A Probabilistic Approach to Modelling Uncertain Logical Arguments

Anthony Hunter
Department of Computer Science
University College London
London, WC1E 6BT, UK
(a.hunter@cs.ucl.ac.uk)

August 6, 2012

Abstract

Argumentation can be modelled at an abstract level using a directed graph where each node denotes an argument and each arc denotes an attack by one argument on another. Since arguments are often uncertain, it can be useful to quantify the uncertainty associated with each argument. Recently, there have been proposals to extend abstract argumentation to take this uncertainty into account. This assigns a probability value for each argument that represents the degree to which the argument is believed to hold, and this is then used to generate a probability distribution over the full subgraphs of the argument graph, which in turn can be used to determine the probability that a set of arguments is admissible or an extension. In order to more fully understand uncertainty in argumentation, in this paper, we extend this idea by considering logic-based argumentation with uncertain arguments. This is based on a probability distribution over models of the language, which can then be used to give a probability distribution over arguments that are constructed using classical logic. We show how this formalization of uncertainty of logical arguments relates to uncertainty of abstract arguments, and we consider a number of interesting classes of probability assignments.

1 Introduction

Argumentation is a vital aspect of intelligent behaviour by humans for dealing with the myriad of situations in everyday life where there are conflicting opinions and options, and where the only available information about these is incomplete and inconsistent. Computational models of argument aim to reflect how human argumentation uses incomplete and inconsistent information to construct and analyze arguments about the conflicting opinions and options. For reviews of some of the options available for computational models of argument, see [BCD07, BH08, RS09].

We start by considering abstract argumentation as proposed by Dung [Dun95]. In the spirit of generality, abstract argumentation is vague about what constitutes an argument and what constitutes an attack. It is assumed that these can be identified in a meaningful way. Since arguments are often uncertain, it may be useful to quantify the uncertainty associated with each argument. Continuing with the spirit of generality, we may assume that the uncertainty of an argument in abstract argumentation can also be identified in a meaningful way.

Example 1. *Consider arguments A_1 = "Patient has hypertension so prescribe diuretics", A_2 = "Patient has hypertension so prescribe betablockers", and A_3 = "Patient has emphysema which is a contraindication for betablockers". Here, we assume that A_1 and A_2 attack each other because we should only give one treatment and so giving one precludes the other, and we assume that A_3 attacks A_2 because it provides a counterargument to A_2. Hence, we get the following abstract argument graph.*

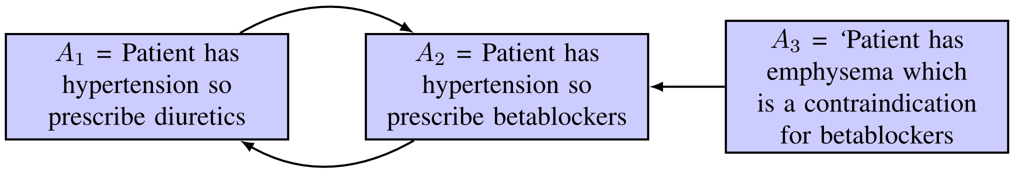

Normally there is uncertainty associated with such arguments. For instance, there may be uncertainty about whether the patient has emphysema, and there may be uncertainty as to whether this implies that it is a contraindication for using betablockers. We may suppose that a medical expert would be in a position to give a probability for this argument based on whether the probability of occurrence, and severity, of the side-effects arising with the use of betablockers on patients with emphysema being of such a degree that the patient should not take the treatment (i.e. that emphysema can be regarded as a contraindication for betablockers).

1.1 Introducing probabilities over arguments into abstract argumentation

Recently, developments of abstract argumentation that take into account the uncertainty of arguments have been presented including [BGW05, RRS$^+$07, DT10, LON11]. These introduce a probability assignment for each argument to represent the degree to which the argument is believed to hold, giving rise to a *probabilistic argument graph* (as proposed in [LON11]). This has provided a means to qualify the uncertainty that often exists concerning arguments in abstract argumentation, and more importantly, opens the way for better modelling of important phenomena in argumentation. By the later, we mean that uncertainty is a very important component in argumentation. Often doubt surrounds individual arguments, and qualifying that uncertainty enables a better modelling of the situation.

What is meant by the probability of an argument holding is an open question. Indeed, there seem to be various ways that we could answer this. In [Hun12], the *justification perspective* views the probability as indicating the degree to which the argument belongs to the graph (i.e. the probability that the argument is justified in appearing in the argument graph). We revisit this perspective later under the constellations approach (Sections 3.2 and 3.3). Alternatively, we could regard the probability as indicating the degree to which the argument is true. This in turn requires clarification. One perspective, which we will develop in this paper, is the *premises perspective* on the probability of an argument (being true). In this, the probability is based on the degree to which the premises of the argument are true, or believed to be true.

We can assume each argument in an argument graph has a proponent. This is someone, or some software system, that has put the argument into the graph. It may be that one person proposes all arguments in the argument graph, and annotates each argument with a probability that it is true. For instance, this could be when a clinician is in a discussion with a patient (as illustrated in Example 1) or when a clinician is making a record of a patient consultation (as illustrated in Example 2). Even for arguments that have a low probability, it may be useful to present it in the graph so that all the key information is present and so that risks can be taken into account.

Example 2. *Consider a doctor updating a patient's medical records after a consultation with the patient. First she writes "From his symptoms, the patient most likely has a cold.", and then she writes "However, there is a small possibility that he has influenza since it is currently very common". Let us represent the first sentence by argument A_1 and the second sentence by argument A_2 as follows.*

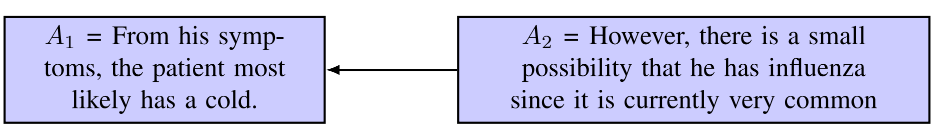

This representation hides the fact that the first argument is much more likely to be true than the second. If we use dialectical semantics to the above graph, A_1 is defeated by A_2. A better solution may be to translate the arguments to the following arguments that have the uncertainty removed from the textual descriptions and then express the uncertainty in the probability function over the arguments such that $p(A'_1) = 0.9$ and $p(A'_2) = 0.1$.

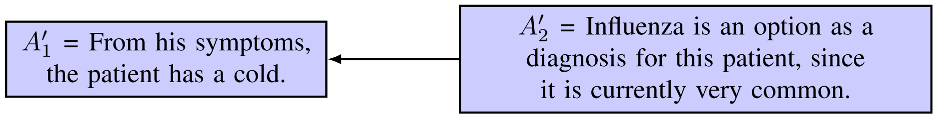

It may be the case that arguments come from different proponents. For example, in a political discussion on television, a sequence of arguments are presented by the various participants in the discussion. A member of the audience can listen to these arguments, and mentally put them into an argument graph, and at the same time annotate each argument with a probability of being true. So the person assigning the probability is separate from the proponents of the arguments. We illustrate this scenario in Figure 1.

Another scenario to consider is a meeting where the participants are a set of experts co-operating on making some joint decisions. In this case, when a participant posits an argument, s/he may also provide a probability value for that argument (i.e. the probability that s/he believes the argument is true), and then the other participants will accept the probability value as the probability that they use for the argument.

For instance, in multi-disciplinary meetings for care of cancer patients, different healthcare professionals in a hospital meet to discuss the current set of patients in their care in order to decide actions to take. So for each patient, different professionals will make arguments about diagnosis, prognosis, and options for interventions. Often, arguments will be presented based on expertise. For instance, if a surgeon is presenting an argument about a surgical option, the surgeon may qualify the argument with a belief in it being true. Other members of the meeting, particularly those who are not surgeons may well then accept that as the probability of it being true.

So whether the arguments and probability values are identified by the same person, or by different people, it is possible to estimate for each argument the probability that it is true. This gives the flexibility for arguments to be posited for which there is some doubt in whether they are true. It also allows for easier aggregation and analysis of arguments from multiple sources, since it allows the agent gathering and analyzing the arguments to identify their own probability function over the arguments. This is particularly useful when judging argumentation by others such as in debates and discussions. The probability value represents important meta-level information about the arguments, and for situations, where we need to deconstruct and analyze the information, it is valuable to represent this information explicitly.

1.2 Towards an understanding of probabilities over arguments

On the one hand, the generality of the definitions for probabilistic argument graphs as proposed in [LON11] is appealing, and as we have seen above, it appears possible to relate real-world argumentation to the probabilistic argument graphs. However, on the other hand, it is not so clear what these probabilities mean. For instance, what is the formal relationship between the uncertainty of the premises of an argument, and the probability of the argument holding? And what constitutes an acceptable, in the general sense of the word, probability distribution over arguments?

If we are more specific about the nature of the arguments, then we can be more precise about what the probability is. For instance, if we think of an argument as involving some premises (i.e. the support of the argument) and some conclusion derived from the premises (i.e. the claim of the argument), then the probability of the argument being true is a function of the probability of the premises being true, and the probability that the claim follows from those premises. So the probability of the argument being true is based on the uncertainty inherent in the information used and the uncertainty inherent in the reasoning used. In this paper, our focus is on the uncertainty inherent in the premises used in argumentation rather than the reasoning (i.e. the premises perspective).

Unfortunately, there is a gap in our understanding of how we can go from the purely probabilistic understanding of the evidence to this argumentation-theoretic representation. We need to fill this gap if we are to better understand probabilistic argument graphs, and if we are to be able to construct such argument graphs from data. In order to bridge this gap, we can consider logical arguments. Giving a probability assignment over premises, then raises the question of what this means, how can we check that it is consistent, and what can we do it is not consistent? Furthermore, how can we then use this information to get an evaluation of belief in the arguments, and how can we then use the probabilities over arguments to interpret the probabilities over the extensions. By adopting an established approach to representing belief in propositional formulae, for modelling uncertainty in arguments, we can address these questions in a

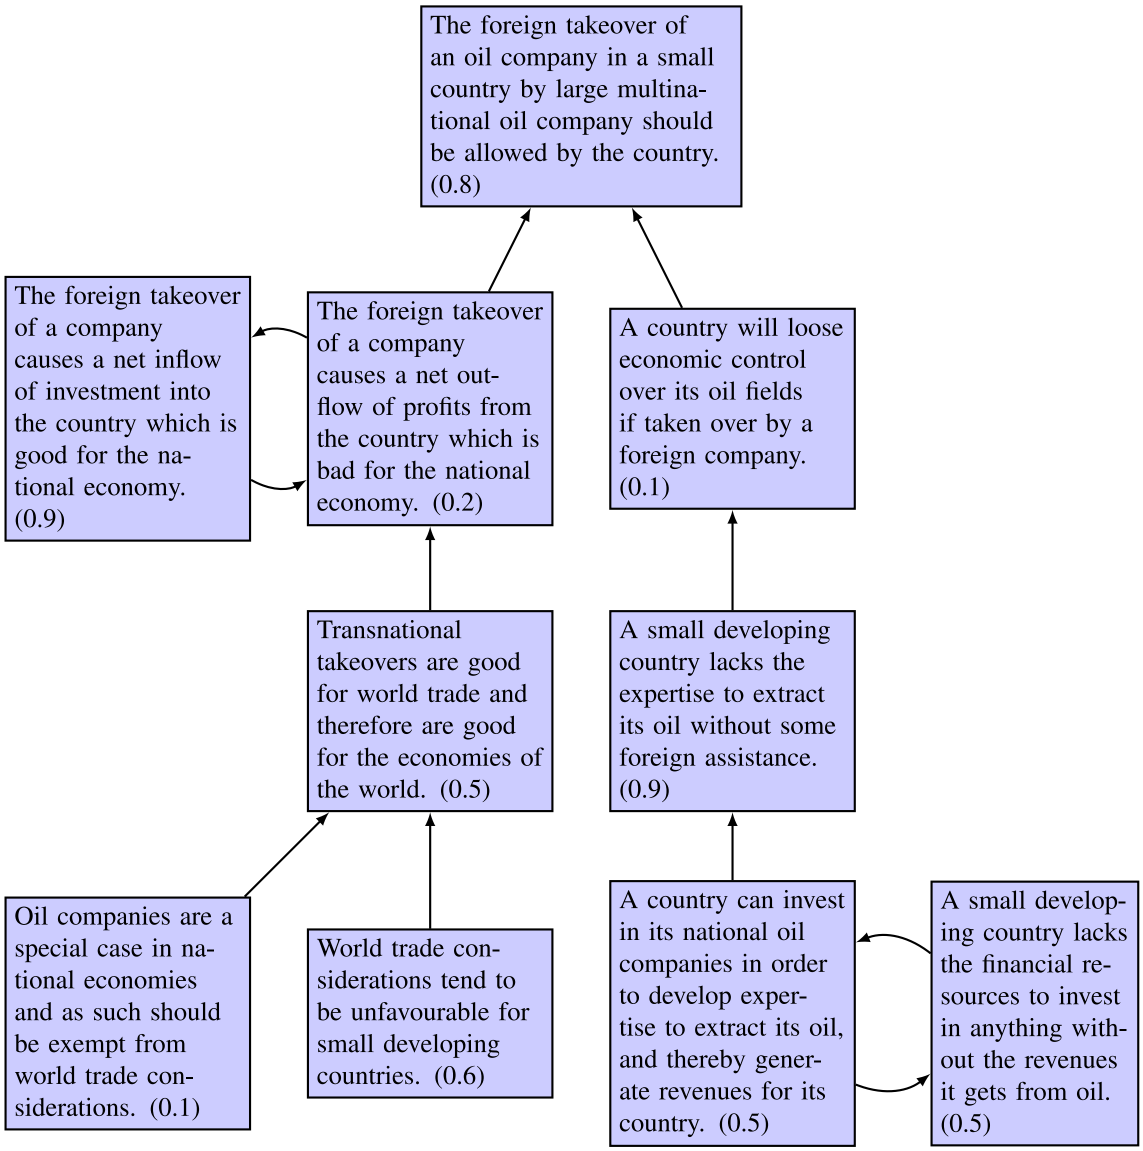

Figure 1: Consider a member of the audience listening to a radio documentary about the takeover of oil production companies in small developing countries by large multinationals. The documentary may be exploring the question of whether such small countries should permit these foreign takeovers. To explore the question, the documentary includes a number of interviews with experts from small developing countries, from multinational oil companies, and from financial institutes. Suppose the member of the audience records ten arguments, and puts them into the argument graph shown. For someone who is reasonably optimistic about multinational oil companies playing a beneficial role in developing countries, the probability value given for each argument (given in brackets in each box) may reflect their belief in each argument.

systematic way. We give an overview of how we do this in the next subsection (Section 1.3) and then formally develop the framework in Sections 2 onwards.

Dung and Thang [DT10] have also extended Dung's approach to abstract argumentation by introducing a probability distribution over sets of arguments, which they use with a version of assumption-based argumentation in which a subset of the rules are probability rules: Each probability rule has an associated probability value, and then through a form of specificity criterion, the probability of an argument can be determined. However, the approach does not provide a direct link between the belief in the set of all premises and the belief in the argument, rather it appears to provide a rule-based mechanism for qualifying the belief in some inferences.

In another rule-based system for argumentation, Riveret *et al* [RRS$^+$07], the belief in the premises of an argument is used to calculate the belief in the argument. For this, there is an assumption that the premises are independent. So for an argument with support $\{\phi_1, ..., \phi_n\}$, where the belief in ϕ_i is given by $P(\phi_i)$, then the belief in the argument is given by $P(\phi_1) \times ... \times P(\phi_n)$. The proposal does not investigate further the nature of this assignment, in particular, it is not used to qualify the probability of extensions in Dung's approach, but goes on to explore its use in dialogue.

For logical arguments, a probability distribution on models has been used by Haenni *et al* [Hae98, HKL00, Hae01] for a notion of probabilistic argumentation for diagnosis. The formulae are divided into those that represent observations and those that represent explanations. Then arguments are constructed for and against particular diagnoses (i.e. only arguments and counterarguments that rebut each other are considered). The approach provides valuable insights into how argumentation can be applied to diagnosis, and viable algorithms and tools have been developed for it. However, they restrict consideration to rebuttals and do not consider undercuts, and by recursion, undercuts to undercuts. More significantly, they do not consider the framework of abstract argumentation, and as a result there is no relationships between the proposal by Haenni *et al* and Dung's proposals for dialectical semantics [Dun95] or probabilistic argument graphs (as defined in [DT10, LON11]) have been established.

In the LA system, another logic-based framework for argumentation, probabilities are also introduced into the rules, and these probabilities are propagated by the inference rules so that arguments are qualified by probabilities (such as via labels such as "likely", "very likely", etc). However, there is no consideration of how the qualitative probabilities relate to some underlying semantics for the language, and there is no consideration of how this probabilistic information relates to Dung's proposals for dialectical semantics [EGKF93, FD00].

Whilst using weights on arguments (such as discussed in [BGW05]), allow for a notion of uncertainty to be represented, our understanding is incomplete for using such weights in a way that conforms with established theories of quantitative uncertainty. Preferences over arguments have been harnessed in argumentation theory (see for example [AC98, AC02]) in order to decide on a pairwise basis whether one argument defeats another argument. In some situations, this is a useful solution. However, it is not always clear where these preferences come from, or what they mean. Often, the preferences seem to be based on the relative strength of belief in the arguments. But, the notion has not yet been adequately explored.

Therefore there is a need to develop a deeper understanding of the role of probability theory in argumentation. In particular, we need a clearer understanding of the uncertainty of the premises of arguments relate to the uncertainty of the arguments, and how those arguments are evaluated. Of course probability theory is only one way of capturing uncertainty about arguments, and indeed, some interesting proposals have been made for using possibility theory in argumentation (see for example [AP04, ACGS08a, ACGS08b]).

1.3 Overview of the probabilistic approach used in this paper

To address the shortcomings raised above, the aim of this paper is to develop a framework for using more detailed information about uncertainty in arguments. For this, we will use classical logic for representing the arguments and for defining the attacks relations. We consider two types of probability distribution, and we investigate what they mean, how they relate to each other, and how they can be used.

- A probability distribution over a language $\mathcal{L}$, or equivalently over the classical (Boolean) models $\mathcal{M}^{\mathcal{L}}$ of the langauge, which gives the probability of each formula in the language.

- A probability distribution over the set of arguments $\mathcal{A}^{\mathcal{L}}$ in a language $\mathcal{L}$, where each argument $\langle \Phi, \psi \rangle$ is such that $\Phi \subseteq \mathcal{L}$ is a set of formulae, $\psi \in \mathcal{L}$ is a formula, and Φ entails ψ.

Given a probability distribution P over a language $\mathcal{L}$, we let the probability of an argument $A = \langle \Phi, \psi \rangle$ be the probability of the conjunction of its support (i.e. if Φ is $\{\phi_1, ..., \phi_n\}$, then the probability of the argument A is $P(\phi_1 \wedge .. \wedge \phi_n)$), as illustrated in Example 3. In this way, there is a simple and clear relationship between belief in the premises of an argument, and the belief in the argument, when we have a consistent probability distribution over a language $\mathcal{L}$ (i.e. when $\sum_{m \in \mathcal{M}^{\mathcal{L}}} P(m) = 1$).

Example 3. *Consider the question of whether a particular new treatment should be used by a patient. Suppose the following are the key points that the doctor and patient are discussing and that the evidence for these points come from the only published clinical trial for this new treatment.*

- *80% of patients taking drug X live for 5 or more years versus 40% of patients taking the placebo*
- *10% of patients taking drug X have a stroke versus 1% of patients taking the placebo*

Whilst this analysis is informative to clincians, it is unlikely to be the way that it would be explained to a patient. For patients it is normal to consider the treatment in terms of reasons for taking it, and for not taking it. For instance, if there is a good reason for taking it, such as there is strong evidence that it will have a positive effect on the patient, then this will be given as an argument. Then if there are any reasons to doubt this, for instance if there is some risk of side-effects, then this may be given as a counterargument. This is summarized in the following argument graph.

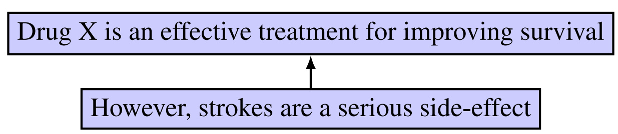

There are various ways we could represent this information in logic. We will adopt a simple approach based on propositional logic using the propositional atoms p = *"problematic treatment",* s = *"effective treatment for improving survival", and* r = *"strokes are a serious side-effect". Using these, we can construct two arguments. The first reflects the reasoning that the treatment will substantially improve survival, and the second reflects the reasoning that there is a risk of stroke as a side-effect, thereby undermining the assumption that the treatment is unproblematic.*

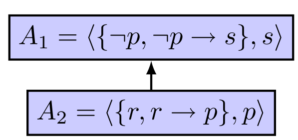

We can consider the models for the support of each argument as tabulated below. Then we can assign probabilities to the models to denote the belief that the model is the correct model.

Model	*p*	*s*	*r*	*Probability of model*
m_1	*false*	*true*	*true*	*0.35*
m_2	*false*	*true*	*false*	*0.35*
m_3	*true*	*true*	*true*	*0.12*
m_4	*true*	*false*	*true*	*0.18*

Now we can we go further and express belief in arguments. So the belief in an argument is the sum of the belief in the models of its support. The models of A_1 *are* m_1 *and* m_2*, and so the belief in the argument* A_1 *is* 0.7*. The models of* A_2 *are* m_3 *and* m_4*, and so the belief in the argument* A_2 *is* 0.3*.*

In practice, it is not always possible to have enough information to get a probability distribution over a language $\mathcal{L}$. We may for instance start with some subset of arguments $\mathcal{A} \subset \mathcal{A}^{\mathcal{L}}$ for which there is a probability assignment (i.e. there is a probability function $P : \mathcal{A} :\mapsto [0, 1]$) but this may be insufficient to determine the probability distribution over $\mathcal{L}$ or $\mathcal{A}^{\mathcal{L}}$. More challengingly is that the probability assignment to $\mathcal{A}$ may be such that there is no consistent probability distribution over $\mathcal{L}$ or $\mathcal{A}^{\mathcal{L}}$. By consistent, we mean that the sum of the probability assigned to the models sums to 1 (as we will fomally define in Section 2). Since argumentation is about dealing with inconsistent information (e.g. conflicting beliefs about what the world), we should be able to handle inconsistent probability distributions within computational models of argument. For this, we investigate some of the options available to us.

A simple situation where inconsistency arises is when someone is pooling evidence. For this, the agent may have some belief in the evidence given by the source of the evidence, and there is may be some conflicts in the beliefs about the sources. We illustrate this in the next example.

Example 4. *Consider a situation where conflicting evidence is being aggregated by a detective investigating a robbery. Suppose the detective uses an argument-based approach in the process, and identifies the following abstract arguments in an argument graph.*

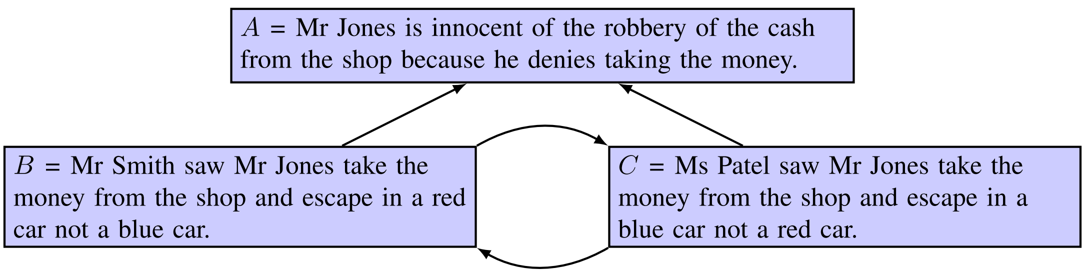

Also suppose, the detective thinks that both Mr Smith and Ms Patel seem to be reliable witnesses. The detective may use the probability function P to reflect this by letting $P(A) = 0.1$, $P(B) = 0.9$, and $P(C) = 0.9$.

Now, we can analyse this in more detail by representing the arguments by the following logical arguments in an argument graph where i denotes "Mr Jones is innocent", s denotes "Mr Jones stole the money", t denotes "Mr Jones took the money from the shop", r denotes "Mr Jones escaped in a red car", and b denotes "Mr Jones escaped in a blue car"

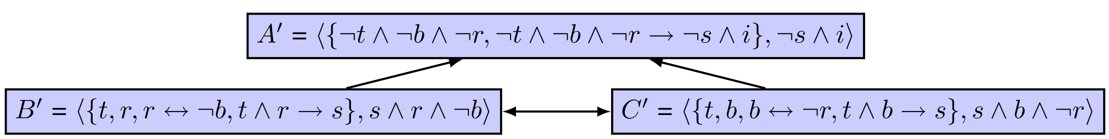

We can consider the models for the support of each of the arguments, and from the probabilities assigned to arguments, assign probabilities to the models. So m_1 is the model for the support of A, m_2 is the model for the support of B, and m_3 is the model for the support of C. Suppose the remaining models are assigned probability of 0.

Model	*i*	*s*	*t*	*r*	*b*	*Probability of model*
m_1	*true*	*false*	*false*	*false*	*false*	*0.1*
m_2	*false*	*true*	*true*	*true*	*false*	*0.9*
m_3	*false*	*true*	*true*	*false*	*true*	*0.9*

Here the sum of the probability assigned to the models is greater than 1. Hence, the probability distribution over the models is inconsistent.

In our approach, we decouple the logic of an argument from the uncertainty of an argument. In this way, we can harness existing proposals for logical argument systems based on classical logic, and in effect extend them with the ability to harness meta-level information about the quality of the formulae. Another reason that we decouple the logic and the uncertainty is that we want to look at the interplay of them, as illustrated in the next example.

Example 5. *Consider the following argument tree. Here, the undercut has explicit uncertainty (70% of cases). This would suggest that the undercut is a strong argument against the root.*

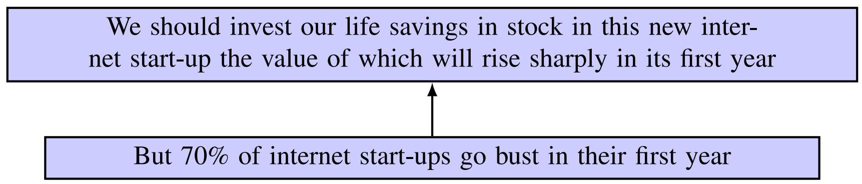

Now consider the following argument tree. Here, the undercut has explicit uncertainty (1% of cases). This would suggest that the undercut is a much weaker argument against the root than above. So we see that there is an interplay between the strength of belief in the argument and the strength of belief in the undercut. In particular, there appears to be a tipping point X for the value in the statement "But X% of internet start-ups go bust in their first year" below which we can ignore the undercut, and above which the undercut defeats the root argument.

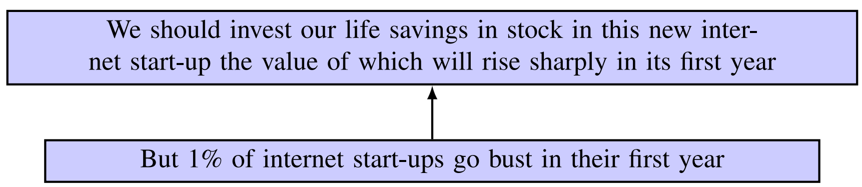

By decoupling the uncertainty from the logic, we can see what the reasons for the argument and counterargument are, and then identify the meta-level information concerning the uncertainty each of them.

Therefore, to develop our understanding of logic and probability theory in argumentation, we have three points to connect, namely "argumentation theory", "probability theory", and "formal logic". For "argumentation theory", we are assuming abstract argumentation proposed by Dung [Dun95]. This is a cornerstone of argumentation theory. It has been extensively studied formally, and it offers considerable insights into the nature of argumentation. For "probability theory", we can assume the usual presentation of it (e.g. [Hac01]), without at this stage committing to any of the particular philosophical perspectives on probability theory. Then for "formal logic", there is a myriad of options developed in philosophy, mathematics, and computer science that we could consider.

Now to connect the three points of "argumentation theory", "probability theory", and "formal logic", we are drawing on three lines of research. To connect "argumentation theory" and "probability theory", we are using proposals by Dung [DT10] and by Li *et al* [LON11], to introduce a probability function on arguments, which are the first proposals to reason with argument graphs using probability theory. To connect "argumentation theory" and "formal logic", the two main contenders, in terms of established research are defeasible logic and classical logic. Here, we use the later (drawing on the line of research the investigates the generation and and comparison of arguments and counterarguments in classical logic starting with Pollock [Pol89, Pol95], and then developed by a number of researchers including Cayrol [Cay95], Besnard and Hunter [BH01], and Gorogiannis and Hunter [GH11], for relating abstract argumentation with classical logic-based argumentation. The reason we use classical logic comes from the needs for linking "probability theory" and "formal logic". Much research has been undertaken on this topic, and it is not possible

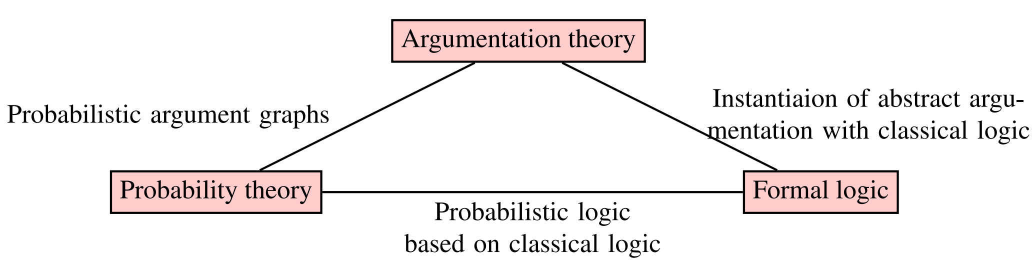

Figure 2: Summary of influences (given as labels on the arcs) used to combine the the topics of argumentation theory, formal logic, and probability theory

to do review of this literature in this paper (for some good reviews of aspects of this complex topic see [Car50, Bac90, Ada98, Wil02]). Many proposals consider how some aspects of probabilistic reasoning can be incorporated in classical logic. A simple and well-studied example is the proposal by Paris [Par94] for using a probability distribution over the classical models of a language as the starting point for probabilistic reasoning with logical formulae. By using the proposal by Paris for the connection between "probability theory and "formal logic", we then use classical logic-based argumentation for connecting "argumentation theory" and "formal logic". We summarize this situation in Figure 2.

We will proceed as follows: (Section 2) provides a review of established notions of abstract argumentation, logical argumentation, and probabilistic belief in propositional formulae; (Section 3) provides a review of probabilistic argument graphs plus gives a new proposal called the epistemic approach; (Section 4) introduces logical arguments into probabilistic argument graphs, with a focus on consistent probability distributions; (Section 5) investigates classes of probability distributions and the resulting argumentation; (Section 6) extends the framework to deal with inconsistent probability distributions; (Section 7) extends the framework to deal with probability distributions coming from multiple sources; (Section 8) provides a discussion of the framework; and (Section 9) provides a discussion of future work.

2 Preliminaries

In this section, we review three established areas of the literature that we require for our investigations: (1) Abstract argumentation for representing and analysing arguments in the form of a directed graph; (2) Logical argumentation for using logic to make explicit the premises and claim of each argument; and (3) Probabilistic belief in propositional formulae by using a probability function over a language to represent the belief in each formula of the language.

2.1 Abstract argumentation

In this section, we review the proposal for abstract argumentation by Dung [Dun95]. Essentially, a collection of arguments can be formalized as a directed binary graph.

Definition 1. *An* **argument graph** *is a pair* $(\mathcal{A}, \mathcal{R})$ *where* $\mathcal{A}$ *is a set and* $\mathcal{R}$ *is a binary relation over* $\mathcal{A}$ *(i.e.,* $\mathcal{R} \subseteq \mathcal{A} \times \mathcal{A}$*).*

Each element $A \in \mathcal{A}$ is called an **argument** and $(A, B) \in \mathcal{R}$ means that A **attacks** B (accordingly, A is said to be an **attacker** of B). So A is a **counterargument** for B when $(A, B) \in \mathcal{R}$ holds.

Arguments can work together as a coalition by attacking other arguments and by defending their members from attack as follows.

Definition 2. *Let* $\Gamma \subseteq \mathcal{A}$ *be a set of arguments.* Γ **attacks** $B \in \mathcal{A}$ *iff there is an argument* $A \in \Gamma$ *such that* A *attacks* B*.* Γ **defends** $A \in \Gamma$ *iff for each argument* $B \in \mathcal{A}$*, if* B *attacks* A *then* Γ *attacks* B*.*

The following gives a requirement that should hold for a coalition of arguments to make sense. If it holds, it means that the arguments in the set offer a consistent view on the topic of the argument graph.

Definition 3. *A set* $\Gamma \subseteq \mathcal{A}$ *of arguments is* **conflictfree** *iff there are no* A*,* B *in* Γ *such that* A *attacks* B*.*

Now, we consider how we can find an acceptable set of arguments from an abstract argument graph. The simplest case of arguments that can be accepted is as follows.

Definition 4. *A set* $\Gamma \subseteq \mathcal{A}$ *of arguments is* **admissible** *iff* Γ *is conflictfree and defends all its elements.*

The intuition here is that for a set of arguments to be accepted, we require that, if any one of them is challenged by a counterargument, then they offer grounds to challenge, in turn, the counterargument. There always exists at least one admissible set: The empty set is always admissible. Clearly, the notion of admissible sets of arguments is the minimum requirement for a set of arguments to be accepted. In this paper, we will focus on the following classes of acceptable arguments proposed by Dung [Dun95].

Definition 5. *Let* Γ *be a conflictfree set of arguments, and let* $\mathsf{Defended} : \wp(\mathcal{A}) \mapsto \wp(\mathcal{A})$ *be a function such that* $\mathsf{Defended}(\Gamma) = \{A \mid \Gamma \text{ defends } A\}$.

1. Γ *is a* **complete extension** *iff* $\Gamma = \mathsf{Defended}(\Gamma)$
2. Γ *is a* **grounded extension** *iff it is the minimal (w.r.t. set inclusion) complete extension.*
3. Γ *is a* **preferred extension** *iff it is a maximal (w.r.t. set inclusion) complete extension.*
4. Γ *is a* **stable extension** *iff it is a preferred extension that attacks all arguments in* $\mathcal{A} \setminus \Gamma$.

In general, the grounded extension provides a skeptical view on which arguments can be accepted, whereas each preferred extension take a credulous view on which arguments can be accepted.

Example 6. *Consider the argument graph below. The conflictfree sets are* $\{\}$, $\{A\}$, $\{B\}$, $\{C\}$, *and* $\{A, C\}$*; The admissible sets are* $\{\}$, $\{A\}$, $\{C\}$, *and* $\{A, C\}$*; And the only complete set is* $\{A, C\}$*, and so this set is grounded and preferred. Also this set is stable.*

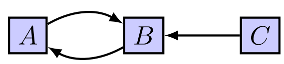

Whilst the focus of this paper is on Dung's definitions for extensions, it would appear that the ideas would generalize to other definitions for extensions such as semi-stable semantics [Cam06] and ideal semantics [DMT07].

2.2 Logical argumentation

In general, we use $\mathcal{L}$ to denote the set of propositional formulae of a language that can formed from the logical connectives of $\vee$, $\wedge$, $\neg$ and $\rightarrow$. The **classical consequence relation**, denoted $\vdash$, is the usual classical consequence relation: For $\Delta \subseteq \mathcal{L}$, if α is a classical inference from Δ, then $\Delta \vdash \alpha$.

The classical consequence relation has been investigated in a number of proposals for logical argument systems including [Cay95, AC98, BH01]. In these proposals, an argument is defined as follows: For $\Phi \subseteq \mathcal{L}$, and a formula $\alpha \in \mathcal{L}$, $\langle \Phi, \alpha \rangle$ is an **argument** iff $\Phi \vdash \alpha$, $\Phi \not\vdash \bot$ and there is no proper subset Φ' of Φ such that $\Phi' \vdash \alpha$. Let $\mathcal{A}^{\mathcal{L}}$ denote the set of arguments that can be formed from language $\mathcal{L}$. For an argument $A = \langle \Phi, \alpha \rangle$, we call Φ the **support**, and α the **claim**, of the argument. Also let $\mathsf{Support}(A) = \Phi$ and $\mathsf{Claim}(A) = \alpha$. For defining the attacks relation, we have a number of options including those given in the next definition.

Definition 6. *Let* A *and* B *be two arguments. We define the following types of* **attack**.

- A *is a* **defeater** *of* B *if* $\mathsf{Claim}(A) \vdash \neg \bigwedge \mathsf{Support}(B)$.
- A *is a* **direct defeater** *of* B *if there is* $\phi \in \mathsf{Support}(B)$ *such that* $\mathsf{Claim}(A) \vdash \neg\phi$.
- A *is a* **undercut** *of* B *if there is* $\Psi \subseteq \mathsf{Support}(B)$ *such that* $\mathsf{Claim}(A) \equiv \neg \bigwedge \Psi$.
- A *is a* **direct undercut** *of* B *if there is* $\phi \in \mathsf{Support}(B)$ *such that* $\mathsf{Claim}(A) \equiv \neg\phi$.
- A *is a* **canonical undercut** *of* B *if* $\mathsf{Claim}(A) \equiv \neg \bigwedge \mathsf{Support}(B)$.
- A *is a* **rebuttal** *of* B *if* $\mathsf{Claim}(A) \equiv \neg\mathsf{Claim}(B)$.
- A *is a* **defeating rebuttal** *of* B *if* $\mathsf{Claim}(A) \vdash \neg\mathsf{Claim}(B)$.

The concepts behind these notions of counterargument have been very widely employed in the literature, so citing originating papers with exactness is difficult. Rebuttals appear in [Pol87] and also [Pol92]. Direct undercuts appear in [EGKF93, EGH95, Cay95]. Undercuts and canonical undercuts were proposed in the above form and studied extensively in [BH01]. Note that canonical undercuts were originally defined using the notion of *maximal conservativeness* but for simplicity we use the above equivalent definition.

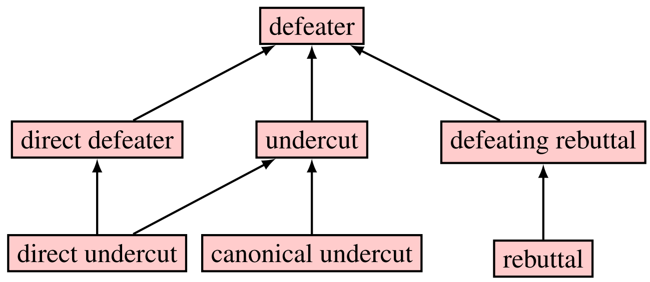

Figure 3: We can represent the containment between the attack relations as above where an arrow from R_1 to R_2 indicates that $R_1 \subseteq R_2$. For proofs, see [BH01, GH11].

Example 7. *Let* $\Delta = \{a \vee b, a \leftrightarrow b, \neg a, c \rightarrow a, \neg a \wedge \neg b, a, b, c, a \rightarrow b, \neg a, \neg b, \neg c\}$

$\langle\{a \vee b, a \leftrightarrow b\}, a \wedge b\rangle$ *is a defeater of* $\langle\{\neg a, c \rightarrow a\}, \neg c\rangle$
$\langle\{a \vee b, a \leftrightarrow b\}, a\rangle$ *is a direct defeater of* $\langle\{\neg a, c \rightarrow a\}, \neg c\rangle$
$\langle\{\neg a \wedge \neg b\}, \neg(a \wedge b)\rangle$ *is a undercut of* $\langle\{a, b, c\}, a \wedge b \wedge c\rangle$
$\langle\{\neg a \wedge \neg b\}, \neg a\rangle$ *is a direct undercut of* $\langle\{a, b, c\}, a \wedge b \wedge c\rangle$
$\langle\{\neg a \wedge \neg b\}, \neg(a \wedge b \wedge c)\rangle$ *is a canonical undercut of* $\langle\{a, b, c\}, a \wedge b \wedge c\rangle$
$\langle\{a, a \rightarrow b\}, b \vee c\rangle$ *is a rebuttal of* $\langle\{\neg b, \neg c\}, \neg(b \vee c)\rangle$
$\langle\{a, a \rightarrow b\}, b\rangle$ *is a defeating rebuttal of* $\langle\{\neg b, \neg c\}, \neg(b \vee c)\rangle$

We often restrict ourselves to arguments over a given knowledegebase $K \subseteq \mathcal{L}$. So for an argument $\langle\Phi, \alpha\rangle$, $\Phi \subseteq K$ has to hold for the support but the claim $\alpha \in \mathcal{L}$ is not necessarily contained in K. Let $\mathcal{A}^K$ be the set of arguments that are generated from K.

2.3 Probabilistic belief in propositional formulae

We use an established proposal for capturing probabilistic belief in propositional formulae [Par94], though we will change some of the nomenclature for our purposes, and we will introduce some further definitions to enable us to develop our proposal.

We assume that the propositional language $\mathcal{L}$ is finite. Given a language $\mathcal{L}$, the set of models (i.e. interpretations) of the language is denoted $\mathcal{M}^{\mathcal{L}}$. Each **model** in $\mathcal{L}$ is an assignment of true or false to the formulae of the language defined in the usual way for classical logic. For $\phi \in \mathcal{L}$, $\mathsf{Models}(\phi)$ denotes the set of models of ϕ (i.e. $\mathsf{Models}(\phi) = \{m \in \mathcal{M}^{\mathcal{L}} \mid m \models \phi\}$), and for $\Delta \subseteq \mathcal{L}$, $\mathsf{Models}(\Delta)$ denotes the set of models of Δ (i.e. if $\Delta = \{\phi_1, ..., \phi_n\}$, then $\mathsf{Models}(\Delta) = \mathsf{Models}(\phi_1) \cap .. \cap \mathsf{Models}(\phi_n)$).

As a simple way to represent models $\mathcal{M}^{\mathcal{L}}$ of the language $\mathcal{L}$, we first declare a **signature**, denoted $\mathcal{S}^{\mathcal{L}}$, which is the atoms of the language $\mathcal{L}$ given in a sequence $(a_1, ..., a_n)$, and then each model is given as a binary number $b_1, ..., b_n$ where for each digit b_i, if b_i is 1, then a_i is true in the model, otherwise a_i is false in the model.

A model can also be represented by a conjunction of literals. Let $\mathcal{S}^{\mathcal{L}} = (s_1, .., s_n)$ be the signature, and let $m \in \mathcal{M}^{\mathcal{L}}$ be represented by the binary number $k_1, ..., k_n$. Then the conjunction of literals $l_1 \wedge ... \wedge l_n$ represents m when for each literal l_i, if k_i is 1, then l_i is s_1, otherwise l_i is $\neg s_1$.

Example 8. *Let the atoms of* $\mathcal{L}$ *be* $\{a, b, c\}$*, and so* $\mathcal{L}$ *contains the usual propositional formulae that can be formed from these three atoms. Let the signature* $\mathcal{S}^{\mathcal{L}}$ *be* (a, b, c)*, and so the models* $\mathcal{M}^{\mathcal{L}}$ *are* $\{111, 110, 101, 100, 011, 010, 001, 000\}$*. Consider* $m = 101$ *which means that* a *is true,* b *is false, and* c *is true. This can equivalently be represented by the conjunction of literals* $a \wedge \neg b \wedge c$*.*

In order to represent uncertainty we use a belief function defined next.

Definition 7. *Let* $\mathcal{L}$ *be a propositional language, and let* $\Phi \subseteq \mathcal{L}$*. A* **belief function** *on* Φ *is a function* $P : \mathcal{L} \rightarrow [0, 1]$*. If* P *is a belief function on* $\mathcal{L}$ *(i.e.* Φ *is* $\mathcal{L}$*), then* P *is* **complete***, otherwise* P *is* **incomplete***.*

A belief function is a measure of how much belief holds for each formula in a set of formulae. For a formula ϕ, the higher $P(\phi)$ is, the higher the belief in ϕ. To provide a useful constraint on belief functions, we consider the special case of a finitely additive function, which we will call a probability function, as defined next.

Definition 8. *Let $\mathcal{L}$ be a propositional language, and let $\mathcal{M}^{\mathcal{L}}$ be the models of the language $\mathcal{L}$. A belief function P on $\mathcal{L}$ is a* **probability function** *on $\mathcal{L}$ iff for each $\phi \in \mathcal{L}$*

$$P(\phi) = \sum_{m \in \mathsf{Models}(\phi)} P(m)$$

The above definition assumes that each model is represented by a conjunction of literals (i.e. a formula in $\mathcal{L}$). Therefore, a probability distribution on $\mathcal{M}^{\mathcal{L}}$ is a probability distribution on $\mathcal{L}$ and vice versa.

Example 9. *Let the atoms of $\mathcal{L}$ be $\{a, b\}$, and so $\mathcal{L}$ contains the usual propositional formulae that can be formed from these two atoms. Let the signature $\mathcal{S}^{\mathcal{L}}$ be (a, b), and so the models $\mathcal{M}^{\mathcal{L}}$ are $\{11, 10, 01, 00\}$. Now suppose $P(11) = 0.8$ and $P(10) = 0.2$. Hence, $P(a) = 1$, $P(a \wedge b) = 0.8$, $P(b \vee \neg b) = 1$, $P(a \vee \neg b) = 0.2$, etc.*

Since we are concerned with conflict in argumentation, we want to consider whether or not a probability distribution on $\mathcal{L}$ makes sense. For this, we use the following definition of a probability distribution being consistent.

Definition 9. *Let $\mathcal{L}$ be a propositional language and let $\mathcal{M}^{\mathcal{L}}$ be the models of the language. A probability function P on $\mathcal{L}$ is* **consistent** *iff*

$$\sum_{m \in \mathcal{M}^{\mathcal{L}}} P(m) = 1$$

What we call a "consistent probability function" is what is normally called a "probability function". The reason we have used the term "consistent" is that we also want to consider an "inconsistent probability function" (i.e. a finitely additive function that does not sum to 1).

Example 10. *Let the atoms of $\mathcal{L}$ be $\{a, b\}$, and so $\mathcal{L}$ contains the usual propositional formulae that can be formed from these two atoms. Let the signature $\mathcal{S}^{\mathcal{L}}$ be (a, b), and so the models $\mathcal{M}^{\mathcal{L}}$ are $\{11, 10, 01, 00\}$. Now consider the belief function $P(a) = 0.8$, $P(\neg a \vee \neg b) = 0.8$, and $P(b) = 0.8$. This belief function provides constraints on the models.*

- *From $P(a) = 0.8$, we have $P(10) + P(11) = 0.8$.*
- *From $P(\neg a \vee \neg b) = 0.8$, we have $P(10) + P(00) + P(01) = 0.8$.*
- *From $P(b) = 0.8$, we have $P(11) + P(01) = 0.8$.*

There is no probability function P on $\mathcal{L}$ such that the above constraints hold and P is consistent. An example of an inconsistent probability distribution that satisfies the constraints is $P(11) = 0.4$, $P(10) = 0.4$, and $P(01) = 0.4$.

Not every belief function on $\mathcal{L}$ is a probability function on $\mathcal{L}$, since the constraints on the models imposed by P for it to be a probability function may not be satisfiable, as illustrated by the next example.

Example 11. *Let the atoms of $\mathcal{L}$ be $\{a, b\}$, and so $\mathcal{L}$ contains the usual propositional formulae that can be formed from these two atoms. Let the signature $\mathcal{S}^{\mathcal{L}}$ be (a, b), and so the models $\mathcal{M}^{\mathcal{L}}$ are $\{11, 10, 01, 00\}$. Now consider the belief function $P(a) = 0.5$, $P(a \wedge b) = 1$, and $P(b) = 0.8$. This provides the following constraints on the models.*

- *From $P(a) = 0.5$, we have $P(10) + P(11) = 0.5$.*
- *From $P(a \wedge b) = 1$, we have $P(11) = 1$.*
- *From $P(b) = 0.8$, we have $P(11) + P(01) = 0.8$.*

There is no probability function on $\mathcal{M}^{\mathcal{L}}$, even an inconsistent probability function, such that the above constraints hold, and therefore, the belief function on $\mathcal{L}$ is not a probability function on $\mathcal{L}$.

When a probability distribution is consistent then the following properties hold. These properties are sometimes used to give a definition for a consistent probability distribution (see [Par94] for more details on this definition).

Proposition 1. *A probability function P on $\mathcal{L}$ is* **consistent** *iff*

- *if* $\models \alpha$, *then* $P(\alpha) = 1$
- *if* $\models \neg(\alpha \wedge \beta)$, *then* $P(\alpha \vee \beta) = P(\alpha) + P(\beta)$

A consistent probability distribution on $\mathcal{L}$ has an intuitive meaning and numerous desirable properties. However, as we discussed in Section 1, we are concerned with argumentation with conflicting information, we also need to consider inconsistent probability distributions. In the following sections, we will investigate both consistent and inconsistent probability distributions on $\mathcal{L}$.

3 Probabilistic argument graphs

We can qualify each argument in an argument graph by a probability value that indicates the belief that the argument is true. So each argument A is assigned a value $P(A)$ in the unit interval. This then gives us the following notion of a probabilistic argument graph as proposed in [LON11].

Definition 10. *A* **probabilistic argument graph** *is a tuple $(\mathcal{A}, \mathcal{R}, P)$ where $(\mathcal{A}, \mathcal{R})$ is an argument graph and $P : \mathcal{A} \mapsto [0, 1]$.*

When capturing informal arguments, there is often the question of how to handle the explicit or implicit uncertainty that has been expressed. Having the extra expressibility of probabilistic argument graphs is valuable in this process.

Example 12. *Continuing Example 1, if there is some uncertainty about whether a given patient has emphysema, we could let the probabilities be $P(A_1) = 1$, $P(A_2) = 1$, and $P(A_3) = 0.7$. So this probability represents that there is quite strong belief that the patient has emphysema and that this implies a contraindication for betablockers. It also represents that there is a degree of doubt about whether the patient has emphysema and/or whether this is a contraindication for betablockers.*

In general, there are no further constraints on the probability assignment beyond Definition 10. So, for example, it is possible for every argument in a graph to be assigned a probability of 1, in which case we will return to Dung's original proposal. Similarly, it is possible for any or every argument in a graph to be assigned a probability of 0, in which case, we will in effect have a graph with no nodes. In the rest of this paper, we will explore some of the options we have for choosing and using probability assignments.

In order to harness the notion of a probabilistic argument graph, we require some subsidiary definitions and notation that we will use in the rest of the paper. Let $G = (\mathcal{A}, \mathcal{R}, P)$ be an probabilistic argument graph, and let $\mathcal{A}' \subseteq \mathcal{A}$. The marginalization of $\mathcal{R}$ to $\mathcal{A}'$, denoted $\mathcal{R} \otimes \mathcal{A}'$, is the subset of $\mathcal{R}$ involving just the arguments in $\mathcal{A}'$ (i.e. $\mathcal{R} \otimes \mathcal{A}' = \{(A, B) \in \mathcal{R} \mid A, B \in \mathcal{A}'\}$). If $G = (\mathcal{A}, \mathcal{R}, P)$ and $G' = (\mathcal{A}', \mathcal{R}', P)$ are probabilistic argument graphs, then G' is a **full subgraph** of G, denoted $G' \sqsubseteq G$, iff $\mathcal{A}' \subseteq \mathcal{A}$ and $\mathcal{R}'$ is $\mathcal{R} \otimes \mathcal{A}'$. This is illustrated in Figure 5.

We will proceed by considering two ways of interpreting the probability over arguments. The first interpretation is the **epistemic approach** in which we are interested in "rational" probability distributions. These are distributions where if an attacker is assigned a high degree of belief, then the attacked argument is assigned a low degree of belief, and vice versa. The second interpretation is the **constellations approach** in which we use the probability distribution over arguments to generate a probability distribution over the full subgraphs. Hence, we can regard each $G' \sqsubseteq G$ as an "interpretation" of G. Such a distribution can then be used to generate a probability distribution over extensions.

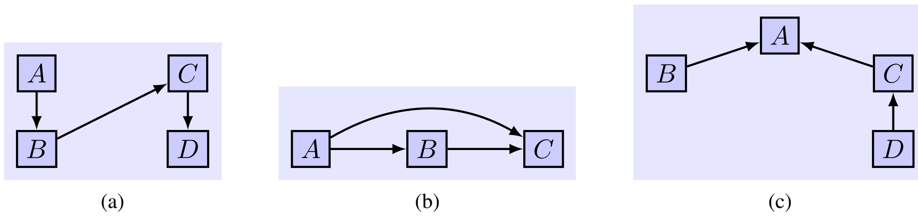

Figure 4: Further examples of argument graphs discussed in Section 3.

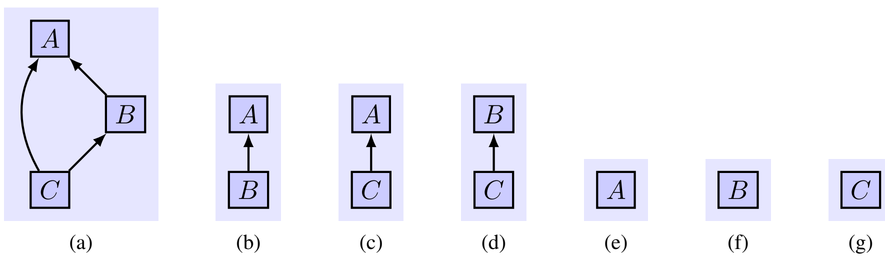

Figure 5: For Graph 5a, the seven non-empty full subgraphs are given (Graphs 5a to 5g).

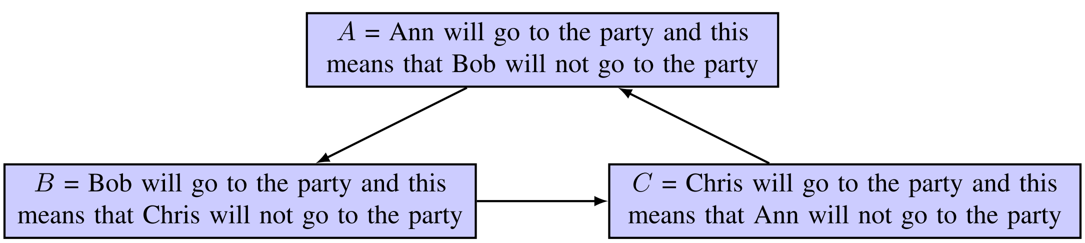

Figure 6: Example of three arguments in a simple cycle.

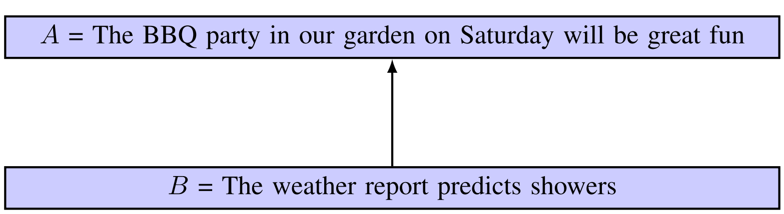

Figure 7: Example of two arguments where one attacks the other. In this example, the degree to which the undercut is believed should affect whether the root argument is believed or not.

3.1 The epistemic approach

In the epistemic approach, which is a novel proposal, the probability distribution over arguments is used directly to identify which arguments are believed. So the higher the probability of the argument, then the more it is believed. This is useful because together with the structure of the argument graph, i.e. information about which arguments are being considered, we can determine what are "rational" beliefs about the arguments. To motivate this approach, we consider a couple of examples.

First consider the graph given in Figure 6. Here, we may believe that say A is true and that B and C are false. In which case, with this extra epistemic information about the arguments, we can resolve the conflict and so take the set $\{A\}$ as the "epistemic" extension. In contrast, there is only one admissible set which is the empty set. So by having this extra epistemic information, we get a more informed extension (in the sense that it has harnessed the extra information in a sensible way).

Now consider another example, given in Figure 7, containing two arguments A and B, where A attacks B. As with many real world arguments, these arguments are enthymemes which are arguments for which the premises and/or claim are implicit. For a discussion of formal modelling of enthymemes, see [Hun07]. Even though we assume that technically B attacks A, we may have very low belief in B. Hence, it would seem reasonable to have very high belief in A. However, for this graph, there is only one non-empty admissible set which is $\{B\}$. Yet, with our epistemic information, it would seem better to obtain $\{A\}$ as the extension. In this way, the probabilistic information over-rides the dialectical information given by the attacks relation.

To address these kinds of example, we can use the probability distribution P to interpret the belief in each argument A. This correspondence between quantitative evaluations and truth valuations arises in various approaches to reasoning with uncertainty including qualitative probabilistic reasoning [Par96] and many-valued logics.

- $P(A) = 0$ represents that A is believed to be false with certainty.
- $P(A) < 0.5$ represents that A is believed to be false to some degree.
- $P(A) = 0.5$ represents that A is neither believed to be true nor believed to be false.
- $P(A) > 0.5$ represents that A is believed to be true to some degree.
- $P(A) = 1$ represents that A is believed to be true with certainty.

So for example, in the example about Ann, Bob and Chris in Figure 6, we may have $P(A) = 0.9$, $P(B) = 0.2$, $P(C) = 0.1$. And for the example about the BBQ in Figure 7, we may have $P(A) = 1$ and $P(B) = 0$. In order, to capture these intuitions, we introduce a new definition of an extension that harnesses this extra epistemic information.

Definition 11. *For a probabilistic argument graph, $(\mathcal{A}, \mathcal{R}, P)$, the* **epistemic extension** *is* $\{A \in \mathcal{A} \mid P(A) > 0.5\}$.

So an epistemic extension is just the set of arguments that are believed to be true to some degree (i.e the arguments with probability greater than 0.5).

Example 13. *For the argument graph in Figure 6, together with probability function P such that $P(A) = 0.1$, $P(B) = 0.8$, and $P(C) = 0.1$, the epistemic extension is $\{B\}$.*

The definition of epistemic extension is very general, and it permits any set of arguments to be an epistemic extension. For instance, given arguments A and B where A attacks B, and a probability function P where $P(A) = 1$ and $P(B) = 1$, then the epistemic extension is $\{A, B\}$, which is not conflictfree. To avoid this kind of situation, we can impose restrictions on the probability function such as given by a rational probability distribution over a set of arguments (defined below). This is a probability distribution that is, in a sense, consistent with the structure of the argument graph. It is consistent in the sense that if the belief in an argument is high, then the belief in the arguments it attacks is low. We call this kind of distribution rational because it would be rational to have such a distribution over a set of arguments. In a sense, if the distribution is not rational, then there is too much belief in an argument and its attacker.

Definition 12. *A probability function P is* **rational** *for an argument graph $(\mathcal{A}, \mathcal{R})$ iff for each $(A, B) \in \mathcal{R}$, if $P(A) > 0.5$, then $P(B) \leq 0.5$. A* **rational argument graph** *is a probabilistic argument graph, $(\mathcal{A}, \mathcal{R}, P)$, where P is rational for $(\mathcal{A}, \mathcal{R})$.*

Perhaps the minimum requirement for a set of arguments to be an extension is that it is conflictfree. As the next result shows, rational extensions coincide with conflictfree sets.

Proposition 2. *Let $(\mathcal{A}, \mathcal{R})$ be an argument graph with $\Gamma \subseteq \mathcal{A}$. Γ is a rational extension iff Γ is conflictfree.*

Proof. $(=>)$ Let Γ be a rational extension. Therefore, there is a rational probability function P for $(\mathcal{A}, \mathcal{R})$, and Γ is the epistemic extension for $(\mathcal{A}, \mathcal{R}, P)$. Therefore, for each $(A, B) \in \mathcal{R}$, if $p(A) > 0.5$, then $p(B) \leq 0.5$. Since Γ is the epistemic extension, if $(A, B) \in \mathcal{R}$, and $p(A) > 0.5$, then $A \in \Gamma$, and $B \notin \Gamma$. Therefore, for each $(A, B) \in \mathcal{R}$, either ($A \in \Gamma$ and $B \notin \Gamma$) or $A \notin \Gamma$. Therefore, Γ is conflictfree.

$(<=)$ Let Γ be conflictfree. Therefore, there is no $A, B \in \Gamma$ such that A attacks B. Therefore, for each $(A, B) \in \mathcal{R}$, either ($A \in \Gamma$ and $B \notin \Gamma$) or $A \notin \Gamma$. Therefore, there is a probability function such that for each $(A, B) \in \mathcal{R}$, if $p(A) > 0.5$, then $p(B) \leq 0.5$. Therefore, there is a rational probability function P for $(\mathcal{A}, \mathcal{R})$, and Γ is the epistemic extension for $(\mathcal{A}, \mathcal{R}, P)$. Therefore, Γ is a rational extension. □

The empty set is a rational extension for any $(\mathcal{A}, \mathcal{R})$ since we can set $P(A) = 0.5$ for every $A \in \mathcal{A}$, and hence get the empty set as the epistemic extension of $(\mathcal{A}, \mathcal{R}, P)$. Furthermore, any admissible set Γ can be obtained by as a rational extension by choosing a probability function P so if $A \in \Gamma$, then $P(A) = 1$, otherwise, $P(A) = 0$.

Before we conclude this section, we will consider a particular kind of rational probability function, called an involutory probability function, that seems to incorporate a natural and useful constraint on the probabilities. The idea for this comes from the proposal for equational semantics in argumentation networks [Gab11]. However, we will see that our definition next is actually of limited use.

Definition 13. *For a probabilistic argument graph $(\mathcal{A}, \mathcal{R}, P)$, the probability function P is* **involutory** *iff $p(A) = 1 - p(B)$ for each $(A, B) \in \mathcal{R}$.*

Example 14. *For an argument graph in Figure 4a, if the assignment for the probability function P is $P(A) = 0.8$, $P(B) = 0.2$, $P(C) = 0.8$, and $P(D) = 0.2$, then P is involutory.*

When the argument graph has odd cycles, there is no probability function that is involutory, apart from a neutral probability function (i.e. where every argument is assigned 0.5). For instance, there is no involutory probability function for the argument graph given in Figure 6 that is not a neutral probability function. Even when the graph is acyclic, it may be the case that there is no involutory probability function (apart from the neutral probability function). Consider for example the argument graph in Figure 4b for which there is no involutory probability function (apart from the neutral probability function). If we restrict consideration to trees, then we are guaranteed to have a probability function that is involutory and not neutral. But even here there are constraints such as siblings have to have the same assignment. For instance, in Figure 4c, it is necessary that $P(A) = P(D)$ and $P(B) = P(C)$ hold for P to be involutory. As a consequence of these limitations, we do not consider involutory functions further.

3.2 The constellations approach

In the constellations approach, proposed in [LON11], we interpret the uncertainty associated with each argument in the probabilistic argument graph as an uncertainty over the argument graph. In other words, we use the probability distribution over arguments to generate a probability distribution over full subgraphs of the original argument graph. Using the full subgraphs, we can then explore the notions of probability distributions over admissible sets, extensions, and inferences.

According to [LON11], for an argument A in a graph G, with a probability assignment P, $P(A)$ is the probability that A is exists in an arbitrary full subgraph of G, and $1 - P(A)$ is the probability that A does not exist in an arbitrary full subgraph of G. We develop this idea further in the justification perspective [Hun12], where $P(A)$ means the probability that A is justified in appearing in the graph.

As the next result shows, we can use the probability assigned to each argument to generate a probability distribution over the subgraphs. So each subgraph can be viewed as an "interpretation" of what

	Full subgraphs	Probability of subgraph	Grounded extension	Preferred extensions
G^1	$A \leftrightarrow B \leftarrow C$	$P(A).P(B).P(C)$	$\{A, C\}$	$\{A, C\}$
G^2	$A \leftrightarrow B$	$P(A).P(B).(1 - P(C))$	$\{\}$	$\{A\}, \{B\}$
G^3	$A \ \ C$	$P(A).(1 - P(B)).P(C)$	$\{A, C\}$	$\{A, C\}$
G^4	$B \leftarrow C$	$(1 - P(A)).P(B).P(C)$	$\{C\}$	$\{C\}$
G^5	A	$P(A).(1 - P(B)).(1 - P(C))$	$\{A\}$	$\{A\}$
G^6	B	$(1 - P(A)).P(B).(1 - P(C))$	$\{B\}$	$\{B\}$
G^7	C	$(1 - P(A)).(1 - P(B)).P(C)$	$\{C\}$	$\{C\}$
G^8		$(1 - P(A)).(1 - P(B)).(1 - P(C))$	$\{\}$	$\{\}$

Table 1: The full subgraphs with their probability and extensions for Example 15.

the argument graph should be. If all the arguments have probability 1, then the argument graph itself has probability 1, and that is the only interpretation we should consider (using the constellations approach). But, if one or more arguments has a probability less than 1, then there will be multiple "interpretations. So for instance, if there is exactly one argument A in the graph G with probability less than one, then there are two interpretations, the first with A, and the second without A. So using the justification perspective, with the constellations approach, we can treat the set of subgraphs of a G as a sample space, where one of the subgraphs is the true argument graph.

Definition 14. *Let $G = (\mathcal{A}, \mathcal{R}, P)$ and $G' = (\mathcal{A}', \mathcal{R}', P)$ be probabilistic argument graphs such that $G' \sqsubseteq G$. The* **probability of subgraph** *G', denoted $p(G')$, is*

$$\left(\prod_{A \in \mathcal{A}'} P(A)\right) \times \left(\prod_{A \in \mathcal{A} \setminus \mathcal{A}'} (1 - P(A))\right)$$

So the probability of a full subgraph captures the degree of certainty that the full subgraph contains exactly the arguments that are regarded as holding.

Proposition 3. *If $G = (\mathcal{A}, \mathcal{R}, p)$ is a probabilistic graph, then $\sum_{G' \sqsubseteq G} p(G') = 1$*

Proof. We show this by induction on the size of $\mathcal{A}$. For the base case, where the graph G_1 contains one argument, there are two full subgraphs, the graph G_1 and the empty graph $G_\emptyset$. Let α be the argument in G_1. So $p(G_1) = p(\alpha)$, and $p(G_\emptyset) = 1 - p(\alpha)$. Hence, $p(G_1) + p(G_\emptyset) = 1$. So in the base case, $\sum_{G' \sqsubseteq G} p(G') = 1$. For the inductive step, assume $\sum_{G'_n \sqsubseteq G_n} p(G'_n) = 1$ holds for any graph G_n containing n arguments. We now show that $\sum_{G'_{n+1} \sqsubseteq G_{n+1}} p(G'_{n+1}) = 1$ holds for any graph G_{n+1} containing $n+1$ arguments. Suppose α is an argument that is not in G_n. So we can extend G_n by adding α to give G_{n+1}. For each subgraph $G'_n \sqsubseteq G_n$, there is the subgraph G'_{n+1} of G_{n+1} where G'_n and G'_{n+1} have the same arguments, in which case $p(G'_{n+1})$ is $p(G'_n) \times (1 - p(\alpha))$, and there is the subgraph G'_{n+1} of G_{n+1} where G'_{n+1} is formed from G'_n by adding the argument α, in which case $p(G'_{n+1})$ is $p(G'_n) \times p(\alpha)$. Therefore, $\sum_{G'_{n+1} \sqsubseteq G_{n+1}} p(G'_{n+1}) = \sum_{G'_n \sqsubseteq G_n} ([p(G'_n) \times p(\alpha)] + [p(G'_n) \times (1 - p(\alpha))])$ $= \sum_{G'_n \sqsubseteq G_n} (p(G'_n) \times (p(\alpha) + (1 - p(\alpha)))) = \sum_{G'_n \sqsubseteq G_n} p(G'_n) = 1$. Therefore, in the inductive case, $\sum_{G' \sqsubseteq G} p(G') = 1$. □

Example 15. *Continuing Example 1, we have each of the eight full subgraphs of the graph in Table 1. For each constellation, we see the probability of the full subgraph (in terms of the probability of the arguments) and we see the grounded extension and preferred extensions that are obtained for each full subgraph.*

If we assume the probabilities $P(A) = 1$, $P(B) = 1$, and $P(C) = 0.8$, then we get the following probability distribution over the full subgraphs.

	G^1	G^2	G^3	G^4	G^5	G^6	G^7	G^8
P	*0.8*	*0.2*	*0*	*0*	*0*	*0*	*0*	*0*

Full subgraph	Admissible sets
G^1	$\{A, B, D\}, \{A, B\}, \{A, D\}, \{B, D\}, \{A\}, \{B\}, \{\}$
G^2	$\{A, B\}, \{A\}, \{B\}, \{\}$
G^3	$\{A, B, D\}, \{A, B\}, \{A, D\}, \{B, D\}, \{A\}, \{B\}, \{D\}, \{\}$
G^4	$\{A, B\}, \{A\}, \{B\}, \{\}$

Table 2: The admissible sets for the full subgraphs considered in Example 16.

Whereas if we use the following probabilities $P(A) = 0.8$, $P(B) = 0.8$, and $P(C) = 0.8$, then we get the following probability distribution over the full subgraphs.

	G^1	G^2	G^3	G^4	G^5	G^6	G^7	G^8
P	*0.512*	*0.128*	*0.128*	*0.128*	*0.032*	*0.032*	*0.032*	*0.008*

If all the arguments in a probabilistic argument graph G have probability of 1, then the only full subgraph of G to have non-zero probability is G, and so it has probability 1. Hence, we will see that we recover Dung's original definitions and results by assuming all arguments have probability 1. At the other extreme, if all the arguments in a probabilistic argument graph G have probability of 0, then the empty graph has probability 1.

For a probabilistic argument graph $G = (\mathcal{A}, \mathcal{R}, P)$, and a set of arguments $\Gamma \subseteq \mathcal{A}$, $G \Vdash_X \Gamma$ denotes that Γ is an X extension of $G = (\mathcal{A}, \mathcal{R})$ where $X = \{\mathsf{ad}, \mathsf{co}, \mathsf{pr}, \mathsf{gr}, \mathsf{st}\}$, and ad denotes admissible semantics, co denotes complete semantics, pr denotes preferred semantics, st denotes stable semantics, and gr denotes grounded semantics. When $G \Vdash_X \Gamma$ holds, we say that G **entails** Γ according to X semantics. The set of full subgraphs that entails a set of arguments Γ according to X semantics, denoted $Q_X(\Gamma)$, is $Q_X(\Gamma) = \{G' \sqsubseteq G \mid G' \Vdash_X \Gamma\}$.

Given a probabilistic argument graph $G = (\mathcal{A}, \mathcal{R}, P)$, and a set of arguments $\Gamma \subseteq \mathcal{A}$, we want to calculate the probability that Γ is an X extension, which we denote by $P(\Gamma^X)$, where $X \in \{\mathsf{ad}, \mathsf{co}, \mathsf{pr}, \mathsf{gr}, \mathsf{st}\}$. For this, we take the sum of the probability of the full subgraphs for which Γ is an X extension.

Definition 15. *Let $G = (\mathcal{A}, \mathcal{R}, P)$ be a probabilistic argument graph and let $\Gamma \subseteq \mathcal{A}$. The probability that Γ is an X extension is*

$$\sum_{G' \in Q_X(\Gamma)} P(G')$$

In the following, we consider examples of using this definition, and discuss its utility. First we analyse the argument graph presented in Figure 8, and then we return to the example raised in Figure 6.

Example 16. *Consider the example given in Figure 8. Suppose the couple are confident that arguments A and B are true, but they are really not sure about the truth or falsity of C and D, this uncertainty can be represented by $P(A) = 1$, $P(B) = 1$, $P(C) = 0.5$, and $P(D) = 0.5$. For this, as presented in Figure 9, there are four full subgraphs G^1, G^2, G^3, and G^4 with non-zero probability. Each full subgraph has probability of 1/4. As a result, there are eight admissible sets with non-zero probability to consider: $P(\emptyset^{\mathsf{ad}}) = 1$, $P(\{A\}^{\mathsf{ad}}) = 1$, $P(\{B\}^{\mathsf{ad}}) = 1$, $P(\{D\}^{\mathsf{ad}}) = 1/4$, $P(\{A, B\}^{\mathsf{ad}}) = 1$, $P(\{A, D\}^{\mathsf{ad}}) = 1/2$, $P(\{B, D\}^{\mathsf{ad}}) = 1/2$, and $p(\{A, B, D\}^{\mathsf{ad}}) = 1/2$. Hence, any admissible set containing C has zero probability, and any admissible set containing D together with either A or B has probability of 0.5. So after the couple consider the argument graph, C is rejected, and D is neither rejected nor accepted.*

$$G^1 \Vdash_{\mathsf{co}} \{A, B, D\} \quad G^2 \Vdash_{\mathsf{co}} \{A, B\} \quad G^3 \Vdash_{\mathsf{co}} \{A, B, D\} \quad G^4 \Vdash_{\mathsf{co}} \{A, B\}$$

As a result, there are two extensions with non-zero probability to consider for $X \in \{\mathsf{co}, \mathsf{pr}, \mathsf{gr}, \mathsf{st}\}$. These are $p(\{A, B\}^X) = 1/2$, and $p(\{A, B, D\}^X) = 1/2$.

Example 17. *For Figure 6, imagine a scenario where a friend of Ann is certain that argument A is true, whereas a friend of Bob does not know whether B is true, and a friend of Chris does not know whether C*

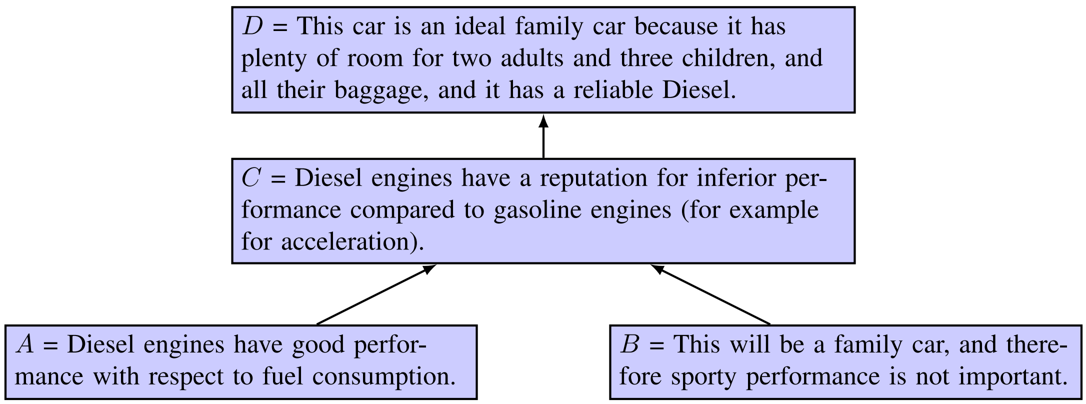

Figure 8: The argument graph for Example 16. The scenario involves a couple deciding on whether to buy a particular car as their family car. A salesperson in the car showroom has given argument D as the key argument in the sales pitch. Then, the salesperson, has presented counterargument C as a rhetorical move. The salesperson is expressing this counterargument since it is quite possible something that the customers might be thinking. By doing this, the salesperson can then provide arguments A and B as counterarguments to C.

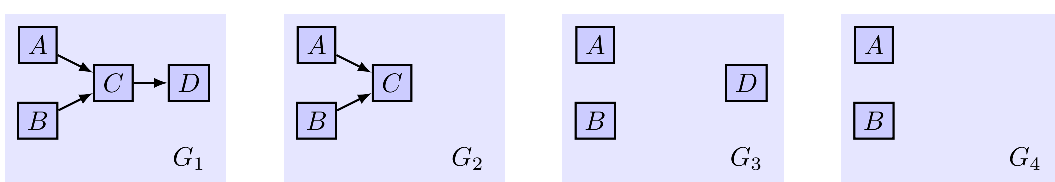

(a) (b) (c) (d)

Figure 9: Full subgraphs for Example 16.

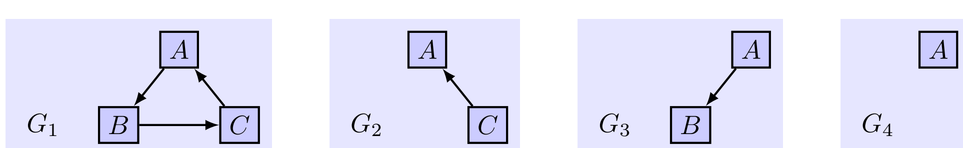

(a) (b) (c) (d)

Figure 10: Full subgraphs for Example 17

is true. Suppose we represent this by $P(A) = 1$, $P(B) = 0.5$, and $P(C) = 0.5$. For this there are four full subgraphs G^1, G^2, G^3, and G^4 with non-zero probability (given in Figure 10). Each full subgraph has probability of 1/4. The complete extensions for these subgraphs are the following.

$$G^1 \Vdash_{\mathsf{co}} \{\} \quad G^2 \Vdash_{\mathsf{co}} \{C\} \quad G^3 \Vdash_{\mathsf{co}} \{A\} \quad G^4 \Vdash_{\mathsf{co}} \{A\}$$

As a result, there are three extensions with non-zero probability to consider for $X \in \{\mathsf{co}, \mathsf{pr}, \mathsf{gr}, \mathsf{st}\}$. These are $p(\emptyset^X) = 1/4$, $P(\{A\}^X) = 1/2$, and $P(\{C\}^X) = 1/4$. Hence, $\{A\}$ is the most likely extension.

In this subsection, we have defined the probability of an extension as the sum of the probabilities of the full subgraphs that entail the extension. Next, we compare this approach with the epistemic approach given in the previous section.

3.3 Comparison of the epistemic and constellations approaches

In this paper, we assume that we start with a set of arguments, as these are important constructs that humans use for dealing with conflicting information (for example in discussions and debates), and then we make a belief assignment to each argument. This gives us a probabilistic argument graph.

Our primary way for dealing with probabilistic argument graphs will be our epistemic approach. The epistemic approach uses the attacks relation in an argument graph in a way that is quite different to the way it is used in Dung's dialectical semantics (i.e. Dung's notions of extensions). Yet if we have the extra information that comes with a probability assignment to arguments, the epistemic approach offers a natural way of viewing the argument graph, as discussed in Section 3.1.

The emphasis in the epistemic approach is to find a probability distribution that is rational with respect to the argument graph. This approach can be used when it can be decided which arguments are believed, and which are disbelieved, in a way that respects the structure of the argument graph (i.e. if the belief in an attacking argument is high, then the belief in the attacked argument is low). In this way, the probability distribution over arguments, and the attack relation between arguments, provide two dimensions of the uncertainty concerning the arguments. As we will see in subsequent sections, when we consider logical arguments, there are useful correspondences between epistemic extensions and certain kinds of probability assignment to the arguments.

Our secondary way for dealing with probabilistic argument graphs is the constellations approach proposed in [DT10, LON11]. The constellations approach is useful when there is a failure to find a probability distribution that respects the structure of the argument graph. For example, consider an argument graph containing arguments A and B where B attacks A. Suppose agent 1 regards A as certain, and agent 2 regards B as certain. So the probability distribution is $P(A) = 1$ and $P(B) = 1$. However, this probability distribution does not respect the structure of the graph (i.e. this probability distribution is not rational according to Definition 12). We will investigate ways that we relax and/or redistribute mass so that we can have a consistent probability distribution over the models, and thereby be able to use the epistemic approach, but it is also useful to consider the constellations approach to deal with this situation.

Note, in the constellations approach, there is an issue with respect to the assumption of independence of arguments when calculating the probability distribution over the full subgraphs (i.e. Definition 14). The proposals for probabilistic argument graphs [LON11] do not address this issue and so we attempt to address this now. We start by considering logical arguments. If we have a knowledgebase containing just two formulae $\{a, \neg a\}$, we can construct arguments $A_1 = \langle\{\alpha\}, \alpha\rangle$, and $A_2 = \langle\{\neg\alpha\}, \neg\alpha\rangle$. The rebuttal relation holds so that the arguments attack each other. In terms of classical logic, it is not possible for both arguments to be true, but each of them is a justified point (i.e. each is a self-contained, and internally valid, contribution given the knowledgebase). So even though logically A_1 and A_2 are not independent (in the sense that if one is known to be true, then the other is known to be false), they are independent as justified points (i.e. knowing that one is a justified point does not affect whether or not the other is a justified point). This means we can construct an argument graph with both arguments appearing.

Continuing the example, we can consider a probability assignment to each argument as introduced in Definition 10. This may perhaps be based on the confidence assigned by the source of each formula. We can treat the probability value assigned as reflecting the confidence that the argument makes a justified point. If each of A_1 and A_2 is assigned 1, then we return to the situation in the previous paragraph. But if the assignment is less than 1 for either argument, then there is some explicit doubt that the argument is a justified point, and therefore there is some doubt that it should appear in the argument graph.

So for the constellations approach, we use the **justification perspective** on the probability of an argument [Hun12]: For an argument A in a graph G, with a probability assignment P, $P(A)$ is treated as the probability that A is a justified point and therefore should appear in the graph, and $1 - P(A)$ is the probability that A is not a justified point and so should not appear in the graph. This means, in the justification perspective, we will use the probability assignment to each argument as a proxy for the probability that the argument is a justified point. The probabilities of the arguments being justified are independent (i.e. knowing that one argument is a justified point does not affect the probability that another is a justified point). Therefore, for the constellations approach, we use the probability assignment to give a finer grained notion of justified point than we get with just logical arguments (i.e. logical arguments without a probability assignment).

4 Probability distributions over logical arguments

We now consider how we can harness probabilistic logic for logical arguments. Essentially, we use a probability function over the language (as reviewed in the Preliminaries section) to represent the uncertainty in the premises of each argument. We then use this as a belief in the argument of the argument as a whole. In this way, we can instantiate probabilistic argument graphs with logical arguments together with the attacks relation being one of those given in Definition 6.

Definition 16. *If P be a probability function on $\mathcal{L}$, then P is a probability function on $\mathcal{A}^{\mathcal{L}}$, where the* **probability of an argument** $\langle \Phi, \alpha \rangle \in \mathcal{A}^{\mathcal{L}}$*, denoted $P(\langle \Phi, \alpha \rangle)$, is $P(\wedge \Phi)$.*

So if P is a probability function on $\mathcal{L}$, then for a set of arguments $\mathcal{A} \subseteq \mathcal{A}^{\mathcal{L}}$, we can form a probabilistic argument graph $(\mathcal{A}, \mathcal{R}, P)$ where $\mathcal{A}$ is a set of logical arguments, $\mathcal{R}$ is attack relation such as undercut, direct undercut, canonical undercut, etc., and P is a probability function on $\mathcal{A}$.

Example 18. *Let P be a probability function on $\mathcal{L}$. Let the signature be (p, q). So the models of $\mathcal{M}^{\mathcal{L}}$ are $\{11, 10, 01, 00\}$. Let $\mathcal{A} = \{A_1, A_2\}$ be defined as follows with $\mathcal{R}$ being the undercut relation. So $P(A_1) = P(11)$, and $P(A_2) = P(01) + P(00)$.*

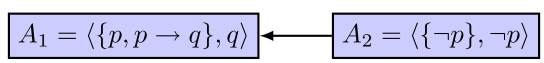

So if P is a probability distribution over $\mathcal{L}$, and $\mathcal{A} \subseteq \mathcal{A}^{\mathcal{L}}$ is a set of arguments, then the probability of an argument $A \in \mathcal{A}$ is uniquely determined by P.

Example 19. *Let P be a probability function on $\mathcal{L}$. Let the signature be (p, q, r). So the models $\mathcal{M}^{\mathcal{L}}$ are $\{111, 110, ..., 000\}$. Let $\mathcal{A} = \{A_1, A_2, A_3\}$ be defined as follows with $\mathcal{R}$ being the undercut relation. So $P(A_1) = P(111) + P(110)$, $P(A_2) = P(101) + P(001)$, and $P(A_3) = P(100) + P(000)$. Here, the undercuts do not share a model.*

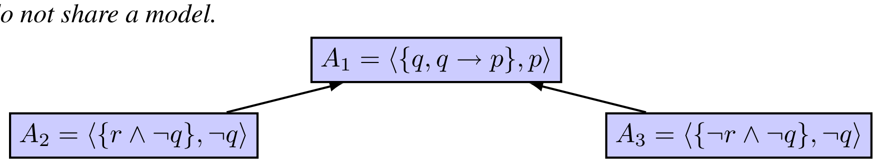

Example 20. *Let P be a probability function on $\mathcal{L}$. Let the signature be (p, q, r, s). So the models $\mathcal{M}^{\mathcal{L}}$ are $\{1111, 1110, ..., 0000\}$. Let $\mathcal{A} = \{A_1, A_2, A_3\}$ be defined as follows with $\mathcal{R}$ being the undercut relation. So $P(A_1) = P(1111) + P(1110) + P(1101) + P(1100)$, $P(A_2) = P(1011) + P(0011) + P(1010) + P(0010)$, and $P(A_3) = P(1101) + P(1001) + P(0101) + P(0001)$.*

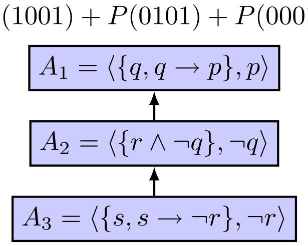

A consistent probability distribution on models means that we have substantial information about the uncertainty of the premises of the arguments. Furthermore, we have the following intuitive properties concerning the uncertainty of logical arguments which we explain as follows: (1) If the claim of an argument is a tautology, then the premises of it are the empty set, and hence, the probability of the argument is 1; (2) If two formulae are mutually inconsistent (i.e. $\neg(\phi \wedge \psi)$ is a tautology), and there is an argument with claim ϕ that is certain to be true, then any argument with claim ψ is certain to be false; (3) If two arguments have semantically equivalent support, then they have the same probability; and (4) If one argument has a subset of the premises of another, then its probability is not lower than the probability of the other argument.

Proposition 4. *Let P be a consistent probability function on $\mathcal{L}$. Let $\phi, \psi \in \mathcal{L}$ and let $\Phi, \Psi \subseteq \mathcal{L}$. For all arguments $\langle \Phi, \phi \rangle$ and $\langle \Psi, \psi \rangle$.*

1. *If $\models \phi$, then $P(\langle \{\}, \phi \rangle) = 1$.*
2. *If $\models \neg(\phi \wedge \psi)$, and $P(\langle \Phi, \phi \rangle) = 1$, then $P(\langle \Psi, \psi \rangle) = 0$.*
3. *If $\models (\bigwedge \Phi) \leftrightarrow (\wedge \Psi)$, then $P(\langle \Phi, \phi \rangle) = P(\langle \Psi, \psi \rangle)$.*
4. *If $\Phi \subseteq \Psi$, then $P(\langle \Phi, \phi \rangle) \geq P(\langle \Psi, \psi \rangle)$.*

Proof. (1) For the emptyset, $\mathsf{Models}(\emptyset) = \mathcal{M}^{\mathcal{L}}$. Therefore, if P is a consistent probability function on $\mathcal{L}$, then $P(\langle \{\}, \phi \rangle) = 1$. (2) If $\models \neg(\phi \wedge \psi)$, then $\mathsf{Models}(\phi) \cap \mathsf{Models}(\psi) = \emptyset$. So for arguments, $A_i = \langle \Phi, \phi \rangle$ and $A_j = \langle \Psi, \psi \rangle$, $\mathsf{Models}(\mathsf{Support}(A_i)) \cap \mathsf{Models}(\mathsf{Support}(A_j)) = \emptyset$. So if $\sum_{m \in \mathsf{Models}(\mathsf{Support}(A_i))} P(m) = 1$, then $\sum_{m \in \mathsf{Models}(\mathsf{Support}(A_j))} P(m) = 0$. (3) If $\models (\bigwedge \Phi) \leftrightarrow (\wedge \Psi)$, then $\mathsf{Models}(\Phi) = \mathsf{Models}(\Psi)$. Hence, $P(\langle \Phi, \phi \rangle) = P(\langle \Psi, \psi \rangle)$. (4) If $\Phi \subseteq \Psi$, then $\mathsf{Models}(\Psi) \subseteq \mathsf{Models}(\Phi)$. Hence, $P(\langle \Psi, \psi \rangle) \leq P(\langle \Phi, \phi \rangle)$. □

For any probability distribution on models, the probability of the premises is always less than or equal to the probability of any individual premise, and the probability of the premises is always less than or equal to the probability of the claim.

Proposition 5. *Let P be a consistent probability function on $\mathcal{L}$. For any argument $\langle \Phi, \phi \rangle \in \mathcal{A}^{\mathcal{L}}$, and any $\psi \in \Phi$, $P(\langle \Phi, \phi \rangle) \leq P(\psi)$, and $P(\langle \Phi, \phi \rangle) \leq P(\phi)$ hold.*

Proof. If $\psi \in \Phi$, then $\mathsf{Models}(\Phi) \subseteq \mathsf{Models}(\psi)$. Hence, $P(\langle \Phi, \phi \rangle) \leq P(\psi)$. And if $\Phi \vdash \phi$, $\mathsf{Models}(\Phi) \subseteq \mathsf{Models}(\phi)$. Hence, $P(\langle \Phi, \phi \rangle) \leq P(\phi)$. □

To analyse the probabilistic argument graphs that we have instantiated with logical arguments, we use either the epistemic approach or the constellations approach discussed in Section 3 and illustrated in the next example.

Example 21. *Let the signature $\mathcal{S}^{\mathcal{L}}$ be (p, q). So the models of $\mathcal{M}^{\mathcal{L}}$ are $\{11, 10, 01, 00\}$. Consider the following argument graph where A_2 is a canonical undercut of A_1.*

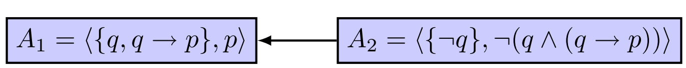

Suppose $P_1(11) = 0.8$ and $P_2(00) = P_2(10) = 0.1$. Therefore $P_1(A_1) = 0.8$, and $P_2(A_2) = 0.2$. For the epistemic approach, the epistemic extension is $\{A_1\}$. For the constellations approach, there are four full subgraphs to consider:

- *G^1 is the original graph above containing A_1 and A_2 where $P(G^1)$ = 0.16.*
- *G^2 is the graph containing just A_1 where $P(G^2)$ = 0.64.*
- *G^3 is the graph containing just A_2 where $P(G^3)$ = 0.04.*
- *G^4 is the empty graph, i.e. containing no arguments, where $P(G^4)$ = 0.16.*

Hence, for each dialectical semantics, the extension $\{A_1\}$ has probability 0.64, the extension $\{A_2\}$ has probability 0.2, and the extension $\{\}$ has probability 0.16.

In this section, we have started with a probability distribution over $\mathcal{L}$, or equivalently a probability distribution over the models of $\mathcal{L}$ (i.e. $\mathcal{M}^{\mathcal{L}}$), and then shown that this uniquely determines a probability distribution over the arguments of the language (i.e. $\mathcal{A}^{\mathcal{L}}$). The probability distribution over $\mathcal{A}^{\mathcal{L}}$ can then be used in a probabilistic argument graph. Starting with a consistent probability distribution over $\mathcal{L}$, means that we have a clear understanding of the probability distribution over any set of arguments $\mathcal{A} \subseteq \mathcal{A}^{\mathcal{L}}$, and hence of the resulting probabilistic argument graph.

Now that we have a definition of the probability of an argument in terms of the probability distribution over the language of its premises, we can explore how different choices of probability distribution over the language can affect the probability distribution over the arguments and vice versa. This will lead us to consider inconsistent probability distributions over $\mathcal{L}$ and ways to handle them. If we start with a probability distribution over a set of arguments $\mathcal{A} \subseteq \mathcal{A}^{\mathcal{L}}$, then we are not guaranteed to have a consistent probability distribution over $\mathcal{L}$. It may be that the probability distribution over the arguments underconstrains and/or overconstrains the probability distribution over the models of $\mathcal{L}$.

Example 22. *Let P be a probability function on $\mathcal{L}$. Let the signature be (p, q, r). So the models $\mathcal{M}^{\mathcal{L}}$ are $\{111, 110, ..., 000\}$. Let $\mathcal{A} = \{A_1, A_2\}$ be defined as follows with $\mathcal{R}$ being the rebut relation. Suppose each argument comes from a different agent, and each proponent is certain about in their argument. Therefore, $P(A_1) = 1$ and $P(A_2) = 1$. Since $P(A_1) = P(111) + P(110)$, and $P(A_2) = P(011) + P(001)$, the following is a probability distribution over $\mathcal{L}$: For each $m \in \{111, 110, 011, 001\}$, $P(m) = 0.5$.*

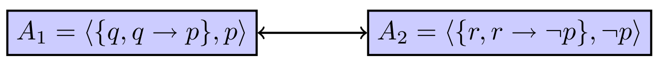

In order to structure this investigation, we introduce in the next section (Section 5), the notion of a probability distribution on arguments being coherent. We investigate coherent distributions in Section 5, and then we investigate incoherent distributions in Section 6, .

5 Logical arguments with coherent uncertainty

Given a probability distribution P on a language $\mathcal{L}$, we may want to know whether P makes sense with respect to the arguments that can be generated from $\mathcal{L}$. For this, we use the following definition for a probability distribution being coherent for a set of arguments $\mathcal{A} \subseteq \mathcal{A}^{\mathcal{L}}$.

Definition 17. *Let P be a probability function on $\mathcal{L}$. For $\mathcal{A} \subseteq \mathcal{A}^{\mathcal{L}}$, P is* **coherent** *on $\mathcal{A}$ iff for all $A_i, A_j \in \mathcal{A}$, if A_i is a defeater of A_j, then $P(A_i) + P(A_j) \leq 1$. Also, P is* **incoherent** *iff P is not coherent.*

Recall from Section 2.2, all of the logical attacks considered in this paper are defeaters (and so include direct defeaters, undercuts, direct undercuts, canonical undercuts, rebuttals, and defeating rebuttals).

Example 23. *Let P be a probability function on $\mathcal{L}$. Let the signature be (p, q, r). So the models $\mathcal{M}^{\mathcal{L}}$ are $\{111, 110, ..., 000\}$. Let $\mathcal{A} = \{A_1, A_2\}$ be defined as follows with $\mathcal{R}$ being the rebut relation. So $P(A_1) = P(111) + P(110)$, and $P(A_2) = p(011) + P(001)$. If we let $P(A_1) = 0.8$ and $P(A_2) = 0.2$, then P is coherent, whereas if we let $P(A_1) = 0.8$ and $P(A_2) = 0.4$, then P is not coherent.*

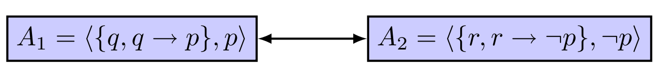

Coherence is a desirable feature of a probabilistic argument graph. For arguments A_i and A_j, where A_i attacks A_j, if A_i is believed to some degree (i.e. $P(A_i) > 0.5$), then A_j should be disbelieved to some degree (i.e. $P(A_j) \leq 0.5$), and if A_j is believed to some degree, then A_j should be disbelieved to some degree.

Proposition 6. *Let P be a probability function on $\mathcal{L}$. For $\mathcal{A} \subseteq \mathcal{A}^{\mathcal{L}}$, if P is consistent on $\mathcal{L}$, then P is coherent on $\mathcal{A}$.*

Proof. Assume P is consistent on $\mathcal{L}$. For any $A_i, A_j \in \mathcal{A}$, if A_i is a defeater of A_j, then $\mathsf{Support}(A_i) \cup \mathsf{Support}(A_j) \vdash \bot$. Let $\mathsf{Models}(A) = \{m \in \mathcal{M}^{\mathcal{L}} \mid m \models \bigwedge \mathsf{Support}(A)\}$. So $\mathsf{Models}(A_i) \cap \mathsf{Models}(A_j) = \emptyset$. Hence,

$$\left(\sum_{m \in \mathsf{Models(A_i)}} P(m) + \sum_{m \in \mathsf{Models(A_j)}} P(m) \right) \leq \left(\sum_{m \in \mathcal{M}^{\mathcal{L}}} P(m) \right) = 1$$

Therefore, $P(A_i) + P(A_j) \leq 1$, and so P is coherent on $\mathcal{A}$. □

So taking the contrapositive of the above result, we see clearly why coherence is useful for linking the probability of an argument to the probability of the models of its support: If P is not coherent, then there is no consistent probability distribution over the models of its premises.

5.1 Coherent uncertainty with commitment

Coherence can come by being skeptical. For instance, if each argument in a probabilistic argument graph is assigned 0.5, or even 0, then the probability function is coherent. However, this may be too skeptical. Next we give a subsidiary definition that allows us to consider the probability distributions that make more commitment to believing the arguments in a given set.

Definition 18. *Let P and P' be probability functions on $\mathcal{L}$. For $\mathcal{A} \subseteq \mathcal{A}^{\mathcal{L}}$, P is* **more committed** *on $\mathcal{A}$ than P' iff for all $A \in \mathcal{A}$, $P(A) \leq P'(A)$.*

Definition 19. *Let P be a probability function on $\mathcal{L}$. For $\mathcal{A} \subseteq \mathcal{A}^{\mathcal{L}}$, P is* **maximally coherent** *on $\mathcal{A}$ iff P is coherent on $\mathcal{A}$ and for all P', if P' is more committed on $\mathcal{A}$ than P, then P' is not coherent.*

Example 24. *Let P be a probability function on $\mathcal{L}$. Let the signature be (p, q, r). So the models $\mathcal{M}^{\mathcal{L}}$ are $\{111, 110, ..., 000\}$. Let $\mathcal{A} = \{A_1, A_2\}$ be defined as follows with $\mathcal{R}$ being the rebut relation. So $P(A_1) = P(111) + P(110)$, and $P(A_2) = p(011) + P(001)$.*

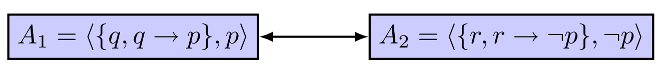

If we let $P(A_1) = 0.3$ and $P(A_2) = 0.1$, then P is coherent on $\mathcal{A}$, though not maximally coherent, whereas if we let $P(A_1) = 0.8$ and $P(A_2) = 0.2$, then P is maximally coherent on $\mathcal{A}$.

With the notion of maximally coherent, we can relate probability distributions on models with those on arguments, as follows.

Proposition 7. *For a probability function P, P is consistent on $\mathcal{L}$ iff P is maximally coherent on $\mathcal{A}^{\mathcal{L}}$.*

Proof. ($\Rightarrow$) Assume P is consistent on $\mathcal{L}$. From Proposition 6, P is coherent on $\mathcal{A}^{\mathcal{L}}$. Furthermore, for any $X \subset \mathcal{M}^{\mathcal{L}}$, $\sum_{m \in X} P(m) + \sum_{m \in \mathcal{M}^{\mathcal{L}} \setminus X} P(m) = 1$. Therefore, for any $X \subset \mathcal{M}^{\mathcal{L}}$, there are arguments $A_i, A_j \in \mathcal{A}^{\mathcal{L}}$, such that A_i is a defeater of A_j and $\mathsf{Models}(A_i) = X$ and $\mathsf{Models}(A_j) = \mathcal{M}^{\mathcal{L}} \setminus X$. Therefore, $P(A_i) + P(A_j) = 1$. Hence, P is maximally coherent on $\mathcal{A}^{\mathcal{L}}$. ($\Leftarrow$) Assume P is maximally coherent on $\mathcal{A}^{\mathcal{L}}$. Therefore, for all $\alpha \in \mathcal{L}$, there are arguments $A_i, A_j \in \mathcal{A}^{\mathcal{L}}$ such that $\mathsf{Support}(A_i) = \{\alpha\}$ and $\mathsf{Support}(A_j) = \{\neg\alpha\}$ and A_i is a defeater of A_j. So $\mathsf{Models}(A_i) \cup \mathsf{Models}(A_i) = \mathcal{M}^{\mathcal{L}}$ and $\mathsf{Models}(A_i) \cap \mathsf{Models}(A_i) = \emptyset$. Therefore, by appropriate choice of α, any bipartition of $\mathcal{M}^{\mathcal{L}}$ is obtainable (i.e. $\mathcal{M}^{\mathcal{L}}_{\alpha} \cup \mathcal{M}^{\mathcal{L}}_{\neg\alpha} = \mathcal{M}^{\mathcal{L}}$ and $\mathcal{M}^{\mathcal{L}}_{\alpha} \cap \mathcal{M}^{\mathcal{L}}_{\neg\alpha} = \emptyset$). Furthermore, because P is maximally coherent on $\mathcal{A}^{\mathcal{L}}$, $P(A_i) + P(A_j) = 1$. Therefore, for all $X \subset \mathcal{M}^{\mathcal{L}}$, $\sum_{m \in X} P(m) + \sum_{m \in \mathcal{M}^{\mathcal{L}} \setminus X} P(m) = 1$ Therefore, $\sum_{m \in \mathcal{M}^{\mathcal{L}}} P(m) = 1$, and so P is consistent on $\mathcal{L}$. □

Note, the above result is when P is maximally coherent on $\mathcal{A}^{\mathcal{L}}$ rather than a subset of it. A ramification of it is that we can identify maximally coherent distributions by given a consistent distribution on the models of the language.

5.2 Coherent uncertainty with rational extensions

As long as we have a consistent probability distribution on $\mathcal{L}$, then the resulting probabilistic argument graph is rational, as shown in the next result. Later we shall see that the converse does not hold.

Proposition 8. *Let $\mathcal{A} \subseteq \mathcal{A}^{\mathcal{L}}$ and let $\mathcal{R}$ be a defeater relation. For a probability function P on $\mathcal{L}$, if P is consistent, then P is rational for $(\mathcal{A}, \mathcal{R})$, and hence $(\mathcal{A}, \mathcal{R}, P)$ is a rational argument graph.*

Proof. Assume P is a consistent probability function P on $\mathcal{L}$. Then by Proposition 6, P is coherent on $\mathcal{A}$. Therefore, for all $A_i, A_j \in \mathcal{A}$, if A_i is a defeater of A_j, then $P(A_i) + P(A_j) \leq 1$. Therefore, if $P(A_i) > 0.5$, then $P(A_j) < 0.5$. Therefore, P is rational for $(\mathcal{A}, \mathcal{R})$. □

It is straightforward to show that any complete extension can be obtained as a rational extension by appropriate choice of probability function. However, it is not always possible to find a consistent probability function that gives this rational extension.

Proposition 9. *Let $\mathcal{A} \subseteq \mathcal{A}^{\mathcal{L}}$. If E is a complete extension of $(\mathcal{A}, \mathcal{R})$, then there is a probability distribution P on $\mathcal{A}$ such that E is a rational extension of $(\mathcal{A}, \mathcal{R}, P)$. However, it is not guaranteed that there exists a probability distribution such that P is consistent on $\mathcal{L}$ and E is a rational extension of $(\mathcal{A}, \mathcal{R}, P)$.*

Proof. Let E be a complete extension of $(\mathcal{A}, \mathcal{R})$. For every $A \in E$, let $P(A)$ be greater than 0.5, and for every $A \in \mathcal{A} \setminus E$, let $P(A)$ be less than or equal to 0.5. Therefore, E is a rational extension of $(\mathcal{A}, \mathcal{R}, P)$. To see that it is not always possible to find a a probability distribution such that P is consistent on $\mathcal{L}$ and E is a rational extension of $(\mathcal{A}, \mathcal{R}, P)$, consider the following argument tree.

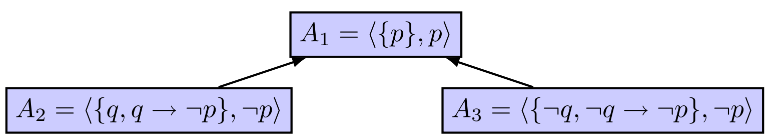

Here $E = \{A_2, A_3\}$ is a complete extension, but because $\bigcup_{A_i \in E} \mathsf{Support}(A_i)$ is inconsistent, there is no consistent probability distribution on $\mathcal{L}$, such that $P(A_2) > 0.5$ and $P(A_3) > 0.5$. □

Earlier we argued that a rational extension is desirable in that it provides a preference for a specific extension that is compliant with the structure of the argument graph. So if the probability distribution is consistent on $\mathcal{L}$, then the probability distribution on $\mathcal{A} \subseteq \mathcal{A}^{\mathcal{L}}$ is rational. However, a probability distribution on $\mathcal{L}$ does not need to be consistent for the probability distribution on $\mathcal{A} \subseteq \mathcal{A}^{\mathcal{L}}$ to be rational.

5.3 Coherent uncertainty with consistent extensions

In the context of rule-based systems, Caminada and Amgoud proposed that the union of the supports of the logical arguments should be consistent [CA05]. This work has been developed for classical logic argumentation by Amgoud and Besnard [AB09] and by Gorogiannis and Hunter [GH11]. Following this lead, we investigate the consistency of the union of the supports of arguments in epistemic extensions.

We start with a focus on argument graphs that are exhaustive in the set of logical arguments and attacks. In other words, given a knowledgebase $K \subseteq \mathcal{L}$, and a probability distribution over models P, we form the probabilistic argument graph $(\mathcal{A}^K, \mathcal{R}, P)$, where $\mathcal{A}^K = \{\langle X, \phi \rangle \in \mathcal{A}^{\mathcal{L}} \mid X \subseteq K\}$, and the attack relation $\mathcal{R}$ is one of the options presented in Definition 6 such as direct undercut, canonical undercut or rebuttal. We illustrate an exhaustive set of arguments and attacks in Figures 11 and 12.

For this section, we assume the following subsidiary definition where E is a set of arguments in an extension.

$$\mathsf{Support}(E) = \bigcup_{A \in E} \mathsf{Support}(A)$$

Because the empty set is a rational extension of any argument graph, it is straightforward to show that the support of the arguments in an epistemic extension is consistent.

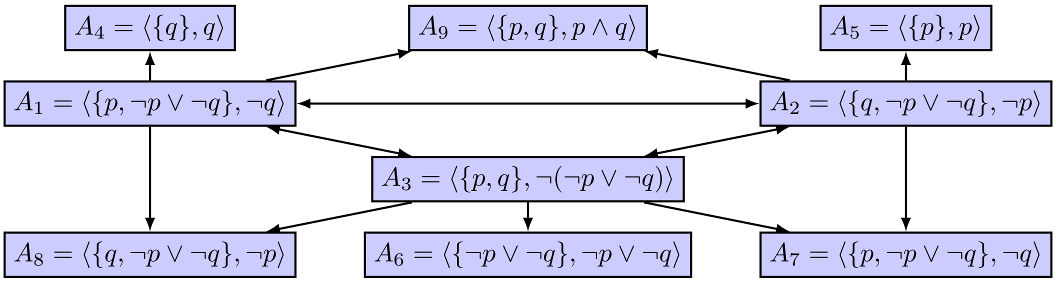

Figure 11: An example of an exhaustive argument graph (i.e. it contains all arguments that can be generated from the knowledgebase) where the knowledgebase K is $\{p, q, \neg p \vee \neg q\}$ and the attack relation is direct undercut. Let the signature $\mathcal{S}^{\mathcal{L}}$ be (p, q), and so the models $\mathcal{M}^{\mathcal{L}}$ are $\{11, 10, 01, 00\}$. Suppose we have $P(11) = 0.6$, $P(10) = 0.2$, and $P(01) = 0.2$. So $P(A_1) = 0.2$, $P(A_2) = 0.2$, $P(A_3) = 0.6$, $P(A_4) = 0.8$, $P(A_5) = 0.8$, $P(A_6) = 0.4$, $P(A_7) = 0.2$, $P(A_8) = 0.2$, and $P(A_9) = 0.6$. Hence, the epistemic extension is $\{A_3, A_4, A_5, A_9\}$. Furthermore, $\mathsf{Support}(A_3) \cup \mathsf{Support}(A_4) \cup \mathsf{Support}(A_5) \cup \mathsf{Support}(A_9)$ is $\{p, q\}$, which is consistent. Now suppose we have a rational probability distribution P' where $P'(A_1) = 0.8$, $P'(A_2) = 0.1$, $P'(A_3) = 0.1$, $P'(A_4) = 0.2$, $P'(A_5) = 0.9$, $P'(A_6) = 0.9$, $P'(A_7) = 0.8$, $P'(A_8) = 0.1$, and $P'(A_9) = 0.1$. So the epistemic extension is $\{A_1, A_5, A_6, A_7\}$. Furthermore, $\mathsf{Support}(A_1) \cup \mathsf{Support}(A_5) \cup \mathsf{Support}(A_6) \cup \mathsf{Support}(A_7)$ is $\{p, \neg p \vee \neg q\}$, which is consistent.

Proposition 10. *Let $K \subseteq \mathcal{L}$. If $\mathcal{E}$ is defined as follows, then $\mathsf{Support}(\mathcal{E}) \not\vdash \bot$ holds.*

$$\mathcal{E} = \bigcap \{E \text{ is an epistemic extension of } (\mathcal{A}^K, \mathcal{R}, P) \mid P \text{ is a rational probability distribution for } (\mathcal{A}^K, \mathcal{R})\}$$

Proof. For $(\mathcal{A}^K, \mathcal{R})$, let $P'(A) = 0.5$ for all $A \in \mathcal{A}$. So P' is a rational probability distribution for $(\mathcal{A}^K, \mathcal{R})$. Therefore, if E' is an epistemic extension of $(\mathcal{A}^K, \mathcal{R}, P')$, then $E' = \emptyset$. Since $E' \in \mathcal{E}$, then $\mathcal{E} = \emptyset$. Therefore, $\mathsf{Support}(\mathcal{E}) \not\vdash \bot$. □

If we consider individual epistemic extensions, then ensuring consistency is more challenging. For an exhaustive argument graph formed from a knowledgebase K with direct undercut, then we get consistent supports in the rational extension. We illustrate this in the example in Figure 11.

Proposition 11. *Let $K \subseteq \mathcal{L}$ and let P be a rational probability function over $\mathcal{A}^K$. If E is the rational extension of $(\mathcal{A}^K, \mathcal{R}, P)$, and $\mathcal{R}$ is direct attack, then $\mathsf{Support}(E) \not\vdash \bot$.*

Proof. Let $\mathsf{MI}(K)$ denote the minimal inconsistent subsets of K. For each $\langle \Phi, \alpha \rangle \in E$, the direct undercuts of it are formed from $\mathsf{MI}(K)$ as follows: For each $\Gamma \in \mathsf{MI}(K)$, if $\psi \in \Gamma \cap \Phi$, then $\langle \Gamma \setminus \{\alpha\}, \neg\psi \rangle$ is a direct undercut of $\langle \Phi, \alpha \rangle \in E$. So each direct undercut of $\langle \Phi, \alpha \rangle$ that can be formed from K is in $\mathcal{A}^K$. Since E is rational, there is a rational probability distribution for $(\mathcal{A}^K, \mathcal{R})$ such that for each argument $A \in E$, $P(A) > 0.5$, and for each direct undercut B of A, $P(B) \leq 0.5$. So for each argument $A \in E$, and for each direct undercut B of A, $B \in \mathcal{A}^K \setminus E$. Therefore, for each argument $A \in E$, and for each $\Gamma \in \mathsf{MI}(K)$, if $\mathsf{Support}(A) \subset \Gamma$, there are no $B \in E$, such that $\mathsf{Support}(A) \cup \mathsf{Support}(B) = \Gamma$. Therefore, $\mathsf{Support}(E) \not\vdash \bot$. □

If we consider some of the alternatives to direct undercut, such as undercut, canonical undercut or rebuttal, then we are not guaranteed to get a consistent set of formulae from the supports in the rational extension as illustrated by the example in Figure 12.

Given that the consistency of the premises used in the arguments in an extension is a desirable property, we explore this feature further using the following definition.

Definition 20. *For a probabilistic argument graph $(\mathcal{A}, \mathcal{R}, P)$, P is* **cohesive** *for $(\mathcal{A}, \mathcal{R})$ iff E is an epistemic extension of $(\mathcal{A}, \mathcal{R}, P)$ and $\mathsf{Support}(E) \not\vdash \bot$.*

Proposition 12. *For a probabilistic argument graph $(\mathcal{A}, \mathcal{R}, P)$, if P is cohesive for $(\mathcal{A}, \mathcal{R})$, then P is rational for $(\mathcal{A}, \mathcal{R})$.*

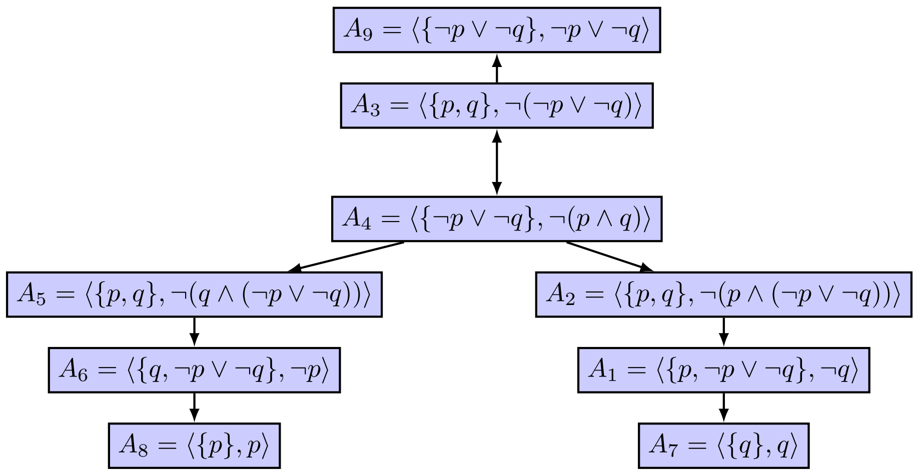

Figure 12: An exhaustive argument graph for $K = \{p, q, \neg p \vee \neg q\}$, where $\mathcal{R}$ is canonical undercut. Let $\mathcal{S}^{\mathcal{L}}$ be (p, q) and let the probability distribution be such that $P(11) = 1/3$, $P(10) = 1/3$, and $P(01) = 1/3$. Therefore, the epistemic extension E is $\{A_4, A_7, A_8, A_9\}$ since for each of these arguments, the probability is $2/3$. However, $\mathsf{Support}(E) \vdash \bot$.

Proof. Let P be cohesive for $(\mathcal{A}, \mathcal{R})$. Therefore, there is an E such that E is an epistemic extension of $(\mathcal{A}, \mathcal{R}, P)$ and $\mathsf{Support}(E) \not\vdash \bot$. Therefore, for all $A, B \in \mathcal{A}$, if A is a defeater of B, then $A \notin E$ or $B \notin E$. So $P(A) \leq 0.5$ or $P(B) \leq 0.5$. Hence, if $P(A) > 0.5$, then $P(B) \leq 0.5$. Therefore, P is rational for $(\mathcal{A}, \mathcal{R})$. $\square$

The converse does not hold as illustrated by the following example.

Example 25. *Consider the following argument graph. Let signature $\mathcal{S}^{\mathcal{L}}$ be (p, q). So the models $\mathcal{M}^{\mathcal{L}}$ are $\{11, 10, 01, 00\}$. Let $P(A_1) = 0$, $P(A_2) = 1$, $P(A_3) = 1$, and $P(A_4) = 1$. So P is rational. Furthermore, $E = \{A_2, A_3, A_4\}$ is the epistemic extension, but $\mathsf{Support}(E)$ is inconsistent. Hence P is not cohesive.*

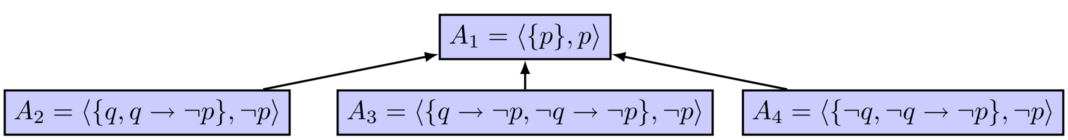

Clearly, there is a close relationship between consistency and cohesion. But we have seen that consistency on $\mathcal{L}$ does not imply cohesion, and the next example shows that cohesion does not imply consistency on $\mathcal{L}$.

Example 26. *Let $K = \{a, \neg a\}$ and let $\mathcal{S}^{\mathcal{K}} = (a, b)$. So $\mathcal{M}^{\mathcal{L}} = \{11, 10, 01, 00\}$. Let $A = \langle\{a\}, a\rangle$, and $\mathcal{A} = \{A\}$. Also let $P(11) = 0.3$, $P(10) = 0.3$, $P(01) = 0$, and $P(00) = 0$. So $P(A) = 0.6$. Hence, the epistemic extension is $\{A\}$, and so $(\mathcal{A}, \mathcal{R}, P)$ is cohesive. But P is not consistent.*

An interesting special case of a probability distribution being cohesive is when each argument in the epistemic extension has a probability 1.

Definition 21. *For a probabilistic argument graph $(\mathcal{A}, \mathcal{R}, P)$, where E be the epistemic extension of $(\mathcal{A}, \mathcal{R}, P)$ P is* **strongly cohesive** *for $(\mathcal{A}, \mathcal{R})$ iff P is cohesive for $(\mathcal{A}, \mathcal{R})$ and for all $A \in E$, $P(A) = 1$.*

Proposition 13. *For a probabilistic argument graph $(\mathcal{A}^{\mathcal{L}}, \mathcal{R}, P)$, if P is strongly cohesive for $(\mathcal{A}^{\mathcal{L}}, \mathcal{R})$, then P is consistent on $\mathcal{L}$.*

Proof. Assume P is strongly cohesive for $(\mathcal{A}^{\mathcal{L}}, \mathcal{R})$. Let E be the epistemic extension of $(\mathcal{A}^{\mathcal{L}}, \mathcal{R}, P)$. Therefore P is cohesive for $(\mathcal{A}, \mathcal{R})$ and for all $A \in E$, $P(A) = 1$. So, for all $A \in \mathcal{A}^{\mathcal{L}}$, $P(A) = 0$ or $P(A) = 1$. Hence, for all $m \in \mathsf{Models}(\mathsf{Support}(E))$, $m > 0$, and for all $m \in \mathcal{M}^{\mathcal{L}} \setminus \mathsf{Models}(\mathsf{Support}(E))$, $m = 0$. Since for all $A \in E$, $P(A) = 1$, there is exactly one model $m \in \mathsf{Models}(\mathsf{Support}(E))$, and this has unit assignment (i.e. $P(m) = 1$). Therefore, P is consistent on $\mathcal{L}$. □

So strong cohesion means that the belief in the arguments in the epistemic extension is certain. It also means that there is just one model with non-zero belief, and the premises of all arguments are based on this one model.

6 Logical arguments with incoherent uncertainty

A probabilistic argument system $(\mathcal{A}, \mathcal{R}, P)$ is incoherent when P is incoherent on $\mathcal{A}$, and this implies P is inconsistent on $\mathcal{L}$. We can use the sum of the probability assigned to the models as a measure of inconsistency, and thereby as a measure of incoherence. So the more $\sum_{m \in \mathcal{M}^{\mathcal{L}}} P(m)$ is greater than 1, the more inconsistent it is. We illustrate this situation in the following examples, where we see maximal inconsistency being $|\mathcal{M}^{\mathcal{L}}|$.

Example 27. *Consider arguments A_1 and A_2 below, where $P(A_1) = 1$ and $P(A_2) = 1$. Since $\mathcal{M}^{\mathcal{L}} = \{1, 0\}$ with the signature (q). So $P(1) = 1$ and $P(0) = 1$, and hence $\sum_{m \in \mathcal{M}^{\mathcal{L}}} P(m) = 2$.*

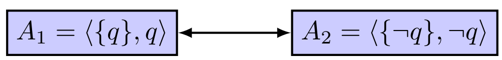

Example 28. *Consider arguments A_1 and A_2 below, where $P(A_1) = 1$ and $P(A_2) = 1$. Since $\mathcal{M}^{\mathcal{L}} = \{11, 10, 01, 00\}$ with the signature (p, q). So $P(11) = 1$ and $P(10) = 1$, and hence $\sum_{m \in \mathcal{M}^{\mathcal{L}}} P(m) = 2$.*

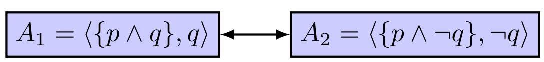

Example 29. *Consider arguments A_1 and A_2 below, where $P(A_1) = 1$, $P(A_2) = 1$, $P(A_3) = 1$, and $P(A_4) = 1$. Since $\mathcal{M}^{\mathcal{L}} = \{11, 10, 01, 00\}$ with the signature (p, q). So $P(11) = 1$, $P(10) = 1$, $P(01) = 1$, and $P(00) = 1$, and hence $\sum_{m \in \mathcal{M}^{\mathcal{L}}} P(m) = 4$.*

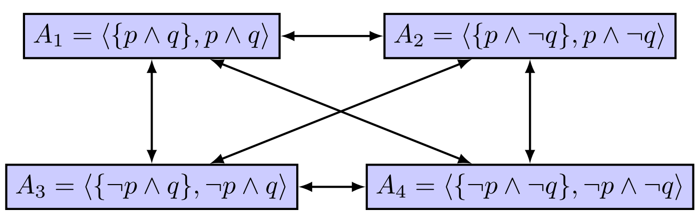

Now consider the following distribution with the same arguments, $P(A_1) = 0.5$, $P(A_2) = 0.5$, $P(A_3) = 0.5$, and $P(A_4) = 0.5$. So $P(11) = 0.5$, $P(10) = 0.5$, $P(01) = 0.5$, and $P(00) = 0.5$, and hence $\sum_{m \in \mathcal{M}^{\mathcal{L}}} P(m) = 2$.

We consider two ways of dealing with inconsistent distributions. The first is based on the constellations approach to probabilistic argument graphs (which we illustrate next), and the second is based on the epistemic approach.

Example 30. *Consider the following graph G_1 with $P(A_1) = 0.5$ and $P(A_2) = 0.8$. So the probabilistic argument graph is incoherent.*

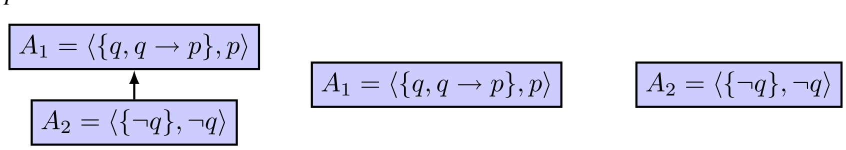

G_1 G_2 G_3

We can apply the constellations approach to the probabilistic argument graph, and so there are four full subgraphs to consider G_1 to G_4 where G_4 is the empty graph. For this, $P(G_1) = 0.4$, $P(G_2) = 0.1$, $P(G_3) = 0.4$, and $P(G_4) = 0.1$. Also, $G_1 \Vdash_{\mathsf{gr}} \{A_2\}$, $G_2 \Vdash_{\mathsf{gr}} \{A_1\}$, $G_3 \Vdash_{\mathsf{gr}} \{A_2\}$, and $G_4 \Vdash_{\mathsf{gr}} \{\}$. Therefore, $P(\{A_1\}^{\mathsf{gr}}) = 0.1$, $P(\{A_2\}^{\mathsf{gr}}) = 0.8$, and $P(\{\}^{\mathsf{gr}}) = 0.1$. Hence, the most likely interpretation of the probabilistic argument graph is that the grounded extension is $\{A_2\}$ with a probability of 0.8.

Whilst using the constellations approach is a reasonable way of directly using inconsistent probability distributions, an alternative is to identify a consistent probability distribution that reflects the information in the inconsistent probability distribution. For this, we consider how to "discount" a distribution in the next subsection, and we consider how to identify a "representative distribution" in the subsequent subsection.

6.1 Using discounting of inconsistent distributions

It is possible to rectify the inconsistency in a probability distribution, in the case when the sum of the models is greater than 1, by adopting a less committed probability distribution. Hence, there is always a probability distribution on arguments in this case, such that the probability distribution on models is consistent. Moreover, for any probability distribution P in this case, there is a consistent probability distribution P' such that P is more committed than P', To obtain this consistent probability distribution involves relaxing and/or redistributing the original probability distribution. However, this is not unique and there is the problem of choosing the appropriate relaxation and/or redistribution. For this, we consider the following notion of discounting which results in a consistent probability distribution.

Definition 22. *Let P be a probability distribution on $\mathcal{L}$ such that $\sum_{m \in \mathcal{M}^{\mathcal{L}}} P(m) > 1$. A probability distribution P' is* **discounted** *for P iff for all $m \in \mathcal{M}^{\mathcal{L}}$,*

$$P'(m) = \frac{P(m)}{\sum_{m \in \mathcal{M}^{\mathcal{L}}} P(m)}$$

Example 31. *Continuing Example 29, we get a discounted probability distribution P' for P where $P'(11) = 0.25$, $P'(10) = 0.25$, $P'(01) = 0.25$, and $P'(00) = 0.25$.*

Example 32. *We return to the example of prescriptions given in the introduction. A = "Patient has hypertension so prescribe diuretics", B = "Patient has hypertension so prescribe betablockers", and C = "Patient has emphysema which is a contraindication for betablockers". We can represent these by the following atoms h = "Patient has hypertension", b = "Prescribe betablockers", d = "Prescribe diuretics", x = "Patient is contraindicated for betablockers", y = "Patient is contraindicated for diuretics", and e = "Patient has emphysema". Hence, we can construct the following argument graph where the attack relation is the defeat relation.*

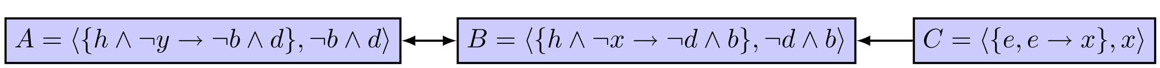

Let $\mathcal{S}^{\mathcal{L}}$ be (h, b, d, x, y, e). Suppose, we consider the arguments first, and we give a distribution as follows: $P(A) = 0.8$, $P(B) = 0.4$, and $P(C) = 0.8$. Hence, the epistemic extension is $\{A, C\}$, but P is not consistent. Suppose P, as defined over models with non-zero assignment, is

$$P(101101) = 0.8 \qquad P(101100) = 0.2 \qquad P(110000) = 0.2 \qquad P(110001) = 0.2$$

By using the discounted probability function P', we have the following non-zero assignments.

$$P(101101) = 0.58 \qquad P(101100) = 0.14 \qquad P(110000) = 0.14 \qquad P(110001) = 0.14$$

Therefore, $P'(A) = 0.72$, $P'(B) = 0.28$, and $P'(C) = 0.58$

The following result ensures that discounting gives a less committed but consistent probability distribution on $\mathcal{L}$.

Proposition 14. *For all probability distributions P on $\mathcal{L}$, if $\sum_{m\in\mathcal{M}^{\mathcal{L}}} P(m) > 1$, then the discounted probability distribution P' on $\mathcal{L}$, is such that P is more committed than P', and $\sum_{m\in\mathcal{M}^{\mathcal{L}}} P'(m) = 1$.*

Proof. Let $k = \sum_{m\in\mathcal{M}^{\mathcal{L}}} P(m)$. So $P'(m) = P(m)/k$. Hence, $\sum_{m\in\mathcal{M}^{\mathcal{L}}} P'(m) = \sum_{m\in\mathcal{M}^{\mathcal{L}}}(P(m)/k) = k/k = 1$. For each $A \in \mathcal{A}^{\mathcal{L}}$, $P(A) = \sum_{m\in\mathsf{Models}(\mathsf{Support}(A))} P(m)$ and $P'(A) = \sum_{m\in\mathsf{Models}(\mathsf{Support}(A))} P'(m)$ = $\sum_{m\in\mathsf{Models}(\mathsf{Support}(A))}(P(m)/k)$. Therefore, $P'(A) = P(A)/k$, and since $k > 1$, $P(A) > P'(A)$. Therefore, P is more committed than P', □

Obviously by relaxing a probability function, the epistemic extension will be decreased.

Proposition 15. *Let P is be an inconsistent probability function $\mathcal{L}$, and let E be the epistemic extension of $(\mathcal{A}, \mathcal{R}, P)$ where $\mathcal{A} \subseteq \mathcal{A}^{\mathcal{L}}$. If P' is discounted for P, and E' is the epistemic extension of $(\mathcal{A}, \mathcal{R}, P')$, then $E' \subseteq E$.*

Proof. For all $m \in \mathcal{M}^{\mathcal{L}}$, $P'(m) \leq P(m)$. Therefore, for all arguments $A \in \mathcal{A}^{\mathcal{L}}$,

$$\left(\sum_{m\in\mathsf{Models}(\mathsf{Support}(A))} P'(m)\right) \leq \left(\sum_{m\in\mathsf{Models}(\mathsf{Support}(A))} P(m)\right)$$

Therefore, for all arguments $A \in \mathcal{A}^{\mathcal{L}}$, $P'(A) \leq P(A)$. Therefore, for all arguments $A \in \mathcal{A}^{\mathcal{L}}$, if $P'(A) > 0.5$, then $P(A) > 0.5$. Therefore, $E' \subseteq E$. □

We can view the assignment of probabilities to extensions as providing a ranking over extensions. For instance, if grounded extension E_1 has a probability 0.8, and grounded extension E_2 has a probability 0.1, then E_1 is more probable than E_2. Discounting allows us to get a consistent probability distribution on $\mathcal{L}$ whilst at the same time not loosing the interpretation of a probabilistic argument graph that comes with the constellations approach.

Proposition 16. *Let P be a probability distribution on $\mathcal{L}$ such that $\sum_{m\in\mathcal{M}^{\mathcal{L}}} P(m) > 1$, and let P' be the discounted probability distribution for P. Let $\mathcal{A} \subseteq \mathcal{A}^{\mathcal{L}}$, and let $\mathcal{R}$ be a defeater relation. For all $E \subseteq \mathcal{A}$, if $P(E) > 0$ then $P'(E) > 0$.*

Proof. Consider $E \subseteq \mathcal{A}$. For all $G' \sqsubseteq G$, if $P(G') > 0$, then $P'(G') > 0$. So, if $(\sum_{G'\in Q_X(E)} P(G')) > 0$, then $(\sum_{G'\in Q_X(E)} P'(G')) > 0$ Therefore, if $P(E) > 0$, then $P'(E) > 0$. □

Example 33. *Let $\mathcal{M}^{\mathcal{L}} = \{111, 110, 101, ..., 000\}$ with the signature $\mathcal{S}^{\mathcal{L}} = (p, q, r)$. Suppose $P(111) = 0.4$, $P(011) = 0.4$, $P(101) = 0.4$, and for the remaining models $m \in \mathcal{M}^{\mathcal{L}}$, $P(m) = 0$, and hence $\sum_{m\in\mathcal{M}^{\mathcal{L}}} P(m) = 1.2$. So for arguments A_1 and A_2 below, $P(A_1) = 0.8$ and $P(A_2) = 0.4$. Using the constellations approach with $(\mathcal{A}, \mathcal{R}, P)$, we get the grounded extension $\{A_1\}$ with probability 0.48, the grounded extension $\{A_2\}$ with probability 0.4, and the grounded extension $\{\}$ with probability 0.12.*

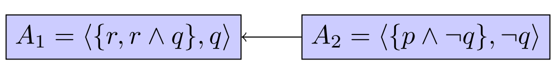

By discounting, we get $P'(111) = 1/3$, $P'(011) = 1/3$, $P'(101) = 1/3$, and for the remaining models $m \in \mathcal{M}^{\mathcal{L}}$, $P(m) = 0$, and hence $\sum_{m\in\mathcal{M}^{\mathcal{L}}} P'(m) = 1$. So for arguments A_1 and A_2 above, $P'(A_1) = 2/3$ and $P'(A_2) = 1/3$. Using the constellations approach with $(\mathcal{A}, \mathcal{R}, P')$, we get the grounded extension $\{A_1\}$ with probability $4/9$, the grounded extension $\{A_2\}$ with probability $3/9$, and the grounded extension $\{\}$ with probability $2/9$.

Example 34. *Let $\mathcal{M}^{\mathcal{L}} = \{111, 110, 101, ..., 000\}$ with the signature $\mathcal{S}^{\mathcal{L}} = (p, q, r)$. Suppose $P(100) = 1$, $P(010) = 1$, $P(001) = 1$, and for the remaining models $m \in \mathcal{M}^{\mathcal{L}}$, $P(m) = 0$, and hence $\sum_{m\in\mathcal{M}^{\mathcal{L}}} P(m)$ = 3. Also $P(A_1) = 1$, $P(A_2) = 1$, and $P(A_3) = 1$. Using the constellations approach with $(\mathcal{A}, \mathcal{R}, P)$, we get the grounded extension $\{\}$ with probability 1.*

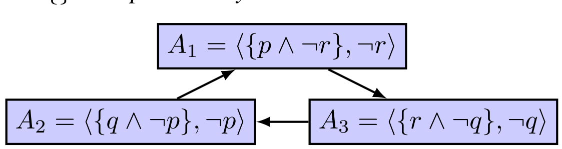

By discounting, we get $P'(100) = 1/3$*,* $P'(010) = 1/3$*,* $P'(001) = 1/3$*, and for the remaining models* $m \in \mathcal{M}^{\mathcal{L}}$*,* $P(m) = 0$*, and hence* $\sum_{m \in \mathcal{M}^{\mathcal{L}}} P'(m) = 1$*. So for arguments* $P(A_1) = 1/3$*,* $P(A_2) = 1/3$*, and* $P(A_3) = 1/3$*. Using the constellations approach with* $(\mathcal{A}, \mathcal{R}, P')$*, we get the following grounded extensions*

Extension	$\{\}$	$\{A_1\}$	$\{A_2\}$	$\{A_3\}$	$\{A_1, A_2\}$	$\{A_1, A_3\}$	$\{A_2, A_3\}$	$\{A_1, A_2, A_3\}$
Probability	*9/27*	*4/27*	*4/27*	*4/27*	*2/27*	*2/27*	*2/27*	*0/27*

Discounting is an obvious way of relaxing a probability function so that it is consistent. With the epistemic approach, it can reduce the membership of the epistemic extension, but with the constellations approach, all the extensions with non-zero probability remain as extensions with non-zero probability after discounting.

6.2 Using a representative distribution that is consistent

In some situations, for a rational argument graph $(\mathcal{A}, \mathcal{R}, P)$, it is possible to find a consistent probability distribution P' on $\mathcal{L}$ such that the rational extensions of $(\mathcal{A}, \mathcal{R}, P)$ and $(\mathcal{A}, \mathcal{R}, P')$ are the same. We illustrate this in the following example.

Example 35. *Consider* $P(A_1) = 0.5$ *and* $P(A_2) = 0.8$*. Let the signature* $\mathcal{S}^{\mathcal{L}}$ *be* (p, q)*. So* P *is a rational assignment, but not consistent on* $\mathcal{L}$*. Now consider* P' *where* $P'(A_1) = 0.2$ *and* $P'(A_2) = 0.8$*. So* P' *is a rational and consistent on* $\mathcal{L}$ *(since* $P'(11) = 0.2$ *and* $P'(10) + P'(01) = 0.8$*), and* P' *and* P *have the same rational extension.*

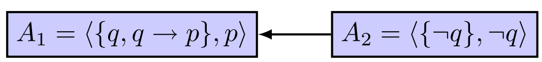

However, in general, for a rational argument graph $(\mathcal{A}, \mathcal{R}, P)$, it is not guaranteed that there is a consistent probability distribution P' on $\mathcal{L}$ such that the rational extensions of $(\mathcal{A}, \mathcal{R}, P)$ and $(\mathcal{A}, \mathcal{R}, P')$ are the same.

Example 36. *Consider the following argument graph with the probability distribution* $P(A_1)$ = 0*,* $P(A_2)$ = 1*, and* $P(A_3)$ = 1*. So* P *is a rational probability distribution, and the epistemic extension is* $\{A_2, A_3\}$*. But because* $\mathsf{Support}(A_2) \cup \mathsf{Support}(A_3) \vdash \bot$*, there is no consistent probability distribution* P' *such that the rational extension is* $\{A_2, A_3\}$*.*

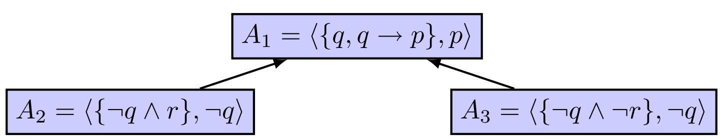

As a special case of inconsistency for the probability distribution, we can assume that we have an argument graph $(\mathcal{A}, \mathcal{R})$ with a cohesive probability distribution on the arguments. Then we are guaranteed to be able to find a consistent probability distribution on the models and this distribution agrees with the rational probability distribution. So finding a probability distribution that agrees may involve relaxing and/or redistributing the probability distribution. This then raises the question of how can we effectively choose an appropriate relaxation or redistribution. For this, we will consider the following definition of a representative distribution.

Definition 23. *Let* $(\mathcal{A}, \mathcal{R}, P)$ *be a probabilistic argument graph where* P *is a rational probability distribution on* $\mathcal{L}$ *and* $\sum_{m \in \mathcal{M}^{\mathcal{L}}} P(m) > 1$ *and* P *is cohesive for* $(\mathcal{A}, \mathcal{R})$*. Let* E *be the epistemic extension of* $(\mathcal{A}, \mathcal{R}, P)$ *and let the cardinality of* $\mathsf{Models}(\mathsf{Support}(E))$ *be* κ*. A probability distribution* P' *is a* **representative** *for* $(\mathcal{A}, \mathcal{R}, P)$ *iff*

$$P'(m) = \begin{cases} 1/\kappa & \text{if } m \in \mathsf{Models}(\mathsf{Support}(E)) \\ 0 & \text{otherwise} \end{cases}$$

Example 37. *Continuing with the probabilistic argument graph $(\mathcal{A}, \mathcal{R}, P)$ given in Example 35, suppose $P(11) = 0.5$, $P(10) = 0.4$, and $P(00) = 0.4$, and for the remaining models $m \in \mathcal{M}^{\mathcal{L}}$, $P(m) = 0$. So P is a rational assignment, but not consistent. For this graph,* $\mathsf{Models}(\mathsf{Support}(E)) = \{10, 00\}$*, and so the probability distribution P' that is representative for $(\mathcal{A}, \mathcal{R}, P)$ is given by $P'(10) = 0.5$, $P'(00) = 0.5$, and for the remaining models $m \in \mathcal{M}^{\mathcal{L}}$, $P(m) = 0$. So $P(A_2) = 1$.*

Proposition 17. *Let $\mathcal{A} \subseteq \mathcal{A}^{\mathcal{L}}$. If P is cohesive for $(\mathcal{A}, \mathcal{R})$ and P' is representative for $(\mathcal{A}, \mathcal{R}, P)$, then P' is consistent.*

Proof. Assume P is cohesive. So $\mathsf{Support}(E) \not\vdash \bot$. So $\mathsf{Models}(\mathsf{Support}(E)) \neq \emptyset$. So for all models $m \in \mathsf{Models}(\mathsf{Support}(E))$, $P'(m) = 1/\kappa$, where the cardinality of $\mathsf{Models}(\mathsf{Support}(E))$ is κ. Therefore, $\sum_{m \in \mathsf{Models}(\mathsf{Support}(E))} P'(m) = 1$, and $\sum_{m \in \mathcal{M}^{\mathcal{L}} \setminus \mathsf{Models}(\mathsf{Support}(E))} P'(m) = 0$. So P' is consistent. □

For our next result, we consider the important case of a probabilistic argument graph that ensures all arguments not in the epistemic extension are involved in a conflict with an argument in the epistemic extension.

Definition 24. *Let $(\mathcal{A}, \mathcal{R}, P)$ be a probabilistic argument graph, and let E be the epistemic extension of $(\mathcal{A}, \mathcal{R}, P)$. $(\mathcal{A}, \mathcal{R}, P)$ is* **bellicose** *iff for each $A \in \mathcal{A} \setminus E$, there is a $A' \in E$ such that $(A, A') \in \mathcal{R}$ or $(A', A') \in \mathcal{R}$.*

Proposition 18. *Let $(\mathcal{A}, \mathcal{R}, P)$ be a bellicose probabilistic argument graph where P is cohesive and rational for $(\mathcal{A}, \mathcal{R})$. Let E be the epistemic extension for $(\mathcal{A}, \mathcal{R}, P)$. If P' is the representative probability distribution for P, and E' is the epistemic extension for $(\mathcal{A}, \mathcal{R}, P')$, then $E = E'$.*

Proof. To show that $E = E'$ holds, we consider two cases. Case 1 is to show that $E \subseteq E'$. Since P is cohesive for $(\mathcal{A}, \mathcal{R})$, $\mathsf{Support}(E) \not\vdash \bot$. Therefore, $\mathsf{Models}(\mathsf{Support}(E)) \neq \emptyset$. Furthermore, for all $m \in \mathsf{Models}(\mathsf{Support}(E))$, $P'(m) = 1/k$. Also, for all $A \in E$, $\mathsf{Models}(\mathsf{Support}(E)) \subseteq \mathsf{Models}(\mathsf{Support}(A))$. Let $\mathsf{Assigned}(A) = \{m \in \mathsf{Models}(\mathsf{Support}(A)) \mid P'(m) > 0\}$ for each $A \in E$. So, for each $A \in E$, $\mathsf{Assigned}(A) = \mathsf{Models}(\mathsf{Support}(E))$. Therefore, for each $A \in E$, $P'(A) = 1$. Therefore, for all $A \in E$, we have $A \in E'$. Therefore, $E \subseteq E'$ holds. Case 2 is to show that $E' \subseteq E$. We assume that $(\mathcal{A}, \mathcal{R}, P)$ is a bellicose probabilistic argument graph. So for each $A \in \mathcal{A} \setminus E$, there is an $A' \in E$ such that $(A, A') \in \mathcal{R}$ or $(A', A') \in \mathcal{R}$. Therefore, for each $A \in \mathcal{A} \setminus E$, there is an $A' \in E$ such that A is a defeater of A' or A' is a defeater of A. Therefore, $\mathsf{Support}(A) \cup \mathsf{Support}(A') \vdash \bot$. Therefore, $\mathsf{Models}(\mathsf{Support}(A)) \cap \mathsf{Models}(\mathsf{Support}(E)) = \emptyset$. Therefore, $\mathsf{Assigned}(A) = \emptyset$. So $P'(A) = 0$. Hence, $A \notin E'$. So, we have $E' \subseteq E$. From cases 1 and case 2 holding, we have $E = E'$. □

Example 38. *Consider the following argument graph $(\mathcal{A}, \mathcal{R})$ with the probability function $P(A_1) = 0.8$, $P(A_2) = 0.4$, and $P(A_3) = 0.7$. Let $\mathcal{S}^{\mathcal{L}}$ be (a, b, c, d). Also, let P be such that $P(1101) = 0.7$, $P(1111) = 0.1$, $P(1010) = 0.2$, $P(1011) = 0.2$, and $P(m) = 0$ for the remaining models m. So P is rational but not consistent. Also $(\mathcal{A}, \mathcal{R}, P)$ is bellicose. The epistemic extension is $\{A_1, A_3\}$. The representative probability distribution is $P'(1101) = 1$, and $P'(m) = 0$ for the remaining models m. Hence, $P'(A_1) = 1$, $P'(A_2) = 0$, and $P'(A_3) = 1$.*

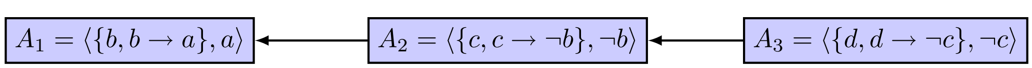

Example 39. *Consider the following argument graph $(\mathcal{A}, \mathcal{R})$ where $\mathcal{S}^{\mathcal{L}}$ is (a, b, c, d), and the probability function is $P(10) = 1/2$, and $P(01) = 1/2$. So $P(a) = 1/2$, $P(b) = 1/2$, $P(a \vee b) = 1$, and $P(a \wedge b) = 0$. Therefore the epistemic extension of $(\mathcal{A}, \mathcal{R}, P)$ is $\{A_3\}$. So $(\mathcal{A}, \mathcal{R}, P)$ is not bellicose.*

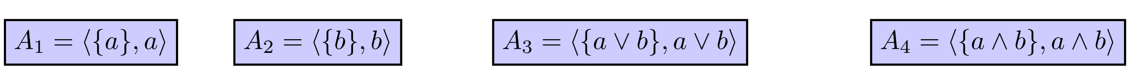

Now consider the representative probability distribution $P'(10) = 1/3$, $P'(11) = 1/3$, and $P'(01) = 1/3$. Therefore, $P'(a) = 2/3$, $P'(b) = 2/3$, $P'(a \vee b) = 1$, and $P'(a \wedge b) = 1/3$. Therefore the epistemic extension of $(\mathcal{A}, \mathcal{R}, P')$ is $\{A_1, A_2, A_3\}$.

So by using a representative probability distribution, with a probabilistic argument graph that is bellicose, we get a consistent probability distribution without affecting the epistemic extension.

Definition 25. *Let $(\mathcal{A}, \mathcal{R}, P)$ be a probabilistic argument graph and let E be the epistemic extension of $(\mathcal{A}, \mathcal{R}, P)$. P is* **maximally cohesive** *for $(\mathcal{A}, \mathcal{R})$ iff P is cohesive for $(\mathcal{A}, \mathcal{R})$ and there is no P' such that P' is cohesive for $(\mathcal{A}, \mathcal{R})$ and $E \subset E'$ where E' is the epistemic extension of $(\mathcal{A}, \mathcal{R}, P')$*

Proposition 19. *If $(\mathcal{A}, \mathcal{R}, P)$ is bellicose and P is cohesive for $(\mathcal{A}, \mathcal{R})$, then P is maximally cohesive for $(\mathcal{A}, \mathcal{R})$.*

Proof. Assume $(\mathcal{A}, \mathcal{R}, P)$ is bellicose and P is cohesive for $(\mathcal{A}, \mathcal{R})$. Let E be the epistemic extension of $(\mathcal{A}, \mathcal{R}, P)$. So for all $A \in \mathcal{A} \setminus E$, there is an $A' \in E$ such that $(A, A') \in \mathcal{R}$ or $(A', A) \in \mathcal{R}$. So for all $A \in \mathcal{A} \setminus E$, there is an $A' \in E$ such that $\mathsf{Support}(\{A, A'\}) \vdash \bot$. So for all $A \in \mathcal{A} \setminus E$, $\mathsf{Support}(E \cup \{A\}) \vdash \bot$. So for all $E' \subseteq \mathcal{A}$, if $E \subset E'$, then $\mathsf{Support}(E') \vdash \bot$ Therefore, there is no P' such that E' be the epistemic extension of $(\mathcal{A}, \mathcal{R}, P')$ and $\mathsf{Support}(E') \not\vdash \bot$. Therefore, there is no P' such that P' is cohesive for $(\mathcal{A}, \mathcal{R})$ and $E \subset E'$ where E' is the epistemic extension of $(\mathcal{A}, \mathcal{R}, P')$. Therefore, P is maximally cohesive for $(\mathcal{A}, \mathcal{R})$. □

Example 40. *We return to the example in Figure 6 of three arguments A = "Ann will go to the party and this means that Bob will not go to the party", B = "Bob will go to the party and this means that Chris will not go to the party", and C = "Chris will go to the party and this means that Ann will not go to the party", which are in a cycle. We can represent these by the following atoms a = "Ann will go to the party", b = "Bob will go to the party", and c = "Chris will go to the party". Hence, we can construct the following argument graph.*

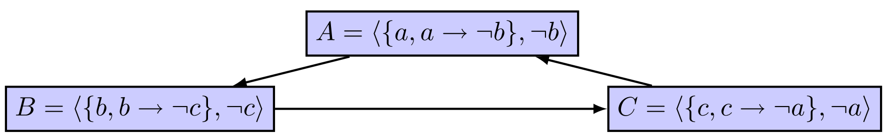

Let $\mathcal{S}^{\mathcal{L}}$ be (a, b, c). Suppose, we consider the arguments first, and we give a distribution as follows: $P(A) = 0.8$, $P(B) = 0.3$, and $P(C) = 0.1$. So we are quite confident that Ann will go, and the others will not. Hence, the epistemic extension is $\{A\}$. Also, P is bellicose for $(\mathcal{A}, \mathcal{R})$ but P is not consistent, since $P(A) = P(101) + P(100) = 0.8$, $P(B) = P(110) + P(010) = 0.3$, and $P(C) = P(011) + P(001) = 0.1$. By using the representative probability function P', we have $P'(101) = 1/2$ and $P'(100) = 1/2$. Therefore, $P'(A) = 1$, $P'(B) = 0$, and $P'(C) = 0$

A representative distribution for a probabilistic argument graph $(\mathcal{A}, \mathcal{R}, P)$ is consistent when $(\mathcal{A}, \mathcal{R}, P)$ is cohesive. Furthermore, in practice, it is often the case that a probabilistic argument graph $(\mathcal{A}, \mathcal{R}, P)$ is bellicose. This means a representative distribution P' would give the same epistemic extension as P. Therefore, a representative distribution is an effective approximation for reflecting the beliefs of an agent with a consistent probability distribution.

7 Multiple agent incoherency

In this section, we consider how arguments may come from different agents. Each agent has their own probability distribution on $\mathcal{L}$, and therefore over the models $\mathcal{M}^{\mathcal{L}}$. This means that the probability distribution over arguments $\mathcal{A}^{\mathcal{L}}$ may disagree, and therefore may lead to incoherence.

For instance, it possible that there are arguments A and A' such that A defeats A', and agent 1 with probability distribution P_1 and agent 2 with probability distribution P_2 such that

$$P_1(A) = 1 \qquad P_1(A') = 0 \qquad P_2(A) = 0 \qquad P_2(A') = 1$$

This kind of situation is common in argumentation. It reflects a strong conflict in the beliefs of the two agents. In order to deal with this situation, we first consider in Section 7.1 how to measure the difference between the two probability distributions on $\mathcal{L}$, we then consider in Section 7.2 how it may be possible to find agreement between the agents, and finally we consider how we can combine probability distributions by each agent into a combined probability distribution over the arguments in Section 7.3.

7.1 Measuring divergence between agents' probability distributions

Given a pair of agents each with their own probability distribution over the language $\mathcal{L}$, we can evaluate how much they disagree by using the following measure. This is the total variation distance which we use on grounds of simplicity. But potentially there are various alternatives that may be considered such the Kullback-Leibler divergence, or the class of f-divergences.

Definition 26. *Let P_i and P_j be probability distributions on $\mathcal{L}$. The* **divergence** *of P_i from P_j, denoted $|P_i - P_j|$, is defined as follows, where $|P_i(m) - P_j(m)|$ is the larger of $P_i(m)$ and $P_j(m)$ minus the smaller.*

$$|P_i - P_j| = \frac{1}{2} \sum_{m \in \mathcal{M}^{\mathcal{L}}} |P_i(m) - P_j(m)|$$

So for all consistent probability distributions, P_i and P_j, on $\mathcal{L}$, we have that $0 \leq |P_i - P_j| \leq 1$. If $|P_i - P_j| = 0$, then P_i and P_j fully agree on all models. If $|P_i - P_j| = 1$, and P_i and P_j are each consistent, then P_i and P_j completely disagree on all models, and we call them fully divergent. Obviously, the divergence of a probability distribution from itself is zero (i.e. $|P_i - P_i| = 0$), and divergence is symmetrical (i.e. $|P_i - P_j| = |P_j - P_i|$).

Example 41. *Let the signature $\mathcal{S}^{\mathcal{L}}$ be (p). So the models of $\mathcal{M}^{\mathcal{L}}$ are $\{1, 0\}$. Consider the following argument graph. Suppose agent 1 has $P_1(A_1) = 1$, and agent 2 has $P_2(A_2) = 1$. Then $P_1(1) = 1$, and $P_2(0) = 1$. So, the divergence $|P_1 - P_2| = 1$ is maximal.*

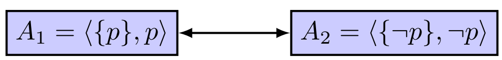

Example 42. *Let the signature $\mathcal{S}^{\mathcal{L}}$ be (p, q). So the models of $\mathcal{M}^{\mathcal{L}}$ are $\{11, 10, 01, 00\}$. Consider the following argument graph. Suppose agent 1 has $P_1(11) = 1$ and agent 2 has $P_2(00) = P_2(10) = 0.5$. So agent 1 has $P_1(A_1) = 1$, and agent 2 has $P_2(A_2) = 1$. So, the divergence $|P_1 - P_2|$ is maximal.*

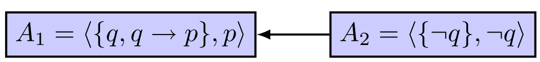

Example 43. *Let the signature be (p, q). So the models of $\mathcal{M}^{\mathcal{L}}$ are $\{11, 10, 01, 00\}$. Consider the following argument graph using canonical undercut as the attack relation. Suppose agent 1 has $P_1(00) = 1$ and agent 2 has $P_2(11) = P_2(10) = P_2(01) = P_2(00) = 0.25$. Therefore, $|P_1 - P_2| = 1/2(3/4 + 1/4 + 1/4 + 1/4) = 1/2 \times 6/4 = 3/4$. This is less than 1 because there is a model to which both assign some mass.*

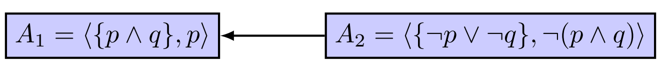

As shown next, the divergence is maximal when the agents assign their belief to different models. In other words, the set of models to which one agent assigns non-zero belief is disjoint from the set of models to which other agent assigns non-zero belief.

Proposition 20. *Let P_i and P_j each be consistent probability distributions on $\mathcal{L}$.*

$$|P_i - P_j| = 1 \textit{ iff for each } m \in \mathcal{M}^{\mathcal{L}}, \textit{ if } P_i(m) > 0, \textit{ then } P_j(m) = 0$$

Proof. $|P_i - P_j| = 1$ iff $\frac{1}{2}\sum_{m \in \mathcal{M}^{\mathcal{L}}} |P_i(m) - P_j(m)| = 1$ iff $\sum_{m \in \mathcal{M}^{\mathcal{L}}} |P_i(m) - P_j(m)| = 2$ iff $\sum_{m \in \mathcal{M}^{\mathcal{L}}} |P_i(m) - P_j(m)| = \sum_{m \in \mathcal{M}^{\mathcal{L}}} P_i(m) + \sum_{m \in \mathcal{M}^{\mathcal{L}}} P_j(m)$ iff there are subsets $M_i \subset \mathcal{M}^{\mathcal{L}}$, and $M_j \subset \mathcal{M}^{\mathcal{L}}$, such that for each $m_i \in M_i$, $P_i(m_i) > 0$ and $P_i(m_j) = 0$, and for each $m_j \in M_j$, $P_i(m_j) = 0$ and $P_j(m_j) > 0$ iff for each $m \in \mathcal{M}^{\mathcal{L}}$, if $P_i(m) > 0$, then $P_j(m) = 0$. □

The following result can be explained as follows: (1) If the divergence is less than 1, then agents agree on some of the models to which some non-zero belief is assigned; (2) If the divergence is zero, then the agents agree on the probability of each argument; and (3) If the agents agree on the probability of each argument, then the divergence is less than one.

Proposition 21. *Let P_i and P_j be consistent probability distributions on $\mathcal{L}$.*

1. *If $|P_i - P_j| < 1$, then there is an $m \in M$ such that $P_i(m) > 0$ and $P_j(m) > 0$.*

2. *If $|P_i - P_j| = 0$, then for all $A \in \mathcal{A}$, $P_i(A) = P_j(A)$*

3. *If for all $A \in \mathcal{A}$, $P_i(A) = P_j(A)$, then $|P_i - P_j| < 1$*

Proof. (1) Assume $|P_i - P_j| < 1$. So $\frac{1}{2}\sum_{m \in \mathcal{M}^{\mathcal{L}}} |P_i(m) - P_j(m)| < 1$. So $\sum_{m \in \mathcal{M}^{\mathcal{L}}} |P_i(m) - P_j(m)| < 2$. Hence there is an $m \in M$ such that $P_i(m) > 0$ and $P_j(m) > 0$. (2) Assume $|P_i - P_j| = 0$. Therefore, $\frac{1}{2}\sum_{m \in \mathcal{M}^{\mathcal{L}}} |P_i(m) - P_j(m)| = 0$. Therefore, for all $m \in \mathcal{M}^{\mathcal{L}}$, $P_i(m) - P_j(m) = 0$. Therefore, for all $m \in \mathcal{M}^{\mathcal{L}}$, $P_i(m) = P_j(m)$. Therefore, for all $A \in \mathcal{A}$, $P_i(A) = P_j(A)$. (3) Assume for all $A \in \mathcal{A}$, $P_i(A) = P_j(A)$. Therefore, there is an $m \in M$ such that $P_i(m) = P_j(m)$. Therefore, there is an $m \in M$ such that $P_i(m) - P_j(m) = 0$. Therefore, $\frac{1}{2}\sum_{m \in \mathcal{M}^{\mathcal{L}}} |P_i(m) - P_j(m)| < 1$. Therefore, $|P_i - P_j| < 1$. □

In the next section, we consider how we may seek some agreement between the agents, and then in the subsequent section we consider how we may combine the probability functions.

7.2 Finding agreement between agents

When a pair of agents have divergent probability distributions, knowing about the differences between the probability distributions may help us in deciding how the agents could act. For exampe, if one agent has an inconsistent probability function, and the other agent has a consistent probability function, then the first agent is at a disadvantage to the second agent since it does not have a consistent view on the information under consideration by the agents. As another example, if each agent is prepared to compromise, they should aim to redistribute mass so that the divergence is reduced. However, if each agent has a consistent probability function, and when the respective probability functions are applied to the argument graph, the resulting probabilistic extensions are the same, then the differences between the two agents might be regarded as negligible.

Example 44. *Consider the following example, where A_1 is an argument by an employer to an employee and A_2 is the reply by the employee. We may model this by an argument graph as follows.*

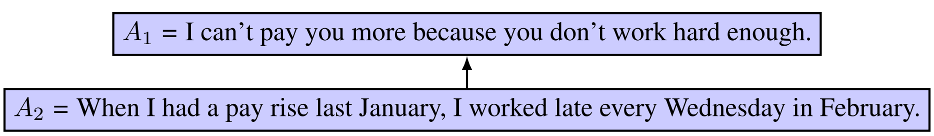

For A_1 there is uncertainty about whether or not the employer can pay more (since we don't know enough about the business, its profitability, etc), and we don't know whether the employee works hard or not. For A_2, there is also some uncertainty about the fact the employee had a pay rise in January, and there is some uncertainty about whether she worked late every Wednesday in February). Probably she will not lie about the fact that she did indeed have a pay rise in January, but she may be being economical with the truth with regard to working late every *Wednesday, and of course there is doubt as to what "working late" means anyway. Is it that she stayed 10 minutes after her normal time, or is that she worked 6 hours extra?*

Let us suppose that we know the employer and the employee, and so when we hear this exchange of arguments, we can assign a probability for each argument. But, this still leaves the question of whether the attack is justified. In other words, suppose both arguments are true, it may be that we do not believe that A_2 is a counterargument to A_1. Just because the employee works late every Wednesday for a month, does that imply that the employee works hard?

Another way of looking at this is if we represent the arguments in propositional logic, as follows, where x is a "Employer can pay employee more", y is "Employee does not work hard enough", and v is "Employee had pay rise in January", and w is "Employee worked late every Wednesday in February". From these propositional atoms, we can encode A_1 as B_1, but for A_2 we have a choice of either B_2 or B_2'. So there is uncertainty as to the precise logical nature of the counterargument, and hence there is uncertainty as to whether the attack holds.

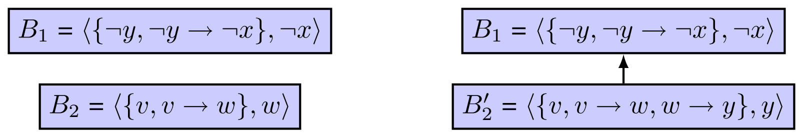

Let the signature $\mathcal{S}^{\mathcal{L}}$ be (v, w, x, y). So $\mathsf{Models}(B_1) = \{1100, 1000, 0100, 0000\}$, $\mathsf{Models}(B_2) = \{1111, 1110, 1101, 1100\}$, and $\mathsf{Models}(B_2') = \{1101, 1111\}$. Now we can represent the perspectives by the employee and employer as follows. Perhaps we may interpret the employee's statements as reflecting certainty in her position. If P_e denotes the employee's probability distribution, then $P_e(1111) = 1$. In contrast, the employer might be inclined to think the worker shouldn't have a pay rise, but may have some uncertainty about his beliefs. If P_b denotes the employer's probability distribution, (with the index b for boss), then $P_b(1100) = 0.6$, $P_b(1000) = 0.3$, and $P_e(1111) = 0.1$. So $P_e(B_1) = 0$, $P_e(B_2) = 1$, and $P_e(B_2') = 1$. Also, $P_b(B_1) = 0.9$, $P_b(B_2) = 0.7$, and $P_b(B_2') = 0.1$. Therefore, the divergence between P_e and P_b is 0.9 and so near maximal. The extra information from the probability distributions shows how far apart the two agents are. Furthermore, the employee may then proceed by trying to get the employer to shift mass to the model 1111, whereas the employer may try to proceed by getting the employee to shift mass to models 1100 and 1000.

For the probabilistic argument graphs $(\mathcal{A}^{\mathcal{L}}, \mathcal{R}, P)$ and $(\mathcal{A}^{\mathcal{L}}, \mathcal{R}, P')$, if they have the same epistemic extensions, then the divergence is limited to a maximum of 1/2. Recall that $\mathcal{A}^{\mathcal{L}}$ is the set of all arguments that can be formed from language $\mathcal{L}$.

Proposition 22. *Let P and P' be consistent probability distributions on $\mathcal{L}$. Let E be the epistemic extension of $(\mathcal{A}^{\mathcal{L}}, \mathcal{R}, P)$, and let E' be the epistemic extension of $(\mathcal{A}^{\mathcal{L}}, \mathcal{R}, P')$. If $E = E'$, then $|P - P'| \leq 0.5$.*

Proof. Assume $E = E'$. So for all $A \in \mathcal{A}^{\mathcal{L}}$, $A \in E$ iff $A' \in E'$. So for all $A \in \mathcal{A}^{\mathcal{L}}$, $P(A) > 0.5$ iff $P'(A) > 0.5$. So for all $\phi \in \mathcal{L}$, $P(\phi) > 0.5$ iff $P'(\phi) > 0.5$. So for all $m_1, .., m_i \in \mathcal{M}^{\mathcal{L}}$, $P(m_1) + ... + P(m_i) > 0.5$ iff $P'(m_1) + ... + P'(m_i) > 0.5$. Since P and P' are each consistent, we have that for all $m_1, .., m_i \in \mathsf{Models}(\phi)$,

$$|(P(m_1) + ... + P(m_i)) - (P'(m_1) + ... + P'(m_i))| \leq 1$$

Hence, $\sum_{m \in \mathcal{M}^{\mathcal{L}}} |P(m) - P'(m)| \leq 1$. Therefore, $|P - P'| \leq 0.5$. □

Example 45. *Let the signature $\mathcal{S}^{\mathcal{L}}$ be (p, q). So $\mathcal{M}^{\mathcal{L}}$ is $\{11, 10, 01, 00\}$. Suppose $\mathcal{R}$ is the empty set. Also suppose we have the following consistent probability functions, and so $|P - P'|$ is 1/3.*

$$\begin{array}{lll} P(10) = 1/3 & P(11) = 1/3 & P(01) = 1/3 \\ P'(10) = 1/4 & P'(11) = 1/2 & P'(01) = 1/4 \end{array}$$

The epistemic extension of both $(\mathcal{A}, \mathcal{R}, P)$ and $(\mathcal{A}, \mathcal{R}, P')$ includes the following arguments

$$\langle\{p \vee q\}, p \vee q\rangle, \langle\{\neg p \vee \neg q\}, p \vee q\rangle, \langle\{p\}, p\rangle, \langle\{q\}, q\rangle$$

Whereas the epistemic extension of both $(\mathcal{A}, \mathcal{R}, P)$ and $(\mathcal{A}, \mathcal{R}, P')$ excludes the following arguments

$$\langle\{p \wedge q\}, p \wedge q\rangle, \langle\{\neg p \wedge \neg q\}, p \wedge q\rangle, \langle\{\neg p\}, \neg p\rangle, \langle\{\neg q\}, \neg q\rangle$$

However, for the probabilistic argument graphs where the arguments are a subset of $\mathcal{A}^{\mathcal{L}}$, then they are not guaranteed to have limited divergence when they have the same epistemic extension. We illustrate this in the next example.

Example 46. *Let the signature $\mathcal{S}^{\mathcal{L}}$ be (p, q). So $\mathcal{M}^{\mathcal{L}}$ is $\{11, 10, 01, 00\}$. Consider $A_1 = \langle\{p \vee q\}, p \vee q\rangle$. Suppose $\mathcal{A}$ contains just A_1 and $\mathcal{R}$ is the empty set. Also suppose $P(10) = 1$, and $P'(01) = 1$. Therefore the epistemic extension of both $(\mathcal{A}, \mathcal{R}, P)$ and $(\mathcal{A}, \mathcal{R}, P')$ is $\{A_1\}$. However, $|P - P'|$ is 1.*

For the probabilistic argument graphs $(\mathcal{A}^{\mathcal{L}}, \mathcal{R}, P)$ and $(\mathcal{A}^{\mathcal{L}}, \mathcal{R}, P')$, if they have the same probabilistic assignment for each extension, then they agree on the arguments.

Proposition 23. *Let P_i and P_j be consistent probability distributions on $\mathcal{L}$. For any semantics $X \in \{\mathsf{ad}, \mathsf{co}, \mathsf{pr}, \mathsf{gr}, \mathsf{st}\}$,*

$$\text{for all } E \subseteq \mathcal{A}, P_i(E^X) = P_j(E^X) \text{ iff for all } A \in \mathcal{A}, P_i(A) = P_j(A)$$

Proof. Let $G = (\mathcal{A}, \mathcal{R})$. (<=) Assume that for all $A \in \mathcal{A}$, $P_i(A) = P_j(A)$. Therefore, for all $G' \sqsubseteq G$, $P_i(G') = P_j(G')$. Therefore, for all $E \subseteq \mathcal{A}$, $P_i(E^x) = P_j(E^x)$. (=>) For a proof by contradiction, assume that it is not the case that $P_i(A) = P_j(A)$ for all $A \in \mathcal{A}$. Therefore, there is an argument $A \in \mathcal{A}$ such that $P_i(A) \neq P_j(A)$. Suppose, without loss of generality, $P_i(A) > P_j(A)$. Therefore, for each graph G' containing A, $P_i(G') > P_j(G')$, and for each graph G'' not containing A, $P_i(G'') < P_j(G'')$. For each $E \subseteq \mathcal{A}$, $P_i(E^X) = \sum_{G' \in Q_x(E)} P_i(G')$ and $P_j(E^X) = \sum_{G' \in Q_x(E)} P_j(G')$. For each $E \subseteq \mathcal{A}$, if $A \in E$, then $P_i(E^X) > P_j(E^X)$ and if $A \notin E$, then $P_i(E^X) < P_j(E^X)$. Therefore, it is not the case that for all $E \subseteq \mathcal{A}$, $P_i(E^X) = P_j(E^X)$. □

So if there is divergence between agents, then there does not necessarily mean that this is manifested in differences in the probabilities over arguments and/or extensions. As the above result shows, they can behave in the same way with respect to arguments and extensions, and yet be somewhat divergent.

7.3 Combined probability distributions

Given two conflicting probability assignments, we can form a merged probability assignment as follows. We use the max function to identify the maximum conflict between agents regarding each argument.

Definition 27. *Let P_i and P_j be probability functions over $\mathcal{A}^{\mathcal{L}}$. The combined probability function over $\mathcal{A}^{\mathcal{L}}$ is defined as follows, where* max *returns the larger of $P_i(A)$ and $P_j(A)$.*

$$P_{i \oplus j}(A) = \mathsf{max}(P_i(A), P_j(A))$$

Obviously, the combination is commutative (i.e. $P_{i \oplus j}(A) = P_{j \oplus i}(A)$), associative (i.e. $P_{i \oplus (j \oplus k)}(A) = P_{(i \oplus j) \oplus k}(A)$), and idempotent (i.e. $P_{i \oplus i)}(A) = P_i(A)$).

Example 47. *Let the signature $\mathcal{S}^{\mathcal{L}}$ be (p, q). So the models of $\mathcal{M}^{\mathcal{L}}$ are $\{11, 10, 01, 00\}$. Consider the following argument graph. Suppose agent 1 has $P_1(11) = 1$ and agent 2 has $P_2(00) = P_2(10) = 0.5$. So agent 1 has $P_1(A_1) = 1$, and agent 2 has $P_2(A_2) = 1$. Therefore, the combined probability distribution is $P_{1 \oplus 2}(A_1) = 1$ and $P_{1 \oplus 2}(A_2) = 1$.*

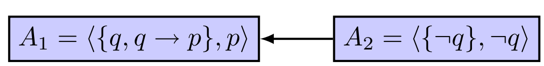

Example 48. *Suppose $P_1(1101) = 1$ and $P_2(1011) = P_2(0011) = P_2(1010) = P_2(0010) = 1/4$. So $P_1(A_1) = 1$, $P_1(A_2) = 0$, $P_1(A_3) = 1$, $P_2(A_1) = 0$, $P_2(A_2) = 1$ and $P_2(A_3) = 0$. Therefore, $P_{1 \oplus 2}(A_1) = 1$, $P_{1 \oplus 2}(A_2) = 1$, and $P_{1 \oplus 2}(A_3) = 1$.*

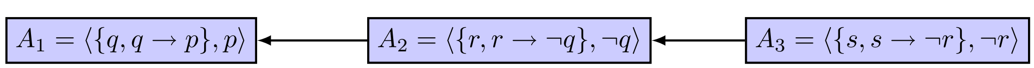

Even if the agents have consistent probability distributions on $\mathcal{L}$, the combined distribution will be not consistent on $\mathcal{L}$ when they diverge.

Proposition 24. *Let P_i and P_j be consistent probability distributions on $\mathcal{L}$. $P_{i \oplus j}$ is consistent on $\mathcal{L}$ iff $|P_i - P_j| = 0$*

Proof. Let P_i and P_j be consistent probability distributions on $\mathcal{L}$. Since $P_{i \oplus j}(A) = \mathsf{max}(P_i(A), P_j(A))$, $P_{i \oplus j}$ is not consistent iff $\sum_{m \in \mathcal{M}^{\mathcal{L}}} P_{i \oplus j}(m) > 1$ iff there is an argument $A \in \mathcal{A}^{\mathcal{L}}$ where $P_i(A) > P_j(A)$ and there is an argument $A' \in \mathcal{A}^{\mathcal{L}}$ where $P_i(A') < P_j(A')$ iff there is an argument $A \in \mathcal{A}^{\mathcal{L}}$ where $\mathsf{Models}(A) = \{m_1, .., m_k\}$ and $P_i(m_i) + ... + P_i(m_k) > P_j(m_i) + ... + P_j(m_k)$, and there is an argument $A' \in \mathcal{A}^{\mathcal{L}}$ where $\mathsf{Models}(A') = \{m'_1, .., m'_k\}$ and $P_i(m'_i) + ... + P_i(m'_k) < P_j(m'_i) + ... + P_j(m'_k)$, iff $|P_i - P_j| > 0$. So we have shown that $P_{i \oplus j}$ is not consistent iff iff $|P_i - P_j| > 0$. Therefore, $P_{i \oplus j}$ is consistent iff iff $|P_i - P_j| = 0$. □

Normally $P_{1\oplus 2}$ is not rational. This is because the combined probability reflects the conflict between the two with respect to the arguments. A solution is to use the constellations approach to evaluate the resulting probabilistic argument graph $(\mathcal{A}, \mathcal{R}, P_{1\oplus 2})$, as illustrated by the following example.

Example 49. *Consider the following argument graph. Let the signature* $\mathcal{S}^{\mathcal{L}}$ *be* (p, q, r, s)*. Suppose,* $P_1(1010) = 0.2$*,* $P_1(1110) = 0.8$*,* $P_2(1001) = 0.8$*, and* $P_2(1101) = 0.2$*. Therefore,* $P_1(A_1) = 0.8$*,* $P_1(A_2) = 0.2$ *and* $P_1(A_3) = 0$*. Also,* $P_2(A_1) = 0.2$*,* $P_2(A_2) = 0$ *and* $P_2(A_3) = 0.8$*.*

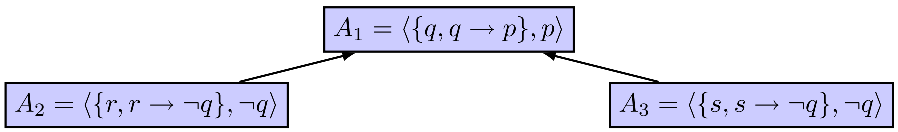

The combined probability distribution is $P_{1\oplus 2}(A_1) = 0.8$*,* $P_{1\oplus 2}(A_2) = 0.2$*, and* $P_{1\oplus 2}(A_3) = 0.8$*. This is rational but not consistent. If we consider the constellations approach, then we get the extensions* $\{A_1\}$ *with probability 0.128,* $\{A_2\}$ *with probability 0.04,* $\{A_3\}$ *with probability 0.64,* $\{A_2, A_3\}$ *with probability 0.16, and* $\emptyset$ *with probability 0.032.*

The combined probability distribution reflects the conflicting views of the two agents. Normally, the combined probability distribution is not consistent. Also, normally, the probabilistic argument graph with a combined probability function is not cohesive. However, if the probabilistic argument graph is bellicose, then using Proposition 17, does mean that the representative probability distribution can be used to give the same epistemic extension.

Acknowledgements

The author is very grateful to the anonymous reviewers for their valuable comments for improving the paper.

8 Discussion

For formal modelling of argumentation, logical inconsistency in a set of beliefs $K \subseteq \mathcal{L}$ is central issue. Logical inconsistency means that there are conflicts between arguments based on the beliefs. If we have no inconsistency, then there are no conflicts. So in a sense, information without logical inconsistency is better than information with logical inconsistency. When we augment our logical argumentation with the probabilistic approach presented in this paper, we have a finer grained differentiation of our beliefs. Obviously, if we have a consistent probability distribution over $\mathcal{L}$, it does not mean that our knowledgebase K is consistent. However, when K is logically inconsistent, then a consisistent probability distribution $\mathcal{L}$ is better than an inconsistent probability distribution over $\mathcal{L}$. So we can rank these scenarios from best to worst, from the point of view of "information quality".

- $K \subseteq \mathcal{L}$ is logically consistent and the probability distribution over $\mathcal{L}$ is consistent.
- $K \subseteq \mathcal{L}$ is logically inconsistent and the probability distribution over $\mathcal{L}$ is consistent.
- $K \subseteq \mathcal{L}$ is logically inconsistent and the probability distribution over $\mathcal{L}$ is inconsistent.
- $K \subseteq \mathcal{L}$ is logically inconsistent and there is no probability distribution over $\mathcal{L}$.

In this paper, we have seen how we can start with a probability distribution over a language $\mathcal{L}$, or equivalently over the set of models of the language $\mathcal{M}^{\mathcal{L}}$, and use this to generate a probability distribution over the logical arguments $\mathcal{A}^{\mathcal{L}}$ based on that language. This gives a precise meaning to probabilistic arguments and probabilistic argument graphs. We have introduced the epistemic approach on probabilistic argument graphs, and we have shown how the probability distribution on a language allows for a clearer understanding of the nature of the epistemic and constellations approach on probabilistic argument graphs. We summarise the key results of this framework as follow:

- If P is consistent on $\mathcal{L}$, then P is coherent on $\mathcal{A} \subseteq \mathcal{A}^{\mathcal{L}}$.
- P is consistent on $\mathcal{L}$ iff P is maximally coherent on $\mathcal{A}^{\mathcal{L}}$.
- If P is consistent on $\mathcal{L}$, then P is rational for $(\mathcal{A}, \mathcal{R})$ where $\mathcal{A} \subseteq \mathcal{A}^{\mathcal{L}}$.
- If P is consistent on $\mathcal{L}$, then P is cohesive for $(\mathcal{A}, \mathcal{R})$ where $\mathcal{A} \subseteq \mathcal{A}^{\mathcal{L}}$.
- If P is cohesive for $(\mathcal{A}, \mathcal{R})$, then P is rational for $(\mathcal{A}, \mathcal{R})$.
- If P is strongly cohesive for $(\mathcal{A}^{\mathcal{L}}, \mathcal{R})$, then P is consistent on $\mathcal{L}$.
- If P' is a discounted probability distribution for $(\mathcal{A}, \mathcal{R}, P)$, then P' is consistent on $\mathcal{L}$.
- If P' is a representative probability distribution for $(\mathcal{A}, \mathcal{R}, P)$, and P is cohesive for$(\mathcal{A}, \mathcal{R})$, then P' is consistent on $\mathcal{L}$.
- If P' is a representative probability distribution for $(\mathcal{A}, \mathcal{R}, P)$, and P is cohesive and rational for $(\mathcal{A}, \mathcal{R})$, then $(\mathcal{A}, \mathcal{R}, P)$ and $(\mathcal{A}, \mathcal{R}, P')$ have the same epistemic extension.

We do not necessarily start with a probability distribution over a language $\mathcal{L}$. We may start with a probability distribution over a set of arguments $\mathcal{A} \subseteq \mathcal{A}^{\mathcal{L}}$. In this case, it is possible to make a probability assignment to the arguments, and this may mean that the corresponding probability distribution over $\mathcal{L}$ is inconsistent. We have investigated various kinds of inconsistency that arise when obtaining that distribution via a probability distribution over arguments. Inconsistency can come when the distribution is incoherent, so when the distribution does not respect the structure of the argument graph (i.e. when the sum of the probability of an argument and an attacker of it is greater than 1). This can arise when an agent is not sure what probabilities can be assigned to arguments, or when two agents disagree over which arguments to believe. We have provided various ways of dealing with inconsistent probability distributions.

As we suggested in Section 1, the framework in this paper is novel. There have been previous proposals for introducing probabilities into abstract argumentation and into rule-based argumentation, but there has been no approach to consider how a probability distribution over models relates to a probability distribution over arguments in instantiations of abstract argumentation. Furthermore, this paper provides a comprehensive analysis of different types of inconsistency that can naturally arise in probability distributions in the context of argumentation. This paper therefore provides a clearer and more robust framework for modelling uncertainty in arguments and counterarguments.

The proposal in this paper fits with the proposal to extend abstract argumentation with probabilities by Li et al [LON11]. In [LON11], probability assignments are given to both arguments and attacks. This probabilistic information is then used to generate all possible argument graphs that contain at least all the arguments and attacks with probability 1. However, the proposal by [LON11] does not consider how these arguments relate to logical arguments, nor how the probabilities over the arguments could be generated from probabilistic information about the premises of the arguments. Furthermore, in this paper in Section 3.1, we have extended the proposal by Li et al [LON11], by introducing epistemic extensions, and rational probability functions, which are useful for considering consistent probability distributions, and for modelling a form of rational analysis of inconsistent information.

This paper also complements the ABEL framework where reasoning with propositional information is augmented with probabilistic information so that individual arguments are qualified by a probability value [Hae98, HKL00]. However, there is no consideration in ABEL of how this probabilistic information relates to Dung's proposals, or how it could be used to decide which arguments are acceptable according to Dung's dialectical semantics. The emphasis is on generating pros and cons for diagnosis, and there is a particular language introduced for this purpose. Nonetheless, there is consideration of computational issues of how probabilities can be computed which may be useful for implementing the proposal in this paper.

Probabilistic reasoning with logical statements has also been considered by Pollock [Pol95]. However, the approach taken is to assign probabilities to formulae without considering the meaning of this in terms of models. Various issues arising from an assignment based on frequency that a consequent holds when the antecedent holds are considered, as well as how such an assignment could be used for statistical syllogism. The emphasis of the work is therefore different as it does not consider what would be acceptable probability assignments for a language, and it does not consider how a probabilistic perspective relates to abstract argumentation.

In addition to the other proposals reviewed in Section 1, probabilities have also been proposed in the qualitative probabilistic reasoning (QPR) framework, which is a qualitative form of Bayesian network, and whilst this can be regarded as a form of argumentation, it does not incorporate the dialectical reasoning seen with the generation of arguments and counterarguments, and so again there is no consideration of how this probabilistic information relates to Dung's proposals [Par96, Par98, Par01]. In another approach based on Bayesian networks, defeasible reasoning is simulated with such a network, so that conditional probabilities represent defeasible rules, and the probabilities are used to decide which inference is propagated. By harnessing existing software for Bayesian networks, it has been shown to be a computationally viable approach, though again there is no consideration of how this probabilistic information relates to Dung's proposals [Vre05]. Finally, argumentation has also been used for merging multiple Bayesian networks, but then the Bayesian networks are the subject of argumentation, rather than probabilistic information being used to quantify the uncertainty of arguments in general [NP06].

9 Future work

In the proposal in this paper, we have used a very simple representation of uncertainty. By starting with a probability distribution over the models of the language, the probability of an argument is the sum of the probability assigned to the models of the premises of the argument. Whilst this has brought benefits to our model of argumentation, we may wish to bring a richer formulism of uncertainty into argumentation, such as a probabilistic logic that allows us to reason in the object language about the probabilities. There have been a number of proposals for probabilistic logics including [Nil86, Bac90, Hal90]. Here we focus on the approach in [Hal90]. Consider the following two statements. The first is a probabilistic belief, whereas the second is a statistical assertion.

- The probability that tweety (a particular bird) flies is greater than 0.9.
- The probability that a randomly chosen flies is greater than 0.9.

These two kinds of statement require different approaches to their semantics and proof theory. For a probabilistic belief, possible worlds (a generalization of the semantics for classical logic) can be used to capture the semantics of a statistical assertion. Each possible world can be regarded as a classical interepretation. A possible worlds interpretation is then a multiset of classical interpretations. Each world is a possibility. If the proportion of worlds in the possible worlds interpretation, where $\texttt{flies}(\texttt{tweety})$ is true is greater than 0.9 implies $\texttt{P}(\texttt{flies}(\texttt{tweety})) > 0.9$ is true.

The second statement, (i.e. "The probability that a randomly chosen bird will fly is greater than 0.9.") represents a chance set-up. It is the result of doing some experiment or trial, and so it is a statistical statement. We may wish to represent this as follows.

$$\texttt{P}(\forall \texttt{X}\ \texttt{bird}(\texttt{X}) \rightarrow \texttt{fly}(\texttt{X})) > 0.9$$

However, the possible worlds approach is not adequate for a semantics for the above kind of statement, since in each possible world there may be a bird that does not fly. A better alternative is to use a single classical world plus a probability distribution over birds. If we randomly choose a bird from the domain of the model, then the probability of the relation fly holding for the randomly chosen bird is greater than 0.9. Now we can extend classical predicate logic with the new syntax for formulae, where $\texttt{X}$ is a randomly chosen element of the domain, and the formula is a statement that is either true or false in this extended classical model.

$$\texttt{P}_{\texttt{X}}(\texttt{fly}(\texttt{X}) \mid \texttt{bird}(\texttt{X})) > 0.9$$

This kind of formula can be used in more complicated classical formulae in the usual way. As illustrated by the following examples.

- $\texttt{P}_{\texttt{X}}(\texttt{son}(\texttt{X},\texttt{Y}))$ which denotes the probability that a randomly chosen $\texttt{X}$ is the son of $\texttt{Y}$.

- $\mathtt{P_Y(son(X,Y))}$ which denotes the probability that X is the son of a randomly chosen Y.
- $\mathtt{P_{X,Y}(son(X,Y))}$ which denotes the probability that a randomly chosen pair (X, Y) will have the property that X is the son of Y.
- $\mathtt{P(\exists X.malfunction(car,X)) > 0.5}$ which denotes the car probably has some type of malfunction.
- $\mathtt{P(malfunction(car,electrical)) > 2.P(malfunction(car,carburettor))}$ which denotes that it is more than twice as likely that the car malfunction is electrical rather than the carburettor.

So in probabilistic logic, first-order classical logic is extended with syntax and semantics to represent and reason with both beliefs and statistical assertions in the object language. Harnessing this in a model of argumentation will enable richer statements about uncertainty of premises and claims to be captured. We leave it to future work to extend the proposal in this paper to using probabilistic logics.

Another area for future work is investigating how to manipulate probability distributions over models, or over arguments. For example, if we want a specific rational extension, how do we manipulate the probabilities to get it? This may have implications for strategies for persuasion. As another example, how robust is a rational extension to changes in the probability distribution over models? If it is very sensitive, then perhaps we would be less confident in the extension than if it remained unchanged for all but large swings in the probability assignment.

One of the key motivations for developing a probabilistic approach to argumentation is handle enthymemes. Most arguments in natural language are enthymemes, which means that they do not explicitly present all their premises and/or claims. For example, for a claim that you need an umbrella today, a husband may give his wife the premise the weather report predicts rain. Clearly, the premise does not entail the claim, but it is easy for the wife to identify the common knowledge used by the husband in order to reconstruct the intended argument correctly. Abduction is used to determine what common knowledge should be used to reconstruct the intended argument, and often there are multiple possibilities for what to abduce.

With the incompleteness inherent in enthymemes, it is difficult to be certain whether one argument attacks another. If a counterargument has an explicit claim, there may be uncertainty as to whether the attacked argument has the premise that the attacker has contradicted. And if a counterargument has an implicit claim, there may be further uncertainty as to what is being contradicted. To illustrate, we consider an example from Sperber and Wilson [SW95]. Suppose in response to an offer of coffee by a host of a dinner party, a guest says "Coffee would keep me awake", the process of deconstructing this enthymeme is problematical. The claim of the argument by the guest is either "Yes, I would like some coffee" or "No, I would not like some coffee". If the former is the claim, then the premises of the argument might be "I need to drive home", "Coffee will keep me awake", and "If I need to drive home, and coffee will keep me awake, then yes I would like some coffee". Whereas if the later is the claim, then the premises of the argument might be "I need to sleep", "Coffee will keep me awake", and "If I need to sleep, and coffee will keep me awake, then no I would not like some coffee".

Whilst humans are constantly handling enthymemes, the logical formalization that characterizes the process remains underdeveloped. Therefore, we need to investigate enthymemes because of their ubiquity in the real world, and because of the difficulties they raise for formalizing and automating argumentation. In [Hun07, BH12], proposals have been made for how common knowledge can be used to construct enthymemes from logical arguments (so that a proponent can send an enthymeme to an intended recipient by removing common knowledge) and deconstruct enthymemes (so that a recipient can rebuild the intended logical argument from the enthymeme by reintroducing the common knowledge). The proposal is based on logical argumentation, with the deconstruction process being based on abduction. In future work, it would be desirable to enhance these proposals with the framework for probabilistic qualification of uncertainty in arguments as presented in this paper.

Another application area for probabilistic arguments is in decision making. For this, we could investigate extending the approach of decision theoretic argumentation of Amgoud and Prade [AP09] with implicit and explicit uncertainty, with the problem of generating intentions in the context of uncertainty [ADM11], and/or the optimization of decision making using assumption-based argumentation of Matt *et al* [MTV10].

Finally, developing an implementation of the proposal for using logical arguments in probabilistic argument graphs would be valuable future work. We have already developed a viable system for argumentation with classical logic [EH11]. A number of proposals have been made for algorithms for abstract argumentation (e.g. [BG02, CDM03, DMT07]). A good starting point would be the ASP-based ASPARTIX system [EGW08] as this would appear to support development of algorithms as ASP programs, which in turn would appear appropriate for also reasoning with the full subgraphs. Furthermore, given a probabilistic function over a language $\mathcal{L}$, the problem of determining whether it is consistent is the PSAT problem [Boo54]. This problem can be viably addressed by algorithms based on SAT solvers [FB11]. Hence, given a belief function over some set of formulae, it can be determined whether it corresponds to a consistent probabilistic function. In addition, algorithms for generating probabilistic arguments have been proposed for the ABEL system [HKL00, Hae01], and these could be adapted as an alternative approach to implementing the framework in this paper.

References

[AB09] L. Amgoud and Ph Besnard. Bridging the gap between abstract argumentation systems and logic. In *Proceedings of the Third International Conference on Scalable Uncertainty Management (SUM'09)*, volume 5785 of *Lecture Notes in Computer Science*, pages 12–27. Springer, 2009.

[AC98] L. Amgoud and C. Cayrol. On the acceptability of arguments in preference-based argumentation. In G. Cooper and S. Moral, editors, *Proceedings of the 14th Conference on Uncertainty in Artificial Intelligence (UAI 1998)*, pages 1–7. Morgan Kaufmann, 1998.

[AC02] L. Amgoud and C. Cayrol. A reasoning model based on the production of acceptable arguments. *Annals of Mathematics and Artificial Intelligence*, 34:197–216, 2002.

[ACGS08a] T. Alsinet, C.Chesñevar, L. Godo, and G. Simari. Formalizing argumentative reasoning in a possibilistic logic programming setting with fuzzy unification. *International Journal of Approximate Reasoning*, 48:711–729, 2008.

[ACGS08b] T. Alsinet, C. Chesñevar, L. Godo, and G. Simari. A logic programming framework for possibilistic argumentation: Formalization and logical properties. *Fuzzy Sets and Systems*, 159(10):1208–1228, 2008.

[Ada98] E. Adams. *A Primer of Probability Logic*. CSLI Publications (Chicago University Press), 1998.

[ADM11] L. Amgoud, C. Devred, and M-C Lagasquie-Schiex. Generating possible intentions with constrained argumentation systems. *International Journal of Approximate Reasoning*, 52(9):1363–1391, 2011.

[AP04] L. Amgoud and H. Prade. Reaching agreement through argumentation: A possibilistic approach. In *Proceedings of the 9th International Conference on Principles of Knowledge Representation ans Reasonning (KR'04)*, pages 175–182. AAAI Press, 2004.

[AP09] L. Amgoud and H. Prade. Using arguments for making and explaining decisions. *Artificial Intelligence*, 173:413–436, 2009.

[Bac90] F. Bacchus. *Representing and Reasoning with Probabilisitic Knowledge: A Logical Approach to Probabilities*. MIT Press, 1990.

[BCD07] T. Bench-Capon and P. Dunne. Argumentation in artificial intelligence. *Artificial Intelligence*, 171(10-15):619–641, 2007.

[BG02] P. Baroni and M. Giacomin. Argumentation through a distributed self-stabilizing approach. *Journal of Experimental and Theoretical Artificial Intelligence*, 14(4):273–301, 2002.

[BGW05] H. Barringer, D. Gabbay, and J. Woods. Temporal dynamics of support and attack networks: From argumentation to zoology. In *Mechanizing Mathematical Reasoning*, volume 2605 of *Lecture Notes in Computer Science*, pages 59–98. Springer, 2005.

[BH01] Ph. Besnard and A. Hunter. A logic-based theory of deductive arguments. *Artificial Intelligence*, 128:203–235, 2001.

[BH08] Ph. Besnard and A. Hunter. *Elements of Argumentation*. MIT Press, 2008.

[BH12] E. Black and A. Hunter. A relevance-theoretic framework for constructing and deconstructing enthymemes. *Journal of Logic and Computation.*, 22(1):55–78, 2012. in press.

[Boo54] G. Boole. *An Investigation on the Laws of Thought*. Macmillan, 1854. (available at www.gutenberg.org).

[CA05] M. Caminada and L. Amgoud. An axiomatic account of formal argumentation. In *Proceedings of the 20th National Conference on Artificial Intelligence (AAAI'05)*, pages 608–613. AAAI Press, 2005.

[Cam06] M. Caminada. Semi-stable semantics. In *Computational Models of Argument (COMMA'06)*, pages 121–130, 2006.

[Car50] R. Carnap. *Logical Foundations of Probability*. University of Chicago Press, 1950.

[Cay95] C. Cayrol. On the relationship between argumentation and non-monotonic coherence-based entailment. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI'95)*, pages 1443–1448, 1995.

[CDM03] C. Cayrol, S. Doutre, and J. Mengin. On the decision problems related to the preferred semantics for argumentation frameworks. *Journal of Logic and Computation*, 13(3):377–403, 2003.

[DMT07] P. Dung, P. Mancarella, and F. Toni. Computing ideal skeptical argumentation. *Artificial Intelligence*, 171:642–674, 2007.

[DT10] P. Dung and P. Thang. Towards (probabilistic) argumentation for jury-based dispute resolution. In *Computational Models of Argument (COMMA'10)*, pages 171–182. IOS Press, 2010.

[Dun95] P. 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.

[EGH95] M. Elvang-Gøransson and A. Hunter. Argumentative logics: Reasoning with classically inconsistent information. *Data & Knowledge Engineering*, 16(2):125–145, 1995.

[EGKF93] M. Elvang-Gøransson, P. Krause, and J. Fox. Acceptability of arguments as 'logical uncertainty'. In *Proceedings of the European Conference on Symbolic and Quantitative Approaches to Reasoning and Uncertainty (ECSQARU'93)*, pages 85–90. Springer-Verlag, 1993.

[EGW08] U. Egly, S. Gaggl, and S. Woltran. Aspartix: Implementing argumentation frameworks using answer-set programming. In *Proceedings of the Twenty-Fourth International Conference on Logic Programming (ICLP'08),*, volume 5366 of *LNCS*, pages 734–738. Springer, 2008.

[EH11] V. Efstathiou and A. Hunter. Algorithms for generating arguments and counterarguments in propositional logic. *International Journal of Approximate Reasoning*, 52:672–704, 2011.

[FB11] M. Finger and G. De Bona. Probabilistic satisfiability: Logic-based algorithms and phase transitions. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI'11)*, 2011.

[FD00] J. Fox and S. Das. *Safe and Sound: Artificial Intelligence in Hazardous Applications*. MIT Press, 2000.

[Gab11] D. Gabbay. Introducing equational semantics for argumentation networks. In *Symbolic and Quantitative Approaches to Reasoning with Uncertainty (ECSQARU'11)*, volume 6717 of *LNCS*, pages 19–35. Springer, 2011.

[GH11] N. Gorogiannis and A. Hunter. Instantiating abstract argumentation with classical logic arguments: Postulates and properties. *Artificial Intelligence*, 175(9-10):1479–1497, 2011.

[Hac01] I. Hacking. *An Introduction to Probability and Inductive Logic*. Cambridge University Press, 2001.

[Hae98] R. Haenni. Modelling uncertainty with propositional assumptions-based systems. In A. Hunter and S. Parsons, editors, *Applications of Uncertainty Formalisms*, pages 446–470. Springer, 1998.

[Hae01] R. Haenni. Cost-bounded argumentation. *International Journal of Approximate Reasoning*, 26(2):101–127, 2001.

[Hal90] J. Halpern. An analysis of first-order logics of probability. *Artificial Intelligence*, 46:311–350, 1990.

[HKL00] R. Haenni, J. Kohlas, and N. Lehmann. Probabilistic argumentation systems. In D. Gabbay and Ph Smets, editors, *Handbook of Defeasible Reasoning and Uncertainty Management Systems, Volume 5*, pages 221–288. Kluwer, 2000.

[Hun07] A. Hunter. Real arguments are approximate arguments. In *Proceedings of the Twenty-Second AAAI Conference on Artificial Intelligence (AAAI'07)*, 2007.

[Hun12] A. Hunter. Some foundations for probabilistic argumentation. In *Proceedings of the International Conference on Computational Models of Argument (COMMA'12)*, 2012.

[LON11] H. Li, N. Oren, and T. Norman. Probabilistic argumentation frameworks. In *Proceedings of the First International Workshop on the Theory and Applications of Formal Argumentation (TAFA'11)*, 2011.

[MTV10] P. Matt, F. Toni, and J Vaccari. Dominant decisions by argumentation agents. In *Argumentation in Multi-Agent Systems (ArgMAS'09)*, volume 6057 of *LNCS*, pages 42–59. Springer, 2010.

[Nil86] N. Nilsson. Probabilistic logic. *Artificial Intelligence*, 28:71–87, 1986.

[NP06] S. Neilson and S. Parsons. An application of formal argumentation: Fusing Bayes nets in MAS. In *Computational Models of Argument (COMMA'06)*, pages 33–46. IOS Press, 2006.

[Par94] J. Paris. *The Uncertain Reasoner's Companion: A Methematical Perspective*. Cambridge University Press, 1994.

[Par96] S. Parsons. Defining normative systems for qualitative argumentation. In *Proceedings of the International Conference on Formal and Applied Practical Reasoning*, volume 1085 of *LNCS*, pages 449–463, 1996.

[Par98] S. Parsons. A proof theoretic approach to qualitative probabilistic reasoning. *International Journal of Approximate Reasoning*, 19:265–297, 1998.

[Par01] S. Parsons. *Qualitative Approaches to Reasoning Under Uncertainty*. MIT Press, 2001.

[Pol87] J. Pollock. Defeasible reasoning. *Cognitive Science*, 11(4):481–518, 1987.

[Pol89] J. Pollock. *How to Build a Person*. MIT Press, 1989.

[Pol92] J. Pollock. How to reason defeasibly. *Artificial Intelligence*, 57(1):1–42, 1992.

[Pol95] J. Pollock. *Cognitive Carpentary*. MIT Press, 1995.

[RRS+07] R. Riveret, A. Rotolo, G. Sartor, H. Prakken, and B. Roth. Success chances in argument games: a probabilistic approach to legal disputes. In *Legal Knowledge and Information Systems (JURIX'07)*, pages 99–108. IOS Press, 2007.

[RS09] I. Rahwan and G. Simari, editors. *Argumentation in Artifical Intelligence*. Springer, 2009.

[SW95] D. Sperber and D. Wilson. *Relevance: Communication and Cognition*. Blackwell, 2nd edition, 1995.

[Vre05] G. Vreeswijk. Argumentation in Bayesian belief networks. In *Argumentation in Multi-agent Systems (ArgMAS'04)*, volume 3366 of *LNAI*. Springer, 2005.

[Wil02] J. Williamson. Probability logic. In D.Gabbay, R. Johnson, and J. Woods, editors, *Handbook of the Logic of Argumentation and Inference*, pages 397–424. Elsevier, 2002.