

Probabilistic Abstract Argumentation based on SCC Decomposability

Tjitze Rienstra

University of Koblenz-Landau
Germany

Matthias Thimm

University of Koblenz-Landau
Germany

Beishui Liao

Zhejiang University
China

Leendert van der Torre

University of Luxembourg
Luxembourg

Abstract

In this paper we introduce a new set of general principles for probabilistic abstract argumentation. The main principle is a probabilistic analogue of SCC decomposability, which ensures that the probabilistic evaluation of an argumentation framework complies with the probabilistic (in)dependencies implied by the graph topology. We introduce various examples of probabilistic semantics and determine which principles they satisfy. Our work also provides new insights into the relationship between abstract argumentation and the theory of Bayesian networks.

1 Introduction

Abstract argumentation deals with *argumentation frameworks* (AFs) and their *semantics* (Dung 1995). The former are sets of arguments together with an attack relation, and the latter are methods to determine their extensions (i. e., sets of acceptable arguments) or labelings (i.e., functions mapping arguments to an acceptance status) (Caminada 2006). Abstract argumentation has found many applications in the field of AI. Beyond its original role as an abstraction of various defeasible reasoning formalisms (Dung 1995), abstract argumentation now forms a cornerstone of the interdisciplinary field of *formal argumentation* (Baroni et al. 2018).

An essential aspect of argumentation is uncertainty, which appears, for example, in opponent or audience models in strategic argumentation or persuasion (Rienstra, Thimm, and Oren 2013; Grossi and van der Hoek 2013; Hunter 2015). It therefore makes sense to consider probabilistic extensions. One approach is to assume that arguments and attacks of an AF are associated with probabilities (Li, Oren, and Norman 2011; Rienstra 2012; Hunter 2013; Hunter and Thimm 2016b). Another approach is to use probabilities to represent degrees of belief in whether arguments are accepted (Thimm 2012; Hunter 2013; Baroni, Giacomin, and Vicig 2014; Hunter and Thimm 2016a; 2017). These two approaches are often referred to, respectively, as the *constellations* and *epistemic* approach (a distinction introduced in (Hunter 2013)).

In this paper we present a new perspective on the epistemic approach. Our main focus is the question of how the topology of an AF determines probabilistic

(in)dependencies that the probabilistic evaluation of an AF should obey. The following example demonstrates this idea.




Figure 1: The eating out example

Example 1 (Eating out). The AF shown in Figure 1 represents the decision making of an agent planning to eat out. He will eat meat or fish (m or f) and take a taxi or drive himself (t or d). He drinks red wine (r) but not with fish or when driving (f and d attack r). Finally, he drinks either cola or water (c or w), but no cola if he drinks red wine (r attacks c).

The direction of the attacks imply that the agent first chooses independently between m and f and between t and d . Then he determines the status of r , which depends on f and d . Finally he chooses between c and w , which depends on r . Note that we can, of course, imagine different scenarios, but this would involve different directions of attack. E.g., if the decision about r came *before* the decision between t and d , then the attack of d on r would be reversed.

Suppose we have probabilistic beliefs about what the agent will do. That is, each argument is associated with a probability that it is accepted. What happens if we learn that the agent accepts m or f ? Clearly, this may affect the probability of r and hence c and w . The preceding discussion suggests, however, that t and d are unaffected. That is, m and f are *probabilistically independent* of t and d .

Now suppose we learn that the agent accepts m or f , but *after* having learnt that he rejects r . Since the effect of m and f on c and w is mediated by r (whose status is now fixed) the arguments c and w are not affected this time. In other words, m and f are *probabilistically independent* of c and w given r . On the other hand, if the agent accepts m , we can infer that he accepts d , since this is now the only way to account for rejection of r . Thus, after learning r is rejected, it no longer holds that m and f are independent of t and d .

This example shows that the topology of an AF quite naturally translates into probabilistic (in)dependence rela-

tionships if the AF is evaluated probabilistically. So far, this aspect has mostly been neglected. One reason is that the epistemic approach as described in e.g. (Hunter 2013) deals mostly with marginal probabilities without addressing the question of what the underlying probability distribution looks like. Therefore conditional probabilities are undefined and questions of probabilistic independence do not arise.

Our approach consists of three steps. We first define the general notion of a *probabilistic semantics*, which is a function mapping each AF to a probability distribution over its labelings. This will be the object of our study. We can think of this as the probabilistic counterpart of a regular (labeling-based) semantics, which maps each AF to a set of labelings.

Secondly, we propose some general principles that a probabilistic semantics should satisfy. Here, a central principle is a probabilistic analogue of *SCC decomposability*, which states that each SCC (strongly connected component) is evaluated independently given the status of its attackers (Baroni et al. 2014). The (in)dependencies in Example 1 (where SCCs are enclosed in rectangles) can be understood as arising from this principle. Another principle is *SCC factorability*, which connects our approach to the theory of Bayesian networks. This connection is already apparent in Example 1, where attacks represent—like edges in a Bayesian network—a specific kind of influence between variables.

Third, we discuss a number of concrete schemes to define a probabilistic semantics, which we then evaluate using the principles we proposed. Two schemes are based on uniform distributions over the choices that can be made in selecting labelings. Two more schemes are generalisations capable of taking into account probabilistic strength of arguments.

The overview of this paper: Section 2 deals with basics of argumentation and probability theory. In Section 3 we define the general notion of a probabilistic semantics and introduce general principles that it should satisfy. Sections 4 and 5 deal with concrete schemes to define a probabilistic semantics, which we evaluate using the principles from Section 3. In Section 6 we discuss open issues and related work.

2 Preliminaries

2.1 Abstract Argumentation

The basic notion in abstract argumentation is an *argumentation framework* (AF for short) (Dung 1995). An AF consists of a set of arguments and a binary attack relation between arguments. We restrict our attention to finite AFs.

Definition 1. Let \mathcal{A} be a set called the *universe of arguments*. An argumentation framework (AF) is a pair $F = (A_F, \rightsquigarrow_F)$ where $A_F \subseteq \mathcal{A}$ is a finite set of arguments, and $\rightsquigarrow_F \subseteq A_F \times A_F$ the *attack relation*.

A *semantics* determines, given an AF, rational points of view on argument acceptability. We use the three-valued labeling-based semantics introduced in (Caminada 2006), where a *labeling* maps each argument to a label **I** (*in* or *accepted*), **O** (*out* or *rejected*), or **U** (*undecided*).

Definition 2. A *labeling* of a set A is a function $\mathbf{L} : A \rightarrow \{\mathbf{I}, \mathbf{O}, \mathbf{U}\}$. We denote by $\mathcal{L}(A)$ the set of labelings of A and, given an AF F , by $\mathcal{L}(F)$ the set of labelings of A_F . We

denote by $\mathbf{I}(\mathbf{L})$, $\mathbf{U}(\mathbf{L})$ and $\mathbf{O}(\mathbf{L})$ the set of arguments labeled **I**, **U** and **O** by \mathbf{L} , respectively.

A basic condition for a labeling to be regarded as a rational point of view is *completeness* (Caminada 2006).

Definition 3. Let F be an AF. A labeling $\mathbf{L} \in \mathcal{L}(F)$ is a *complete* labeling of F iff, for all $x \in A_F$:

1. $\mathbf{L}(x) = \mathbf{O}$ iff $\exists y \in A_F$ s.t. $y \rightsquigarrow_F x$ and $\mathbf{L}(y) = \mathbf{I}$.
2. $\mathbf{L}(x) = \mathbf{I}$ iff $\forall y \in A_F$ s.t. $y \rightsquigarrow_F x$, $\mathbf{L}(y) = \mathbf{O}$.

A semantics σ maps each AF F to a set of labelings of F and is called *universally defined* if it maps each AF to a *nonempty* set of labelings.

Definition 4. A *semantics* σ associates each AF F with a set $\mathcal{L}_\sigma(F) \subseteq \mathcal{L}(F)$ of labelings. A semantics σ is *universally defined* iff for each AF F , $\mathcal{L}_\sigma(F) \neq \emptyset$.

We limit our attention in this paper to universally defined semantics and will not make this explicit every time. Thus, if we speak of “a semantics σ ” we mean “a universally defined semantics σ ”. This implies that our results do not apply to the *stable* semantics, which is not universally defined (Caminada 2006). We return to this issue in Section 6.

In this paper we focus on the *complete semantics*, which yields complete labelings, and the *preferred*, *grounded* and *semi-stable* semantics, which yield the **I**-maximal, **I**-minimal and **U**-minimal complete labelings (Caminada 2006). These semantics are all universally defined. Note that an AF always has one grounded labeling but may have multiple complete, preferred or semi-stable labelings.

Definition 5. The **co** (*complete*), **pr** (*preferred*), **gr** (*grounded*) and **ss** (*semi-stable*) semantics are defined by

$$\begin{aligned} \mathcal{L}_{\text{co}}(F) &= \{\mathbf{L} \in \mathcal{L}(F) \mid \mathbf{L} \text{ is a complete labeling of } F\}, \\ \mathcal{L}_{\text{pr}}(F) &= \{\mathbf{L} \in \mathcal{L}_{\text{co}}(F) \mid \nexists \mathbf{L}' \in \mathcal{L}_{\text{co}}(F), \mathbf{I}(\mathbf{L}) \subset \mathbf{I}(\mathbf{L}')\}, \\ \mathcal{L}_{\text{gr}}(F) &= \{\mathbf{L} \in \mathcal{L}_{\text{co}}(F) \mid \nexists \mathbf{L}' \in \mathcal{L}_{\text{co}}(F), \mathbf{I}(\mathbf{L}) \supset \mathbf{I}(\mathbf{L}')\}, \\ \mathcal{L}_{\text{ss}}(F) &= \{\mathbf{L} \in \mathcal{L}_{\text{co}}(F) \mid \nexists \mathbf{L}' \in \mathcal{L}_{\text{co}}(F), \mathbf{U}(\mathbf{L}) \supset \mathbf{U}(\mathbf{L}')\}. \end{aligned}$$

2.2 Probability Theory

In this section we present the necessary basics concerning probability theory. We start with the general notion of a probability distribution (Koller and Friedman 2009).

Definition 6. Let Ω be a finite set of elements called *possible worlds*. A *probability distribution* over Ω is a function $P : \Omega \rightarrow \mathbb{R}$ such that for all $w \in \Omega$, $P(w) \geq 0$ and $\sum_{w \in \Omega} P(w) = 1$. P is extended to a function over *events* (i. e., subsets of Ω) by defining $P(X) = \sum_{w \in X} P(w)$, for all $X \subseteq \Omega$. If $X, Y \subseteq \Omega$ and if $P(Y) > 0$, then the probability of X *conditional on* Y , denoted $P(X \mid Y)$, is defined by $P(X \mid Y) = P(X \cap Y) / P(Y)$.

A set of possible worlds may be associated with a set \mathcal{X} of variables. For simplicity we assume that every variable has the same (finite) domain denoted by *Dom*.

Definition 7. A *valuation* of a set \mathcal{X} of variables is a function \mathbf{V} from \mathcal{X} to *Dom*. If \mathbf{V} is a valuation of \mathcal{X} and $B \subseteq \mathcal{X}$ then $\mathbf{V} \downarrow B$ denotes the restriction of \mathbf{V} to B . We say that Ω is *determined* by \mathcal{X} iff Ω consists of all valuations of \mathcal{X} .

If Ω is determined by \mathcal{X} and if $x \in \mathcal{X}$ and $v \in \text{Dom}$, we denote by $x =_{\Omega} v$ the set $\{\mathbf{V} \in \Omega \mid \mathbf{V}(x) = v\}$, omitting the subscript Ω if clear from context. Furthermore, a valuation \mathbf{V} of a set $B \subseteq \mathcal{X}$ will also be used (abusing notation) to denote the event $\{\mathbf{V}' \in \Omega \mid \mathbf{V}' \downarrow B = \mathbf{V}\}$.

Given three sets U, V, W of variables, U is said to be *independent of V given W* if, once W is fixed, learning U does not affect the probability of V and vice versa (Koller and Friedman 2009). By setting $W = \emptyset$, the following also defines unconditional independence between U and V .

Definition 8. Let P be a probability distribution over a set Ω determined by \mathcal{X} and let $U, V, W \subseteq \mathcal{X}$. We say that U is *independent of V given W (with respect to P)* iff, for all valuations \mathbf{U}, \mathbf{V} and \mathbf{W} of U, V and W , respectively, we have $P(\mathbf{U} \mid \mathbf{V} \cap \mathbf{W}) = P(\mathbf{U} \mid \mathbf{W})$ if $P(\mathbf{V} \cap \mathbf{W}) > 0$.

Remark 1. The following fact will be useful in what follows: If U is independent of V given W then, for any $U' \subseteq U$ and $V' \subseteq V$, we have that U' is independent of V' given W . This is called the *decomposition* property of conditional independence (Koller and Friedman 2009).

We apply the notions defined here to abstract argumentation by interpreting arguments as variables with domain $\{\mathbf{I}, \mathbf{U}, \mathbf{O}\}$. Given an AF F , a probability distribution over $\mathcal{L}(F)$ represents a probabilistic evaluation of F , and an event like $a = \mathbf{I}$ represents the event that a is accepted.

3 Probabilistic Semantics

We now define a probabilistic counterpart of a regular semantics. A *probabilistic semantics* associates each AF with a probability distribution over its labelings and represents a particular type of probabilistic evaluation for any AF.

Definition 9. A *probabilistic semantics* π maps each AF F to a probability distribution over $\mathcal{L}(F)$ denoted by $P_{\pi}(F)$.

A probabilistic semantics yields absolute probabilities (e.g. $P_{\pi}(F)(r = \mathbf{I})$ for the probability that r is accepted) as well as joint and conditional probabilities (e.g. $P_{\pi}(F)(r = \mathbf{I} \mid f = \mathbf{O} \cap d = \mathbf{O})$ for the probability that r is accepted if we learn that f and d are rejected).

We now discuss some general principles for a probabilistic semantics. Like the principles studied in standard abstract argumentation (van der Torre and Vesic 2017) these principles can be used to check whether a semantics behaves as desired or act as guiding principles in defining one.

3.1 σ -compatibility

The first principle states that a probabilistic semantics π is compatible with a semantics σ in the sense that a labeling receives nonzero probability if and only if it is a σ labeling.

Definition 10. A probabilistic semantics π is *σ -compatible* iff for all F , $P_{\pi}(F)(\mathbf{L}) > 0$ iff $\mathbf{L} \in \mathcal{L}_{\sigma}(F)$.

Note that a probabilistic semantics π cannot be σ -compatible if σ is not universally defined, as this implies that $P_{\pi}(F)$ is not a probability distribution if F has no labelings.

3.2 Probabilistic SCC Decomposability

We now describe a probabilistic analogue of a principle called *decomposability w.r.t. SCC partitioning* (Baroni et al. 2014). Here we call it *SCC decomposability* and we simplify its definition somewhat. We first define the notion of SCC formally, as well as the notion of an *outparent* of a set of arguments. An outparent of a set S of arguments is any argument attacking S but not a member of S .

Definition 11. Let F be an AF. The set of *SCCs* (strongly connected components) of F , denoted $\text{SCC}(F)$, contains all equivalence classes induced by the *path equivalence relation* \sim_F over A_F defined by $x \sim_F y$ iff $x = y$ or there is a directed path from x to y and y to x . An argument $x \in A_F$ is an *outparent* of a set $S \subseteq A_F$ iff there is a $y \in S$ such that $x \rightsquigarrow_F y$, and $x \notin S$. The set of outparents of S is denoted by $OP_F(S)$.

We also use the following auxiliary notions.

Definition 12. Let F be an AF and let $S \subseteq A_F$. The *context* of S is the set $S \cup OP_F(S)$ and is denoted by $C_F(S)$. S is an *unattacked set* iff it has no outparents. The *restriction of F to S* is the AF $(S, \rightsquigarrow_F \cap S \times S)$ and is denoted by $F \downarrow S$.

SCC decomposability states that the labelings of an AF are determined independently for *each SCC S given the outparents of S* by means of a *local function*. For the following definition, note that the set of outparents of an SCC S of an AF F is always an unattacked set of the AF $F \downarrow C_F(S)$.

Definition 13. (Baroni et al. 2014) A *local function* \mathbb{L} is a function that assigns to each AF F , unattacked set I of F , and labeling $\mathbf{L}_I \in \mathcal{L}(I)$, a set $\mathbb{L}(F, \mathbf{L}_I) \subseteq \mathcal{L}(A_F \setminus I)$. We say that a local function \mathbb{L} *represents* a semantics σ iff, for each AF F , we have that $\mathbf{L} \in \mathcal{L}_{\sigma}(F)$ iff

$$\forall S \in \text{SCC}(F), \mathbf{L} \downarrow S \in \mathbb{L}(F \downarrow C_F(S), \mathbf{L} \downarrow OP_F(S)).$$

We say that a semantics σ is *SCC decomposable* iff there exists a local function \mathbb{L} that represents σ .

To prove that a semantics σ is SCC decomposable, a local function must be defined and it must be shown that it represents σ . An important example of a local function is the one that mimics the effect of a labeling \mathbf{L}_I of the outparents I of an SCC S by (1) attacking \mathbf{O} -labeled outparents; (2) making \mathbf{U} -labeled outparents self-attacking; and (3) ignoring attacks between outparents (Baroni et al. 2014). We denote this local function by \mathbb{L}_{σ} .

Definition 14. Given a semantics σ we define the local function \mathbb{L}_{σ} by $\mathbb{L}_{\sigma}(F, \mathbf{L}_I) = \mathcal{L}_{\sigma}(F \downarrow_{(I, \mathbf{L}_I)} \downarrow (A_F \setminus I))$, where $F \downarrow_{(I, \mathbf{L}_I)} = (A_F \cup \{x' \mid x \in I, L(x) = \mathbf{O}\}, (\rightsquigarrow_F \setminus (I \times I)) \cup \{(x', x) \mid x \in I, L(x) = \mathbf{O}\} \cup \{(x, x) \mid x \in I, L(x) = \mathbf{U}\})$.

It was shown in (Baroni et al. 2014) that, for all $\sigma \in \{\mathbf{co}, \mathbf{pr}, \mathbf{gr}\}$, σ is SCC decomposable because \mathbb{L}_{σ} represents σ . However, there is no local function that represents \mathbf{ss} and hence \mathbf{ss} is not SCC decomposable.

We now define a probabilistic analogue of SCC decomposability. We say that π satisfies *probabilistic SCC decomposability* if for each AF F , $P_{\pi}(F)$ is decomposable into the product of local probabilities, determined by a *local probability function*, for each SCC given its outparents.

Definition 15. A *local probability function* \mathbb{P} assigns to each AF F , unattacked set I of F , and labeling \mathbf{L}_I of I , a probability distribution over $\mathcal{L}(A_F \setminus I)$ denoted by $\mathbb{P}(F, \mathbf{L}_I)$. We say that \mathbb{P} represents π iff, for each AF F ,

$$P_\pi(F)(\mathbf{L}) = \prod_{S \in \text{SCC}(F)} \mathbb{P}(F \downarrow C_F(S), \mathbf{L} \downarrow OP_F(S))(\mathbf{L} \downarrow S).$$

We say that π satisfies *probabilistic SCC decomposability* iff there exists a local probability function \mathbb{P} that represents π .

The following proposition establishes a link between the two notions of SCC decomposability. Note that the only-if direction does not hold (this is demonstrated in Section 4.1).

Proposition 1. *Let π be σ -compatible. If π is probabilistically SCC decomposable then σ is SCC decomposable.*

Proof. Suppose π is σ -compatible and π satisfies probabilistic SCC decomposability. Let \mathbb{P} represent π . Define \mathbb{L} by $\mathbf{L} \in \mathbb{L}(F, \mathbf{L}_I)$ iff $\mathbb{P}(F, \mathbf{L}_I)(\mathbf{L}) > 0$. It can be checked that \mathbb{L} represents σ . Hence σ is SCC decomposable. \square

Note that probabilistic SCC decomposability is not just a principle that a probabilistic semantics may or may not satisfy, but also a way to *define* a probabilistically SCC decomposable semantics, i.e., via a local probability function. We use this strategy in Section 4 and 5. Proposition 1 implies, however, that such a semantics cannot be compatible with a non-SCC decomposable semantics like the *ss* semantics.

The main benefit of probabilistic SCC decomposability is that it ensures that the evaluation of an AF complies with the probabilistic independencies implied by the topology of the AF. To see why, we need to consider two additional principles, both of which follow from probabilistic SCC decomposability. These are the topic of the next two sections.

3.3 SCC Factorability

The next principle, called *SCC factorability*, precisely states how the topology of an AF translates into probabilistic independencies when the AF is evaluated probabilistically. It connects the well-known factorisation condition for a Bayesian network with the probabilistic SCC decomposability principle discussed above. We discuss the relationship with Bayesian networks in more detail in Section 6.

We say that a probability distribution is *SCC factorable* if it is decomposable into the product of conditional probabilities corresponding to *each SCC given its outparents*.

Definition 16. Let F be an AF. A probability distribution P over $\mathcal{L}(F)$ is *SCC-factorable* w.r.t. F iff, for all $\mathbf{L} \in \mathcal{L}(F)$,

$$P(\mathbf{L}) = \prod_{S \in \text{SCC}(F)} P(\mathbf{L} \downarrow S \mid \mathbf{L} \downarrow OP_F(S)). \quad (1)$$

SCC factorability can be expressed equivalently by a set of probabilistic independence statements. For this we need to define the notion of (non)descendant. This is defined like a (non)descendant of a variable in a Bayesian network, except that we also consider nondescendants of *sets* of variables.

Definition 17. Let F be an AF, $x, y \in A_F$ and $S \subseteq A_F$. x is a *descendant* of y iff $x = y$ or there is a directed path from y to x . x is a *descendant* of S iff x is a descendant of some $y \in S$. x is *nondescendant* of S iff x is not a descendant of S . The set of nondescendants of S is denoted by $ND_F(S)$.

SCC factorability is equivalent to stating that *every SCC S is independent of its nondescendants given its outparents*. To see why this makes sense, note that the influence—if any—of a nondescendant of an SCC S on S is always mediated by the outparents of S . That is, any directed path of attacks from a nondescendant of S to S must go through an outparent of S . Intuitively, then, the nondescendants of S should not influence S if we fix the status of the outparents of S .

Proposition 2. *Let F be an AF. A probability distribution P over $\mathcal{L}(F)$ is SCC-factorable w.r.t. F iff for all $S \in \text{SCC}(F)$, S is independent of $ND_F(S)$ given $OP_F(S)$.*

Proof. (Only if) Suppose P is SCC-factorable w.r.t. F and let $S \in \text{SCC}(F)$. Let \mathbf{S} , \mathbf{OP}_S and \mathbf{ND}_S be valuations of S , $OP_F(S)$ and $ND_F(S)$, respectively. Assume $P(\mathbf{OP}_S \cap \mathbf{ND}_S) > 0$. Then $P(\mathbf{S} \mid \mathbf{ND}_S) = P(\mathbf{S} \cap \mathbf{ND}_S) / P(\mathbf{ND}_S)$. (1) implies that $P(\mathbf{S} \cap \mathbf{ND}_S) = P(\mathbf{S} \mid \mathbf{OP}_S)P(\mathbf{ND}_S)$. Hence $P(\mathbf{S} \mid \mathbf{ND}_S) = P(\mathbf{S} \mid \mathbf{OP}_S)$. Since $OP_F(S) \subseteq ND_F(S)$ and $P(\mathbf{OP}_S \cap \mathbf{ND}_S) > 0$ we have $\mathbf{ND}_S \subseteq \mathbf{OP}_S$ and thus $P(\mathbf{S} \mid \mathbf{OP}_S \cap \mathbf{ND}_S) = P(\mathbf{S} \mid \mathbf{OP}_S)$. Hence S is independent of $ND_F(S)$ given $OP_F(S)$. (If) Suppose each $S \in \text{SCC}(F)$ is independent of $ND_F(S)$ given $OP_F(S)$. Let S_1, \dots, S_n be an ordering of $\text{SCC}(F)$ s.t., if S_j is a descendant of S_i , then $i < j$. The chain rule yields, for all $\mathbf{L} \in \mathcal{L}(F)$, $P(\mathbf{L}) = \prod_{S_i} P(\mathbf{L} \downarrow S_i \mid \mathbf{L} \downarrow S_1 \cup \dots \cup S_{i-1})$ where, for all S_i , $OP_F(S_i) \subseteq S_1 \cup \dots \cup S_{i-1} \subseteq ND_F(S_i)$. Because each S_i is independent of $ND_F(S_i)$ given $OP_F(S_i)$ it follows that, for all $\mathbf{L} \in \mathcal{L}(F)$, $P(\mathbf{L}) = \prod_{S_i} P(\mathbf{L} \downarrow S_i \mid \mathbf{L} \downarrow OP_F(S_i))$. Thus P is SCC factorable w.r.t. F . \square

Example 2. Let F be the AF from Figure 1. Suppose P is SCC factorable w.r.t. F . Consider the SCC $\{m, f\}$, which has nondescendants t and d but no outparents. Hence $\{m, f\}$ is independent of $\{t, d\}$. It also follows that the individual members of these sets are independent (cf. Remark 1). For example, m is independent of t . It similarly follows that c and w are independent of m, f, t and d conditional on r . On the other hand it is not necessarily true that m and f are independent of t and d conditional on r . All this is in line with what we argued for in example 1.

Any probabilistically SCC decomposable semantics produces SCC factorable probability distributions.

Proposition 3. *If π satisfies probabilistic SCC decomposability then for each AF F , $P_\pi(F)$ is SCC factorable w.r.t. F .*

Proof (Sketch). Suppose π satisfies probabilistic SCC decomposability. Let \mathbb{P} represent π , F be an AF, $S \in \text{SCC}(F)$ and $\mathbf{L} \in \mathcal{L}(F)$. Rewriting $P_\pi(F)(\mathbf{L} \downarrow S \mid \mathbf{L} \downarrow OP_F(S))$ using definitions 6 and 15 and taking out common factors yields $\mathbb{P}(F \downarrow C_F(S), \mathbf{L} \downarrow OP_F(S))(\mathbf{L} \downarrow S)$. Via definitions 15 and 16 it follows that $P_\pi(F)$ is SCC factorable w.r.t. F . \square

3.4 Probabilistic Directionality

We now introduce a probabilistic version of *directionality* (Baroni and Giacomini 2007). This principle states that, since the members of an unattacked set B of an AF F do not depend on arguments outside B , evaluating B in isolation (i.e., evaluating F restricted to B) is equivalent to evaluating F , as far as the members of B is concerned. While the **co**, **gr** and **pr** semantics satisfy directionality, the **ss** semantics does not (Baroni and Giacomini 2007). Below we denote by $\mathcal{L}_\sigma(F)\downarrow B$ the set $\{\mathbf{L}\downarrow B \mid \mathbf{L} \in \mathcal{L}_\sigma(F)\}$.

Definition 18. σ satisfies *directionality* iff, for each AF F and unattacked set B of F we have $\mathcal{L}_\sigma(F)\downarrow B = \mathcal{L}_\sigma(F\downarrow B)$.

We define the probabilistic version as follows.

Definition 19. π satisfies *probabilistic directionality* iff, for each AF F , unattacked set B of F , and valuation \mathbf{B} of B , we have $P_\pi(F)(\mathbf{B}) = P_\pi(F\downarrow B)(\mathbf{B})$.

Example 3. (Continued from Example 1) Suppose π satisfies probabilistic directionality. Because $\{m, f\}$ is an unattacked set, m and f do not depend on the remaining arguments. Thus, evaluating $F\downarrow\{m, f\}$ is equivalent to evaluating F , as far as m and f are concerned. More precisely, let $F' = F\downarrow\{m, f\}$. We then have, for all $l_m, l_f \in \{\mathbf{I}, \mathbf{U}, \mathbf{O}\}$,

$$P_\pi(F)(m = l_m \cap f = l_f) = P_\pi(F')(m = l_m \cap f = l_f).$$

Any probabilistic semantics satisfying probabilistic SCC decomposability also satisfies probabilistic directionality.

Proposition 4. *If π satisfies probabilistic SCC decomposability then π satisfies probabilistic directionality.*

Proof (Sketch). Suppose π satisfies probabilistic SCC decomposability. Let \mathbb{P} represent π . Let F be a AF, B an unattacked set of F , and $\mathbf{B} \in \mathcal{L}(B)$. Rewriting $P_\pi(F\downarrow B)(\mathbf{B})$ using definition 15 yields $\prod_{S \in \text{SCC}(F), S \subseteq B} \mathbb{P}(F\downarrow C_F(S), \mathbf{B}\downarrow OP_F(S)) (\mathbf{B}\downarrow S)$. This can be rewritten equivalently in terms of a product of probabilities corresponding to *all* SCCs of F , which yields $\sum_{\mathbf{L} \in \mathcal{L}(F), \mathbf{L}\downarrow B = \mathbf{B}} \prod_{S \in \text{SCC}(F)} \mathbb{P}(F\downarrow C_F(S), \mathbf{L}\downarrow OP_F(S)) (\mathbf{L}\downarrow S)$. This implies, via definition 6 and 15, that $P_\pi(F\downarrow B)(\mathbf{B}) = P_\pi(F)(\mathbf{B})$. Hence π satisfies directionality. \square

4 Uniform Probabilistic Semantics

In this section we discuss two schemes that define a probabilistic semantics. They are both based on the *principle of indifference*. According to this principle, if there are n possibilities and no way to distinguish them, then each receives probability $1/n$. In other words, we consider *uniform* distributions over the different choices we can make in selecting labelings. We show, however, that this idea can be applied in two ways. We evaluate the resulting semantics on the basis of the principles discussed in the previous section.

4.1 The σ -uniform semantics

The first scheme simply produces uniform distributions over the labelings of an AF under a given semantics σ . This has been done earlier in (Thimm et al. 2017) and was justified empirically in (Toniolo, Norman, and Oren 2017). We call it the *σ -uniform semantics* and denote it by σ^u .

Definition 20. Let σ be a semantics. The σ^u probabilistic semantics of an AF F is defined by

$$P_{\sigma^u}(F)(\mathbf{L}) = \begin{cases} \frac{1}{|\mathcal{L}_\sigma(F)|}, & \text{if } \mathbf{L} \in \mathcal{L}_\sigma(F) \\ 0, & \text{otherwise.} \end{cases}$$

We first note the following.

Proposition 5. *For any semantics σ , σ^u is σ -compatible.*

Does a semantics defined by the σ -uniform scheme satisfy probabilistic SCC decomposability? The **gr**^u semantics does. However, this is a degenerate case: **gr**^u always assigns probability 1 to the grounded labeling. In general, the answer is no. For the semantics we consider, we have:

Proposition 6. *The **co**^u, **pr**^u and **ss**^u semantics do not satisfy probabilistic SCC decomposability.*

This is demonstrated by the following example. It is based on the **pr**^u semantics but also applies to **co**^u and **ss**^u.

Example 4. (Continued from Example 1) The AF F shown in Figure 1 has seven **pr** labelings. Thus, according to $P_{\text{pr}^u}(F)$ each receives probability $1/7$ (see Table 1). First note that m is not independent of t :

$P_{\text{pr}^u}(F)(m = \mathbf{I}) = 3/7$ $P_{\text{pr}^u}(F)(m = \mathbf{I} \mid t = \mathbf{I}) = 1/3$.
It follows that $\{m, f\}$ is not independent of $\{t, d\}$ (cf. Remark 1) and hence that $P_{\text{pr}^u}(F)$ is not SCC factorable. Furthermore, **pr**^u fails probabilistic directionality. To see why, consider the unattacked set $B = \{m, f\}$ and note that

$$P_{\text{pr}^u}(F\downarrow B)(m = \mathbf{I}) = P_{\text{pr}^u}(F\downarrow B)(f = \mathbf{I}) = 1/2.$$

On the other hand we have

$$P_{\text{pr}^u}(F)(m = \mathbf{I}) = 3/7 \text{ and } P_{\text{pr}^u}(F)(f = \mathbf{I}) = 4/7.$$

	m	f	t	d	r	c	w	P
L_1	\mathbf{I}	\mathbf{O}	\mathbf{I}	\mathbf{O}	\mathbf{I}	\mathbf{O}	\mathbf{I}	$1/7$
L_2	\mathbf{I}	\mathbf{O}	\mathbf{O}	\mathbf{I}	\mathbf{O}	\mathbf{I}	\mathbf{O}	$1/7$
L_3	\mathbf{I}	\mathbf{O}	\mathbf{O}	\mathbf{I}	\mathbf{O}	\mathbf{O}	\mathbf{I}	$1/7$
L_4	\mathbf{O}	\mathbf{I}	\mathbf{I}	\mathbf{O}	\mathbf{O}	\mathbf{I}	\mathbf{O}	$1/7$
L_5	\mathbf{O}	\mathbf{I}	\mathbf{I}	\mathbf{O}	\mathbf{O}	\mathbf{O}	\mathbf{I}	$1/7$
L_6	\mathbf{O}	\mathbf{I}	\mathbf{O}	\mathbf{I}	\mathbf{O}	\mathbf{I}	\mathbf{O}	$1/7$
L_7	\mathbf{O}	\mathbf{I}	\mathbf{O}	\mathbf{I}	\mathbf{O}	\mathbf{O}	\mathbf{I}	$1/7$

Table 1: The distribution $P_{\text{pr}^u}(F)$ for F shown in Figure 1

Thus, the straightforward application of the principle of indifference in this scheme does not generally lead to well-behaved semantics with respect to the principles discussed in Section 3. The semantics we consider next fixes this.

4.2 The σ -SCC-uniform semantics

The next scheme is, like the σ -uniform scheme, based on the principle of indifference, but now applied to each SCC separately. We call it the *σ -SCC-uniform* scheme and define it in terms of a local probability function that, in turn, is defined in terms of \mathbb{L}_σ (see definition 14).

Definition 21. Let σ be a semantics. We denote by \mathbb{P}_σ^{su} the local probability function defined by

$$\mathbb{P}_\sigma^{su}(F, \mathbf{L}_I)(\mathbf{L}) = \begin{cases} \frac{1}{|\mathbb{L}_\sigma(F, \mathbf{L}_I)|}, & \text{if } \mathbf{L} \in \mathbb{L}_\sigma(F, \mathbf{L}_I) \\ 0, & \text{otherwise.} \end{cases}$$

We call the semantics represented by \mathbb{P}_σ^{su} the *σ -SCC-uniform semantics* and denote it by σ^{su} .

Earlier we saw that, for any semantics σ , σ^u is σ -compatible. The same does not apply to σ^{su} . The following example demonstrates that ss^{su} is not ss -compatible.

Example 5. Define F by $A_F = \{a, b, c\}$ and $a \rightsquigarrow_F b$, $b \rightsquigarrow_F a$, $b \rightsquigarrow_F c$, and $c \rightsquigarrow_F c$. Define $\mathbf{L} \in \mathcal{L}(F)$ by $\mathbf{L}(a) = \mathbf{I}$, $\mathbf{L}(b) = \mathbf{O}$ and $\mathbf{L}(c) = \mathbf{U}$. While $\mathbf{L} \notin \mathcal{L}_{\text{ss}}(F)$, we do have $P_{\text{ss}^{su}}(F)(\mathbf{L}) = 1/2$.

However, we do have that compatibility is guaranteed if we use an SCC decomposable semantics:

Proposition 7. *If σ is SCC decomposable then σ^{su} is σ -compatible.*

Proof (Sketch). Definitions 13 and 21 imply that, if σ is SCC decomposable, then we have $\mathbf{L} \in \mathcal{L}_\sigma(F)$ iff for all $S \in \text{SCC}(F)$, $\mathbb{P}_\sigma^{su}(F \downarrow C_F(S), \mathbf{L} \downarrow OP_F(S)) > 0$. Hence $\mathbf{L} \in \mathcal{L}_\sigma(F)$ iff $P_\sigma^{su}(F)(\mathbf{L}) > 0$. \square

Unlike the σ -uniform scheme, any semantics defined by the σ -SCC-uniform scheme satisfies probabilistic SCC decomposability. This follows directly from how it is defined.

Proposition 8. *For any semantics σ , σ^{su} satisfies probabilistic SCC decomposability.*

The following example demonstrates how the pr^{su} semantics deals with our running example.

	m	f	t	d	r	c	w	P
\mathbf{L}_1	\mathbf{I}	\mathbf{O}	\mathbf{I}	\mathbf{O}	\mathbf{I}	\mathbf{O}	\mathbf{I}	$1/4$
\mathbf{L}_2	\mathbf{I}	\mathbf{O}	\mathbf{O}	\mathbf{I}	\mathbf{O}	\mathbf{I}	\mathbf{O}	$1/8$
\mathbf{L}_3	\mathbf{I}	\mathbf{O}	\mathbf{O}	\mathbf{I}	\mathbf{O}	\mathbf{O}	\mathbf{I}	$1/8$
\mathbf{L}_4	\mathbf{O}	\mathbf{I}	\mathbf{I}	\mathbf{O}	\mathbf{O}	\mathbf{I}	\mathbf{O}	$1/8$
\mathbf{L}_5	\mathbf{O}	\mathbf{I}	\mathbf{I}	\mathbf{O}	\mathbf{O}	\mathbf{O}	\mathbf{I}	$1/8$
\mathbf{L}_6	\mathbf{O}	\mathbf{I}	\mathbf{O}	\mathbf{I}	\mathbf{O}	\mathbf{I}	\mathbf{O}	$1/8$
\mathbf{L}_7	\mathbf{O}	\mathbf{I}	\mathbf{O}	\mathbf{I}	\mathbf{O}	\mathbf{O}	\mathbf{I}	$1/8$

Table 2: The distribution $P_{\text{pr}^{su}}(F)$ for F shown in Figure 1




Figure 2: An AF with local conditional probabilities

Example 6. (Continued from Example 1) Table 2 shows the distribution $P_{\text{pr}^{su}}(F)$. To see how this distribution is determined, consider Figure 2. Here, each SCC S is annotated with a table showing, for all the relevant labelings \mathbf{L}_I of the outparents of S , the distribution $\mathbb{P}_{\text{pr}}^{su}(F \downarrow C_F(S), \mathbf{L}_I)$ over S . Columns containing labelings of the outparents of S (if any) are gray. We can use these tables to verify the probability of a labeling by multiplying entries according to Definition 15. For example $P_{\text{pr}^{su}}(F)(\mathbf{L}_1) = 1/2 \times 1/2 \times 1 \times 1 = 1/4$.

Since pr^{su} satisfies probabilistic SCC decomposability it satisfies probabilistic directionality. Indeed, we now have

$$P_{\text{pr}^{su}}(F)(m = \mathbf{I}) = P_{\text{pr}^{su}}(F)(f = \mathbf{I}) = 1/2.$$

It also follows that $P_{\text{pr}^{su}}(F)$ is SCC factorable. Indeed, m is now, unlike in Example 4, independent of t :

$$P_{\text{pr}^{su}}(F)(m = \mathbf{I}) = P_{\text{pr}^{su}}(F)(m = \mathbf{I} \mid t = \mathbf{I}) = 1/2.$$

On the other hand m is *not* independent of t given r , since we have $P_{\text{pr}^{su}}(F)(m = \mathbf{I} \mid r = \mathbf{O}) = 1/3$ and $P_{\text{pr}^{su}}(F)(m = \mathbf{I} \mid r = \mathbf{O} \cap t = \mathbf{I}) = 0$. However in Example 1 we argued that this is to be expected.

In sum, the σ -SCC uniform scheme defines probabilistic semantics that are well-behaved with respect to the principles discussed in Section 3. However, σ^{su} is σ -compatible only if σ is SCC decomposable. In this sense, the σ -SCC uniform scheme cannot be combined with a non-SCC decomposable semantics like the ss semantics.

5 Parameterised Probabilistic Semantics

So far we only considered purely abstract AFs, where each argument is treated equally. We now consider semantics for AFs where each argument x is associated with a probability $p(x)$. This probability can be thought of as being associated with the event that x is a valid argument. A probabilistic semantics should, in this setting, take these probabilities into account. To motivate this, consider the following example.

Example 7 (Alcohol and antibiotics). Let F be the AF shown in Figure 3. The scenario described here is similar to the one in Example 1, except that it only deals with choice between m and f , which affects r , which in turn affects c and w . Furthermore we believe that the agent might have taken antibiotics (a), implying that he cannot have alcohol and therefore does not drink red wine (a attacks r). However, we are uncertain about this; a is valid only with probability $3/4$. All the other arguments are valid with probability 1. In evaluating this AF, similar considerations apply as in Example 1. For example, learning that the agent chose m or f should not affect our belief about whether he took antibiotics: $\{m, f\}$ is independent of a .




Figure 3: The alcohol and antibiotics example

The AF in Figure 3 is an example of what (Li, Oren, and Norman 2011) call a *probabilistic AF*. Formally, this is a pair (F, p) where F is an AF and p assigns to each $x \in A_F$ a probability. However, to obtain a compatible notion of probabilistic semantics, we model it differently. First of all, we take a *probability assignment* to be a function assigning a probability to each member of the universe of arguments.

Definition 22. A *probability assignment* (PA) is a function $p : \mathcal{A} \rightarrow [0, 1]$.

Instead of pairing p with F , we use p as a parameter of a *parameterised* probabilistic semantics. Thus, a parameterised probabilistic semantics π^p (i. e., π parameterised by p) is a probabilistic semantics in the sense of definition 9. This does introduce some redundancy, since p must assign probabilities to arguments that we may not be interested in. To solve this we assume that $p(x) = 1$ whenever $p(x)$ is not specified. If one insists on thinking in terms of probabilistic AFs, one can substitute $P_{\pi^p}(F)$ with $P_{\pi}((F, p))$ on the understanding that $p(x)$ is specified for all $x \in A_F$.

5.1 The σ -Constellations Semantics

We now describe a scheme for defining a parameterised probabilistic semantics that generalises the σ -uniform scheme. We call this the σ -constellations scheme. It is based on the *constellations* approach (Li, Oren, and Norman 2011; Hunter 2013). Here, the idea is that the probabilities of the arguments of an AF F induce a probability distribution over the subsets of A_F . In this context we refer to these subsets as *framework states*. Intuitively, a framework state B represents the situation that the members of B are valid and that the members of $A_F \setminus B$ are not. Given a PA p , the distribution over the framework states of F (denoted by $J_{(A_F, p)}$) is determined by assuming that events of different arguments being valid are independent:

$$J_{(A, p)}(B) = \left(\prod_{x \in B} p(x) \right) \times \left(\prod_{x \in A \setminus B} (1 - p(x)) \right).$$

Each framework state B of F gives rise to a sub-graph $F \downarrow B$ and hence (given a semantics σ) to the set $\mathcal{L}_{\sigma}(F \downarrow B)$ of labelings. We assume that invalid arguments are labeled \mathbf{O} , which leads to the following definition.

Definition 23. Let σ be a semantics, let F an AF and let $B \subseteq A_F$. The set of σ/B *labelings* of F is denoted by $\mathcal{L}_{\sigma/B}(F)$ and consists of all $\mathbf{L} \in \mathcal{L}(F)$ satisfying

$$\mathbf{L} \downarrow B \in \mathcal{L}_{\sigma}(F \downarrow B) \text{ and } \forall x \in A_F \setminus B, \mathbf{L}(x) = \mathbf{O}.$$

If we follow (Li, Oren, and Norman 2011) then the probability of a labeling \mathbf{L} of F is the sum of the probabilities of all framework states that produce \mathbf{L} :

$$P(\mathbf{L}) = \sum_{B \subseteq A_F, \mathbf{L} \in \mathcal{L}_{\sigma/B}(F)} J_{(A_F, p)}(B).$$

But if we use this formula, the probabilities of the labelings of F may sum up to more than 1, since the same labeling may be produced by more than one framework state. Instead, we uniformly distribute the probability of a framework state over the set of labelings it produces. Given a PA p , the σ -constellations semantics (denoted σ_p^c) is therefore defined as follows.

Definition 24. Let σ be a semantics and p a PA. The σ -constellation semantics σ_p^c of an AF F is defined by

$$P_{\sigma_p^c}(F)(\mathbf{L}) = \sum_{B \subseteq A_F, \mathbf{L} \in \mathcal{L}_{\sigma/B}(F)} \frac{J_{(A_F, p)}(B)}{|\mathcal{L}_{\sigma/B}(F)|}.$$

The σ -constellation scheme generalises the σ -uniform scheme in the following sense.

Proposition 9. *If σ is a semantics, F an AF, and p a PA where, for all $x \in A_F$, $p(x) = 1$, then $P_{\sigma_p^c}(F) = P_{\sigma^u}(F)$.*

Proof (Sketch). If for all $x \in A_F$, $p(x) = 1$ then $J_{(A_F, p)}(B) = 1$ iff $B = A_F$. Hence $P_{\sigma_p^c}(F) = P_{\sigma^u}(F)$. \square

Obviously, if p assigns probability less than one to an argument, then σ_p^c may assign nonzero probability to a labeling that is not a σ labeling. In this case, σ_p^c is not σ -compatible. What remains is to check whether a semantics defined by the σ -constellations scheme satisfies probabilistic SCC decomposability. For the grounded semantics it does:

Proposition 10. *For each PA p , \mathbf{gr}_p^c satisfies probabilistic SCC decomposability.*

Proof (Sketch). Using the fact that the grounded semantics is SCC decomposable and that the grounded labeling is unique, we can show that, for all F , $P_{\mathbf{gr}_p^c}(F)$ equals $P_{\mathbf{gr}_p^{cs}}(F)$ as defined in definition 26. Proposition 12 establishes that $P_{\mathbf{gr}_p^{cs}}$ is probabilistically SCC decomposable. \square

However, propositions 6 and 9 imply that this does not hold in general. Here we demonstrate failure of probabilistic SCC decomposability under the \mathbf{pr}_p^c semantics.

	m	f	a	r	c	w	P
L_1	\mathbf{I}	\mathbf{O}	\mathbf{O}	\mathbf{I}	\mathbf{O}	\mathbf{I}	1/12
L_2	\mathbf{O}	\mathbf{I}	\mathbf{O}	\mathbf{O}	\mathbf{I}	\mathbf{O}	1/12
L_3	\mathbf{O}	\mathbf{I}	\mathbf{O}	\mathbf{O}	\mathbf{O}	\mathbf{I}	1/12
L_4	\mathbf{I}	\mathbf{O}	\mathbf{I}	\mathbf{O}	\mathbf{O}	\mathbf{I}	3/16
L_5	\mathbf{I}	\mathbf{O}	\mathbf{I}	\mathbf{O}	\mathbf{I}	\mathbf{O}	3/16
L_6	\mathbf{O}	\mathbf{I}	\mathbf{I}	\mathbf{O}	\mathbf{O}	\mathbf{I}	3/16
L_7	\mathbf{O}	\mathbf{I}	\mathbf{I}	\mathbf{O}	\mathbf{I}	\mathbf{O}	3/16

Table 3: The distribution $P_{\mathbf{pr}_p^c}(F)$ for F shown in Figure 3

Example 8. (Continued from Example 7) Let F be the AF shown in Figure 3. Let p be a PA that assigns probability $3/4$ to a and probability 1 to all other arguments. The probability distribution $P_{\mathbf{pr}_p^c}(F)$ is shown in Table 3. Explanation: the two framework states with nonzero probability are $B_1 = \{m, f, r, c, w\}$ and $B_2 = \{m, f, r, c, w, a\}$ with $J_{(A_F, p)}(B_1) = 1/4$ and $J_{(A_F, p)}(B_2) = 3/4$. We have $\mathcal{L}_{\mathbf{pr}/B_1}(F) = \{L_1, L_2, L_3\}$ and $\mathcal{L}_{\mathbf{pr}/B_2}(F) = \{L_4, L_5, L_6, L_7\}$. Thus we distribute $1/4$ uniformly over L_1, L_2 and L_3 and $3/4$ uniformly over L_4, L_5, L_6 and L_7 .

Note that m is not independent of a :

$$P_{\mathbf{pr}_p^c}(F)(m = \mathbf{I}) = 11/24 \quad P_{\mathbf{pr}_p^c}(F)(m = \mathbf{I} \mid a = \mathbf{O}) = 1/3.$$

This is a violation of SCC factorability. Therefore \mathbf{pr}_p^c does not satisfy probabilistic SCC decomposability.

5.2 The σ -SCC-Constellations Semantics

We now consider a scheme that adapts the σ -constellations scheme similarly to how the σ -SCC-uniform scheme adapts the σ -uniform scheme. Like the σ -constellations scheme, it is based on the constellations approach, but now applied to each SCC separately. We first define the following function,

which we use to determine the labelings of an SCC given a valuation of its outparents and a “local framework state”.

Definition 25. Given a semantics σ we denote by \mathbb{C}_σ the function that assigns, to each AF F , unattached set I of F , labeling $\mathbf{L}_I \in \mathcal{L}(I)$ and set $B \subseteq A_F \setminus I$, a set $\mathbb{C}_\sigma(F, \mathbf{L}_I, B)$ containing all $\mathbf{L} \in \mathcal{L}(A_F \setminus I)$ such that $\mathbf{L} \downarrow B \in \mathbb{L}_\sigma(F \downarrow B \cup I, \mathbf{L}_I)$ and $\forall x \in A_F \setminus (B \cup I), \mathbf{L}(x) = \mathbf{O}$.

We define the σ -SCC-constellations semantics in terms of a local probability function parameterised by a PA p . Note the similarity with Definition 24.

Definition 26. Let σ be a semantics and p a PA. We define the local probability function $\mathbb{P}_{(\sigma,p)}^{cs}$ by

$$\mathbb{P}_{(\sigma,p)}^{cs}(F, \mathbf{L}_I)(\mathbf{L}) = \sum_{\substack{B \subseteq A_F \setminus I, \\ \mathbf{L} \in \mathbb{C}_\sigma(F, \mathbf{L}_I, B)}} \frac{J_{(A_F \setminus I, p)}(B)}{|\mathbb{C}_\sigma(F, \mathbf{L}_I, B)|}.$$

We call the semantics represented by $\mathbb{P}_{(\sigma,p)}^{cs}$ the σ -SCC-constellations semantics and denote it by σ_p^{cs} .

The σ -SCC-constellation semantics generalises the σ -SCC-uniform semantics in the following sense.

Proposition 11. If σ is a semantics, F an AF, and p a PA where, for all $x \in A_F, p(x) = 1$, then $P_{\sigma_p^{cs}}(F) = P_{\sigma^{su}}(F)$.

Proof (Sketch). If for all $x \in A_F, p(x) = 1$ then $J_{(A_F, p)}(B) = 1$ iff $B = A_F$. Hence $\mathbb{P}_{(\sigma,p)}^{cs}(F) = \mathbb{P}_{\sigma^{su}}(F)$. \square

Since σ_p^{cs} is represented by $\mathbb{P}_{(\sigma,p)}^{cs}$, we have:

Proposition 12. For any semantics σ and PA p, σ_p^{cs} satisfies probabilistic SCC decomposability.

The following example shows how this approach deals with our running example.

	m	f	a	r	c	w	P
\mathbf{L}_1	\mathbf{I}	\mathbf{O}	\mathbf{O}	\mathbf{I}	\mathbf{O}	\mathbf{I}	$1/8$
\mathbf{L}_2	\mathbf{O}	\mathbf{I}	\mathbf{O}	\mathbf{O}	\mathbf{I}	\mathbf{O}	$1/16$
\mathbf{L}_3	\mathbf{O}	\mathbf{I}	\mathbf{O}	\mathbf{O}	\mathbf{O}	\mathbf{I}	$1/16$
\mathbf{L}_4	\mathbf{I}	\mathbf{O}	\mathbf{I}	\mathbf{O}	\mathbf{O}	\mathbf{I}	$3/16$
\mathbf{L}_5	\mathbf{I}	\mathbf{O}	\mathbf{I}	\mathbf{O}	\mathbf{I}	\mathbf{O}	$3/16$
\mathbf{L}_6	\mathbf{O}	\mathbf{I}	\mathbf{I}	\mathbf{O}	\mathbf{O}	\mathbf{I}	$3/16$
\mathbf{L}_7	\mathbf{O}	\mathbf{I}	\mathbf{I}	\mathbf{O}	\mathbf{I}	\mathbf{O}	$3/16$

Table 4: The distribution $P_{\mathbf{pr}_p^{cs}}(F)$ for F shown in Figure 3




Figure 4: Local conditional probabilities for Example 9

Example 9. (Continued from Example 8) The probability distribution $P_{\mathbf{pr}_p^{cs}}(F)$ is shown in Table 4. Figure 4 shows, for each SCC S of F and all the relevant input labelings \mathbf{L}_I of S , a table depicting the probability distribution $\mathbb{P}_{(\mathbf{pr}, p)}^{cs}(F \downarrow C_F(S), \mathbf{L}_I)$. Like in Example 6, we can use these tables to verify the probabilities shown in Table 4 by multiplying entries according to Definition 15.

Since \mathbf{pr}_p^{cs} satisfies probabilistic SCC decomposability it also satisfies probabilistic directionality. Thus, we now have

$$P_{\mathbf{pr}_p^{cs}}(F)(m = \mathbf{I}) = P_{\mathbf{pr}_p^{cs}}(F)(f = \mathbf{I}) = 1/2.$$

Furthermore, $P_{\mathbf{pr}_p^{cs}}(F)$ is SCC factorable, implying that $\{c, a\}$ is independent of a . Indeed we have, for example

$$P_{\mathbf{pr}_p^{cs}}(F)(m = \mathbf{I} \mid a = \mathbf{O}) = P_{\mathbf{pr}_p^{cs}}(F)(m = \mathbf{I}) = 1/2.$$

Thus, unlike the σ -uniform scheme, the σ -SCC uniform defines parameterised probabilistic semantics that are well-behaved with respect to the principles discussed in Section 3.

6 Discussion and Related Work

In this section we discuss a number of open questions, possible directions for future work, and connections with related approaches in the literature.

6.1 Bayesian Networks

In Section 3.3 we mentioned that the notion of SCC factorisation connects our approach to the theory of Bayesian networks. Let us first recall the definition. A Bayesian network over a set \mathcal{X} of variables is a directed acyclic graph $G = (\mathcal{X}, \rightarrow)$ (Koller and Friedman 2009). A probability distribution over \mathcal{X} is *factorable* with respect to a G if P can be decomposed into the product of conditional probabilities corresponding to each variable given its parents:

Definition 27. Let \mathcal{X} be a set of variables. A *Bayesian network* over \mathcal{X} is a directed acyclic graph $G = (\mathcal{X}, \rightarrow)$. A probability distribution P over the set Ω determined by \mathcal{X} is *factorable* with respect to G iff for all $\mathbf{V} \in \Omega$,

$$P(\mathbf{V}) = \prod_{x \in \mathcal{X}} P(x = \mathbf{V}(x) \mid \mathbf{V} \downarrow Pa_G(x)) \quad (2)$$

where $Pa_G(x)$ denotes the parents of x in G .

A crucial difference between AFs and Bayesian networks is that AFs may contain cycles. However, if we consider acyclic AFs, the two notions of factorability coincide. In this sense factorability is a special case of SCC factorability.

Proposition 13. If F is acyclic then a distribution P over $\mathcal{L}(F)$ is SCC-factorable w.r.t. F iff it is factorable w.r.t. F .

Moreover, if F is acyclic, Proposition 2 implies that F gets interpreted as a Bayesian network: each argument is independent of its nondescendants given its parents.

Finally, a probabilistically SCC decomposable semantics describes the evaluation of an AF similarly to how a Bayesian network plus a set of conditional probability tables (CPTs) describe a distribution. This can be seen in Figures 2 and 4, which include tables playing the role of the CPTs, but defined over sets of variables rather than single variables.

We believe that abstract argumentation theory can profit from the extensive theory of Bayesian networks. The insights obtained here establish a connection between the two formalisms on a formal level and can be understood as laying the foundations for future investigations in this direction. One example is to look at hybrid approaches, based on graphs that combine arguments and attacks (an “AF part”) as well as random variables whose relation is described using CPTs (a “BN part”). This is an example of *fibring* as considered in (Gabbay 2009).

Other work combining argumentation with Bayesian networks includes approaches to do argumentation *based on* Bayesian networks (Timmer et al. 2017; Vreeswijk 2004) or fusing Bayesian networks using argumentation (Nielsen and Parsons 2007). All this is quite different from what we do. An exception is (Gabbay and Rodrigues 2016), which deals with translating Bayesian nets into a kind of numerical AFs, but leaves handling of cycles to future work. A qualitative version of SCC factorability called *SCC stratification* was considered before in (Rienstra and Thimm 2018) in the context of ranking-based semantics for argumentation.

6.2 Decomposition based on SCC Recursiveness

Suppose we extend our earlier example as shown in Figure 5. We now have two SCCs. Hence, SCC factorability no longer implies independence between the SCCs of the original AF. However, these independencies arguably should still hold, because the argument y , which is responsible for making the original SCCs interdependent, is “neutralised” by x .




Figure 5: An AF consisting of two SCCs

We can deal with this example using the *SCC recursive* scheme (Baroni, Giacomin, and Guida 2005). This means that we first evaluate the SCC $\{x\}$ and assign to x the label **I**. Then we remove any argument attacked by arguments previously labeled **I**. Thus, we remove y , thereby eliminating the dependencies introduced by y , ending up again with the original SCCs $\{m, f\}$, $\{t, d\}$, $\{r\}$ and $\{c, w\}$. In future work we plan to address the question of how the SCC recursive scheme can be applied in a probabilistic setting.

6.3 Properties Based on Marginal Probabilities

Various rationality postulates for the epistemic approach were discussed in (Hunter and Thimm 2017). These postulates also apply to the our setting but deal purely with *marginal* probabilities of arguments. In our setting, the marginal probability of an argument x is simply the probability of $x = \mathbf{I}$. Consistent with their notation we write, given an argument x and distribution P , $P(x)$ as shorthand for $P(x = \mathbf{I})$. We consider only the most basic postulates considered in (Hunter and Thimm 2017). Let F be an AF and let P be a probability distribution over $\mathcal{L}(F)$. Then:

FOU P is *founded* wrt. F if $P(x) = 1$ for every $x \in A_F$ s.t. there is no $y \in A_F$ s.t. $y \rightsquigarrow_F x$.

RAT P is *rational* wrt. F if for every $x, y \in A_F$, if $x \rightsquigarrow_F y$ then $P(x) > 0.5$ implies $P(y) \leq 0.5$.

COH P is *coherent* wrt. F if for every $x, y \in A_F$, if $x \rightsquigarrow_F y$ then $P(x) \leq 1 - P(y)$.

The intuition behind these properties is as follows. **FOU**: if x is not attacked, then x is believed without doubt (i. e. $P(x) = 1$). **RAT**: if x attacks y , and x is somewhat believed (i. e. $P(x) > 0.5$), then y is somewhat not believed (i. e. $P(y) \leq 0.5$). **COH**: if x attacks y , then the degree to which x is believed caps the degree to which y can be believed.

These postulates were already shown to hold for uniform distributions over labelings in (Thimm et al. 2017). These results carry over to the σ -uniform semantics:

Proposition 14. *Let F be an AF. If $\sigma \in \{\mathbf{gr}, \mathbf{co}, \mathbf{pr}, \mathbf{ss}\}$ then $P_{\sigma^u}(F)$ is founded, rational, and coherent.*

We now turn to the σ -SCC-uniform semantics. For that, we actually get the same result as above.

Proposition 15. *Let F be an AF. If $\sigma \in \{\mathbf{gr}, \mathbf{co}, \mathbf{pr}, \mathbf{ss}\}$ then $P_{\sigma^{su}}(F)$ is founded, rational, and coherent.*

Thus, our approach satisfies at least some basic postulates from the epistemic approach. Additional postulates from this approach will be considered in future work.

6.4 Further Semantics and Principles

While in this paper we focus on the complete, grounded, preferred and semi-stable semantics, our approach does not depend on this restriction. What matters in our approach is whether a semantics is SCC decomposable or not. Apart from SCC decomposability, many other principles for semantics of abstract AFs have been studied (van der Torre and Vesic 2017). What these principles mean in a probabilistic setting is a topic for future work. Finally, we have considered only universally-defined semantics, ruling out the stable semantics. This restriction can be lifted if we require that a principle is valid only for AFs admitting at least one labeling, as done in (Baroni, Dunne, and Giacomin 2011).

7 Conclusion

We formalised the requirement that the probabilistic evaluation of an AF should comply with probabilistic independence relationships implied by the topology of the AF. Our formalisation is based on probabilistic analogues of the SCC decomposability and directionality principles known from abstract argumentation. Probabilistic SCC decomposability can be used to determine whether a given probabilistic semantics is well-behaved but can also be used to define a (well-behaved) probabilistic semantics. A related principle called SCC factorability furthermore provides new insights into the relationship between abstract argumentation and the theory of Bayesian networks. We believe that these insights can guide future investigations into the connection between these formalisms.

Acknowledgements This work was partially supported by the Deutsche Forschungsgemeinschaft (grant KE 1686/3-1).

References

- Baroni, P., and Giacomin, M. 2007. On principle-based evaluation of extension-based argumentation semantics. *Artificial Intelligence* 171(10-15):675–700.
- Baroni, P.; Boella, G.; Cerutti, F.; Giacomin, M.; van der Torre, L.; and Villata, S. 2014. On the input/output behavior of argumentation frameworks. *Artificial Intelligence* 217:144–197.
- Baroni, P.; Gabbay, D.; Giacomin, M.; and van der Torre, L., eds. 2018. *Handbook of Formal Argumentation*. College Publications.
- Baroni, P.; Dunne, P. E.; and Giacomin, M. 2011. On the resolution-based family of abstract argumentation semantics and its grounded instance. *Artif. Intell.* 175(3-4):791–813.
- Baroni, P.; Giacomin, M.; and Guida, G. 2005. Secrecursiveness: a general schema for argumentation semantics. *Artificial Intelligence* 168(1-2):162–210.
- Baroni, P.; Giacomin, M.; and Vicig, P. 2014. On rationality conditions for epistemic probabilities in abstract argumentation. In *Proceedings of the Fifth International Conference on Computational Models of Argumentation (COMMA'14)*.
- Caminada, M. 2006. On the issue of reinstatement in argumentation. In *Logics in Artificial Intelligence, 10th European Conference, JELIA 2006, Liverpool, UK, September 13-15, 2006, Proceedings*, volume 4160 of *LNCS*, 111–123. Springer.
- Dung, P. M. 1995. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. *Artificial Intelligence* 77(2):321–358.
- Gabbay, D. M., and Rodrigues, O. 2016. Introducing bayesian argumentation networks. *The IfColog Journal of Logics and their Applications* 3(2):241–278.
- Gabbay, D. M. 2009. Fibring argumentation frames. *Studia Logica* 93(2-3):231–295.
- Grossi, D., and van der Hoek, W. 2013. Audience-based uncertainty in abstract argument games. In *IJCAI*, 143–149. IJCAI/AAAI.
- Hunter, A., and Thimm, M. 2016a. On partial information and contradictions in probabilistic abstract argumentation. In Baral, C.; Delgrande, J. P.; and Wolter, F., eds., *Principles of Knowledge Representation and Reasoning: Proceedings of the Fifteenth International Conference, KR 2016, Cape Town, South Africa, April 25-29, 2016.*, 53–62. AAAI Press.
- Hunter, A., and Thimm, M. 2016b. Optimization of dialectical outcomes in dialogical argumentation. *International Journal of Approximate Reasoning* 78:73–102.
- Hunter, A., and Thimm, M. 2017. Probabilistic reasoning with abstract argumentation frameworks. *Journal of Artificial Intelligence Research* 59:565–611.
- Hunter, A. 2013. A probabilistic approach to modelling uncertain logical arguments. *International Journal of Approximate Reasoning* 54(1):47–81.
- Hunter, A. 2015. Modelling the persuadee in asymmetric argumentation dialogues for persuasion. In Yang, Q., and Wooldridge, M., eds., *Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence, IJCAI 2015, Buenos Aires, Argentina, July 25-31, 2015*, 3055–3061. AAAI Press.
- Koller, D., and Friedman, N. 2009. *Probabilistic Graphical Models - Principles and Techniques*. MIT Press.
- Li, H.; Oren, N.; and Norman, T. J. 2011. Probabilistic argumentation frameworks. In *Theory and Applications of Formal Argumentation - First International Workshop, TFAFA 2011, Barcelona, Spain, July 16-17, 2011, Revised Selected Papers*, volume 7132 of *LNCS*, 1–16. Springer.
- Nielsen, S. H., and Parsons, S. 2007. An application of formal argumentation: Fusing bayesian networks in multi-agent systems. *Artif. Intell.* 171(10-15):754–775.
- Rienstra, T., and Thimm, M. 2018. Ranking functions over labelings. In *Proceedings of the Seventh International Conference on Computational Models of Argument (COMMA-2018), September 11-14, 2018, Warsaw, Poland*. IOS Press.
- Rienstra, T.; Thimm, M.; and Oren, N. 2013. Opponent models with uncertainty for strategic argumentation. In *IJCAI*, 332–338. IJCAI/AAAI.
- Rienstra, T. 2012. Towards a probabilistic Dung-style argumentation system. In *The First International Conference on Agreement Technologies, AT 2012, Dubrovnik, Croatia, October 15-16, 2012*, volume 918 of *CEUR Workshop Proceedings*, 138–152. CEUR-WS.org.
- Thimm, M.; Baroni, P.; Giacomin, M.; and Vicig, P. 2017. Probabilities on extensions in abstract argumentation. In *Proceedings of the 2017 International Workshop on Theory and Applications of Formal Argument (TFAFA'17)*.
- Thimm, M. 2012. A probabilistic semantics for abstract argumentation. In *ECAI 2012 - 20th European Conference on Artificial Intelligence, Montpellier, France, August 27-31, 2012*, volume 242 of *Frontiers in Artificial Intelligence and Applications*, 750–755. IOS Press.
- Timmer, S. T.; Meyer, J.-J. C.; Prakken, H.; Renooij, S.; and Verheij, B. 2017. A two-phase method for extracting explanatory arguments from bayesian networks. *International Journal of Approximate Reasoning* 80:475–494.
- Toniolo, A.; Norman, T. J.; and Oren, N. 2017. Enumerating preferred extensions: A case study of human reasoning. In *Proceedings of the 2017 International Workshop on Theory and Applications of Formal Argument (TFAFA'17)*.
- van der Torre, L., and Vesic, S. 2017. The principle-based approach to abstract argumentation semantics. *The IfColog Journal of Logics and their Applications*.
- Vreeswijk, G. 2004. Argumentation in bayesian belief networks. In *ArgMAS*, volume 3366 of *Lecture Notes in Computer Science*, 111–129. Springer.