

Epistemic Argumentation with Conditional Probabilities and Labeling Constraints

Glauber de Bona

Victor Hugo Nascimento Rocha

Fabio Gagliardi Cozman

Escola Politécnica, Universidade de São Paulo, SP – Brazil

GLAUBER.BONA@USP.BR

VICTOR.HUGO.ROCHA@USP.BR

FGCOZMAN@USP.BR

Abstract

We extend epistemic graphs, a powerful representation language employed in argumentation theory, first, by allowing conditional probabilities in that language. We also offer a new way of interpreting the graph as a set of restrictions based on a selected semantics for the abstract argumentation frameworks. The resulting semantics for epistemic graphs are given by credal sets that we characterize through inequalities. We illustrate the main issues in our proposals by resorting to arguments related to climate change.

Keywords: Argumentation, Epistemic graphs, Probabilistic logic, Conditional probabilities.

1. Introduction

Human argumentation is a remarkably complex activity with a vast array of applications, for instance in negotiation, in political debate and in court proceedings [25]. During the last few decades several computational techniques have been proposed to support human argumentation and to enable machine argumentation. In particular, since 1995 there has been significant activity around *argumentation frameworks* as defined by Dung [7]. Argumentation frameworks take arguments in a very abstract form, looking at ways to decide which arguments prevail and which fail without examining the internal structure of each argument. Some of the ensuing literature studies arguments handled by a single agent, while others look at sets of agents; some allow for specification of preferences or values amongst arguments, while others take beliefs and uncertainties into account; some aim at convincing disagreeing parties, while others are primarily concerned with building consensus around decisions. In this paper we work within this research program, and we equate “argumentation” with “argumentation as it appears in work aimed at artificial intelligence”.

There are many *probabilistic* argumentation schemes that attach to an argument A a probability $P(A)$. Several such theories resort to inequalities over probabilities. An example can be found in Thimm’s proposal for probabilistic argumentation frameworks [24]: there the statement that argument A is attacked by arguments B_1, \dots, B_n is translated

into the inequalities

$$1 - \sum_i P(B_i) \leq P(A) \leq 1 - \max_i B_i.$$

The intuition here is that $P(A)$ is bounded from above by each one of $1 - P(B_i)$, that is by the belief in the “complement” of its attacker B_i , and also $P(A)$ is bounded from below by one minus the the sum of beliefs in all attackers of A . Yet another example is to be found in *epistemic graphs* [14, 16] as those allow for sentences such as

$$(P(A) \geq 0.7) \wedge (P(A) < 0.9).$$

Epistemic graphs combine the graph-theoretical tools that have been central to argumentation theory in the past two decades with rather general probabilistic assessments; the usual semantics is based on attaching a credal set to any given epistemic graph.

In this paper we expand the language of epistemic graphs, first, by adding assessments based on conditional probability. For instance, we allow constraints such as

$$P(A|B) \geq 0.6 \quad \text{or} \quad P(A|B \wedge C) \leq P(A|B).$$

We show that many relevant inferences can be reduced to linear integer programming and related optimization techniques, and we examine the consequences of assessments based on our proposed constraints. We also introduce a novel semantics for the nodes and edges that appear in an epistemic graph, by introducing novel constraints based on labelings (we refer to them as *labeling constraints*). For instance, one might require that only admissible labelings should get positive probability. Thus the new semantics introduces a possible way of interpreting the epistemic graph structure. Throughout we use examples from debates on connections between climate changes and oceans, a topic that has produced countless arguments in recent years.

In Section 2 we review relevant concepts about probabilistic argumentation. Section 3 looks at possible ways to employ conditional probabilities in epistemic graphs and their resulting semantics. In Section 4 we examine how constraints affect edges in graphs, and in Section 5 we look at novel constraints that can be extracted from graphs. Section 6 presents a few thoughts on difficulties still faced by probabilistic argumentation theory.

2. Argumentation and Probabilities

In this paper we suppose that agents are trying to reach decisions as to which arguments, from a given pool of arguments, are acceptable and which are not. Given that arguments can support or oppose each other in defeasible ways [11, 25], the direct construction of a state space associated with probabilities and utilities may not be trivial. Thus it makes sense to resort to argumentation graphs and other devices of argumentation theory. On the other hand, it seems also sensible to adopt probability theory as a key decision support tool when it comes to uncertainty — even though there are other proposals based on weights or possibility measures in the literature [1]. (Note that we do not employ utilities in this paper, leaving the discussion about values and actual decision making to future work.)

There are two main approaches to probabilistic argumentation in the literature, often referred to as the *constellation* and the *epistemic* approaches [14]. Before we describe them, we summarize the main concepts employed in argumentation frameworks [5, 7].

An *abstract argumentation framework (AAF)*, as originally proposed by Dung [7], is a tuple $(\mathcal{A}, \rightarrow)$ where \mathcal{A} is a set whose elements are abstract arguments and \rightarrow is a binary relation in $\mathcal{A} \times \mathcal{A}$. The intended interpretation of $A \rightarrow B$ is that A attacks B ; that is, if argument A is accepted then argument B is to be rejected. These days the literature considers not only attacks but also supports; we might write $A \bar{\rightarrow} B$ to indicate that A attacks B and $A \overset{+}{\rightarrow} B$ to indicate that A supports B . Moreover, as noted previously, there are extensions to Dung’s AAFs where weights can be attached to arguments, attacks, supports. Figures 1 and 2 illustrate these relations in a particular domain, namely, climate change in the Blue Amazon (a vast region of the Southern Atlantic Ocean). A directed graph where each node is an argument and each edge corresponds to an attack or support is an *argumentation graph*.

The semantics of an AAF is established by associating each argument to a label, either in, out, or undecided, respectively indicating whether it is accepted, rejected, or remains neither accepted nor rejected. A *labeling* $L: \mathcal{A} \rightarrow \{\text{in, out, undecided}\}$ assigns a label to each argument (Dung defined semantics using *extensions*, but labelings are equivalent [4]). A labeling is *conflict-free* if for no A, B that are in we have $A \rightarrow B$; a conflict-free labeling is *admissible* when: first, for each argument A that is *rejected*, there is an argument B that is in and that attacks A ; second, for each argument A that is in and that attacks B that attack A are out. And an admissible labeling is *complete* when for each undecided argument A there is no in argument B that attacks A and there is some argument C that attacks A and that is not out. A complete labeling is said *stable* when there is no argument labeled undecided; it is *semi-stable* when the set of arguments labeled undecided is minimal.

There are many other semantics besides the stable and the semi-stable ones [4].

Now suppose we take \mathcal{A} as a sample space, and specify a probability measure over it. What could these probabilities mean? One possibility is to interpret $P(A)$ as the probability that A is in fact included in the argumentation graph. An event (a set of arguments) is then mapped to a subgraph of the original argumentation graph containing all arguments. Thus the probability measure over \mathcal{A} induces a probability measure over a constellation of argumentation graphs. We may then compute the probability that some argument A is labeled in: it is the probability of all those subgraphs where A is in. Likewise, we can compute the probability that A is out or left undecided. Still, the meaning of $P(A)$ may be a bit mysterious in the constellations approach: if an event with arguments A and B is realized, then one should expect both of these arguments to hold, thus it does not seem reasonable to question whether they should be accepted or rejected. Such issues have been discussed by Hunter, who has proposed that $P(A)$ should be treated as the probability that A is a *justified* contribution to the argumentation graph [13]. We note that even if this perspective is embraced, there are still issues with assumptions of independence usually made within the constellation approach [10, 18].

The epistemic approach is different. It assumes that all arguments and relations are given in a single argumentation graph, and that probabilities reflect whether arguments are believed or not. One of the first proposals related to epistemic argumentation, by Thimm [24], explicitly indicated that $P(A)$ captures the degree of belief of A and then extracted labelings from those probabilities. Similarly, Hunter and Thimm take probabilities to denote “the belief that an agent has that an argument is justifiable” (that is, premises and associated reasoning are valid), and then extract labelings from them [14]. This seems to be the perspective in the recent work on epistemic graphs [14, 16] that we use as starting point in our investigation.

An epistemic graph combines nodes and edges respectively standing for arguments and their relations, plus inequalities connecting probabilities of arguments. Figure 3 depicts the structure of an epistemic graph constructed using arguments about whether an investment in the preservation of the Blue Amazon is appropriate.

In many ways, epistemic graphs offers an extremely powerful language that subsumes and extends most of what has been done within epistemic argumentation. Notably, so far epistemic graphs do not mention conditional probabilities; this is a gap we address in the next section.

3. Epistemic Graphs and Conditional Probabilities

In this section we introduce epistemic graphs that are endowed with conditional probabilities, thus extending the

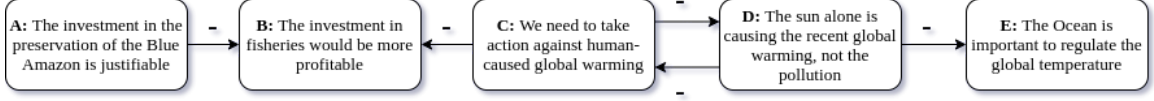


Figure 1: Attack relation between arguments.




Figure 2: Support relation between arguments.

epistemic graphs in the literature [14, 16]. Instead of first presenting the current language of epistemic graphs and then inserting conditional probabilities into that language, we present our complete language at once, commenting on how it differs from past proposals.

Note that we introduce epistemic graphs with the power to represent conditional probabilities in two incremental steps. Firstly, we use linear restrictions over unconditional probabilities. These include constraints like $P(A \wedge B) - 0.2P(B) \geq 0$ that represent $P(A|B) \geq 0.2$.¹ Then we expand the language to allow for linear restrictions directly over conditional probabilities such as $P(A|B) - P(B|C) \geq 0$.

Epistemic graphs will be introduced as pairs consisting of argumentation graphs and probabilistic constraints over these arguments. The semantics of the probabilistic constraints will be precisely defined in this section, but the semantics of the graph, thus the semantics of the whole epistemic graph, will be discussed in Sections 4 and 5.

3.1. The Language \mathcal{L}_{linear}

To formally introduce our languages for probabilistic constraints, let (\mathcal{G}, ℓ) be a *labelled graph*, where $\mathcal{G} = (V, R)$ is a directed graph, with V being a finite set of arguments (nodes) and $R \subseteq V \times V$ is the set of arcs. The function $\ell : R \rightarrow 2^\Omega$ associates a set of labels from $\Omega = \{+, -, *\}$ to each arc, which are read as support, attack or dependency, respectively. If $\ell(r)$ is a singleton for every $r \in R$, the graph is *uni-labelled*; if additionally $\ell(r) \in \{+, -\}$ for all $r \in R$, the graph is *bipolar*.

In the first language we present, which is basically the one proposed by Fagin, Halpern and Meggido [9], probabilistic linear constraints are built over the probability of propositional formulas (*terms*) over the arguments, like $\phi = A \vee \neg B$. Formally, we define an *epistemic atom* as

$$\sum_{i=1}^n \alpha_i P(\phi_i) \geq \alpha_0$$

1. We assume that conditional probabilities are defined only when the conditioning proposition has positive probability.

where, for all i , ϕ_i is a term and $\alpha_i \in \mathbb{Q}$. The language \mathcal{L}_{linear} is then defined as the set of Boolean formulas formed with epistemic atoms, called here *epistemic formulas* [9]. Adopting the graph shown in Figure 3, the following would be examples of some epistemic formulas in \mathcal{L}_{linear} :

$$\begin{aligned} P(A \wedge B) - 0.5P(B) &\leq 0, \\ P(A \wedge C) - 0.7P(C) &\geq 0, \\ P(B) > 0.5 \vee P(D) > 0.5. \end{aligned} \quad (1)$$

The semantics is then given by a *belief distribution* $\pi : 2^V \rightarrow [0, 1]$. Each $\Gamma \subseteq V$ can be seen as a possible world or valuation, where $\Gamma \models A$ iff $A \in \Gamma$. As usual, this satisfaction relation can be extended via $\Gamma \models \neg\phi$ iff $\Gamma \not\models \phi$; $\Gamma \models \phi_1 \wedge \phi_2$ iff $\Gamma \models \phi_1$ and $\Gamma \models \phi_2$; and $\Gamma \models \phi_1 \vee \phi_2$ iff $\Gamma \models \phi_1$ or $\Gamma \models \phi_2$. A belief distribution π is assumed to behave as a probability distribution, so that for each term ϕ :

$$P_\pi(\phi) = \sum \{\pi(\Gamma) | \Gamma \models \phi\}.$$

A belief distribution π satisfies an epistemic atom $\psi = \sum_{i=1}^n \alpha_i P(\phi_i) \geq \alpha_0$ iff $\sum_{i=1}^n \alpha_i P_\pi(\phi_i) \geq \alpha_0$, which we denote by $\pi \models \psi$. The classical semantics apply to epistemic formulas: $\pi \models \neg\psi$ iff $\pi \not\models \psi$; $\pi \models \psi_1 \wedge \psi_2$ iff $\pi \models \psi_1$ and $\pi \models \psi_2$; $\pi \models \psi_1 \vee \psi_2$ iff $\pi \models \psi_1$ or $\pi \models \psi_2$. If $\mathcal{C} \subseteq \mathcal{L}_{linear}$ is a set of formulas, $\pi \models \mathcal{C}$ iff $\pi \models \psi$ for every $\psi \in \mathcal{C}$. We also use $\mathcal{C} \models \phi$ to denote entailment between a set of formulas \mathcal{C} and a formula $\phi \in \mathcal{L}_{linear}$, which means that $\pi \models \phi$ for every belief distribution π such that $\pi \models \mathcal{C}$. We use \mathcal{C}^* to denote the closure of \mathcal{C} , which is the set of all formulas $\phi \in \mathcal{L}_{linear}$ entailed by it.

For simplicity, the only mathematical relation allowed in epistemic atoms is \geq . Nevertheless, given the semantics, we can multiply all α_i by -1 to represent a \leq -inequality and negate an atom to emulate strict inequalities ($<$, $>$). Equality and its negation can be represented via conjunctions of atoms. Thus, we might use $<$, $>$, $=$, \neq in atoms to denote these formulas. Furthermore, we use $a \bowtie b$, with $\bowtie \in \{\geq, \leq, >, <, =, \neq\}$, to denote a probabilistic formula $a - b \bowtie 0$.

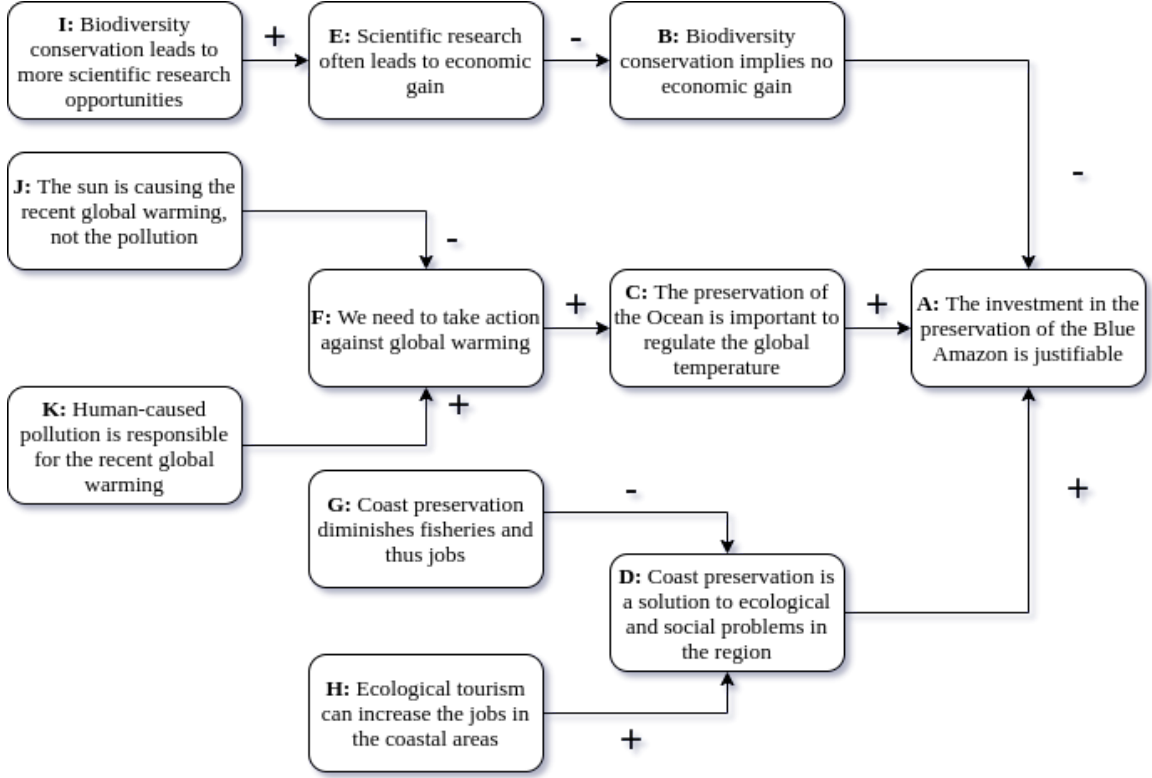


Figure 3: Nodes and edges in epistemic graph.

The epistemic language proposed by Hunter, Polberg and Thimm [16] is a proper subset of \mathcal{L}_{linear} , where the coefficients are restricted to $\alpha_i \in \{-1, 1\}$, for $i \geq 1$, and additionally $\alpha_0 \times \alpha_1 \in [0, 1]$ ². The related computational complexity is the same in both languages, though. For instance, the decision problem of finding a belief distribution satisfying a set of epistemic formulas is NP-complete in both cases. On one hand, both decision problems generalize the Boolean Satisfiability problem (SAT), thus being NP-hard. On the other hand, the satisfiability problem for the language \mathcal{L}_{linear} (the most general of the two) is known to be in NP [9]. Furthermore, both satisfiability problems can be viewed as instances of the Classical Generalized Probabilistic Satisfiability problem (GGenPSAT) [6], also in NP. The latter can be solved in practice for instance via a reduction to Satisfiability Modulo Theory (SMT) with respect to the theory of Quantifier-Free Linear Integer and Real Arithmetic (QF-LIRA) [6].

The key advantage of \mathcal{L}_{linear} is that conditional probability assessments in the form of $P(B|A) \geq q$ can be encoded through epistemic atoms via $P(B \wedge A) - qP(A) \geq 0$. When $P_\pi(A) > 0$, this atom is satisfied exactly by those belief distributions π with $P_\pi(B \wedge A)/P_\pi(A) = q$. For instance, in

2. In the language proposed by Hunter, Polberg and Thimm [16], $\alpha_1 = 1$ and $\alpha_0 \in [0, 1]$ and \leq is in the language. We are allowing also $\alpha_1 = -1$ and $\alpha_0 \in [-1, 0]$ to emulate \leq .

our example represented by Figure 3 and the epistemic formulas 1, the first formula $P(A \wedge B) - 0.5P(B) \geq 0$ encodes $P(A|B) \geq 0.5$.

Combining a labelled graph (\mathcal{G}, ℓ) with a set \mathcal{C} of epistemic formulas, also called constraints, we have an *epistemic graph* $(\mathcal{G}, \ell, \mathcal{C})$.

3.2. The language \mathcal{L}_{cond}

Arcs in an epistemic graph could be given probabilistic semantics as well. The most intuitive, and perhaps the first one we should try, semantics for a bipolar argumentation graph is to allow each arc from A to B to mean either

$$P(B|A) > P(B)$$

when the arc is a supporting one, or

$$P(B|A) < P(B)$$

when the arc is an attacking one. Adopting the ratio formula interpretation ($P(B|A) = P(B \wedge A)/P(A)$), the inequalities above would become non-linear when clearing the denominators. The result cannot be captured by epistemic atoms in \mathcal{L}_{linear} .

To circumvent that, we can enrich our language by allowing conditional probabilities in the epistemic atoms:

$$\sum_{i=1}^n \alpha_i P(\phi_i | \phi'_i) \geq \alpha_0.$$

The set of all Boolean combinations out of these atoms forms the language \mathcal{L}_{cond} .

The semantics of the epistemic atoms is based on the non-linear restrictions obtained when conditional probabilities are interpreted via the ratio formula and the denominators are cleared. We point out that the belief distributions $\pi : 2^V \rightarrow [0, 1]$ are still unconditional. We say a belief distribution π satisfies $\psi = \sum_{i=1}^n \alpha_i P(\phi_i | \phi'_i) \geq \alpha_0$ iff

$$\sum_{i=1}^n \alpha_i \frac{P_\pi(\phi_i \wedge \phi'_i)}{P_\pi(\phi'_i)} \geq \alpha_0$$

or

$$P_\pi(\phi'_i) = 0 \text{ for some } 1 \leq i \leq n.$$

Note that the second condition implies the non-linear restriction resulting from clearing the denominators to be trivially satisfied.

The semantics for epistemic formulas in \mathcal{L}_{cond} is derived again from the semantics of the atoms as usual, and we use the entailment relation \models in the same way. Using a set \mathcal{C} of epistemic formulas in \mathcal{L}_{cond} , we may form also an *epistemic graph* $(\mathcal{G}, \ell, \mathcal{C})$.

By allowing explicitly for non-linearity, \mathcal{L}_{cond} enriches the epistemic approach to argumentation and allows greater expression power. For instance, the epistemic formulas in Expression (2) can be associated with Figure 3 so as to build an epistemic graph:

$$\begin{aligned} P(A|C \wedge D) &\geq P(A|C \vee D), \\ P(A|C) &\geq P(A), \\ P(A|D) &\geq P(A), \\ P(B|E) &\leq P(B), \\ P(E|I) &\geq 0.5, \\ P(C|F) &\geq P(C), \\ P(F|J) &\leq P(F|K), \\ P(D|H) &\geq P(D), \\ P(D|G) &\leq P(D|H). \end{aligned} \quad (2)$$

However, for the corresponding satisfiability problem, the formulas in \mathcal{L}_{cond} are basically polynomial weight formulas from Fagin et al. [9], and they show the satisfiability problem to be PSPACE-complete.

3.3. Recovering Credal and Bayesian Networks

In this section we just point out that, by moving to \mathcal{L}_{cond} , one can capture a large class of probabilistic modeling techniques that rely on independence assumptions. Of course,

this is not too surprising: within that language, one can explicitly impose $P(A|B) - P(A) = 0$ for arguments A and B . So independence is easily asserted; similarly, conditional independence can be readily expressed.

Now, with independence assumptions we can have epistemic graphs that actually specify credal networks [19] or Bayesian networks [21]. To do so, we cannot have any directed cycles in the epistemic graph (that is, we cannot leave one node and, following directed edges, get back to the same node). Note that the assumption of acyclicity of arguments has been made in connection with existing frameworks [8]. Once we have such an acyclic directed graph, we can write down independence relations to the effect that each argument is conditionally independent of its nondescendants given its parents (that is, the conditional probability of each argument given all configurations of nondescendants is not affected by the nondescendants that are not parents). By imposing such independence relations we in fact obtain the usual Markov condition that is adopted both by credal and by Bayesian networks. We then add probability assessments that specify probabilities for each argument given configurations of its parents. If all such probabilities are sharply specified, we obtain a Bayesian network; otherwise, we have a credal network.

Consider our running example on the Blue Amazon. The example was built without any intention to produce a credal network. Still, the graph is acyclic, as it was the natural structure given arguments that we found in public debates. Note that the assessments conveyed by Expression (2) always relate a node with its parents (also the natural constraints to impose). So, it is a typical credal network — minus the independence assumptions that must be imposed separately.

Other kinds of assessments can be used in practice. One possibility is to borrow from *qualitative Bayesian networks* [26]. For instance, in a qualitative Bayesian network we may mark the edge from B to A with label $+$, referred to as a *positive influence*, to mean that

$$P(A|B \wedge (\mathbf{C} = \mathbf{c})) = P(A|B)$$

for all configurations $(\mathbf{C} = \mathbf{c})$ of parents of A that are distinct from B . Finally, if all configurations of arguments and their parents are associated with sharp probabilities, we obtain a Bayesian network.

4. From Probabilistic Constraints to Arcs

An epistemic graph has two separate parts, the graph itself and the probabilistic constraints. In the previous section, we introduced two languages for the probabilistic constraints and defined their semantic via a set of belief distributions. In this section we explore a viewpoint from which the graph structure can be interpreted. One can view, following

Hunter [16], the graph arcs as summarizing some information from the probabilistic constraints, thus verifying if the former are justified by the latter under some reasonable interpretation. We explore this path and adapt it to the presence of conditional probabilities, as these are enabled by the languages we propose.

The semantics of an epistemic graph can be simply defined as the semantics of its probabilistic constraints, which is the credal set satisfying them. From such a point of view, the arcs appearing in the graph should simply reflect properties that hold for the involved arguments. These properties can be analyzed in two steps, first verifying whether an arc is justifiable, then searching for its appropriate labels. One could start with only a set of probabilistic constraints and derive the missing arc. Alternatively, a labeled graph can be given together with the probabilistic constraints and one just checks if the former is in fact corresponding to the latter. Maybe only the arcs are given, and one is searching the correct labels. In any case, the methods are the same, for the problem is always finding whether an arc between a pair of arguments and its label correspond in some precise and justifiable way to the probabilistic constraints.

To illustrate this procedure, let's consider the relation between arguments A and C from the epistemic graph in Figure 3 and the epistemic formulas in Expression (2). Considering just the latter, we have $P(A|C) \geq P(A)$, that could be interpreted as "the probability of A being accepted is greater when C is also accepted". That means a support relation between both the arguments, which corresponds precisely to what is represented in the graph.

Given a set \mathcal{C} of probabilistic constraints in \mathcal{L}_{linear} or \mathcal{L}_{cond} and a labelled graph (\mathcal{G}, ℓ) , with $\mathcal{G} = (V, R)$, one can see each arc $(A, B) \in R$, together with its label $\ell((A, B))$, as representing information already contained in \mathcal{C} . For instance, if A attacks (supports) B , one expects that the constraints somehow entail that a high probability for A implies a low (high) probability for B . This can be encoded in different ways, as Hunter, Polberg and Thimm explore [16].

The basic idea is that an arc (A, B) , regardless its label, means that adding a probability for A to the probabilistic constraints change the set of possible values for the probability of B . This notion is parameterized by a set $F \subset V$ of arguments whose probability are held fixed. In the most flexible definition, not every probabilistic constraint in the set \mathcal{C} is considered, but only a subset of its logical closure \mathcal{C}^* . This is formally captured by the definition below, where the probabilistic constraints can be in \mathcal{L}_{linear} or \mathcal{L}_{cond} :

Definition 1 (By Hunter, Polberg and Thimm [16]) *Let $(\mathcal{G}, \ell, \mathcal{C})$ be an epistemic graph, with $\mathcal{G} = (V, R)$, $Z \subset \mathcal{C}^*$ be set of probabilistic constraints and $F = \{C_1, \dots, C_m\} \subseteq V \setminus \{A, B\}$ be arguments. Then $(A, B) \in V \times V$ is **semi-effective w.r.t** (Z, F) if there are*

$x, y, q_1, \dots, q_m \in [0, 1]$ such that

$$D \doteq Z \cup \{P(C_i) = q_i | 1 \leq i \leq m\} \cup \{P(A) = y\}$$

is consistent and at least one of the following conditions hold:

- $D \not\models P(B) \neq x$ and $D \cup \{P(A) = y\} \models P(B) \neq x$;
- $D \models P(B) \neq x$ and $D \cup \{P(A) = y\} \not\models P(B) \neq x$.

For instance, suppose we have a the set probabilistic constraints $\mathcal{C} = \{P(C) > 0.5 \rightarrow P(A) + P(B) = 1\}$. Fixing value $q \in [0.5, 1]$ for $P(C)$ entails $P(B) = 1 - P(A)$. Without restricting the value of $P(A)$, in principle $P(B)$ can have any value in the interval $[0, 1]$. When we further assume $P(A) = x$, $P(B)$ must be precisely $1 - x$, therefore (A, B) is semi-effective w.r.t. $(\mathcal{C}, \{C\})$. Note that (A, B) is not semi-effective w.r.t. (\mathcal{C}, \emptyset) , for there is no restriction on $P(B)$ when $P(C) \leq 0.5$.

If an arc (A, B) is semi-effective w.r.t to some pair (Z, F) , it can be labelled as according to the type of effect $P(A) > 0.5$ has on $P(B)$. Details on how to derive consistent labels to arcs from the the probabilistic constraints can be found in the work of Hunter, Polberg and Thimm [16]. It is worth noting though that a given arc can be supporting and attacking for different pairs (Z, F) , justifying the use of sets for labels and showing the flexibility of the whole framework.

The main intuition behind this interpretation is that an arc (A, B) indicates that the probability of A affects the probability B in a given context, where the context is captured by the pair (Z, F) . Conditional probabilities, though, enables the arguments in the context to be represented directly in the conditioning proposition. Intuitively, given a $Z \subset \mathcal{C}^*$ and a conjunction of arguments F , assigning a value to $P(A|F)$ could change the possible values for $P(B|F)$. Formally, using the power of \mathcal{L}_{linear} to represent conditional probabilities, we can define conditionally effective arcs:

Definition 2 *Let $(\mathcal{G}, \ell, \mathcal{C})$ be an epistemic graph, with $\mathcal{G} = (V, R)$, $Z \subset \mathcal{C}^*$ be a set of probabilistic constraints and F be conjunction of arguments $C_1, \dots, C_m \in V \setminus \{A, B\}$. Then $(A, B) \in V \times V$ is **conditionally effective w.r.t** (Z, F) , if there are $x, y \in [0, 1]$ such that*

$$D \doteq Z \cup \{P(A \wedge F) = yP(F)\}$$

is consistent and at least one of the following conditions hold:

- $Z \not\models P(B \wedge F) \neq xP(F)$ and $D \models P(B \wedge F) \neq xP(F)$;
- $Z \models P(B \wedge F) \neq xP(F)$ and $D \not\models P(B \wedge F) \neq xP(F)$.

For instance, consider a set of probabilistic constraints

$$\mathcal{C} = \{P(A \wedge C) > 0.5P(C) \rightarrow P(B \wedge C) < 0.5P(C)\}.$$

In principle, $P(B)$ can be assigned any value in $[0, 1]$. Furthermore, there is no value for $P(C)$ and $P(A)$ that together with \mathcal{C} restricts the value of $P(B)$. Formally, (A, B) is not semi-effective w.r.t. (Z, \mathcal{C}) . However, $\mathcal{C} \cup P(A \wedge C) = 0.6P(C)$ entails $P(B \wedge C) < 0.5P(C)$, and (A, B) is conditionally effective w.r.t. (Z, C) .

Once an arc is conditionally effective, we would have still to derive the set labels implied by the probabilistic restrictions. Moving from \mathcal{L}_{linear} to \mathcal{L}_{cond} , we have more expressivity, allowing the interpretation of arcs based on the Bayesian confirmation theory [25], where a set of evidences E is said to (dis)confirm a hypothesis H if $P(H|E) > P(H)$ ($P(H|E) < P(H)$). Accordingly, we can interpret an arc (A, B) as effective and attacking if $P(B|A) - P(B) < 0$ is entailed by the constraints and as effective and supporting if $P(B|A) - P(B) > 0$ is so. More interestingly, we can use a subset $Z \subseteq \mathcal{C}^*$ and a context F to parameterize a definition of effectiveness.

Definition 3 Let $(\mathcal{G}, \ell, \mathcal{C})$ be an epistemic graph, with $\mathcal{G} = (V, R)$, $Z \subseteq \mathcal{C}^*$ be set of constraints and F be conjunction of arguments $C_1, \dots, C_m \in V \setminus \{A, B\}$. Then $(A, B) \in V \times V$ is **confirmationally effective w.r.t. (Z, F)** if Z is consistent and one of the following conditions hold

- $Z \models P(B|F \wedge A) - P(B|F) > 0$, when (A, B) is **confirmationally supporting**;
- $Z \models P(B|F \wedge A) - P(B|F) < 0$, when (A, B) is **confirmationally attacking**.

The definitions of effectiveness we propose need not replace the original one, but may instead offer alternative notions of attack and support, without losing the flexibility of the context. Once a definition is chosen, each arc in the graph can be checked for effectiveness, and missing arcs may then be added.

5. From Arcs to Probabilistic Constraints

In a different view of the role of nodes and edges, one may take arcs as additional constraints on the set of belief distributions induced by the epistemic graph. These extra restrictions from the graph may have the form of epistemic formulas in the language or even express higher-order restrictions that cannot be reduced to probabilistic constraints.

Going beyond the epistemic language, one could explore restrictions over the set of belief distributions, and not over each of them. For instance, the notion of positive/negative monotonicity [16] can be used to translate a supporting/attacking arc (A, B) into the restrictions stating that $P_{\pi_1}(A) > P_{\pi_2}(A)$ should imply $P_{\pi_1}(B) > P_{\pi_2}(B)$ (or $P_{\pi_1}(B) < P_{\pi_2}(B)$), within a given context, for any belief distributions π_1 and π_2 in the semantics of the epistemic graph. In general, any constraint outside the probabilistic

language adopted could increase the representation power of the whole framework (generally followed by higher computational complexity). Here we focus on the former path, to stay within the bounds of optimization problems that can be reduced to constraint satisfaction problems.

Staying within the probabilistic language, either \mathcal{L}_{linear} or \mathcal{L}_{cond} , the most direct approach to associate the labeled graph to a set of probabilistic constraints is arc-wise, assuming arcs are either attacking or supporting. For instance, in \mathcal{L}_{linear} , an attacking (supporting) arc (A, B) can represent the probabilistic formula $P(A) > 0.5 \rightarrow P(B) < 0.5$ ($P(A) > 0.5 \rightarrow P(B) > 0.5$). In \mathcal{L}_{cond} , we can again in the spirit of Bayesian confirmation theory use (A, B) to represent $P(A|B) - P(A) > 0$ (support) or $P(A|B) - P(A) < 0$ (attack). In either case, we would have an epistemic formula directly derivable from each arc, so it is not clear what is gained by representing these restrictions using a graph instead of the probabilistic language, besides the visual analysis.

To properly take advantage of the graph representation, we can look for the information conveyed by it as a whole, not by each arc in isolation. If we permit only attacking arcs, the label-based semantics for Abstract Argumentation Frameworks (AAF) can be employed to restrict the belief distributions satisfying the epistemic graph. By construing the probability of an argument ($P(A)$) as the probability of it being accepted ($P(L(A) = \text{in})$), the semantics of the probabilistic restrictions could be given by probability distributions π over labelings. The probability $P_\pi(A)$ could then be defined as the sum of the probabilities of labelings L where A is accepted ($L(A) = \text{in}$). The graph structure could then restrict the possible labeling according to the semantics assumed for the AAF. For instance, one might require conflict-free, admissible, complete, stable or semi-stable labelings, constraining the domain of the probability distribution. This is proposed for Probabilistic Labeling Frames by Riveret et al [23], where the sample space is defined as the set of labelings with a certain property. By viewing each labeling as its set of its accepted arguments, these restrictions on labelings consequently narrow the domain of the belief distributions $\pi : 2^V \rightarrow [0, 1]$.

Here we adopt an alternative, equivalent semantics for the graph, forcing sets of arguments that do not correspond to labelings with a desired property to have null probability instead of removing them from the sample space. For example, under the stable semantics, $\pi(\Gamma)$ may be positive only if there is a stable labeling accepting exactly the arguments in Γ .

Formally, the semantics of an epistemic graph $(\mathcal{G}, \ell, \mathcal{C})$ is given through belief distributions $\pi : 2^V \rightarrow [0, 1]$ over sets of arguments. For a labeled graph (\mathcal{G}, ℓ) , with $\mathcal{G} = (\mathcal{A}, \rightarrow)$, where every arc is labeled simply as an attack, consider the AAF $\mathcal{G} = (\mathcal{A}, \rightarrow)$. Each labeling $L : \mathcal{A} \rightarrow \{\text{in}, \text{out}, \text{undecided}\}$ yields a set of

accepted arguments $\Gamma_L = \{A \in \mathcal{A} \mid L(A) = \text{in}\}$. Each semantics for AAF associates a set of labelings L to a AAF \mathcal{G} , corresponding to a set of sets Γ_L . For each $S \in \{\text{conflict-free, admissible, complete, stable, semi-stable}\}$, let $S(\mathcal{G})$ denote the set

$$\{\Gamma_L \subseteq \mathcal{A} \mid L: \mathcal{A} \rightarrow \{\text{in, out, undecided}\} \text{ is } S\}.$$

For instance, $\text{semi-stable}(\mathcal{G})$ contains all sets Γ_L of accepted arguments according to a semi-stable labeling. For a set of arguments $\Gamma \subseteq \mathcal{A}$, let ϕ_Γ denote the conjunction $\bigwedge \{A \mid A \in \Gamma\} \wedge \bigwedge \{\neg A \mid A \in \mathcal{A} \setminus \Gamma\}$. For instance, if $\mathcal{A} = \{A, B, C, D\}$ and $\Gamma = \{A, B\}$, then $\phi_\Gamma = A \wedge B \wedge \neg C \wedge \neg D$. Note that $P_\pi(\phi_\Gamma) = \pi(\Gamma)$ for any set $\Gamma \subseteq \mathcal{A}$ and any belief distribution $\pi: 2^\mathcal{A} \rightarrow [0, 1]$. Now, for each $S \in \{\text{conflict-free, admissible, complete, stable, semi-stable}\}$, the labeled graph (\mathcal{G}, ℓ) can be associated with the following set of probabilistic formulas in $\mathcal{L}_{\text{linear}}$, which we call *labeling constraints*:

$$\mathcal{C}_{\mathcal{G}} = \{P(\phi_\Gamma) = 0 \mid \Gamma \in 2^\mathcal{A} \setminus S(\mathcal{G})\}.$$

Equivalently, $\mathcal{C}_{\mathcal{G}}$ can be defined via a single restriction that forces the belief distribution of the sets in $S(\mathcal{G})$ to sum up to one:

$$\mathcal{C}_{\mathcal{G}} = \left\{ \sum_{\Gamma \in S(\mathcal{G})} P(\phi_\Gamma) = 1 \right\}.$$

Fixing a label-based semantics for the AAFs ($S \in \{\text{conflict-free, admissible, complete, stable, semi-stable}\}$), the semantics of an epistemic graph $(\mathcal{G}, \ell, \mathcal{C})$, where all arcs in \mathcal{G} are simply attacks, can be defined as the set of belief distributions satisfying $\mathcal{C} \cup \mathcal{C}_{\mathcal{G}}$.

To illustrate how the graph structure can impose restriction based on a given semantics for abstract argumentation frameworks, consider the AAF $\mathcal{G} = (\mathcal{A}, \rightarrow)$ from Figure 4, formed from the graph in Figure 1 by abstracting the arguments. For this AAF, there are exactly only two stable labelings, shown in Figure 5. Each labeling corresponds to a set of accepted arguments, say $\Psi = \{A, C, E\}$ and $\Psi' = \{A, D\}$, so that $\text{stable}(\mathcal{G}) = \{\Psi, \Psi'\}$. The derived probabilistic restrictions prevent any other set of arguments to have positive probability; namely $\mathcal{C}_{\mathcal{G}} = \{P(\phi_\Gamma) = 0 \mid \Gamma \in 2^\mathcal{A} \setminus \{\Psi, \Psi'\}\}$. Note that $\mathcal{C}_{\mathcal{G}}$ implies $P(A) = 1$, $P(C) = P(E)$ and $P(C) + P(D) = 1$; in fact $\mathcal{C}_{\mathcal{G}}$ is equivalent to $\{P(A \wedge \neg B \wedge C \wedge \neg D \wedge E) + P(A \wedge \neg B \wedge \neg C \wedge D \wedge \neg E) = 1\}$, as mentioned above. An epistemic graph $(\mathcal{G}, \ell, \mathcal{C})$, with $\mathcal{C} = \{P(C) = 0.2\}$, would be satisfied only the belief distribution $\pi: 2^\mathcal{A} \rightarrow [0, 1]$ with $\pi(\{A, C, E\}) = 0.2$ and $\pi(\{A, D\}) = 0.8$, for instance.

Epistemic probabilities were firstly proposed as a set of $\{P(A) = q \mid A \in \mathcal{A}, q \in [0, 1]\}$ of probabilistic restrictions over arguments, together with some desiderata relating them to the graph [14]. By imposing a particular label-based semantics, some of these postulates can be recovered.




Figure 4: An Abstract Argumentation Framework.




Figure 5: AAF with its two different stable labelings.

For instance, requiring conflict-free labelings, we have that an arc (A, B) in \mathcal{G} implies $P(A) + P(B) = 1$ [23], a desirable property known as coherence. Furthermore, forcing complete labelings yields $P(A) = 1$ for every argument A with no attackers [23], when the epistemic probabilities are said to be founded.

6. A Bit of Discussion

The past ten years have seen growing interest in probabilistic argumentation frameworks of various kinds; in this paper we just added a few proposals to that literature. Still, probabilistic argumentation, as explored in artificial intelligence, is filled with difficult conceptual questions.

To begin, is the theory a normative or a descriptive one? Abstract argumentation frameworks usually resort to seemingly intuitive conditions to generate labelings, and their descriptive power is often emphasized, particularly through examples that discuss daily issues such as dental care decisions. On the other hand, probabilistic modeling is a normative endeavor; people do not normally follow exactly the probability calculus [17]. Although credal sets are useful in representing imperfect agents, they are also based on a normative theory [2]. Hence probabilistic argumentation tries to glue a mix of descriptive and normative elements. Even if existing work on empirical evaluation [22] and on processing real utterances as arguments [20] may clarify the extent to which a descriptive viewpoint is tenable, more debate is needed on the exact status of the theory.

Suppose we adopt probability theory: What then is the meaning of the probability of an argument? This may be the most difficult question. As noted at the end of Section 2, some authors take $P(A)$ to capture the belief that A is justified or the belief that A holds (apparently meaning that an argument holds when it applies correct rules to correct premises). This is not the universal wording. Some-

times a probability value is said to carry the belief that an argument is acceptable [3, 15]. It does seem attractive to assign uncertainty to whether or not the agent(s) will accept an argument, rather than trying to clarify what exactly it means for an argument to be justified or to hold. However, taking $P(A)$ as the probability of accepting A is somewhat at odds with the approach adopted in AAFs, where acceptability is derived from attacks and supports rather than taken as a starting point. In fact, sometimes probabilities are themselves used to build labelings that indicate acceptability [24]. It does seem that more debate is needed to clarify the semantics of probabilities in probabilistic argumentation.

As a digression, we note that the very flexibility that is sought in probabilistic argumentation seems to point inevitably to a theory based on credal sets. This observation goes beyond the mere fact that agents are imperfect and must be modeled accordingly. Argumentation aims at a gradual combination of arguments and their associated uncertainty; thus it is natural to work with a formalism that allows for incomplete specifications. In fact it is not surprising that some of the initial work on argumentation within artificial intelligence [12] were closely related to propositional probabilistic logic, typically dealing with probability bounds (in some cases resorting to semantics based on Dempster-Shafer theory).

Another nontrivial issue is the relationship between the graph itself (nodes and edges) and the probability assessments that one may have. Is this graph just a visual tool that can be replaced by suitable constraints? Is the graph built first and then constraints extracted from it, or the other way around? These pressing questions are particularly important in the context of epistemic graphs; we have touched on some of them in this paper. In many cases, one can start either with constraints or with nodes and edges, and move freely between them. However, there are some constraints that are intrinsically tied to the structure of the graph, such as labeling constraints. More work seems to be due to understand the status of all these concepts and their connections.

7. Conclusion

In this paper we have extended epistemic graphs in two directions: first, by allowing for conditional probabilities; second, by extracting novel constraints out of nodes and edges (in particular, labeling constraints). Depending on how conditional probabilities are used, the overall complexity of inference may remain essentially the same, or it may change substantially; there is a trade-off here between expressivity and complexity.

It should be noted, in connection with our proposed labeling constraints, that they offer a novel semantics for epistemic graphs. There the graph imposes strong con-

straints over the probabilities, on top of whatever other constraints may be explicitly specified. Namely, the graph indicates which label-based semantics must be respected by probabilities.

There are several possible avenues for future work. One is to develop algorithms for inference and decision making for the languages we have proposed, in particular for \mathcal{L}_{cond} . Another is to investigate how these languages fare in practical scenarios. A third possibility, perhaps the most difficult one, is to look carefully at the questions discussed in Section 6 so as to build a solid and comprehensive theory of argumentation under uncertainty based on (suitably extended and interpreted) epistemic graphs.

Acknowledgments

This work was carried out at the Center for Artificial Intelligence (C4AI-USP), with support by the São Paulo Research Foundation (FAPESP grant #2019/07665-4) and by the IBM Corporation. The third author is partially supported by Conselho Nacional de Desenvolvimento Científico e Tecnológico (CNPq), grant 312180/2018-7.

References

- [1] L. Amgoud and H. Prade. Reaching agreement through argumentation: A possibilistic approach. In *International Conference on Principles of Knowledge Representation and Reasoning (KR'04)*, page 175–182, 2004.
- [2] Thomas Augustin, Frank P. A. Coolen, Gert De Cooman, and Matthias Troffaes. *Introduction to Imprecise Probabilities*. John Wiley & Sons, 2014.
- [3] P. Baroni, M. Giacomin, and P. Vicig. On rationality conditions for epistemic probabilities in abstract argumentation. In *Computational Models of Argument*, pages 121–132, 2014.
- [4] Pietro Baroni, Martin Caminada, and Massimiliano Giacomin. An introduction to argumentation semantics. *The Knowledge Engineering Review*, page 1–24, 2004.
- [5] Phillippe Besnard and Anthony Hunter. *Elements of Argumentation*. MIT Press, 2008.
- [6] Carlos Caleiro, Filipe Casal, and Andreia Mordido. Classical generalized probabilistic satisfiability. In *International Joint Conference on Artificial Intelligence*, pages 908–914, 2017.
- [7] P. M. Dung. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. *Artificial Intelligence*, 77(2):321–357, 1995.

- [8] Phan Minh Dung and Phan Minh Thang. Analysing product reviews using probabilistic argumentation. In *Computational Models of Argument*, pages 171–182, 2010.
- [9] Ronald Fagin, Joseph Y Halpern, and Nimrod Megiddo. A logic for reasoning about probabilities. *Information and Computation*, 87(1-2):78–128, 1990.
- [10] D. M. Gabbay and O. Rodrigues. Probabilistic argumentation: An equational approach. *Log. Univers.*, 9: 345–382, 2015.
- [11] Alejandro J. Garcia and Guillermo R. Simari. Defeasible logic programming: Delp-servers, contextual queries, and explanations for answers. *Argument and Computation*, 5(1):63–88, 2014.
- [12] R. Haenni, B. Anrig, J. Kohlas, and N. Lehmann. A survey on probabilistic argumentation. In *European Conference on Symbolic and Quantitative Approaches to Reasoning with Uncertainty (ECSQARU)*, pages 19–25, 2001.
- [13] Anthony Hunter. Some foundations for probabilistic abstract argumentation. In *Computational Models of Argument*, pages 117–128, 2012.
- [14] Anthony Hunter and Matthias Thimm. Probabilistic argumentation with epistemic extensions. In *International Workshop on Defeasible and Ampliative Reasoning (DARe’14)*, 2014.
- [15] Anthony Hunter and Matthias Thimm. Probabilistic reasoning with abstract argumentation frameworks. *Journal of Artificial Intelligence Research*, 59:565–611, 2017.
- [16] Anthony Hunter, Sylwia Polberg, and Matthias Thimm. Epistemic graphs for representing and reasoning with positive and negative influences of arguments. *Artificial Intelligence*, 281:103236, 2020.
- [17] Daniel Kahneman, Paul Slovic, and Amos Tversky. *Judgment Under Uncertainty: Heuristics and Biases*. Cambridge University Press, 1982.
- [18] Theofrastos Mantadelis and Stefano Bistarelli. Probabilistic abstract argumentation frameworks, a possible world view. *International Journal of Approximate Reasoning*, 119:204–219, 2020.
- [19] Denis Deratani Mauá and Fabio G. Cozman. Thirty years of credal networks: Specification, algorithms and complexity. *International Journal of Approximate Reasoning*, 126:133–157, 2020.
- [20] Kawsar Noor and Anthony Hunter. Analysing product reviews using probabilistic argumentation. In *Computational Models of Argument*, pages 295–306, 2020.
- [21] Judea Pearl. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufmann, 1988.
- [22] Sylwia Polberg and Anthony Hunter. Empirical evaluation of abstract argumentation: Supporting the need for bipolar and probabilistic approaches. *International Journal of Approximate Reasoning*, 93:487–543, 2018.
- [23] Régis Riveret, Pietro Baroni, Yang Gao, Guido Governatori, Antonino Rotolo, and Giovanni Sartor. A labelling framework for probabilistic argumentation. *Annals of Mathematics and Artificial Intelligence*, 83 (1):21–71, 2018.
- [24] Matthias Thimm. A probabilistic semantics for abstract argumentation. In *European Conference on Artificial Intelligence (ECAI)*, pages 750–755, 2012.
- [25] Frans H. van Eemeren, Bart Garssen, Erik C. W. Krabbe, A. Francisca Snoeck Henkemans, Bart Verheij, and Jean H. M. Wagemans. Argumentation and artificial intelligence. In *Handbook of Argumentation Theory*. Springer, 2014.
- [26] Michael P. Wellman. Qualitative probabilistic networks for planning under uncertainty. In *Uncertainty in Artificial Intelligence*, volume 5, pages 197–208, 1988.