Abstract Argumentation Frameworks with Marginal Probabilities

Bettina Fazzinga[1,2] and **Sergio Flesca**[3] and **Filippo Furfaro**[3]
[1]DICES - University of Calabria, Italy
[2]ICAR - CNR, Italy
[3]DIMES - University of Calabria, Italy
{bettina.fazzinga, sergio.flesca, filippo.furfaro}@unical.it

Abstract

In the context of probabilistic AAFs, we introduce *AAFs with marginal probabilities* (mAAFs) requiring only marginal probabilities of arguments/attacks to be specified and not relying on the independence assumption. Reasoning over mAAFs requires taking into account multiple probability distributions over the possible worlds, so that the probability of extensions is not determined by a unique value, but by an interval. We focus on the problems of computing the max and min probabilities of extensions over mAAFs under Dung's semantics, characterize their complexity, and provide closed formulas for polynomial cases.

1 Introduction

Several frameworks based on the probability theory have extended Dung's *Abstract Argumentation Framework* (AAF [Dung, 1995]) to take into account the uncertainty possibly affecting the occurrence of arguments and attacks in the argumentation. In particular, in the *constellations approach* [Hunter, 2012; Hunter, 2014; Dondio, 2014; Doder and Woltran, 2014; Rienstra, 2012; Dung and Thang, 2010; Li *et al.*, 2011; Fazzinga *et al.*, 2015] the dispute is represented with a *probabilistic AAF* (prAAF), that encodes the alternative possible worlds for the argumentation as a set of (deterministic) AAFs, where each AAF is associated with the probability of being the AAF actually occurring. prAAFs can be divided in two categories: those relying on the *independence assumption* (i.e. arguments are independent from one another, and the occurrence of attacks is conditioned only to the occurrence of the related arguments), and those not. The latter allow the analyst to specify any probability distribution function (pdf) over the possible worlds, but, in fact, can be hardly used in complex scenarios: defining such a pdf may require reasoning on a number of possible worlds exponential w.r.t. the number of arguments and attacks, and this in turn may require a strong effort and a deep knowledge of the correlations between arguments and attacks, that often is not available. On the other hand, the prAAFs assuming independence are compact and user-friendly, as they require only the specification of the marginal probabilities of arguments/attacks (which implicitly define a pdf over the possible worlds). In fact, assigning suitable marginal probabilities calls for reasoning over one argument/attack at the time, so is much less burdensome than explicitly describing a pdf over the possible worlds and does not require to have a precise picture of how the arguments/attacks are correlated. For instance, the probability of an argument (resp., attack) can be modeled by looking into statistics about occurrences of single arguments/attacks or by reasoning on the chances of participating to the dispute of the agents who propose the arguments or perceive the attacks. Unfortunately, assuming independence is often inadequate, since the existence of correlations between arguments/attacks cannot be excluded.

In this paper, we introduce a new prAAF, called *AAF with marginal probabilities* (mAAF), that is in between these two categories: it requires to specify no pdf over the possible worlds, but only the marginal probabilities of arguments and attacks, while not assuming independence. This calls for a reasoning paradigm different from the prAAFs in the literature, where a single (explicitly or implicitly encoded) pdf over the possible worlds is considered: in mAAFs several pdfs may be consistent with the marginal probabilities, thus the probability of being extensions cannot measured by a unique value. The following example clarifies this aspect, and gives an insight on why assuming independence when correlations are not known may result in incautious conclusions.

Example 1 *Consider the argumentation graph below, reporting the marginal probabilities of arguments and attacks.*

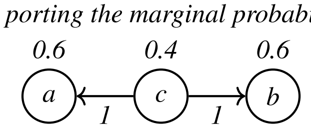

The numbers besides nodes and edges are the marginal probabilities returned by a function $P(\cdot)$

As the three arguments are the only uncertain portions of the argumentation, we have the $2^3 = 8$ *possible worlds reported in the first column of the table below. The other columns report three (out of many other) pdfs consistent with the marginal probabilities. Specifically,* π_1 *corresponds to assuming independence between arguments/attacks, as it assigns to every possible world* ω_i *the product of the marginal probabilities (resp., complements of marginal probabilities) of the arguments/attacks in* ω_i *(resp., not in* ω_i*). As for* π_3*, it corresponds to the case where* a*,* b *are positively correlated and are in mutual exclusion with* c *(so that the only possible scenarios for the argumentation are* ω_4 *and* ω_5*), while* π_2 *to the case where either* a*,* b*,* c *coexist, or at most one between*

a or b occurs. Observe that π_2 (resp., π_3) is a pdf minimizing (resp., maximizing) the probability of ω_5 (0 and 0.6 are the min and max values since no pdf can assign to ω_5 a probability lower than 0 and higher than any of $P(a)$, $P(b)$, $1-P(c)$).

possible world	π_1	π_2	π_3
$\omega_1 = \langle \emptyset, \emptyset \rangle$	*0.096*	*0.2*	*0*
$\omega_2 = \langle \{a\}, \emptyset \rangle$	*0.144*	*0.2*	*0*
$\omega_3 = \langle \{b\}, \emptyset \rangle$	*0.144*	*0.2*	*0*
$\omega_4 = \langle \{c\}, \emptyset \rangle$	*0.064*	*0*	*0.4*
$\omega_5 = \langle \{a, b\}, \emptyset \rangle$	*0.216*	*0*	*0.6*
$\omega_6 = \langle \{a, c\}, \{(c, a)\} \rangle$	*0.096*	*0*	*0*
$\omega_7 = \langle \{b, c\}, \{(c, b)\} \rangle$	*0.096*	*0*	*0*
$\omega_8 = \langle \{a, b, c\}, \{(c, a), (c, b)\} \rangle$	*0.144*	*0.4*	*0*

Let $S = \{a, b\}$. The only ω_i where S is admissible is ω_5. As in prAAFs the probability that a set is an extension is the sum of the probabilities of the possible worlds where the extension's conditions are met, the probability $P(S)$ that S is admissible is the probability of ω_5. Therefore, from what said above on the minimum and maximum probability of ω_5, we conclude that $P(S)$ is in the range $[\pi_2(\omega_5)..\pi_3(\omega_5)] = [0..0.6]$.

Now, suppose that the analyst is a lawyer, and that a, b are arguments possibly claimed by the counterpart's witnesses. If the lawyer trusted the analysis under independence, they would conclude that $\{a,b\}$ is not a robust set of arguments, since $P(S) = \pi_1(\omega_5) = 0.216$ is rather low. Instead, taking into account all the possible pdfs consistent with the marginal probabilities, the lawyer would be aware that $P(S)$ can be rather high, as $P(S)$ may be up to $\pi_3(\omega_5) = 0.6$, thus they can make more cautious decisions regarding the trial strategy.

1.1 Contributions

We introduce mAAFs along with the problems of maximizing and minimizing the probability that a set is an extension over an mAAF. We first focus on MAXP-VER, the decision counterpart of the maximization problem, and provide a thorough complexity analysis under Dung's semantics for extensions. We show that, depending on the semantics, MAXP-VER can be polynomial-time solvable, NP-complete or Σ_2^p-complete. For the PTIME cases, we provide elegant closed formulas returning the maximum probability value. Furthermore, we show the relation between mAAFs and the several prAAFs proposed in the literature. Finally, we also analyze the dual problem MINP-VER referring to the minimum probability, providing a tight characterization under some semantics and showing an interesting asymmetry w.r.t. MAXP-VER.

2 Preliminaries

An *abstract argumentation framework* (*AAF*) is a pair $F = \langle A, D \rangle$, where A is a set of abstract elements, called *arguments*, and D a binary relation over arguments, called *attack relation*. Given $a \in A$ and $S \subseteq A$, we say: "*a is acceptable w.r.t. S*" if for every $(c, a) \in D$ there is some $(s, c) \in D$ with $s \in S$. An attack (a, b) will be also denoted as δ_{ab}.

Several semantics for AAFs have been proposed to identify "reasonable" sets of arguments, called *extensions* [Dung, 1995]. A set $S \subseteq A$ is: a *conflict-free* extension (`cf`) iff there is no attack involving arguments in S; an *admissible extension* (`ad`) iff S is conflict-free and its arguments are acceptable w.r.t. S; a *stable extension* (`st`) iff S is conflict-free and attacks each argument in $A \setminus S$; a *complete extension* (`co`) iff S is admissible and contains all the arguments that are acceptable w.r.t. S; a *grounded extension* (`gr`) iff S is a minimal (w.r.t. $\subseteq$) complete set of arguments; a *preferred extension* (`pr`) iff S is a maximal (w.r.t. $\subseteq$) complete set of arguments. The set of extensions of an AAF F under a semantics σ is denoted as $Ext(F, \sigma)$. The verification problem of deciding if $S \in Ext(F, \sigma)$ and is denoted as VER(F, S, σ).

3 AAFs with Marginal Probabilities (mAAFs)

We consider the case where the arguments and attacks that may occur in the argumentation are known, but the exact composition of the argumentation is not certain, as it is not known which of the "*possible worlds*" (i.e. sets of arguments and attacks) will be the actual argumentation. In particular, we consider the scenario where a probabilistic measure of the uncertainty is available, in terms of the marginal probabilities of the arguments and attacks. Basically, the marginal probability of an argument represents the overall probability of the possible worlds where the argument occurs. The meaning of the marginal probability of an attack is analogous, but its value is conditioned to the occurrence of the involved arguments (e.g. "(a, b) *has probability* 1" means that the attack occurs in all the scenarios where *both* a and b occur). This naturally gives rise to the formal definition below.

Definition 1 (mAAF) *An* AAF with marginal probabilities *(mAAF) is a tuple $\langle A, D, P \rangle$, where $\langle A, D \rangle$ is an AAF and $P : (A \cup D) \rightarrow [0, 1]$ associates arguments and attacks with probabilities. Arguments and attacks are said to be* certain *is they are assigned probability 1 by P,* uncertain *otherwise.*

Formally, given an mAAF $F = \langle A, D, P \rangle$, a possible world of F is any AAF $\omega = \langle A', D' \rangle$ with $A' \subseteq A$, $D' \subseteq (A' \times A') \cap D$. We denote as $PW(F)$ the set of the possible worlds of F. Given two possible worlds $\omega = \langle A, D \rangle$, $\omega' = \langle A', D' \rangle$, we say that ω' *expands* ω if $(A \cup D) \subset (A' \cup D')$. Differently from the probabilistic extensions of AAFs in the literature (see the discussion on *constellations approaches* in Section 5), mAAFs do not rely on a unique *probability distribution function* (pdf) over the possible worlds: different pdfs may be consistent with the marginal probabilities characterizing arguments and attacks, as discussed in Example 1 and below.

Example 2 *Continuing Example 1, it is easy to see that also π_4, that assigns 0.4 to both ω_2 and ω_7, 0.2 to ω_5 and 0 to all the other possible worlds, is consistent with the marginal probabilities. Observe that π_4 maximizes the probability of ω_2 (containing only a): no consistent pdf can assign to ω_2 a value higher than any of $P(a)$, $1 - P(b)$, $1 - P(c)$.*

Given a pdf π over $PW(F)$ and a set of possible worlds $PW' \subseteq PW(F)$, $\pi(PW') = \sum_{\omega \in PW'} \pi(\omega)$ denotes the overall probability assigned by π to the possible worlds in PW'. Then, we say that π *is consistent with* the marginal probability of argument a (resp., attack δ_{ab}), written as $\pi \models P(a)$ (resp., $\pi \models P(\delta_{ab})$) if $P(a) = \sum_{\omega \in PW(F) | a \in \omega} \pi(\omega)$ (resp., $P(\delta_{ab}) =$

$\sum_{\omega \in PW(F) | \delta \in \omega} \pi(\omega) / \sum_{\omega \in PW(F) | a \in \omega \wedge b \in \omega} \pi(\omega))$. In turn, π *is consistent with* P ($\pi \models P$) if it is consistent with the marginal probabilities of F's arguments/attacks. We denote as $\Pi(F)$ the set of pdfs over $PW(F)$ consistent with P.

The presence of multiple possible worlds for the argumentation, along the fact that each possible world may be associated with different probabilities (since different pdfs over $PW(F)$ may be consistent with F), naturally call for revisiting the traditional way of considering extensions. In this spirit, given a pdf π in $\Pi(F)$ and a semantics σ, we define the probability that S is a σ-extension of F according to π as: $\pi(F, S, \sigma) = \sum_{\omega | S \in \text{Ext}(\omega, \sigma)} \pi(w)$, that is the sum of the probabilities assigned by π to the possible worlds where S is a σ-extension. Then, in order to take into account that several probability assignments to the possible worlds can be consistent with the marginal probabilities, we define: $P_{\min}(F, S, \sigma) = \min_{\pi \in \Pi(F)} \pi(F, S, \sigma)$ and $P_{\max}(F, S, \sigma) = \max_{\pi \in \Pi(F)} \pi(F, S, \sigma)$, i.e. the minimum and maximum probabilities that S is a σ-extension.

Proposition 1 states two general properties: 1) there is always a pdf over $PW(F)$ consistent with P, and 2) the set of probabilities that S is a σ-extension of F is a closed interval.

Proposition 1 *Let F be an mAAF, S a set of arguments and $\sigma \in \{\texttt{cf}, \texttt{st}, \texttt{ad}, \texttt{co}, \texttt{gr}, \texttt{pr}\}$. Then:* 1) $\Pi(F) \neq \emptyset$*, and* 2) *for every* $p \in [P_{\min}(F, S, \sigma)..P_{\max}(F, S, \sigma)]$ *there is a pdf* $\pi \in \Pi(F)$ *such that* $p = \pi(F, S, \sigma)$.

1) follows from the existence of the pdf implied by assuming independence, while 2) is a straightforward consequence of the theory of probabilistic logics in [Nilsson, 1986], since the probabilities of sentences entailed by a probabilistic sentence are known to compose a closed interval.

Reasoning over $P_{\min}(F, S, \sigma)$ and $P_{\max}(F, S, \sigma)$ is obviously relevant for an analyst looking into an argumentation modeled via an mAAF F, since this gives insights on the extent to which S can be considered "robust". Thus, we address the decision problems MAXP-VER and MINP-VER, that are natural adaptions of the classical verification problem and that are the decision counterpart of finding $P_{\max}(F, S, \sigma)$ and $P_{\min}(F, S, \sigma)$:

Problem statement: "*Let F be an mAAF F, σ a semantics, S a set of arguments, and p^* a probability value.* MAXP-VER(F, S, σ, p^*) *(resp.,* MINP-VER(F, S, σ, p^*)*) is the problem of deciding if there is a pdf π in $\Pi(F)$ such that $\pi(F, S, \sigma) \geq p^*$ (resp., $\pi(F, S, \sigma) \leq p^*$)*".

4 Characterizing MAXP-VER and MINP-VER

We start with the problem MAXP-VER and show that it can be solved in polynomial time under the conflict-free, admissible, and stable semantics, that it is complete for NP under the complete and grounded semantics, and for Σ_2^p under the preferred semantics. In particular, for the polynomial-time cases, we provide closed formulas allowing an easy computation of $P_{\max}(F, S, \sigma)$. In the following, given a set of arguments X, we denote with $M(X) = \max\{0, 1 - |X| + \sum_{a \in X} P(a)\}$ the minimum probability that, consistently with the marginal probabilities, the arguments in X occur simultaneously.

Theorem 1 *Given an mAAF $F = \langle A, D, P \rangle$ and $S \subseteq A$:*

1) $P_{\max}(F, S, \texttt{cf}) = \min_{s,t \in S} \{(1 - P(\delta_{st})) \cdot \min\{P(s), P(t)\}\}$

2) $P_{\max}(F, S, \texttt{st}) = \min\{P_{\max}(F, S, \texttt{cf}), \min_{a \in A \setminus S}\{1 - P(a) + \sum_{s \in S} P(\delta_{sa}) \cdot M(\{s, a\})\}$.

3) $P_{\max}(F, S, \texttt{ad}) = \min\{P_{\max}(F, S, \texttt{cf}), \min_{a \in A \setminus S}\{(1 - P(a)) + \min_{s \in S}\{(1 - P(\delta_{as})) \cdot M(\{s, a\}) + \sum_{t \in S} P(\delta_{ta}) \cdot M(\{t, a\})\}\}\}$

where $\min \emptyset = 1$ *and* $P(\delta_{st}) = 0$ *if* $\delta_{st} \notin D$.

Proof of Case 1. Without loss of generality, we assume that A and S coincide, as it is easy to see that, denoting as F' the projection of F over S, for each $\pi' \in \Pi(F')$ there is a $\pi \in \Pi(F)$ such that $\pi'(F', S, \texttt{cf}) = \pi(F, S, \texttt{cf})$, and vice versa: π' can be obtained from π by marginalization, while π can be obtained from π' by distributing the marginal probabilities of the arguments and attacks in F but not in F' over the possible worlds of F expanding the possible worlds of F'.

For any pair of (possibly coinciding) arguments $s, t \in S$, let X_{st} be the set of the possible worlds containing s, t, and no attack from s to t, and let Y_{st} be the set of the possible worlds containing s, t. It is straightforward to see that, $\forall s, t \in S$ and $\forall \pi \in \Pi(F)$, the probability $\pi(F, S, \texttt{cf})$ that S is a cf-extension is not greater than the overall probability $\pi(X_{st})$, which, in turn, is equal to $(1 - P(\delta_{st})) \cdot \pi(Y_{st})$. Herein, $\pi(Y_{st}) \leq \min\{P(s), P(t)\}$, since Y_{st} is a subset of the two sets Y_s and Y_t of the possible worlds containing s and t, respectively (where, obviously, $\pi(Y_s) = P(s)$ and $\pi(Y_t) = P(t)$). Hence, we obtain that $\forall \pi \in \Pi(F)$ $\pi(F, S, \texttt{cf}) \leq \min_{s,t \in S} \{(1 - P(\delta_{st})) \cdot \min\{P(s), P(t)\}\}$. Since the right-hand side of this inequality (from now on denoted as *MAX*) coincides with the right-hand side of the formula for $P_{\max}$, it remains to be proved only that there is some $\pi \in \Pi(F)$ for which this inequality holds as an equality. Let $s_1, \ldots, s_k$ be the arguments of S in ascending order of marginal probability. We consider the $k+1$ sets of possible worlds: $PW_0, \ldots, PW_k$ where each PW_i contains the possible worlds containing all the arguments $s_{i+1}, \ldots, s_k$ (if $i < k$), but not $s_1, \ldots, s_i$ (if $i \geq 1$). Observe that PW_k contains only the empty possible world, and that PW_0 contains, among others, the possible world ω_{cf} consisting of S and no attack between its arguments. Given this, in order to make $\pi(F, S, \texttt{cf}) = MAX$, we set $\pi(\omega_{\text{cf}}) = MAX$. Then, in order to guarantee that $\forall s \in S$ $\pi \models P(s)$, it suffices to define π so that 1) $\pi(PW_0) = P(s_1)$, 2) $\forall i \in [2..k]$ $\pi(PW_{i-1}) = P(s_i) - P(s_{i-1})$, 3) $\pi(PW_k) = 1 - P(s_k)$ and 4) $\pi(\omega) = 0$ for each possible world ω not in any PW_i. In particular, for each PW_i, the overall probability $\pi(PW_i)$ is distributed among the possible worlds in PW_i by making π consistent also with the attacks' probabilities. To this aim, for each $s, t \in S$, we first compute $\overline{P(\delta_{st})} = \frac{P(\delta_{st}) \cdot \min\{P(s), P(t)\}}{\min\{P(s), P(t)\} - \pi(\omega_{\text{cf}})}$, if $P(\delta_{st}) \neq 0$, or $\overline{P(\delta_{st})} = 0$, otherwise, that is the (conditioned) marginal probability of δ_{st} "rescaled" against the overall probability of the possible worlds different from ω_{cf} and containing s and t. Observe that, since $\pi(\omega_{\text{cf}}) = MAX$, $\overline{P(\delta_{st})} \leq 1$. Then, $\forall i \in [0..k], \omega \in PW_i$ with $\omega \neq \omega_{\text{cf}}$, we set $\pi(\omega) = X_i \cdot \Pi_{\delta_{st} \in \omega} \overline{P(\delta_{st})} \cdot \Pi_{\delta_{st} \notin \omega} (1 - \overline{P(\delta_{st})})$, where $X_0 = \pi(PW_0) - \pi(\omega_{\text{cf}})$ and $X_i = \pi(PW_i)$, for $i > 0$, which ensures $\pi \models P$.

Proof of Case 2. (Sketch) For every $a \in A \setminus S$, the term $M'(a) = 1 - P(a) + \sum_{s \in S} P(\delta_{sa}) \cdot M(\{s, a\})$ is the probability that a does not occur (as measured by $1 - P(a)$), or is attacked by S, assuming that the co-existence of a with each $s \in S$ has minimum probability. Observe that, in $M'(a)$, the sum over the arguments in S means considering the attacks from different arguments in S to a as alternative, as this maximizes the probability that at least one of them occurs. Finally, taking the minimum between $P_{\max}(F, S, \texttt{cf})$ and $\min_{a \in A \setminus S} M'(a)$ means maximizing the probability that S is conflict-free and, for each $a \in A \setminus S$, either a does not occur or is attacked by S.

Proof of Case 3. (Sketch) $(1 - P(\delta_{as})) \cdot M(\{s, a\})$ is the portion of the probability space occupied by the possible worlds where a does not attack s, when the probability that a and s coexist is minimized. Then, $M'(a) = \min_{s \in S}\{(1 - P(\delta_{as})) \cdot M(\{s, a\})\}$ is the maximum probability that a coexists with all the arguments in S, but a attacks no $s \in S$. Analogously, $M''(a) = \sum_{t \in S} P(\delta_{ta}) \cdot M(\{t, a\})$ is the maximum probability that some $t \in S$ attacks a, provided that the probability that a and t coexist is minimized. Hence, $M'''(a) = (1 - P(a)) + M'(a) + M''(a)$ is the maximum probability that a does not occur or does not attack S or is counterattacked by S. Correspondingly, $\min\{P_{\max}(F, S, \texttt{cf}), \min_{a \in A \setminus S} M'''(a)\}$ is the maximum probability that S is conflict free and every argument outside S either does not occur, or occurs but does not attack S, or occurs and is counterattacked by S. □

Example 3 *Applying the formulas above in Example 1, we obtain* $P_{\max}(F, \{a, b\}, \texttt{cf}) = \min\{(1 - P(\delta_{aa})) \cdot \min\{P(a), P(a)\}, (1 - P(\delta_{bb})) \cdot \min\{P(b), P(b)\}, (1 - P(\delta_{ab})) \cdot \min\{P(a), P(b)\}, (1 - P(\delta_{ba})) \cdot \min\{P(b), P(a)\}\} = 0.6$. *As* $M(\{a, c\}) = M(\{b, c\}) = \max\{0, (0.6 + 0.4) - 1\} = 0$, *the formulas for* $\sigma \in \{\texttt{ad}, \texttt{st}\}$ *simplify to* $P_{\max}(F, \{a, b\}, \texttt{ad}) = \min\{0.6, (1 - P(c))\} = 0.6$ *and* $P_{\max}(F, \{a\}, \texttt{st}) = \min\{P_{\max}(F, \{a\}, \texttt{cf}), \min\{1 - P(c), 1 - P(b)\}\} = 0.4$.

As the formulas for $P_{\max}$ in Theorem 1 can be evaluated in time $O(|A|^2)$, we obtain the following corollary.

Corollary 1 *Under* $\sigma \in \{\texttt{cf}, \texttt{st}, \texttt{ad}\}$, MAXP-VER *is in* P.

Under the other Dungean semantics, MAXP-VER becomes intractable. Interestingly, the intractability does not depend on whether the uncertainty involves the arguments or the attacks, since in both cases the problem becomes NP-complete under $\sigma \in \{\texttt{co}, \texttt{gr}\}$ and Σ_2^p-complete under $\sigma = \texttt{pr}$. We start with showing the NP upper bound under $\sigma \in \{\texttt{co}, \texttt{gr}\}$, that will be proved similarly to what done in [Georgakopoulos *et al.*, 1988] for proving that *probabilistic SAT* is in NP.

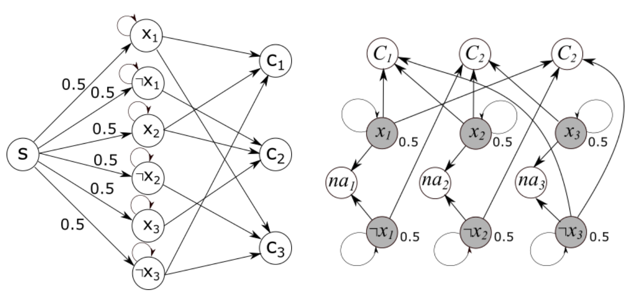

(a) (b)

Figure 1: (a) The mAAF $F(\varphi)$ and (b) the mAAF $F'(\varphi)$ for $\varphi = (x_1 \vee x_2 \vee \neg x_3) \wedge (\neg x_1 \vee x_2 \vee x_3) \wedge (x_1 \vee \neg x_2 \vee \neg x_3)$

Theorem 2 *Given an mAAF* $F = \langle A, D, P \rangle$ *and* $S \subseteq A$, MAXP-VER(F, S, σ, p^*) *is in NP under* $\sigma \in \{\texttt{co}, \texttt{gr}\}$.

Proof. A $\pi \in \Pi(F)$ such that $\pi(F, S, \sigma) \geq p^*$ can be found by solving a system of (in)equalities $\mathcal{S}$ with a positive variable $\pi(\omega_i)$ for each $\omega_i \in PW(F)$. $\mathcal{S}$ contains:
1) the equality $\sum_{\omega_i \in PW(F)} \pi(\omega_i) = 1$ (stating that the overall probability of the possible worlds is 1);
2) $\forall a \in A$ the equality $\sum_{\omega_i \in PW(F) | a \in \omega_i} \pi(\omega_i) = P(a)$ and, $\forall \delta_{ab} \in D$, the equality $\sum_{\omega_i \in PW(F) | \delta_{ab} \in \omega_i} \pi(\omega_i) = P(\delta_{ab}) \cdot \sum_{\omega_i \in PW(F) | a, b \in \omega_i} \pi(\omega_i)$ (imposing $\pi \models P$);
3) the inequality $\sum_{\omega_i \in PW(F) | S \in \text{Ext}(S, \sigma)} \pi(\omega_i) \geq p^*$ (imposing that the overall probability of the possible worlds where S is a σ-extension is not less than p^*).
$\mathcal{S}$ has $m = |A| + |D| + 2$ (in)equalities and $O(2^{|A \cup D|})$ variables, thus, as entailed by linear programming theory, if it has a solution, it has a *basic solution* with at most m non zero values, where the size of each value is polynomially bounded by m and the size of the constants in $\mathcal{S}$. Hence, MAXP-VER(F, S, σ) can be solved by guessing a polynomial-size π over $PW(F)$ assigning a non-zero probability to at most m possible worlds, and then checking if $\pi \models P$ and $\pi(F, S, \sigma) \geq p^*$ (this can be done in polynomial time, as VER is in P for $\sigma \in \{\texttt{co}, \texttt{gr}\}$). □

Theorem 3 states that NP is also a lower bound under $\sigma \in \{\texttt{co}, \texttt{gr}\}$. To prove this result, we rely on some symmetries that arise when reasoning over the pdfs consistent with P, that are exploited to obtain reductions from the NP-complete problem *not-all-equals 3-*SAT *(*NAE3SAT*)*: "*Given a* 3CNF *formula* ϕ*, is there any truth assignment satisfying* ϕ *and such that no clause contains all the three literals set to true?*".

Theorem 3 *Given an mAAF* F *and* $S \subseteq A$*, under* $\sigma \in \{\texttt{co}, \texttt{gr}\}$, MAXP-VER$(F, S, \sigma, p^*)$ *is NP-hard, even if:* 1) *all the arguments are certain,* 2) *all the attacks are certain.*

Proof. In what follows, when referring to CNF formulas, we assume that their form is $C_1 \wedge \cdots \wedge C_m$, where each clause C_j is the disjunction of k_j literals $\bigvee_{h=1}^{k_j} l_h^j$, and each l_h^j is a variable or its negation. We denote the variables as $x_1, \ldots, x_n$. 3CNF denotes formulas where $\forall j \in [1..m]\ k_j = 3$.

Case 1. We consider $\sigma = \texttt{co}$ (an analogous reasoning on the same construction works for $\sigma = \texttt{gr}$). We first introduce the construction of the mAAF $F(\varphi) = \langle A, D, P \rangle$ encoding a generic CNF formula φ. Herein, A consists of an argument s, an argument c_j for each clause C_j, and the arguments $x_i, \neg x_i$ for each variable x_i; D contains, for each $i \in [1..n]$, the defeats (s, x_i), $(s, \neg x_i)$, (x_i, x_i) and $(\neg x_i, \neg x_i)$, as well as, for each clause C_j and literal l_i^j in C_j, the attack (l_i^j, c_j). Finally, P assigns probability $\frac{1}{2}$ to the attacks (s, x_i), $(s, \neg x_i)$ (for $i \in [1..n]$), and 1 to all the other attacks and arguments. Fig. 1(a) depicts an example of $F(\varphi)$.

We show a reduction from NAE3SAT. Given an instance ϕ of NAE3SAT, let $\hat{\phi} = \phi \wedge \overline{\phi} \wedge \phi_X$, where: $\overline{\phi} = \bigwedge_{j=1}^{k}(\neg l_1^j \vee \neg l_2^j \vee \neg l_3^j)$ and $\phi_X = \bigwedge_{i=1}^{n}(x_i \vee \neg x_i)$. We now show the equivalence: "*The instance* ϕ *of* NAE3SAT *is true*" $\Leftrightarrow$ "MAXP-VER$(F(\hat{\phi}, \{s\}, \texttt{co}, 1)$ *is true*".

($\Rightarrow$) As ϕ is a true instance of NAE3SAT, not only there is a truth assignment t satisfying ϕ, but also the complement t' of t satisfies ϕ, and both t and t' set at least one literal in every clause to false (with "complement" we mean: $\forall i \in [1..n]$, $t'(x_i) = \neg t(x_i)$). Hence, t and t' make also $\hat{\phi}$ satisfied. Let $F(\hat{\phi}) = \langle \hat{A}, \hat{D}, P\rangle$ and $\omega = \langle A, D\rangle$, $\omega' = \langle A', D'\rangle$ be the possible worlds of $F(\hat{\phi})$ such that: 1) $A = A' = \hat{A}$, and 2) $D = \hat{D} \setminus \{(s, l_i) \mid (l_i = x_i \wedge t(x_i) = \text{true}) \vee (l_i = \neg x_i \wedge t(x_i) = \text{false})\}$, and 3) $D' = \hat{D} \setminus \{(s, l_i) \mid (l_i = x_i \wedge t'(x_i) = \text{true}) \vee (l_i = \neg x_i \wedge t'(x_i) = \text{false})\}$). Now, $\{s\}$ is an extension in ω and ω', as t and t' make $\hat{\phi}$ satisfied (thus every c_j is attacked by some x_i or $\neg x_i$ not attacked by s). Let π be the pdf over $PW(F(\hat{\phi}))$ assigning probability $\frac{1}{2}$ to both ω and ω', and 0 to the other possible worlds. By construction of ω, ω', it follows that $\pi \models P$. Hence, since $\{s\}$ is a co-extension in ω and ω', and $\pi(\omega) + \pi(\omega') = 1$, MAXP-VER$(F(\hat{\phi}), \{s\}, \texttt{co}, 1)$ is true.

($\Leftarrow$) As MAXP-VER$(F(\hat{\phi}), \{s\}, \texttt{co}, 1)$ is true, there are $k \geq 0$ possible worlds $\omega_1, \dots, \omega_k$ of $F(\hat{\phi})$ such that 1) $\{s\}$ is a complete extension of $\omega_1, \dots, \omega_k$, and 2) there is a pdf $\pi \in \Pi(F(\hat{\phi}))$ assigning non-zero probability only to $\omega_1, \dots, \omega_k$.

We first prove that, for each $i \in [1..n], j \in [1..k]$ ω_j contains exactly one of the attacks (s, x_i) or $(s, \neg x_i)$. Reasoning by contradiction, assume that there is a ω_j containing both (s, x_i) and $(s, \neg x_i)$. This implies that the argument corresponding to the clause $(x_i \vee \neg x_i) \in \hat{\phi}$ is acceptable w.r.t. $\{s\}$ in ω_j, thus contradicting that $\{s\}$ is a complete extension in ω_j. Vice versa, assume that there is an ω_j not containing (s, x_i) and $(s, \neg x_i)$. Since (s, x_i) and $(s, \neg x_i)$ do not occur simultaneously in any possible world in $\omega_1, \dots, \omega_{j-1}, \omega_{j+1}, \dots, \omega_k$, it follows that $P((s, x_i)) + P((s, \neg x_i)) = 1 - \pi(\omega_j)$. But $\pi(\omega_j) > 0$ and $P((s, x_i)) = P((s, \neg x_i)) = \frac{1}{2}$, thus we reach the contradiction $\frac{1}{2} + \frac{1}{2} < 1$.

Given that ω_1 contains, for each variable x_i, exactly one of the attacks (s, x_i) or $(s, \neg x_i)$, the relation t between the variables and the truth values *true*, *false* defined below is a truth assignment over $x_1, \dots, x_n$: $\forall x_i \in \{x_1, \dots, x_n\}$, $t(x_i) =$ *true* if $(s, x_i) \in \omega_1$, and $t(x_i) =$ *false* if $(s, \neg x_i) \in \omega_1$. It is easy to see that t makes $\hat{\phi}$ true, as for each C_j in $\hat{\phi}$ there is at least one argument l_i (of the form x_i or $\neg x_i$) attacking the argument c_j such that (s, l_i) does not appear in ω_1 (otherwise, $\{s\}$ would not be a complete extension in ω_1). This ends the proof, since the satisfiability of $\hat{\phi}$, by construction, implies that ϕ is a true instance of NAE3SAT.

<u>*Case* 2.</u> We reason analogously to Case 1, but use a new construction. For any CNF formula φ, we consider the mAAF $F'(\phi) = \langle A, D, P\rangle$ where: A contains $\forall i \in [1..n]$ the arguments $x_i, \neg x_i, na_i$, and $\forall j \in [1..k]$ an argument c_j; D contains, $\forall i \in [1..n]$, the attacks $(x_i, na_i), (\neg x_i, na_i), (x_i, x_i)$ and $(\neg x_i, \neg x_i)$, and, $\forall j \in [1..k]$, the attacks $(l_1^j, c_j), (l_2^j, c_j)$, (l_3^j, c_j). P assigns 1 to every attack and to $na_1, \dots, na_n, c_1, \dots, c_k$, and $\frac{1}{2}$ to $x_1, \dots, x_n, \neg x_1, \dots, \neg x_n$. An example of $F'(\varphi)$ is in Fig. 1(b). Then, similarly to Case 1, the equivalence "*An instance* ϕ *of* NAE3SAT *is true*" $\Leftrightarrow$ "MAXP-VER$(F'(\phi), \emptyset, \texttt{co}, 1)$ *is true*" can be proved. The difference here is that a truth assignment t can be encoded by a possible world containing, $\forall i \in [1..n]$, either x_i or $\neg x_i$, where $t(x_i) =$ *true* (resp., *false*) is encoded by the presence of x_i (resp., $\neg x_i$). Then, the fact that a clause C_j is made true by having assigned *true* (resp., *false*) to x_i is encoded by the presence of the attack (x_i, c_j) (resp., $(\neg x_i, c_j)$). □

Proposition 2 relates MAXP-VER with the verification problem INCVER over *incomplete AAFs* (iAAF), that are AAFs where some arguments and attacks are marked as uncertain (with no measure of their uncertainty). Herein, INCVER(F^i, S, σ) asks if S is a σ-extension in some possible world of F^i. This proposition states that INCVER can be reduced in polynomial time to MAXP-VER, and will be used to characterize the complexity of MAXP-VER under $\sigma = \texttt{pr}$. However, it is of independent interest, and it will be exploited in the next section, where we discuss the relationship between mAAFs and variants of AAFs dealing with uncertainty.

Proposition 2 *There is a Karp reduction from* INCVER *to* MAXP-VER *adding no new uncertain arguments and attacks.*

Proof. Given INCVER(F^i, S, σ), where A and D (resp., $A^?$ and $D^?$) are the certain (resp., uncertain) arguments and attacks of F^i, let $F = \langle A \cup A^?, D \cup D^?, P\rangle$ be the mAAF where $\forall x \in A \cup D$ $P(x) = 1$ and $\forall x \in A^? \cup D^?$ $P(x) = 1/2$. Let π be the following pdf over $PW(F)$: $\pi(\omega) = \Pi_{a \in A^?} 1/2 \cdot \Pi_{\delta_{ab} \in D^? | a,b \in \omega} 1/2$. Basically, π is the pdf entailed by assuming independence between the arguments and conditioned independence between attacks, thus $\pi \models P$. Let $p^* = \min_{\omega \in PW(F)} \pi(\omega) = \left(\frac{1}{2}\right)^{|A^?|} \times \left(\frac{1}{2}\right)^{|D^?|}$. It is easy to see that the answers of INCVER(F^i, S, σ) and MAXP-VER(F, S, σ, p^*) coincide. □

Theorem 4 MAXP-VER(F, S, σ, p^*) *is* Σ_2^p*-complete under* $\sigma = \texttt{pr}$*, even if all the arguments or all the attacks are certain.*

Proof. The membership proof is analogous to Theorem 2: solving VER over a possible world now requires an *NP* oracle, thus the problem is in NP^{NP}. Proposition 2 and the Σ_2^p-hardness of INCVER (when the arguments or the attacks are certain [Baumeister *et al.*, 2018]) imply the hardness. □

We now turn our attention to MINPVER, the dual problem referring to the minimum probability. Its complexity is characterized by the following theorem.

Theorem 5 *Given an mAAF* $F = \langle A, D, P\rangle$ *and* $S \subseteq A$*,*

1) MINP-VER$(F, S, \texttt{cf}, p^*)$ *is in P, and* $P_{\min}(F, S, \sigma) = \max\{0, M(S) - \sum_{a,b \in S} P(\delta_{ab}) \cdot M(\{a, b\})\}$*;*
2) MINP-VER$(F, S, \texttt{pr}, p^*)$ *is NP-complete;*
3) MINP-VER *is in NP under* $\sigma \in \{\texttt{ad}, \texttt{st}, \texttt{co}, \texttt{gr}\}$.

Proof. 1) and 3) can be obtained reasoning analogously to the proofs of Theorems 1 and 2, respectively. As for $\sigma = \texttt{pr}$, the membership to *NP* follows from modifying the proving strategy of Theorem 2 by replacing inequality I_3 with an I'_3 imposing that the overall probability that S is *not* preferred is not less than $1 - p^*$, and making the guess-phase include also a witness for each of the guessed $\pi(\omega_i)$ involved in I'_3 certifying that S is not preferred in ω_i. A reduction from the complement of $\textsc{ver}(F, S, \mathtt{pr})$ implies the hardness (it suffices to translate F into an mAAF F' with no uncertain arguments/attacks and decide $\textsc{min}$P-$\textsc{Ver}(F', S, \mathtt{pr}, 1/2)$). □

The case $\sigma = \mathtt{pr}$ is particularly interesting, since it shows an asymmetry between MAXP-VER and MINP-VER. We conjecture that an asymmetry holds also under $\sigma \in \{\mathtt{co}, \mathtt{gr}\}$ as replacing the $\max$ operator with $\min$ when moving from MAXP-VER to MINP-VER does not allow the form of exclusive disjunction exploited in our reductions. Thus, overall, we conjecture that the upper bound in point 2) is not tight, and MINP-VER is in P under $\sigma \in \{\mathtt{ad}, \mathtt{st}, \mathtt{co}, \mathtt{gr}\}$.

5 Relationship with Probabilistic AAFs and Further Related Work

mAAFs are in the family of probabilistic AAFs adopting the *constellations approach* (denoted as *prAAFs*), as their semantics (differently from the *epistemic approach* [Hunter and Thimm, 2014; Hunter *et al.*, 2020; Baier *et al.*, 2021; Potyka, 2021], where probabilities are degrees of belief in arguments' acceptance) is based on considering different possible worlds for the argumentation. The prAAFs in the literature typically consider a single pdf over the possible worlds and differ in how this pdf is encoded. *Exhaustive* prAAFs (EX [Dung and Thang, 2010]) enumerate the possible worlds $\omega_1, \ldots, \omega_k$ having a chance to be the actual argumentation, and assign a probability to each of them. Evaluating the probability of extensions over EX reduces to solving VER over every ω_i (with $i \in [1..k]$). Hence, MAXP-VER over EX is in P under $\sigma \in \{\mathtt{cf}, \mathtt{ad}, \mathtt{st}, \mathtt{co}, \mathtt{gr}\}$, for which VER is in P, and is in $P^{||NP}$ under $\sigma \in \{\mathtt{pr}\}$, as in this case it is solvable with k parallel invocations of an NP-oracle solving VER. This seems to mean that MAXP-VER over EX is never more complex than over mAAFs: in particular, it is easier under $\sigma \in \{\mathtt{co}, \mathtt{gr}\}$ (where it is in P over EX and NP-complete over mAAFs) and under $\sigma \in \{\mathtt{pr}\}$ (where it is in $P^{\|NP}$ over EX and NP^{NP}-complete over mAAFs). However, these differences "in favor of" EX follow from the explicit representation of the pdf in the encoding of EX. Thus, the measure of computational complexity over EX benefits from a discount of one exponential level compared with mAAFs, where enumerating the possible worlds has an exponential cost. In fact, such a discount is paid when an EX is defined, as the analyst is due to enumerate all the alternative scenarios, and this may be prohibitive when the possibilities are many. Things change significantly when considering *general* prAAFs (GEN [Fazzinga *et al.*, 2019]), where the pdf over the possible worlds is compactly encoded via *world-sets descriptors*: here, evaluating the probabilities of extensions becomes $\#P$-hard for all the Dungean semantics (while MAXP-VER over mAAFs may be in P, NP-complete or Σ_2^p-complete, depending on σ). Finally, *independence-based* prAAFs (IND [Li *et al.*, 2011; Fazzinga *et al.*, 2015]) share with mAAFs the specification of the marginal probabilities of arguments/attacks, but they assume independence, thus a unique pdf over the possible worlds. Here, evaluating extensions' probabilities can be done in PTIME under $\sigma \in \{\mathtt{cf}, \mathtt{ad}, \mathtt{st}\}$ (the same as MAXP-VER over mAAFs) and is $FP^{\#P}$-complete for the other semantics. Although a precise characterization of the decision problem MAXP-VER over IND is not in the literature, we can draw some conclusions: under $\sigma \in \{\mathtt{co}, \mathtt{gr}\}$, MAXP-VER over IND cannot be simpler than over mAAFs (where it is NP-complete), as the fact that the minimum difference between the probabilities of two possible worlds is known and of polynomial size allows for exploiting a polynomial number of invocations to a MAXP-VER's solver to obtain an extension probability. If MAXP-VER over IND were in NP, this would imply $FP^{\#P} = FP^{NP}$.

mAAFs are also related to iAAFs [Baumeister *et al.*, 2018; Alfano *et al.*, 2022]: Proposition 2 states the reducibility of the verification problem INCVER over iAAFs to MAXP-VER. From our results, it turns out that reasoning over iAAFs is of the same complexity as over mAAFs under $\sigma \in \{\mathtt{cf}, \mathtt{ad}, \mathtt{st}\}$ and $\sigma = \mathtt{pr}$, where both problems are in P and Σ_2^p-complete, respectively, while, under $\sigma \in \{\mathtt{co}, \mathtt{gr}\}$, reasoning over iAAFs is strictly simpler than over mAAFs (as INCVER is in P [Fazzinga *et al.*, 2020] while MAXP-VER is NP-complete). Overall, mAAFs are between iAAFs and IND: introducing measures of uncertainty on top of iAAFs (thus obtaining mAAFs) can increase the computational complexity, as well as introducing the independence assumption on top of mAAFs (thus obtaining IND). Intuitively, a reason for the raise in complexity when moving to IND is that the independence assumption restricts $\Pi(F)$ to a singleton, and this can allow, under some semantics, a form of counting (as the cardinality of a set of possible worlds can be inferred from its overall probability).

Further works related to ours are those introducing probabilistic variants of AAFs generalizations (such as *probabilistic bipolar AAFs* [Fazzinga *et al.*, 2018], *probabilistic Control AFs* [Gaignier *et al.*, 2021], *probabilistic Abstract Dialectical Frameworks* [Polberg and Doder, 2014]) as well as those extending AAFs with preferences [Amgoud and Vesic, 2011], degrees of beliefs [Santini *et al.*, 2018], social values [Bench-Capon, 2003; Atkinson and Bench-Capon, 2016].

6 Conclusions and Future Work

We have introduced *AAFs with marginal probabilities* (mAAFs), where arguments and attacks are probabilistic events whose marginal probability is known (under no independence assumption). We have characterized the complexity of MAXP-VER and MINP-VER, that are the decision counterparts of maximizing and minimizing the probability that a set is an extension. In future work, we plan to prove our conjectures on the complexity of MINP-VER stated in Section 4, and to extend mAAFs with the specification of correlations among arguments/attacks, as done for iAAFs [Fazzinga *et al.*, 2021a; Fazzinga *et al.*, 2021b; Mailly, 2021]. This would exclude pdfs assigning non-zero probability to unrealistic possible worlds from the reasoning. Another interesting direction for future work is the investigation of constraint-programming approaches for the NP-hard cases.

References

[Alfano *et al.*, 2022] Gianvincenzo Alfano, Sergio Greco, Francesco Parisi, and Irina Trubitsyna. Incomplete argumentation frameworks: Properties and complexity. In *Proc. of AAAI*, 2022.

[Amgoud and Vesic, 2011] Leila Amgoud and Srdjan Vesic. A new approach for preference-based argumentation frameworks. *Ann. Math. Artif. Intell.*, 63(2), 2011.

[Atkinson and Bench-Capon, 2016] Katie Atkinson and Trevor Bench-Capon. Value based reasoning and the actions of others. In *Proc. European Conf. on Artificial Intelligence (ECAI)*, pages 680–688, 2016.

[Baier *et al.*, 2021] Christel Baier, Martin Diller, Clemens Dubslaff, Sarah Alice Gaggl, Holger Hermanns, and Nikolai Käfer. Admissibility in probabilistic argumentation. In *Proc, Conf. on Principles of Knowledge Representation and Reasoning, KR*, pages 87–98, 2021.

[Baumeister *et al.*, 2018] Dorothea Baumeister, Daniel Neugebauer, Jörg Rothe, and Hilmar Schadrack. Verification in incomplete argumentation frameworks. *Artif. Intell.*, 264:1–26, 2018.

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

[Doder and Woltran, 2014] Dragan Doder and Stefan Woltran. Probabilistic argumentation frameworks - A logical approach. In *Proc. Conf. Scalable Uncertainty Management (SUM)*, pages 134–147, 2014.

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

[Dung and Thang, 2010] Phan Minh Dung and Phan Minh Thang. Towards (probabilistic) argumentation for jury-based dispute resolution. In *Proc. Computational Models of Argument (COMMA)*, pages 171–182, 2010.

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

[Fazzinga *et al.*, 2015] Bettina Fazzinga, Sergio Flesca, and Francesco Parisi. On the complexity of probabilistic abstract argumentation frameworks. *ACM Trans. Comput. Log.*, 16(3):22:1–22:39, 2015.

[Fazzinga *et al.*, 2018] Bettina Fazzinga, Sergio Flesca, and Filippo Furfaro. Probabilistic bipolar abstract argumentation frameworks: complexity results. In *Proc. Joint Conf. on Artificial Intelligence (IJCAI)*, 2018.

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

[Fazzinga *et al.*, 2020] Bettina Fazzinga, Sergio Flesca, and Filippo Furfaro. Revisiting the notion of extension over incomplete abstract argumentation frameworks. In *Proc. Joint Conf. on Artificial Intelligence (IJCAI)*, 2020.

[Fazzinga *et al.*, 2021a] Bettina Fazzinga, Sergio Flesca, and Filippo Furfaro. Reasoning over argument-incomplete aafs in the presence of correlations. In *Proc. Joint Conf. on Artificial Intelligence (IJCAI)*. ijcai.org, 2021.

[Fazzinga *et al.*, 2021b] Bettina Fazzinga, Sergio Flesca, and Filippo Furfaro. Reasoning over attack-incomplete aafs in the presence of correlations. In *Proc. Principles of Knowledge Representation and Reasoning (KR)*, 2021.

[Gaignier *et al.*, 2021] Fabrice Gaignier, Yannis Dimopoulos, Jean-Guy Mailly, and Pavlos Moraitis. Probabilistic control argumentation frameworks. In *Proc. of Conf. on Autonomous Agents and Multiagent Systems AAMAS*, pages 519–527, 2021.

[Georgakopoulos *et al.*, 1988] George F. Georgakopoulos, Dimitris J. Kavvadias, and Christos H. Papadimitriou. Probabilistic satisfiability. *J. Complex.*, 4(1):1–11, 1988.

[Hunter and Thimm, 2014] Anthony Hunter and Matthias Thimm. Probabilistic argumentation with epistemic extensions. In *Proc. of the International Workshop on Defeasible and Ampliative Reasoning, DARe@ECAI*, 2014.

[Hunter *et al.*, 2020] Anthony Hunter, Sylwia Polberg, and Matthias Thimm. Epistemic graphs for representing and reasoning with positive and negative influences of arguments. *Artif. Intell.*, 281:103236, 2020.

[Hunter, 2012] Anthony Hunter. Some foundations for probabilistic abstract argumentation. In *Proc. Computational Models of Argument (COMMA)*, pages 117–128, 2012.

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

[Li *et al.*, 2011] Hengfei Li, Nir Oren, and Timothy J. Norman. Probabilistic argumentation frameworks. In *Proc. Int. Workshop on Theory and Applications of Formal Argumentation (TAFA)*, pages 1–16, 2011.

[Mailly, 2021] Jean-Guy Mailly. Constrained incomplete argumentation frameworks. In *Proc. of European Conf. on Symbolic and Quantitative Approaches to Reasoning with Uncertainty, ECSQARU*, pages 103–116, 2021.

[Nilsson, 1986] Nils J. Nilsson. Probabilistic logic. *Artif. Intell.*, 28(1):71–87, 1986.

[Polberg and Doder, 2014] Sylwia Polberg and Dragan Doder. Probabilistic abstract dialectical frameworks. In *Proc. European Conf. on Logics in Artificial Intelligence, JELIA*, pages 591–599, 2014.

[Potyka, 2021] Nico Potyka. From probabilistic programming to probabilistic argumentation. In *Proc. of the International Conference on Logic Programming (ICLP)*, 2021.

[Rienstra, 2012] Tjitze Rienstra. Towards a probabilistic dung-style argumentation system. In *Proc. Int. Conf. on Agreement Technologies (AT)*, pages 138–152, 2012.

[Santini *et al.*, 2018] Francesco Santini, Audun Jøsang, and Maria Silvia Pini. Are my arguments trustworthy? abstract argumentation with subjective logic. In *Proc. Int. Conf. on Information Fusion (FUSION)*, pages 1982–1989, 2018.