{45} does not occur, given that a_4 and a_5 occur)* $\times$ 0.7 *(δ_{56} occurs and δ_{65} does not, given that a_5 and a_6 occur), that is* 4.2%. □

7. RELATED WORK

The main state-of-the-art approaches for handling uncertainty in AAFs by relying on probability theory can be classified in two categories, based on the way they interpret the probabilities of the arguments: those adopting the classical *constellations* approach [5, 9, 10, 11, 12, 13, 14, 15, 16] and those adopting the recent *epistemic* one [31, 32, 33]. The former has been widely discussed in the core of the paper (see Section 2), where the sub-categories EX [9, 10, 12] and IND [11, 14, 15, 16] have been analyzed as well. As regards the epistemic approach, probabilities and extensions have a different semantics, compared with the constellations approach. Specifically, the probability of an argument represents the degree of belief in the argument (the higher the probability, the more the argument is believed), and a key concept is the "rational" probability distribution, that requires that if the belief in an argument is high, then the belief in the arguments attacked by it is low. In this approach, epistemic extensions are considered rather than Dung's extensions, where an epistemic extension is the set of arguments that are believed to be true to some degree. The interested reader can find a more detailed comparative description of the two categories in [34].

We now focus our attention on the approaches classified as constellations, as the complexity characterization provided in our work refers to this class of prAAFs, to which obviously our new framework GEN and its subclasses belong too. Among the research works dealing with prAAFs of form EX, [9] addressed the modeling of jury-based dispute resolutions, and proposed a prAAF where uncertainty is taken into account by specifying probability distribution functions (pdfs) over possible AAFs and showing how an instance of the proposed prAAF can be obtained by specifying a probabilistic assumption-based argumentation framework (introduced by themselves). In the same spirit, [10] defined a prAAF as a pdf over the set of possible AAFs, and introduced a probabilistic version of a fragment of the ASPIC framework [35] that can be used to instantiate the proposed prAAF. [36] addresses the problem of computing all the subgraphs of an AAF in which an argument a belongs to the grounded extension, and [12] extends it by focusing on computing the probability that an argument a belongs to the grounded extension of a probabilistic abstract argumentation framework. In particular, [12] assumes to receive a joint probability distribution over the arguments as input. In fact, providing a joint probability distribution usually means specifying the probability values for all the possible correlations, i.e., $P(a)$, $P(a \wedge b)$, $P(a \wedge b \wedge c)$... and so on. This is analogous to providing the probabilities for all the possible AAFs (since defeats are considered as certain). Thus, the prAAF considered in [12] can be viewed as a prAAF of form EX.

Differently from [9] and [10], [14] proposed a prAAF where probabilities are directly associated with arguments and defeats, instead of being associated with possible AAFs, and independence among pairs of arguments/defeats is assumed. After claiming that computing the probability $P^{sem}(S)$ that a set S of arguments is an extension according to sem requires exponential time for every semantics, [14] proposed a Monte-Carlo simulation approach to approximate $P^{sem}(S)$. [11, 15, 16] build upon [14]: [11] characterizes different semantics from the approach of [14] in terms of probabilistic logic, as a first step in the direction of creating a uniform logical formalization for all the proposed AAFs of the literature, in order to understand and compare the different approaches. [15, 16], instead, showed that computing $P^{sem}(S)$ is actually tractable for the *admissible* and *stable* semantics, but it is $FP^{\#P}$-complete for other semantics, including *complete*, *grounded*, *preferred* and *ideal-set*. Furthermore, [37, 38] proposed a Monte-Carlo approach to efficiently estimate $P^{sem}(S)$ based on the polynomiality results of [15, 16]. In [14], as well as in [9] and [10], $P^{sem}(S)$ is defined as the sum of the probabilities of the possible AAFs where S is an extension according to semantics *sem*.

In the above-cited works, probability theory is recognized as a fundamental tool to model uncertainty. However, a deeper understanding of the role of probability theory in abstract argumentation was developed only later in [13, 34], where the *justification* and the *premise* perspectives of probabilities of arguments are introduced. According to the former perspective the probability of an argument indicates the probability that it is justified in appearing in the argumentation system. In contrast, the premise perspective views the probability of an argument as the probability that the argument is true based on the degrees to which the premises supporting the argument are believed to be true. Starting from these perspectives, in [34], a formal framework showing the connection among argumentation theory, classical logic, and probability theory was investigated. Furthermore, qualification of attacks is addressed in [5], where an investigation of the meaning of the uncertainty concerning defeats in probabilistic abstract argumentation is provided.

Besides the approaches that model uncertainty in AAFs by relying on probability theory, many proposals have been made where uncertainty is represented by exploiting weights or preferences on arguments and/or defeats [39, 40, 41, 42, 43, 44, 45]. Another interesting approach to represent uncertainty in argumentation is that based on using possibility theory, as done in [46, 47, 48]. Although the approaches based on weights, preferences, possibilities, or probabilities to model uncertainty have been proved to be effective in different contexts, there is no common agreement on what kind of approach should be used in general. In this regard, [13, 34] observed that the probability-based approaches may take advantage from relying on a well-established and well-founded theory, whereas the approaches based on weights or preferences do not conform to well-established theories yet.

The computational complexity of computing extensions has been thoroughly investigated for classical AAFs [22, 23, 24, 25, 26, 27] with respect to several semantics (a comprehensive overview of argumentation semantics can be found in [49]). In particular, [22] presents a number of results on the complexity of some decision questions for semi-stable semantics, while [24] focuses on ideal semantics; complexity results for preferred semantics can be found in [23]. Complexity results about skeptical and credulous acceptance under admissible, complete, grounded, stable and preferred semantics have been provided in [50, 51, 52], while [24] characterizes the complexity of skeptical and credulous acceptance under ideal and ideal-set. [25] provides complexity results for AAFs in terms of skeptical and credulous acceptance under the semi-stable semantics, while [26] analyses CF2 semantics. The recent work [27] has studied the computational complexity of different decision problems centred on *critical sets* of arguments whose status (i.e., membership to an extension) is sufficient to determine uniquely the status of every other argument. As regard the case of adding weights to AAFs, the computational complexity of computing extensions has been deeply investigated in [43, 44], whereas the complexity for the case of using preferences is studied in [40, 53].

Several systems are available for reasoning in non-probabilistic argumentation frameworks, such as [54, 55, 56, 57]. Furthermore, a system for reasoning in probabilistic argumentation frameworks is presented in [58]. An up-to-date survey of systems for solving reasoning problem in AAFs can be found in [59].

8. CONCLUSION AND FUTURE WORK

A thorough characterization of the complexity of the fundamental problem P-$\textsc{Ext}^{sem}(S)$ and P-$\textsc{Acc}^{sem}(a)$ over probabilistic abstract argumentation frameworks has been provided. The results reported in this paper give an insight on the sensitiveness of the complexity to the semantics adopted for the extensions, the representation paradigm of the pdf, and the forms of correlations between arguments/defeats. This fills a gap in the existing research literature, where the complexity of P-$\textsc{Ext}^{sem}(S)$ and P-$\textsc{Acc}^{sem}(a)$ was studied only in specific cases, such as the case that the terms of the dispute are independent events, or the deterministic case.

This analysis of the two fundamental problems has been allowed by the introduction of a new form of prAAF (namely, GEN), that enables correlations between the terms of the dispute to be easily specified in (and detected from) the encoding of the pdf. The new framework has been shown to exhibit a nice trade-off between expressiveness and compactness, and its use as the core of practical argumentation systems has been discussed and is, in fact, matter of an ongoing research.

There are several interesting directions for future work. On the one hand, the complexity analysis can be easily extended to deal with the *sceptical* variant of the acceptance problem, where an argument is required to belong to all the extensions. For some semantics, extending our results is trivial. For instance, the characterization under the ideal semantics does not change w.r.t. the case studied here, since the ideal extension is unique. Analogously, the case of admissible semantics becomes trivial, since the empty set is always an admissible extension. For the other semantics, very similar reasonings to those used in the proofs of our results can be adopted to provide a characterization under the sceptical variant of P-$\text{ACC}^{sem}(a)$.

On the other hand, other directions for extending our results are less straightforward, such as that of widening the complexity analysis to include other semantics for the extensions (such as *stage* [60], *cf2* [26], and the variants of the ideal semantics defined in [61]).

Finally, we will focus our attention to address some other research problems in the probabilistic frameworks, such as the computation of the most probable extension or the argument having the highest probability of acceptance. The latter can be reduced to invoking a solver P-$\text{ACC}^{sem}(a)$ for each argument, thus its complexity is closely related to that of P-$\text{ACC}^{sem}(a)$. As regards the former, the characterization of its complexity requires further investigation.

Acknowledgements

The authors are grateful to the anonymous reviewers of ECAI 2016, for their insightful suggestions regarding a preliminary version of this paper published in form of poster [62].

References

[1] T. J. M. Bench-Capon, P. E. Dunne, Argumentation in artificial intelligence, Artif. Intell. 171 (10-15) (2007) 619–641.

[2] P. Besnard, A. Hunter (Eds.), Elements Of Argumentation, The MIT Press, 2008.

[3] I. Rahwan, G. R. Simari (Eds.), Argumentation in Artificial Intelligence, Springer, 2009.

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

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

[6] P. M. Dung, P. Mancarella, F. Toni, Computing ideal sceptical argumentation, Artif. Intell. 171 (10-15) (2007) 642–674.

[7] P. Baroni, M. Giacomin, Semantics of abstract argument systems, in: Argumentation in Artificial Intelligence, 2009, pp. 25–44.

[8] M. Caminada, Semi-stable semantics, in: Proc. Int. Conf. Computational Models of Argument (COMMA), 2006, pp. 121–130.

[9] P. M. Dung, P. M. Thang, Towards (probabilistic) argumentation for jury-based dispute resolution, in: Proc. Int. Conf. Computational Models of Argument (COMMA), 2010, pp. 171–182.

[10] T. Rienstra, Towards a probabilistic dung-style argumentation system, in: Proc. Int. Conf. on Agreement Technologies (AT), 2012, pp. 138–152.

[11] D. Doder, S. Woltran, Probabilistic argumentation frameworks - A logical approach, in: Proc. Int. Conf. on Scalable Uncertainty Management (SUM), 2014, pp. 134–147.

[12] P. Dondio, Toward a computational analysis of probabilistic argumentation frameworks, Cybernetics and Systems 45 (3) (2014) 254–278.

[13] A. Hunter, Some foundations for probabilistic abstract argumentation, in: Proc. Int. Conf. Computational Models of Argument (COMMA), 2012, pp. 117–128.

[14] H. Li, N. Oren, T. J. Norman, Probabilistic argumentation frameworks, in: Proc. Int. Workshop on Theorie and Applications of Formal Argumentation (TAFA), 2011.

[15] B. Fazzinga, S. Flesca, F. Parisi, On the complexity of probabilistic abstract argumentation, in: Proc. Int. Joint Conference on Artificial Intelligence (IJ-CAI), 2013.

[16] B. Fazzinga, S. Flesca, F. Parisi, On the complexity of probabilistic abstract argumentation frameworks, ACM Trans. Comput. Log. (TOCL) 16 (3) (2015) 22.

[17] B. Fazzinga, S. Flesca, F. Furfaro, Credulous acceptability in probabilistic abstract argumentation: complexity results, in: Workshop on Advances In Argumentation In Artificial Intelligence (AI^3), 2017, p. (submitted for publication).

[18] L. Antova, T. Jansen, C. Koch, D. Olteanu, Fast and simple relational processing of uncertain data, in: Proc. Int. Conf. on Data Engineering (ICDE), 2008, pp. 983–992.

[19] C. Koch, D. Olteanu, Conditioning probabilistic databases, Proc. of VLDB Endowment (PVLDB) 1 (1) (2008) 313–325.

[20] L. G. Valiant, The complexity of computing the permanent, Theor. Comput. Sci. (TCS) 8 (1979) 189–201.

[21] S. Toda, O. Watanabe, Polynomial Time 1-Turing Reductions from #PH to #P, Theor. Comput. Sci. 100 (1) (1992) 205–221.

[22] P. E. Dunne, M. Caminada, Computational complexity of semi-stable semantics in abstract argumentation frameworks, in: Proc.European Conf. on Logics in Artificial Intelligence (JELIA), 2008, pp. 153–165.

[23] P. E. Dunne, M. Wooldridge, Complexity of abstract argumentation, in: Argumentation in Artificial Intelligence, 2009, pp. 85–104.

[24] P. E. Dunne, The computational complexity of ideal semantics, Artif. Intell. 173 (18).

[25] W. Dvorák, S. Woltran, Complexity of semi-stable and stage semantics in argumentation frameworks, Inf. Process. Lett. 110 (11) (2010) 425–430.

[26] S. A. Gaggl, S. Woltran, The cf2 argumentation semantics revisited, J. Log. Comput. 23 (5) (2013) 925–949.

[27] R. Booth, M. Caminada, P. E. Dunne, M. Podlaszewski, I. Rahwan, Complexity properties of critical sets of arguments, in: Proc. of Computational Models of Argument (COMMA), 2014, pp. 173–184.

[28] World-set decompositions: Expressiveness and efficient algorithms, Theoretical Computer Science (TCS) 403 (23) (2008) 265 – 284.

[29] O. Benjelloun, A. D. Sarma, A. Halevy, J. Widom, Uldbs: Databases with uncertainty and lineage, in: Proc. Int. Conf. on Very Large Data Bases (VLDB), 2006, pp. 953–964.

[30] P. Besnard, S. Doutre, Checking the acceptability of a set of arguments, in: Proc. Int. Workshop on Non-Monotonic Reasoning (NMR), 2004, pp. 59–64.

[31] M. Thimm, A probabilistic semantics for abstract argumentation, in: Proc. European Conf. on Artificial Intelligence (ECAI), 2012, pp. 750–755.

[32] A. Hunter, M. Thimm, Probabilistic argumentation with incomplete information, in: Proc. European Conf. on Artificial Intelligence (ECAI), 2014, pp. 1033–1034.

[33] A. Hunter, M. Thimm, Probabilistic argumentation with epistemic extensions, in: Proc. Int. Workshop on Defeasible and Ampliative Reasoning (DARe@ECAI), 2014.

[34] A. Hunter, A probabilistic approach to modelling uncertain logical arguments, Int. J. Approx. Reasoning 54 (1) (2013) 47–81.

[35] H. Prakken, An abstract framework for argumentation with structured arguments, Argument & Computation 1 (2) (2010) 93–124.

[36] P. Dondio, Computing the grounded semantics in all the subgraphs of an argumentation framework: An empirical evaluation, in: Proc. Int. Workshop Computational Logic in Multi-Agent Systems (CLIMA), 2013, pp. 119–137.

[37] B. Fazzinga, S. Flesca, F. Parisi, Efficiently estimating the probability of extensions in abstract argumentation, in: Proc. Int Conf. on Scalable Uncertainty Management (SUM), 2013, pp. 106–119.

[38] B. Fazzinga, S. Flesca, F. Parisi, On efficiently estimating the probability of extensions in abstract argumentation frameworks, Int. J. Approx. Reasoning 69 (2016) 106–132.

[39] T. J. M. Bench-Capon, Persuasion in practical argument using value-based argumentation frameworks, J. Log. Comput. 13 (3) (2003) 429–448.

[40] L. Amgoud, S. Vesic, A new approach for preference-based argumentation frameworks, Ann. Math. Artif. Intell. 63 (2) (2011) 149–183.

[41] L. Amgoud, C. Cayrol, A reasoning model based on the production of acceptable arguments, Ann. Math. Artif. Intell. 34 (1-3) (2002) 197–215.

[42] S. Modgil, Reasoning about preferences in argumentation frameworks, Artif. Intell. 173 (9-10) (2009) 901–934.

[43] P. E. Dunne, A. Hunter, P. McBurney, S. Parsons, M. Wooldridge, Weighted argument systems: Basic definitions, algorithms, and complexity results, Artif. Intell. 175 (2).

[44] S. Coste-Marquis, S. Konieczny, P. Marquis, M. A. Ouali, Weighted attacks in argumentation frameworks, in: Proc. Int. Conf. on Knowledge Representation and Reasoning (KR), 2012.

[45] G. Brewka, S. Polberg, S. Woltran, Generalizations of dung frameworks and their role in formal argumentation, IEEE Intelligent Systems 29 (1) (2014) 30–38.

[46] L. Amgoud, H. Prade, Reaching agreement through argumentation: A possibilistic approach, in: Proc. Int. Conf. on Knowledge Representation and Reasoning (KR), 2004, pp. 175–182.

[47] T. Alsinet, C. I. Chesñevar, L. Godo, S. Sandri, G. R. Simari, Formalizing argumentative reasoning in a possibilistic logic programming setting with fuzzy unification, Int. J. Approx. Reasoning 48 (3) (2008) 711–729.

[48] T. Alsinet, C. I. Chesñevar, L. Godo, G. R. Simari, A logic programming framework for possibilistic argumentation: Formalization and logical properties, Fuzzy Sets and Systems 159 (10) (2008) 1208–1228.

[49] P. Baroni, M. Caminada, M. Giacomin, An introduction to argumentation semantics, Knowledge Eng. Review 26 (4) (2011) 365–410.

[50] Y. Dimopoulos, A. Torres, Graph theoretical structures in logic programs and default theories, Theor. Comput. Sci. 170 (1-2) (1996) 209–244.

[51] P. E. Dunne, T. J. M. Bench-Capon, Coherence in finite argument systems, Artif. Intell. 141 (1/2) (2002) 187–203.

[52] S. Coste-Marquis, C. Devred, P. Marquis, Symmetric argumentation frameworks, in: Proc. of Symbolic and Quantitative Approaches to Reasoning with Uncertainty (ECSQARU), 2005, pp. 317–328.

[53] E. J. Kim, S. Ordyniak, S. Szeider, Algorithms and complexity results for persuasive argumentation, Artif. Intell. 175 (9-10) (2011) 1722–1736.

[54] M. South, G. Vreeswijk, J. Fox, Dungine: A java dung reasoner, in: Proc. Int. Conf. Computational Models of Argument (COMMA), 2008.

[55] U. Egly, S. A. Gaggl, S. Woltran, ASPARTIX: implementing argumentation frameworks using answer-set programming, in: Proc. Int. Conf on Logic Programming (ICLP), 2008, pp. 734–738.

[56] M. Snaith, C. Reed, TOAST: online aspic$^+$ implementation, in: Proc. Int. Conf. Computational Models of Argument (COMMA), 2012, pp. 509–510.

[57] T. Alsinet, R. Béjar, L. Godo, F. Guitart, Using answer set programming for an scalable implementation of defeasible argumentation, in: Proc. Int. Conf. on Tools with Artificial Intelligence (ICTAI), 2012, pp. 1016–1021.

[58] B. Fazzinga, S. Flesca, F. Parisi, A. Pietramala, PARTY: A mobile system for efficiently assessing the probability of extensions in a debate, in: Proc. Int. Conf. on Database and Expert Systems Applications (DEXA), 2015, pp. 220–235.

[59] G. Charwat, W. Dvorák, S. A. Gaggl, J. P. Wallner, S. Woltran, Methods for solving reasoning problems in abstract argumentation - A survey, Artif. Intell. 220 (2015) 28–63.

[60] M. Caminada, An algorithm for stage semantics, in: Proc. Int. Conf. Computational Models of Argument (COMMA), 2010, pp. 147–158.

[61] P. E. Dunne, W. Dvor'ak, S. Woltran, Parametric properties of ideal semantics, Artif. Intell. 202 (2013) 1–28.

[62] B. Fazzinga, S. Flesca, F. Furfaro, Computing extensions' probabilities in probabilistic abstract argumentation: Beyond independence, in: Proc. European Conf. on Artificial Intelligence (ECAI), 2016, pp. 1588–1589.