

# Belief Revision in Structured Probabilistic Argumentation

Paulo Shakarian<sup>1</sup>, Gerardo I. Simari<sup>2</sup>, and Marcelo A. Falappa<sup>3</sup>

<sup>1</sup> Department of Electrical Engineering and Computer Science  
U.S. Military Academy, West Point, NY, USA  
[paulo@shakarian.net](mailto:paulo@shakarian.net)

<sup>2</sup> Department of Computer Science, University of Oxford, United Kingdom  
[gerardo.simari@cs.ox.ac.uk](mailto:gerardo.simari@cs.ox.ac.uk)

<sup>3</sup> Departamento de Ciencias e Ingeniería de la Computación  
Universidad Nacional del Sur, Bahía Blanca, Argentina  
[mfalappa@cs.uns.edu.ar](mailto:mfalappa@cs.uns.edu.ar)

**Abstract.** In real-world applications, knowledge bases consisting of all the information at hand for a specific domain, along with the current state of affairs, are bound to contain contradictory data coming from different sources, as well as data with varying degrees of uncertainty attached. Likewise, an important aspect of the effort associated with maintaining knowledge bases is deciding what information is no longer useful; pieces of information (such as intelligence reports) may be outdated, may come from sources that have recently been discovered to be of low quality, or abundant evidence may be available that contradicts them. In this paper, we propose a probabilistic structured argumentation framework that arises from the extension of Presumptive Defeasible Logic Programming (PreDeLP) with probabilistic models, and argue that this formalism is capable of addressing the basic issues of handling contradictory and uncertain data. Then, to address the last issue, we focus on the study of non-prioritized belief revision operations over probabilistic PreDeLP programs. We propose a set of rationality postulates – based on well-known ones developed for classical knowledge bases – that characterize how such operations should behave, and study a class of operators along with theoretical relationships with the proposed postulates, including a representation theorem stating the equivalence between this class and the class of operators characterized by the postulates.

## 1 Introduction and Related Work

Decision-support systems that are part of virtually any kind of real-world application must be part of a framework that is rich enough to deal with several basic problems: (i) handling contradictory information; (ii) answering abductive queries; (iii) managing uncertainty; and (iv) updating beliefs. *Presumptions* come into play as key components of answers to abductive queries, and must be maintained as elements of the knowledge base; therefore, whenever candidate answers to these queries are evaluated, the (in)consistency of the knowledge base

together with the presumptions being made needs to be addressed via belief revision operations.

In this paper, we begin by proposing a framework that addresses items (i)–(iii) by extending Presumptive DELP [1] (PreDeLP, for short) with probabilistic models in order to model uncertainty in the application domain; the resulting framework is a general-purpose probabilistic argumentation language that we will refer to as Probabilistic PreDeLP (P-PreDeLP, for short).

In the second part of this paper, we address the problem of updating beliefs – item (iv) above – in P-PreDeLP knowledge bases, focusing on the study of non-prioritized belief revision operations. We propose a set of rationality postulates characterizing how such operations should behave – these postulates are based on the well-known postulates proposed in [2] for non-prioritized belief revision in classical knowledge bases. We then study a class of operators and their theoretical relationships with the proposed postulates, concluding with a representation theorem.

**Related Work.** Belief revision studies changes to knowledge bases as a response to *epistemic inputs*. Traditionally, such knowledge bases can be either belief sets (sets of formulas closed under consequence) [3, 4] or belief bases [5, 2] (which are not closed); since our end goal is to apply the results we obtain to real-world domains, here we focus on belief bases. In particular, as motivated by requirements (i)–(iv) above, our knowledge bases consist of logical formulas over which we apply argumentation-based reasoning and to which we couple a probabilistic model. The connection between belief revision and argumentation was first studied in [6]; since then, the work that is most closely related to our approach is the development of the explanation-based operators of [7].

The study of argumentation systems together with probabilistic reasoning has recently received a lot of attention, though a significant part has been in the combination between the two has been in the form of probabilistic abstract argumentation [8–11]. There have, however, been several approaches that combine structured argumentation with models for reasoning under uncertainty; the first of such approaches to be proposed was [12], and several others followed, such as the possibilistic approach of [13], and the probabilistic logic-based approach of [14]. The main difference between these works and our own is that here we adopt a bipartite knowledge base, where one part models the knowledge that is not inherently probabilistic – uncertain knowledge is modeled separately, thus allowing a clear separation of interests between the two kinds of models. This approach is based on a similar one developed for ontological languages in the Semantic Web (see [15], and references within).

Finally, to the best of our knowledge, this is the first paper in which the combination of structured argumentation, probabilistic models, and belief revision has been addressed in conjunction.

Probabilistic Model (EM)	Analytical Model (AM)
“Malware X was compiled on a system using the English language.”	“Malware X was compiled on a system in English-speaking country Y.”
“Country Y and country Z are currently at war.”	“Country Y has a motive to launch a cyber-attack against country Z
“Malware W and malware X were created in a similar coding style.”	“Malware W and malware X are related.

**Table 1.** Examples of the kind of information that could be represented in the two different models in a cyber-security application domain.

## 2 Preliminaries

The Probabilistic PreDeLP (P-PreDeLP, for short) framework is composed of two separate models of the world. The first is called the *environmental model* (referred to as “EM”), and is used to describe the probabilistic knowledge that we have about the domain. The second one is called the *analytical model* (referred to as “AM”), and is used to analyze competing hypotheses that can account for a given phenomenon – what we will generally call queries. The AM is composed of a classical (that is, non-probabilistic) PreDeLP program in order to allow for contradictory information, giving the system the capability to model competing explanations for a given query.

**Two Kinds of Uncertainty.** In general, the EM contains knowledge such as evidence, uncertain facts, or knowledge about agents and systems. The AM, on the other hand, contains ideas that a user may conclude based on the information in the EM. Table 1 gives some examples of the types of information that could appear in each of the two models in a cyber-security application. Note that a knowledge engineer (or automated system) could assign a probability to statements in the EM column, whereas statements in the AM column can be either true or false depending on a certain combination (or several possible combinations) of statements from the EM. There are thus two kinds of uncertainty that need to be modeled: probabilistic uncertainty and uncertainty arising from defeasible knowledge. As we will see, our model allows both kinds of uncertainty to coexist, and also allows for the combination of the two since defeasible rules and presumptions (that is, defeasible facts) can also be annotated with probabilistic events.

In the rest of this section, we formally describe these two models, as well as how knowledge in the AM can be annotated with information from the EM – these annotations specify the conditions under which the various statements in the AM can potentially be true.

**Basic Language.** We assume sets of variable and constant symbols, denoted with  $\mathbf{V}$  and  $\mathbf{C}$ , respectively. In the rest of this paper, we will use capital letters to represent variables (e.g.,  $X, Y, Z$ ), while lowercase letters represent constants. The next component of the language is a set of  $n$ -ary predicate symbols; the EM and AM use separate sets of predicate symbols, denoted with  $\mathbf{P}_{EM}, \mathbf{P}_{AM}$ ,

respectively – the two models can, however, share variables and constants. As usual, a *term* is composed of either a variable or constant. Given terms  $t_1, \dots, t_n$  and  $n$ -ary predicate symbol  $p$ ,  $p(t_1, \dots, t_n)$  is called an *atom*; if  $t_1, \dots, t_n$  are constants, then the atom is said to be *ground*. The sets of all ground atoms for EM and AM are denoted with  $\mathbf{G}_{EM}$  and  $\mathbf{G}_{AM}$ , respectively.

Given set of ground atoms, a *world* is any subset of atoms – those that belong to the set are said to be *true* in the world, while those that do not are *false*. Therefore, there are  $2^{|\mathbf{G}_{EM}|}$  possible worlds in the EM and  $2^{|\mathbf{G}_{AM}|}$  worlds in the AM. These sets are denoted with  $\mathcal{W}_{EM}$  and  $\mathcal{W}_{AM}$ , respectively. In order to avoid worlds that do not model possible situations given a particular domain, we include *integrity constraints* of the form  $\text{oneOf}(\mathcal{A}')$ , where  $\mathcal{A}'$  is a subset of ground atoms. Intuitively, such a constraint states that any world where more than one of the atoms from set  $\mathcal{A}'$  appears is invalid. We use  $\mathbf{IC}_{EM}$  and  $\mathbf{IC}_{AM}$  to denote the sets of integrity constraints for the EM and AM, respectively, and the sets of worlds that conform to these constraints is denoted with  $\mathcal{W}_{EM}(\mathbf{IC}_{EM})$ ,  $\mathcal{W}_{AM}(\mathbf{IC}_{AM})$ , respectively.

Finally, logical formulas arise from the combination of atoms using the traditional connectives ( $\wedge$ ,  $\vee$ , and  $\neg$ ). As usual, we say a world  $w$  *satisfies* formula ( $f$ ), written  $w \models f$ , iff: (i) If  $f$  is an atom, then  $w \models f$  iff  $f \in w$ ; (ii) if  $f = \neg f'$  then  $w \models f$  iff  $w \not\models f'$ ; (iii) if  $f = f' \wedge f''$  then  $w \models f$  iff  $w \models f'$  and  $w \models f''$ ; and (iv) if  $f = f' \vee f''$  then  $w \models f$  iff  $w \models f'$  or  $w \models f''$ . We use the notation  $\text{form}_{EM}$ ,  $\text{form}_{AM}$  to denote the set of all possible (ground) formulas in the EM and AM, respectively.

## 2.1 Probabilistic Model

The EM or environmental model is largely based on the probabilistic logic of [16], which we now briefly review.

**Definition 1.** *Let  $f$  be a formula over  $\mathbf{P}_{EM}$ ,  $\mathbf{V}$ , and  $\mathbf{C}$ ,  $p \in [0, 1]$ , and  $\epsilon \in [0, \min(p, 1 - p)]$ . A probabilistic formula is of the form  $f : p \pm \epsilon$ . A set  $\mathcal{K}_{EM}$  of probabilistic formulas is called a probabilistic knowledge base.*

In the above definition, the number  $\epsilon$  is referred to as an *error tolerance*. Intuitively, probabilistic formulas are interpreted as “formula  $f$  is true with probability between  $p - \epsilon$  and  $p + \epsilon$ ” – note that there are no further constraints over this interval apart from those imposed by other probabilistic formulas in the knowledge base. The uncertainty regarding the probability values stems from the fact that certain assumptions (such as probabilistic independence) may not be suitable in the environment being modeled.

*Example 1.* Consider the following set  $\mathcal{K}_{EM}$ :

$$\begin{array}{lll} f_1 = a : 0.8 \pm 0.1 & f_4 = d \wedge e & : 0.7 \pm 0.2 & f_7 = k : 1 \pm 0 \\ f_2 = b : 0.2 \pm 0.1 & f_5 = f \wedge g \wedge h & : 0.6 \pm 0.1 & \\ f_3 = c : 0.8 \pm 0.1 & f_6 = i \vee \neg j & : 0.9 \pm 0.1 & \end{array}$$

Throughout the paper, we also use  $\mathcal{K}'_{EM} = \{f_1, f_2, f_3\}$  ■

A set of probabilistic formulas describes a set of possible probability distributions  $Pr$  over the set  $\mathcal{W}_{EM}(\mathbf{IC}_{EM})$ . We say that probability distribution  $Pr$  satisfies probabilistic formula  $f : p \pm \epsilon$  iff:  $p - \epsilon \leq \sum_{w \in \mathcal{W}_{EM}(\mathbf{IC}_{EM})} Pr(w) \leq p + \epsilon$ . We say that a probability distribution over  $\mathcal{W}_{EM}(\mathbf{IC}_{EM})$  satisfies  $\mathcal{K}_{EM}$  iff it satisfies all probabilistic formulas in  $\mathcal{K}_{EM}$ .

Given a probabilistic knowledge base and a (non-probabilistic) formula  $q$ , the *maximum entailment* problem seeks to identify real numbers  $p, \epsilon$  such that all valid probability distributions  $Pr$  that satisfy  $\mathcal{K}_{EM}$  also satisfy  $q : p \pm \epsilon$ , and there does not exist  $p', \epsilon'$  s.t.  $[p - \epsilon, p + \epsilon] \supset [p' - \epsilon', p' + \epsilon']$ , where all probability distributions  $Pr$  that satisfy  $\mathcal{K}_{EM}$  also satisfy  $q : p' \pm \epsilon'$ . In order to solve this problem we must solve the linear program defined below.

**Definition 2.** *Given a knowledge base  $\mathcal{K}_{EM}$  and a formula  $q$ , we have a variable  $x_i$  for each  $w_i \in \mathcal{W}_{EM}(\mathbf{IC}_{EM})$ .*

– For each  $f_j : p_j \pm \epsilon_j \in \mathcal{K}_{EM}$ , there is a constraint of the form:

$$p_j - \epsilon_j \leq \sum_{w_i \in \mathcal{W}_{EM}(\mathbf{IC}_{EM}) \text{ s.t. } w_i \models f_j} x_i \leq p_j + \epsilon_j.$$

– We also have the constraint:  $\sum_{w_i \in \mathcal{W}_{EM}(\mathbf{IC}_{EM})} x_i = 1$ .

– The objective is to minimize the function:  $\sum_{w_i \in \mathcal{W}_{EM}(\mathbf{IC}_{EM}) \text{ s.t. } w_i \models q} x_i$ .

We use the notation  $EP\text{-LP-MIN}(\mathcal{K}_{EM}, q)$  to refer to the value of the objective function in the solution to the  $EM\text{-LP-MIN}$  constraints.

The next step is to solve the linear program a second time, but instead maximizing the objective function (we shall refer to this as  $EM\text{-LP-MAX}$ ) – let  $\ell$  and  $u$  be the results of these operations, respectively. In [16], it is shown that  $\epsilon = \frac{u-\ell}{2}$  and  $p = \ell + \epsilon$  is the solution to the maximum entailment problem. We note that although the above linear program has an exponential number of variables in the worst case (i.e., no integrity constraints), the presence of constraints has the potential to greatly reduce this space. Further, there are also good heuristics (cf. [17,18]) that have been shown to provide highly accurate approximations with a reduced-size linear program.

*Example 2.* Consider KB  $\mathcal{K}'_{EM}$  from Example 1 and a set of ground atoms restricted to those that appear in that program; we have the following worlds:

$$\begin{array}{llll} w_1 = \{a, b, c\} & w_2 = \{a, b\} & w_3 = \{a, c\} & w_4 = \{b, c\} \\ w_5 = \{b\} & w_6 = \{a\} & w_7 = \{c\} & w_8 = \emptyset \end{array}$$

and suppose we wish to compute the probability for formula  $q = a \vee c$ . For each formula in  $\mathcal{K}_{EM}$  we have a constraint, and for each world above we have a variable. An objective function is created based on the worlds that satisfy the query formula (in this case, worlds  $w_1, w_2, w_3, w_4, w_6, w_7$ ). Solving  $EP\text{-LP-MAX}(\mathcal{K}'_{EM}, q)$  and  $EP\text{-LP-MIN}(\mathcal{K}'_{EM}, q)$ , we obtain the solution  $0.9 \pm 0.1$ . ■

### 3 Argumentation Model

For the analytical model (AM), we choose a structured argumentation framework [19] due to several characteristics that make such frameworks highly applicable to many domains. Unlike the EM, which describes probabilistic information about the state of the real world, the AM must allow for competing ideas. Therefore, it must be able to represent contradictory information. The algorithmic approach we shall later describe allows for the creation of *arguments* based on the AM that may “compete” with each other to answer a given query. In this competition – known as a *dialectical process* – one argument may defeat another based on a *comparison criterion* that determines the prevailing argument. Resulting from this process, certain arguments are *warranted* (those that are not *defeated* by other arguments) thereby providing a suitable explanation for the answer to a given query.

The transparency provided by the system can allow knowledge engineers to identify potentially incorrect input information and fine-tune the models or, alternatively, collect more information. In short, argumentation-based reasoning has been studied as a natural way to manage a set of inconsistent information – it is the way humans settle disputes. As we will see, another desirable characteristic of (structured) argumentation frameworks is that, once a conclusion is reached, we are left with an explanation of how we arrived at it and information about why a given argument is warranted; this is very important information for users to have. In the following, we first recall the basics of the underlying argumentation framework used, and then go on to introduce the analytical model (AM).

#### 3.1 Defeasible Logic Programming with Presumptions (PreDeLP)

Defeasible Logic Programming with Presumptions (PreDeLP) [1] is a formalism combining logic programming with defeasible argumentation; it arises as an extension of classical DELP [20] with the possibility of having presumptions, as described below – since this capability is useful in many applications, we adopt this extended version in this paper. In this section, we briefly recall the basics of PreDeLP; we refer the reader to [20, 1] for the complete presentation.

The formalism contains several different constructs: facts, presumptions, strict rules, and defeasible rules. Facts are statements about the analysis that can always be considered to be true, while presumptions are statements that may or may not be true. Strict rules specify logical consequences of a set of facts or presumptions (similar to an implication, though not the same) that must always occur, while defeasible rules specify logical consequences that may be assumed to be true when no contradicting information is present. These building blocks are used in the construction of *arguments*, and are part of a PreDeLP program, which is a set of facts, strict rules, presumptions, and defeasible rules. Formally, we use the notation  $\Pi_{AM} = (\Theta, \Omega, \Phi, \Delta)$  to denote a PreDeLP program, where  $\Omega$  is the set of strict rules,  $\Theta$  is the set of facts,  $\Delta$  is the set of defeasible rules, and  $\Phi$  is the set of presumptions. In Figure 1, we provide an example  $\Pi_{AM}$ . We now define these constructs formally.

$\Theta : \theta_{1a} = p$	$\theta_{1b} = q$	$\theta_2 = r$		
$\Omega : \omega_{1a} = \neg s \leftarrow t$	$\omega_{1b} = \neg t \leftarrow s$	$\omega_{2a} = s \leftarrow p, u, r, v$	$\omega_{2b} = t \leftarrow q, w, x, v$	
$\Phi : \phi_1 = y \prec$	$\phi_2 = v \prec$	$\phi_3 = \neg z \prec$		
$\Delta : \delta_{1a} = s \prec p$	$\delta_{1b} = t \prec q$	$\delta_2 = s \prec u$	$\delta_3 = s \prec r, v$	
$\delta_4 = u \prec y$	$\delta_{5a} = \neg u \prec \neg z$	$\delta_{5b} = \neg w \prec \neg n$		

**Fig. 1.** An example (propositional) argumentation framework.

**Facts** ( $\Theta$ ) are ground literals representing atomic information or its negation, using strong negation “ $\neg$ ”. Note that all of the literals in our framework must be formed with a predicate from the set  $\mathbf{P}_{AM}$ . Note that information in the form of facts cannot be contradicted. We will use the notation  $[\Theta]$  to denote the set of all possible facts.

**Strict Rules** ( $\Omega$ ) represent non-defeasible cause-and-effect information that resembles an implication (though the semantics is different since the contrapositive does not hold) and are of the form  $L_0 \leftarrow L_1, \dots, L_n$ , where  $L_0$  is a ground literal and  $\{L_i\}_{i>0}$  is a set of ground literals. We will use the notation  $[\Omega]$  to denote the set of all possible strict rules.

**Presumptions** ( $\Phi$ ) are ground literals of the same form as facts, except that they are not taken as being true but rather defeasible, which means that they can be contradicted. Presumptions are denoted in the same manner as facts, except that the symbol  $\prec$  is added.

**Defeasible Rules** ( $\Delta$ ) represent tentative knowledge that can be used if nothing can be posed against it. Just as presumptions are the defeasible counterpart of facts, defeasible rules are the defeasible counterpart of strict rules. They are of the form  $L_0 \prec L_1, \dots, L_n$ , where  $L_0$  is a ground literal and  $\{L_i\}_{i>0}$  is a set of ground literals. In both strict and defeasible rules, *strong negation* is allowed in the head of rules, and hence may be used to represent contradictory knowledge.

Even though the above constructs are ground, we allow for schematic versions with variables that are used to represent sets of ground rules. We denote variables with strings starting with an uppercase letter.

**Arguments.** Given a query in the form of a ground atom, the goal is to derive arguments for and against its validity – derivation follows the same mechanism of logic programming [21]. Since rule heads can contain strong negation, it is possible to defeasibly derive contradictory literals from a program. For the treatment of contradictory knowledge, PreDeLP incorporates a defeasible argumentation formalism that allows the identification of the pieces of knowledge that are in conflict and, through the previously mentioned dialectical process, decides which information prevails as warranted. This dialectical process involves

$\langle \mathcal{A}_1, s \rangle$	$\mathcal{A}_1 = \{\theta_{1a}, \delta_{1a}\}$	$\langle \mathcal{A}_2, s \rangle$	$\mathcal{A}_2 = \{\phi_1, \phi_2, \delta_4, \omega_{2a}, \theta_{1a}, \theta_2\}$
$\langle \mathcal{A}_3, s \rangle$	$\mathcal{A}_3 = \{\phi_1, \delta_2, \delta_4\}$	$\langle \mathcal{A}_4, s \rangle$	$\mathcal{A}_4 = \{\phi_2, \delta_3, \theta_2\}$
$\langle \mathcal{A}_5, u \rangle$	$\mathcal{A}_5 = \{\phi_1, \delta_4\}$	$\langle \mathcal{A}_6, \neg s \rangle$	$\mathcal{A}_6 = \{\delta_{1b}, \theta_{1b}, \omega_{1a}\}$
$\langle \mathcal{A}_7, \neg u \rangle$	$\mathcal{A}_7 = \{\phi_3, \delta_{5a}\}$		

**Fig. 2.** Example ground arguments from the framework of Figure 1.

the construction and evaluation of arguments, building a *dialectical tree* in the process. Arguments are formally defined next.

**Definition 3.** An argument  $\langle \mathcal{A}, L \rangle$  for a literal  $L$  is a pair of the literal and a (possibly empty) set of the EM ( $\mathcal{A} \subseteq \Pi_{AM}$ ) that provides a minimal proof for  $L$  meeting the following requirements: (i)  $L$  is defeasibly derived from  $\mathcal{A}$ ; (ii)  $\Omega \cup \Theta \cup \mathcal{A}$  is not contradictory; and (iii)  $\mathcal{A}$  is a minimal subset of  $\Delta \cup \Phi$  satisfying 1 and 2, denoted  $\langle \mathcal{A}, L \rangle$ .

Literal  $L$  is called the conclusion supported by the argument, and  $\mathcal{A}$  is the support of the argument. An argument  $\langle \mathcal{B}, L \rangle$  is a subargument of  $\langle \mathcal{A}, L' \rangle$  iff  $\mathcal{B} \subseteq \mathcal{A}$ . An argument  $\langle \mathcal{A}, L \rangle$  is presumptive iff  $\mathcal{A} \cap \Phi$  is not empty. We will also use  $\Omega(\mathcal{A}) = \mathcal{A} \cap \Omega$ ,  $\Theta(\mathcal{A}) = \mathcal{A} \cap \Theta$ ,  $\Delta(\mathcal{A}) = \mathcal{A} \cap \Delta$ , and  $\Phi(\mathcal{A}) = \mathcal{A} \cap \Phi$ .

Our definition differs slightly from that of [22], where DeLP is introduced, as we include strict rules and facts as part of arguments – the reason for this will become clear in Section 4. Arguments for our scenario are shown next.

*Example 3.* Figure 2 shows example arguments based on the knowledge base from Figure 1. Note that  $\langle \mathcal{A}_5, u \rangle$  is a sub-argument of  $\langle \mathcal{A}_2, s \rangle$  and  $\langle \mathcal{A}_3, s \rangle$ . ■

Given an argument  $\langle \mathcal{A}_1, L_1 \rangle$ , counter-arguments are arguments that contradict it. Argument  $\langle \mathcal{A}_2, L_2 \rangle$  is said to *counterargue* or *attack*  $\langle \mathcal{A}_1, L_1 \rangle$  at a literal  $L'$  iff there exists a subargument  $\langle \mathcal{A}, L'' \rangle$  of  $\langle \mathcal{A}_1, L_1 \rangle$  such that the set  $\Omega(\mathcal{A}_1) \cup \Omega(\mathcal{A}_2) \cup \Theta(\mathcal{A}_1) \cup \Theta(\mathcal{A}_2) \cup \{L_2, L''\}$  is contradictory.

*Example 4.* Consider the arguments from Example 3. The following are some of the attack relationships between them:  $\mathcal{A}_1$ ,  $\mathcal{A}_2$ ,  $\mathcal{A}_3$ , and  $\mathcal{A}_4$  all attack  $\mathcal{A}_6$ ;  $\mathcal{A}_5$  attacks  $\mathcal{A}_7$ ; and  $\mathcal{A}_7$  attacks  $\mathcal{A}_2$ . ■

A *proper defeater* of an argument  $\langle \mathcal{A}, L \rangle$  is a counter-argument that – by some criterion – is considered to be better than  $\langle \mathcal{A}, L \rangle$ ; if the two are incomparable according to this criterion, the counterargument is said to be a *blocking* defeater. An important characteristic of PreDeLP is that the argument comparison criterion is modular, and thus the most appropriate criterion for the domain that is being represented can be selected; the default criterion used in classical defeasible logic programming (from which PreDeLP is derived) is *generalized specificity* [23], though an extension of this criterion is required for arguments using presumptions [1]. We briefly recall this criterion next – the first definition is for generalized specificity, which is subsequently used in the definition of presumption-enabled specificity.



**Definition 4.** Let  $\Pi_{AM} = (\Theta, \Omega, \Phi, \Delta)$  be a PreDeLP program and let  $\mathcal{F}$  be the set of all literals that have a defeasible derivation from  $\Pi_{AM}$ . An argument  $\langle \mathcal{A}_1, L_1 \rangle$  is preferred to  $\langle \mathcal{A}_2, L_2 \rangle$ , denoted with  $\mathcal{A}_1 \succ_{PS} \mathcal{A}_2$  iff:

- (1) For all  $H \subseteq \mathcal{F}$ ,  $\Omega(\mathcal{A}_1) \cup \Omega(\mathcal{A}_2) \cup H$  is non-contradictory: if there is a derivation for  $L_1$  from  $\Omega(\mathcal{A}_2) \cup \Omega(\mathcal{A}_1) \cup \Delta(\mathcal{A}_1) \cup H$ , and there is no derivation for  $L_1$  from  $\Omega(\mathcal{A}_1) \cup \Omega(\mathcal{A}_2) \cup H$ , then there is a derivation for  $L_2$  from  $\Omega(\mathcal{A}_1) \cup \Omega(\mathcal{A}_2) \cup \Delta(\mathcal{A}_2) \cup H$ ; and
- (2) there is at least one set  $H' \subseteq \mathcal{F}$ ,  $\Omega(\mathcal{A}_1) \cup \Omega(\mathcal{A}_2) \cup H'$  is non-contradictory, such that there is a derivation for  $L_2$  from  $\Omega(\mathcal{A}_1) \cup \Omega(\mathcal{A}_2) \cup H' \cup \Delta(\mathcal{A}_2)$ , there is no derivation for  $L_2$  from  $\Omega(\mathcal{A}_1) \cup \Omega(\mathcal{A}_2) \cup H'$ , and there is no derivation for  $L_1$  from  $\Omega(\mathcal{A}_1) \cup \Omega(\mathcal{A}_2) \cup H' \cup \Delta(\mathcal{A}_1)$ .

Intuitively, the principle of specificity says that, in the presence of two conflicting lines of argument about a proposition, the one that uses more of the available information is more convincing. A classic example involves a bird, Tweety, and arguments stating that it both flies (because it is a bird) and doesn't fly (because it is a penguin). The latter argument uses more information about Tweety – it is more specific – and is thus the stronger of the two.

**Definition 5 ([1]).** Let  $\Pi_{AM} = (\Theta, \Omega, \Phi, \Delta)$  be a PreDeLP program. An argument  $\langle \mathcal{A}_1, L_1 \rangle$  is preferred to  $\langle \mathcal{A}_2, L_2 \rangle$ , denoted with  $\mathcal{A}_1 \succ \mathcal{A}_2$  iff any of the following conditions hold:

- (1)  $\langle \mathcal{A}_1, L_1 \rangle$  and  $\langle \mathcal{A}_2, L_2 \rangle$  are both factual arguments and  $\langle \mathcal{A}_1, L_1 \rangle \succ_{PS} \langle \mathcal{A}_2, L_2 \rangle$ .
- (2)  $\langle \mathcal{A}_1, L_1 \rangle$  is a factual argument and  $\langle \mathcal{A}_2, L_2 \rangle$  is a presumptive argument.
- (3)  $\langle \mathcal{A}_1, L_1 \rangle$  and  $\langle \mathcal{A}_2, L_2 \rangle$  are presumptive arguments, and
  - (a)  $\Phi(\mathcal{A}_1) \subsetneq \Phi(\mathcal{A}_2)$  or,
  - (b)  $\Phi(\mathcal{A}_1) = \Phi(\mathcal{A}_2)$  and  $\langle \mathcal{A}_1, L_1 \rangle \succ_{PS} \langle \mathcal{A}_2, L_2 \rangle$ .

Generally, if  $\mathcal{A}, \mathcal{B}$  are arguments with rules  $X$  and  $Y$ , resp., and  $X \subset Y$ , then  $\mathcal{A}$  is stronger than  $\mathcal{B}$ . This also holds when  $\mathcal{A}$  and  $\mathcal{B}$  use presumptions  $P_1$  and  $P_2$ , resp., and  $P_1 \subset P_2$ .

*Example 5.* The following are some relationships between arguments from Example 3, based on Definitions 4 and 5.

- $\mathcal{A}_1$  and  $\mathcal{A}_6$  are incomparable (blocking defeaters);
- $\mathcal{A}_6 \succ \mathcal{A}_2$ , and thus  $\mathcal{A}_6$  defeats  $\mathcal{A}_2$ ;
- $\mathcal{A}_5$  and  $\mathcal{A}_7$  are incomparable (blocking defeaters). ■

A sequence of arguments called an *argumentation line* thus arises from this attack relation, where each argument defeats its predecessor. To avoid undesirable sequences, which may represent circular argumentation lines, in DeLP an *argumentation line* is *acceptable* if it satisfies certain constraints (see [20]). A literal  $L$  is *warranted* if there exists a non-defeated argument  $\mathcal{A}$  supporting  $L$ .

Clearly, there can be more than one defeater for a particular argument  $\langle \mathcal{A}, L \rangle$ . Therefore, many acceptable argumentation lines could arise from  $\langle \mathcal{A}, L \rangle$ , leading to a tree structure. The tree is built from the set of all argumentation lines

rooted in the initial argument. In a dialectical tree, every node (except the root) represents a defeater of its parent, and leaves correspond to undefeated arguments. Each path from the root to a leaf corresponds to a different acceptable argumentation line. A dialectical tree provides a structure for considering all the possible acceptable argumentation lines that can be generated for deciding whether an argument is defeated. We call this tree *dialectical* because it represents an exhaustive dialectical<sup>4</sup> analysis for the argument in its root. For a given argument  $\langle \mathcal{A}, L \rangle$ , we denote the corresponding dialectical tree as  $\mathcal{T}(\langle \mathcal{A}, L \rangle)$ .

Given a literal  $L$  and an argument  $\langle \mathcal{A}, L \rangle$ , in order to decide whether or not a literal  $L$  is warranted, every node in the dialectical tree  $\mathcal{T}(\langle \mathcal{A}, L \rangle)$  is recursively marked as “D” (*defeated*) or “U” (*undefeated*), obtaining a marked dialectical tree  $\mathcal{T}^*(\langle \mathcal{A}, L \rangle)$  as follows:

1. All leaves in  $\mathcal{T}^*(\langle \mathcal{A}, L \rangle)$  are marked as “U”s, and
2. Let  $\langle \mathcal{B}, q \rangle$  be an inner node of  $\mathcal{T}^*(\langle \mathcal{A}, L \rangle)$ . Then  $\langle \mathcal{B}, q \rangle$  will be marked as “U” iff every child of  $\langle \mathcal{B}, q \rangle$  is marked as “D”. The node  $\langle \mathcal{B}, q \rangle$  will be marked as “D” iff it has at least a child marked as “U”.

Given an argument  $\langle \mathcal{A}, L \rangle$  obtained from  $\Pi_{AM}$ , if the root of  $\mathcal{T}^*(\langle \mathcal{A}, L \rangle)$  is marked as “U”, then we will say that  $\mathcal{T}^*(\langle \mathcal{A}, h \rangle)$  *warrants*  $L$  and that  $L$  is *warranted* from  $\Pi_{AM}$ . (Warranted arguments correspond to those in the grounded extension of a Dung argumentation system [24].) There is a further requirement when the arguments in the dialectical tree contains presumptions – the conjunction of all presumptions used in even (respectively, odd) levels of the tree must be consistent. This can give rise to multiple trees for a given literal, as there can potentially be different arguments that make contradictory assumptions.

We can then extend the idea of a dialectical tree to a *dialectical forest*. For a given literal  $L$ , a dialectical forest  $\mathcal{F}(L)$  consists of the set of dialectical trees for all arguments for  $L$ . We shall denote a marked dialectical forest, the set of all marked dialectical trees for arguments for  $L$ , as  $\mathcal{F}^*(L)$ . Hence, for a literal  $L$ , we say it is *warranted* if there is at least one argument for that literal in the dialectical forest  $\mathcal{F}^*(L)$  that is labeled as “U”, *not warranted* if there is at least one argument for the literal  $\neg L$  in the dialectical forest  $\mathcal{F}^*(\neg L)$  that is labeled as “U”, and *undecided* otherwise.

## 4 Probabilistic PreDeLP

Probabilistic PreDeLP arises from the combination of the environmental and analytical models ( $\Pi_{EM}$  and  $\Pi_{AM}$ , respectively). Intuitively, given  $\Pi_{AM}$ , every element of  $\Omega \cup \Theta \cup \Delta \cup \Phi$  might only hold in certain worlds in the set  $\mathcal{W}_{EM}$  – that is, they are subject to probabilistic events. Therefore, we associate elements of  $\Omega \cup \Theta \cup \Delta \cup \Phi$  with a formula from  $form_{EM}$ . For instance, we could associate formula *rainy* to fact *umbrella* to state that the latter only holds when the probabilistic event *rainy* holds; since weather is uncertain in nature, it has been modeled as part of the EM.

<sup>4</sup> In the sense of providing reasons for and against a position.

$af(\theta_{1a}) = af(\theta_{1b}) = k \vee (f \wedge (h \vee (e \wedge l)))$	$af(\phi_3) = b$
$af(\theta_2) = i$	$af(\delta_{1a}) = af(\delta_{1b}) = \text{True}$
$af(\omega_{1a}) = af(\omega_{1b}) = \text{True}$	$af(\delta_2) = \text{True}$
$af(\omega_{2a}) = af(\omega_{2b}) = \text{True}$	$af(\delta_3) = \text{True}$
$af(\phi_1) = c \vee a$	$af(\delta_4) = \text{True}$
$af(\phi_2) = f \wedge m$	$af(\delta_{5a}) = af(\delta_{5b}) = \text{True}$

**Fig. 3.** Example annotation function.

We can then compute the probabilities of subsets of  $\Omega \cup \Theta \cup \Delta \cup \Phi$  using the information contained in  $\Pi_{EM}$ , as we describe shortly. The notion of an *annotation function* associates elements of  $\Omega \cup \Theta \cup \Delta \cup \Phi$  with elements of  $form_{EM}$ .

**Definition 6.** An annotation function is any function  $af : \Omega \cup \Theta \cup \Delta \cup \Phi \rightarrow form_{EM}$ . We shall use  $[af]$  to denote the set of all annotation functions.

We will sometimes denote annotation functions as sets of pairs  $(f, af(f))$  in order to simplify the presentation. Figure 3 shows an example of an annotation function for our running example.

We now have all the components to formally define Probabilistic PreDeLP programs (P-PreDeLP for short).

**Definition 7.** Given environmental model  $\Pi_{EM}$ , analytical model  $\Pi_{AM}$ , and annotation function  $af$ , a probabilistic PreDeLP program is of the form  $\mathcal{I} = (\Pi_{EM}, \Pi_{AM}, af)$ . We use notation  $[\mathcal{I}]$  to denote the set of all possible programs.

Given this setup, we can consider a world-based approach; that is, the defeat relationship among arguments depends on the current state of the (EM) world.

**Definition 8.** Let  $\mathcal{I} = (\Pi_{EM}, \Pi_{AM}, af)$  be a P-PreDeLP program, argument  $\langle \mathcal{A}, L \rangle$  is valid w.r.t. world  $w \in \mathcal{W}_{EM}$  iff  $\forall c \in \mathcal{A}, w \models af(c)$ .

We extend the notion of validity to argumentation lines, dialectical trees, and dialectical forests in the expected way (for instance, an argumentation line is valid w.r.t.  $w$  iff all arguments that comprise that line are valid w.r.t.  $w$ ). We also extend the idea of a dialectical tree w.r.t. worlds; so, for a given world  $w \in \mathcal{W}_{EM}$ , the dialectical (resp., marked dialectical) tree induced by  $w$  is denoted with  $\mathcal{T}_w \langle \mathcal{A}, L \rangle$  (resp.,  $\mathcal{T}_w^* \langle \mathcal{A}, L \rangle$ ). We require that all arguments and defeaters in these trees to be valid with respect to  $w$ . Likewise, we extend the notion of dialectical forests in the same manner (denoted with  $\mathcal{F}_w(L)$  and  $\mathcal{F}_w^*(L)$ , resp.). Based on these concepts we introduce the notion of *warranting scenario*.

**Definition 9.** Let  $\mathcal{I} = (\Pi_{EM}, \Pi_{AM}, af)$  be a P-PreDeLP program and  $L$  be a literal formed with a ground atom from  $\mathbf{G}_{AM}$ ; a world  $w \in \mathcal{W}_{EM}$  is said to be a warranting scenario for  $L$  (denoted  $w \vdash_{war} L$ ) iff there is a dialectical forest  $\mathcal{F}_w^*(L)$  in which  $L$  is warranted and  $\mathcal{F}_w^*(L)$  is valid w.r.t.  $w$ .

Hence, the set of worlds in the EM where a literal  $L$  in the AM *must* be true is exactly the set of warranting scenarios – these are the “necessary” worlds:  $nec(L) = \{w \in \mathcal{W}_{EM} \mid (w \vdash_{\text{war}} L)\}$ . Now, the set of worlds in the EM where AM literal  $L$  *can* be true is the following – these are the “possible” worlds:  $poss(L) = \{w \in \mathcal{W}_{EM} \mid w \not\vdash_{\text{war}} \neg L\}$ . The probability distribution  $Pr$  defined over the worlds in the EM induces an upper and lower bound on the probability of literal  $L$  (denoted  $\mathbf{P}_{L,Pr,\mathcal{I}}$ ) as follows:

$$\ell_{L,Pr,\mathcal{I}} = \sum_{w \in nec(L)} Pr(w), \quad u_{L,Pr,\mathcal{I}} = \sum_{w \in poss(L)} Pr(w)$$

$$\ell_{L,Pr,\mathcal{I}} \leq \mathbf{P}_{L,Pr,\mathcal{I}} \leq u_{L,Pr,\mathcal{I}}$$

Since the EM in general does not define a single probability distribution, the above computations should be done using linear programs EP-LP-MIN and EP-LP-MAX, as described above.

#### 4.1 Sources of Inconsistency

We use the following notion of (classical) consistency of PreDeLP programs:  $\Pi$  is said to be *consistent* if there does not exist ground literal  $a$  s.t.  $\Pi \vdash a$  and  $\Pi \vdash \neg a$ . For P-PreDeLP programs, there are two main kinds of inconsistency that can be present; the first is what we refer to as EM, or Type I, (in)consistency.

**Definition 10.** *Environmental model  $\Pi_{EM}$  is Type I consistent iff there exists a probability distribution  $Pr$  over the set of worlds  $\mathcal{W}_{EM}$  that satisfies  $\Pi_{EM}$ .*

We illustrate this type of consistency in the following example.

*Example 6.* The following formula is a simple example of an EM for which there is no satisfying probability distribution:

$$\begin{aligned} rain \vee hail &: 0.3 \pm 0; \\ rain \wedge hail &: 0.5 \pm 0.1. \end{aligned}$$

A P-PreDeLP program using such an EM gives rise to an example of Type I inconsistency, as it arises from the fact that there is no satisfying interpretation for the EM knowledge base. ■

Assuming a consistent EM, inconsistencies can still arise through the interaction between the annotation function and facts and strict rules. We will refer to this as combined, or Type II, (in)consistency.

**Definition 11.** *A P-PreDeLP program  $\mathcal{I} = (\Pi_{EM}, \Pi_{AM}, af)$ , with  $\Pi_{AM} = \langle \Theta, \Omega, \Phi, \Delta \rangle$ , is Type II consistent iff: given any probability distribution  $Pr$  that satisfies  $\Pi_{EM}$ , if there exists a world  $w \in \mathcal{W}_{EM}$  such that  $\bigcup_{x \in \Theta \cup \Omega \mid w \models af(x)} \{x\}$  is inconsistent, then we have  $Pr(w) = 0$ .*

Thus, any EM world in which the set of associated facts and strict rules are inconsistent (we refer to this as “classical consistency”) must always be assigned a zero probability. The following is an example of this other type of inconsistency.

*Example 7.* Consider the EM knowledge base from Example 1, the AM presented in Figure 1 and the annotation function from Figure 3. Suppose the following fact is added to the argumentation model:

$$\theta_3 = \neg p,$$

and that the annotation function is expanded as follows:

$$af(\theta_3) = \neg k.$$

Clearly, fact  $\theta_3$  is in direct conflict with fact  $\theta_{1a}$  – this does not necessarily mean that there is an inconsistency. For instance, by the annotation function,  $\theta_{1a}$  holds in the world  $\{k\}$  while  $\theta_3$  does not. However, if we consider the world:

$$w = \{f, h\}$$

Note that  $w \models af(\theta_3)$  and  $w \models af(\theta_2)$ , which means that, in this world, two contradictory facts can occur. Since the environmental model indicates that this world can be assigned a non-zero probability, we have a Type II inconsistent program. ■

Another example (perhaps easier to visualize) in the rain/hail scenario discussed above, is as follows: suppose we have facts  $f = umbrella$  and  $g = \neg umbrella$ , and annotation function  $af(f) = rain \vee hail$  and  $af(g) = wind$ . Intuitively, the first fact states that an umbrella should be carried if it either rains or hails, while the second states that an umbrella should not be carried if it is windy. If the EM assigns a non-zero probability to formula  $(rain \vee hail) \wedge wind$ , then we have Type II inconsistency.

In the following, we say that a P-PreDeLP program is **consistent** if and only if it is both Type I and Type II consistent. However, in this paper, we focus on Type II consistency and assume that the program is Type I consistent.

## 4.2 Basic Operations for Restoring Consistency

Given a P-PreDeLP program that is Type II inconsistent, there are two basic strategies that can be used to restore consistency:

*Revise the EM:* the probabilistic model can be changed in order to force the worlds that induce contradicting strict knowledge to have probability zero.

*Revise the annotation function:* The annotations involved in the inconsistency can be changed so that the conflicting information in the AM does not become induced under any possible world.

It may also appear that a third option would be to adjust the AM – this is, however, equivalent to modifying the annotation function. Consider the presence

of two facts in the AM:  $a, \neg a$ . Assuming that this causes an inconsistency (that is, there is at least one world in which they both hold), one way to resolve it would be to remove one of these two literals. Suppose  $\neg a$  is removed; this would be equivalent to setting  $af(\neg a) = \perp$  (where  $\perp$  represents a contradiction in the language of the EM). *In this paper, we often refer to “removing elements of  $\Pi_{AM}$ ” to refer to changes to the annotation function that cause certain elements of the  $\Pi_{AM}$  to not have their annotations satisfied in certain EM worlds.*

Now, suppose that  $\Pi_{EM}$  is consistent, but that the overall program is Type II inconsistent. Then, there must exist a set of worlds in the EM where there is a probability distribution that assigns each of them a non-zero probability. This gives rise to the following result.

**Proposition 1.** *If there exists a probability distribution  $Pr$  that satisfies  $\Pi_{EM}$  s.t. there exists a world  $w \in \mathcal{W}_{EM}$  where  $Pr(w) > 0$  and  $\bigcup_{x \in \Theta \cup \Omega \mid w \models af(x)} \{x\}$  is inconsistent (Type II inconsistency), then any change made in order to resolve this inconsistency by modifying only  $\Pi_{EM}$  yields a new EM  $\Pi'_{EM}$  such that  $(\bigwedge_{a \in w} a \wedge \bigwedge_{a \notin w} \neg a) : 0 \pm 0$  is entailed by  $\Pi'_{EM}$ .*

Proposition 1 seems to imply an easy strategy of adding formulas to  $\Pi_{EM}$  causing certain worlds to have a zero probability. However, this may lead to Type I inconsistencies in the resulting model  $\Pi'_{EM}$ . If we are applying an EM-only strategy to resolve inconsistencies, this would then lead to further adjustments to  $\Pi'_{EM}$  in order to restore Type I consistency. However, such changes could potentially lead to Type II inconsistency in the overall P-PreDeLP program (by either removing elements of  $\Pi'_{EM}$  or loosening probability bounds of the sentences in  $\Pi'_{EM}$ ), which would lead to setting more EM worlds to a probability of zero. It is easy to devise an example of a situation in which the probability mass cannot be accommodated given the constraints imposed by the AM and EM together – in such cases, it would be impossible to restore consistency by only modifying  $\Pi_{EM}$ . We thus arrive at the following observation:

**Observation 1** *Given a Type II inconsistent P-PreDeLP program, consistency cannot always be restored via modifications to  $\Pi_{EM}$  alone.*

Therefore, due to this line of reasoning, in this paper we focus our efforts on modifications to the annotation function only. However, in the future, we intend to explore belief revision operators that consider both the annotation function (which, as we saw, captures changes to the AM) along with changes to the EM, as well as combinations of the two.

## 5 Revising Probabilistic PreDeLP Programs

Given a P-PreDeLP program  $\mathcal{I} = (\Pi_{EM}, \Pi_{AM}, af)$ , with  $\Pi_{AM} = \Omega \cup \Theta \cup \Delta \cup \Phi$ , we are interested in solving the problem of incorporating an epistemic input  $(f, af')$  into  $\mathcal{I}$ , where  $f$  is either an atom or a rule and  $af'$  is equivalent to  $af$ , except for its expansion to include  $f$ . For ease of presentation, we assume

that  $f$  is to be incorporated as a fact or strict rule, since incorporating defeasible knowledge can never lead to inconsistency. As we are only conducting annotation function revisions, for  $\mathcal{I} = (\Pi_{EM}, \Pi_{AM}, af)$  and input  $(f, af')$  we denote the revision as follows:  $\mathcal{I} \bullet (f, af') = (\Pi_{EM}, \Pi'_{AM}, af'')$  where  $\Pi'_{AM} = \Pi_{AM} \cup \{f\}$  and  $af''$  is the revised annotation function.

**Notation.** We use the symbol “ $\bullet$ ” to denote the revision operator. We also slightly abuse notation for the sake of presentation, as well as introduce notation to convert sets of worlds to/from formulas.

- $\mathcal{I} \cup (f, af')$  to denote  $\mathcal{I}' = (\Pi_{EM}, \Pi_{AM} \cup \{f\}, af')$ .
- $(f, af') \in \mathcal{I} = (\Pi_{AM}, \Pi_{EM}, af)$  to denote  $f \in \Pi_{AM}$  and  $af = af'$ .
- $wld(f) = \{w \mid w \models f\}$  – the set of worlds that satisfy formula  $f$ ; and
- $for(w) = \bigwedge_{a \in w} a \wedge \bigwedge_{a \notin w} \neg a$  – the formula that has  $w$  as its only model.
- $\Pi'_{AM}(w) = \{f \in \Theta \cup \Omega \mid w \models af(f)\}$
- $\mathcal{W}_{EM}^0(\mathcal{I}) = \{w \in \mathcal{W}_{EM} \mid \Pi'_{AM}(w) \text{ is inconsistent}\}$
- $\mathcal{W}_{EM}^I(\mathcal{I}) = \{w \in \mathcal{W}_{EM}^0 \mid \exists Pr \text{ s.t. } Pr \models \Pi_{EM} \wedge Pr(w) > 0\}$

Intuitively,  $\Pi'_{AM}(w)$  is the subset of facts and strict rules in  $\Pi_{AM}$  whose annotations are true in EM world  $w$ . The set  $\mathcal{W}_{EM}^0(\mathcal{I})$  contains all the EM worlds for a given program where the corresponding knowledge base in the AM is classically inconsistent and  $\mathcal{W}_{EM}^I(\mathcal{I})$  is a subset of these that can be assigned a non-zero probability – the latter are the worlds where inconsistency in the AM can arise.

### 5.1 Postulates for Revising the Annotation Function

We now analyze the rationality postulates for non-prioritized revision of belief bases first introduced in [2] and later generalized in [25], in the context of P-PreDeLP programs. These postulates are chosen due to the fact that they are well studied in the literature for non-prioritized belief revision.

**Inclusion:** For  $\mathcal{I} \bullet (f, af') = (\Pi_{EM}, \Pi_{AM} \cup \{f\}, af'')$ ,  $\forall g \in \Pi_{AM}$ ,  $wld(af''(g)) \subseteq wld(af'(g))$ .

This postulate states that, for any element in the AM, the worlds that satisfy its annotation after the revision are a subset of the original set of worlds satisfying the annotation for that element.

**Vacuity:** If  $\mathcal{I} \cup (f, af')$  is consistent, then  $\mathcal{I} \bullet (f, af') = \mathcal{I} \cup (f, af')$

**Consistency Preservation:** If  $\mathcal{I}$  is consistent, then  $\mathcal{I} \bullet (f, af')$  is also consistent.

**Weak Success:** If  $\mathcal{I} \cup (f, af')$  is consistent, then  $(f, af') \in \mathcal{I} \bullet (f, af')$ .

Whenever the simple addition of the input doesn't cause inconsistencies to arise, the result will contain the input.

**Core Retainment:** For  $\mathcal{I} \bullet (f, af') = (\Pi_{EM}, \Pi_{AM} \cup \{f\}, af'')$ , for each  $w \in \mathcal{W}_{EM}^I(\mathcal{I} \cup (f, af'))$ , we have  $X_w = \{h \in \Theta \cup \Omega \mid w \models af''(h)\}$ ; for each  $g \in$

$\Pi_{AM}(w) \setminus X_w$  there exists  $Y_w \subseteq X_w \cup \{f\}$  s.t.  $Y_w$  is consistent and  $Y_w \cup \{g\}$  is inconsistent.

For a given EM world, if a portion of the associated AM knowledge base is removed by the operator, then there exists a subset of the remaining knowledge base that is not consistent with the removed element and  $f$ .

**Relevance:** For  $\mathcal{I} \bullet (f, af') = (\Pi_{EM}, \Pi_{AM} \cup \{f\}, af'')$ , for each  $w \in \mathcal{W}_{EM}^I(\mathcal{I} \cup (f, af'))$  ( $(f, af')$ ), we have  $X_w = \{h \in \Theta \cup \Omega \mid w \models af''(h)\}$ ; for each  $g \in \Pi_{AM}(w) \setminus X_w$  there exists  $Y_w \supseteq X_w \cup \{f\}$  s.t.  $Y_w$  is consistent and  $Y_w \cup \{g\}$  is inconsistent.

For a given EM world, if a portion of the associated AM knowledge base is removed by the operator, then there exists a superset of the remaining knowledge base that is not consistent with the removed element and  $f$ .

**Uniformity 1:** Let  $(f, af'_1), (g, af'_2)$  be two inputs where  $\mathcal{W}_{EM}^I(\mathcal{I} \cup (f, af'_1)) = \mathcal{W}_{EM}^I(\mathcal{I} \cup (g, af'_2))$ ; for all  $w \in \mathcal{W}_{EM}^I(\mathcal{I} \cup (f, af'_1))$  and for all  $X \subseteq \Pi_{AM}(w)$ ; if  $\{x \mid x \in X \cup \{f\}, w \models af'_1(x)\}$  is inconsistent iff  $\{x \mid x \in X \cup \{g\}, w \models af'_2(x)\}$  is inconsistent, then for each  $h \in \Pi_{AM}$ , we have that:

$$\begin{aligned} \{w \in \mathcal{W}_{EM}^I(\mathcal{I} \cup (f, af'_1)) \mid w \models af'_1(h) \wedge \neg af'_1(h)\} = \\ \{w \in \mathcal{W}_{EM}^I(\mathcal{I} \cup (g, af'_2)) \mid w \models af'_2(h) \wedge \neg af'_2(h)\}. \end{aligned}$$

If two inputs result in the same set of EM worlds leading to inconsistencies in an AM knowledge base, and the consistency between analogous subsets (when joined with the respective input) are the same, then the models removed from the annotation of a given strict rule or fact are the same for both inputs.

**Uniformity 2:** Let  $(f, af'_1), (g, af'_2)$  be two inputs where  $\mathcal{W}_{EM}^I(\mathcal{I} \cup (f, af'_1)) = \mathcal{W}_{EM}^I(\mathcal{I} \cup (g, af'_2))$ ; for all  $w \in \mathcal{W}_{EM}^I(\mathcal{I} \cup (f, af'_1))$  and for all  $X \subseteq \Pi_{AM}(w)$ ; if  $\{x \mid x \in X \cup \{f\}, w \models af'_1(x)\}$  is inconsistent iff  $\{x \mid x \in X \cup \{g\}, w \models af'_2(x)\}$  is inconsistent, then

$$\begin{aligned} \{w \in \mathcal{W}_{EM}^I(\mathcal{I} \cup (f, af'_1)) \mid w \models af'_1(h) \wedge af'_1(h)\} = \\ \{w \in \mathcal{W}_{EM}^I(\mathcal{I} \cup (g, af'_2)) \mid w \models af'_2(h) \wedge af'_2(h)\}. \end{aligned}$$

If two inputs result in the same set of EM worlds leading to inconsistencies in an AM knowledge base, and the consistency between analogous subsets (when joined with the respective input) are the same, then the models retained in the the annotation of a given strict rule or fact are the same for both inputs.

**Relationships between Postulates.** There are a couple of interesting relationships among the postulates. The first is a sufficient condition for Core Retainment to be implied by Relevance.

**Proposition 2.** *Let  $\bullet$  be an operator such that  $\mathcal{I} \bullet (f, af') = (\Pi_{EM}, \Pi_{AM} \cup \{f\}, af'')$ , where  $\forall w \in \mathcal{W}_{EM}^I(\mathcal{I} \cup (f, af'))$ ,  $\Pi_{AM}^{\mathcal{I} \bullet (f, af')}(w)$  is a maximal consistent subset of  $\Pi_{AM}^{\mathcal{I} \cup (f, af')}(w)$ . If  $\bullet$  satisfies Relevance then it also satisfies Core Retainment.*



Similarly, we can show the equivalence between the two Uniformity postulates under certain conditions.

**Proposition 3.** *Let  $\bullet$  be an operator such that  $\mathcal{I} \bullet (f, af') = (\Pi_{EM}, \Pi_{AM} \cup \{f\}, af'')$  and  $\forall w, \Pi_{AM}^{\mathcal{I} \bullet (f, af')}(w) \subseteq \Pi_{AM}^{\mathcal{I} \cup (f, af')}(w)$ . Operator  $\bullet$  satisfies Uniformity 1 iff it satisfies Uniformity 2.*

Given the results of Propositions 2 and 3, we will not study Core Retainment and Uniformity 2 with respect to the construction of a belief revision operator in the next section.

## 5.2 An Operator for P-PreDeLP Revision

In this section, we introduce an operator for revising a P-PreDeLP program. As stated earlier, any subset of  $\Pi_{AM}$  associated with a world in  $\mathcal{W}_{EM}^I(\mathcal{I} \cup (f, af'))$  must be modified by the operator in order to remain consistent. So, for such a world  $w$ , we introduce a set of candidate replacement programs for  $\Pi_{AM}(w)$  in order to maintain consistency and satisfy the Inclusion postulate.

$$\begin{aligned} \text{candPgm}(w, \mathcal{I}) = \{ & \Pi'_{AM} \mid \Pi'_{AM} \subseteq \Pi_{AM}(w) \text{ s.t. } \Pi'_{AM} \text{ is consistent and} \\ & \nexists \Pi''_{AM} \subseteq \Pi_{AM}(w) \text{ s.t. } \Pi''_{AM} \supset \Pi'_{AM} \text{ s.t. } \Pi''_{AM} \text{ is consistent} \} \end{aligned}$$

Intuitively,  $\text{candPgm}(w, \mathcal{I})$  is the set of maximal consistent subsets of  $\Pi_{AM}(w)$ . Coming back to the rain/hail example presented above, we have:

*Example 8.* Consider the P-PreDeLP program  $\mathcal{I}$  presented right after Example 7, and the following EM knowledge base:

$$\begin{aligned} \text{rain} \vee \text{hail} & : 0.5 \pm 0.1; \\ \text{rain} \wedge \text{hail} & : 0.3 \pm 0.1; \\ \text{wind} & : 0.2 \pm 0. \end{aligned}$$

Given this setup, we have, for instance:

$$\text{candPgm}(\{\text{rain}, \text{hail}, \text{wind}\}, \mathcal{I}) = \left\{ \{ \text{umbrella} \}, \{ \neg \text{umbrella} \} \right\}.$$

Intuitively, this means that, since the world where *rain*, *hail*, and *wind* are all true can be assigned a non-zero probability by the EM, we must choose either *umbrella* or  $\neg$ *umbrella* in order to recover consistency. ■

We now show a series of intermediate results that lead up to the representation theorem (Theorem 1). First, we show how this set plays a role in showing a necessary and sufficient requirement for Inclusion and Consistency Preservation to hold together.

**Lemma 1.** *Given program  $\mathcal{I}$  and input  $(f, af')$ , operator  $\bullet$  satisfies Inclusion and Consistency Preservation iff for  $\mathcal{I} \bullet (f, af') = (\Pi_{EM}, \Pi_{AM}, af'')$ , for all  $w \in \mathcal{W}_{EM}^I(\mathcal{I} \cup (f, af'))$ , there exists an element  $X \in \text{candPgm}(w, \mathcal{I} \cup (f, af'))$  s.t.  $\{h \in \Theta \cup \Omega \cup \{f\} \mid w \models af''(h)\} \subseteq X$ .*

Next, we investigate the role that the set *candPgm* plays in showing the necessary and sufficient requirement for satisfying Inclusion, Consistency Preservation, and Relevance all at once.

**Lemma 2.** *Given program  $\mathcal{I}$  and input  $(f, af')$ , operator  $\bullet$  satisfies Inclusion, Consistency Preservation, and Relevance iff for  $\mathcal{I} \bullet (f, af') = (\Pi_{EM}, \Pi_{AM}, af'')$ , for all  $w \in \mathcal{W}_{EM}^I(\mathcal{I} \cup (f, af'))$  we have  $\{h \in \Theta \cup \Omega \cup \{f\} \mid w \models af''(h)\} \in \text{candPgm}(w, \mathcal{I} \cup (f, af'))$ .*

The last of the intermediate results shows that if there is a consistent program where two inputs cause inconsistencies to arise in the same way, then for each world the set of candidate replacement programs (minus the added AM formula) is the same. This result will be used as a support of the satisfaction of the first Uniformity postulate.

**Lemma 3.** *Let  $\mathcal{I} = (\Pi_{EM}, \Pi_{AM}, af)$  be a consistent program,  $(f_1, af'_1), (f_2, af'_2)$  be two inputs, and  $\mathcal{I}_i = (\Pi_{EM}, \Pi_{AM} \cup \{f_i\}, af'_i)$ . If  $\mathcal{W}_{EM}^I(\mathcal{I}_1) = \mathcal{W}_{EM}^I(\mathcal{I}_2)$ , then for all  $w \in \mathcal{W}_{EM}^I(\mathcal{I}_1)$  and all  $X \subseteq \Pi_{AM}(w)$  we have that:*

1. *If  $\{x \mid x \in X \cup \{f_1\}, w \models af'_1(x)\}$  is inconsistent  $\Leftrightarrow \{x \mid x \in X \cup \{f_2\}, w \models af'_2(x)\}$  is inconsistent, then  $\{X \setminus \{f_1\} \mid X \in \text{candPgm}(w, \mathcal{I}_1)\} = \{X \setminus \{f_2\} \mid X \in \text{candPgm}(w, \mathcal{I}_2)\}$ .*
2. *If  $\{X \setminus \{f_1\} \mid X \in \text{candPgm}(w, \mathcal{I}_1)\} = \{X \setminus \{f_2\} \mid X \in \text{candPgm}(w, \mathcal{I}_2)\}$  then  $\{x \mid x \in X \cup \{f_1\}, w \models af'_1(x)\}$  is inconsistent  $\Leftrightarrow \{x \mid x \in X \cup \{f_2\}, w \models af'_2(x)\}$  is inconsistent.*

We now have the necessary tools to present the construction of our non-prioritized belief revision operator.

**Construction.** Before introducing the construction, we define some preliminary notation. Let  $\Phi : \mathcal{W}_{EM} \rightarrow 2^{[\Theta] \cup [\Omega]}$ . For each  $h$  there is a formula in  $\Pi_{AM} \cup \{f\}$ , where  $f$  is part of the input. Given these elements, we define:

$$\text{newFor}(h, \Phi, \mathcal{I}, (f, af')) = af'(h) \wedge \bigwedge_{w \in \mathcal{W}_{EM}^I(\mathcal{I} \cup (f, af')) \mid h \notin \Phi(w)} \neg \text{for}(w_i)$$

The following definition then characterizes the class of operators called **AFO** (annotation function-based operators).

**Definition 12 (AF-based Operators).** *A belief revision operator  $\bullet$  is an “annotation function-based” (or af-based) operator ( $\bullet \in \mathbf{AFO}$ ) iff given program  $\mathcal{I} = (\Pi_{EM}, \Pi_{AM}, af)$  and input  $(f, af')$ , the revision is defined as  $\mathcal{I} \bullet (f, af') = (\Pi_{EM}, \Pi_{AM} \cup \{f\}, af'')$ , where:*

$$\forall h, af''(h) = \text{newFor}(h, \Phi, \mathcal{I}, (f, af'))$$

where  $\forall w \in \mathcal{W}_{EM}, \Phi(w) \in \text{CandPgm}_{af}(w, \mathcal{I} \cup (f, af'))$ .

As the main result of the paper, we now show that satisfying a key set of postulates is a necessary and sufficient condition for membership in **AFO**.

**Theorem 1 (Representation Theorem).** *An operator  $\bullet$  belongs to class **AFO** iff it satisfies Inclusion, Vacuity, Consistency Preservation, Weak Success, Relevance, and Uniformity 1.*

*Proof.* (Sketch) (If) By the fact that formulas associated with worlds in the set  $\mathcal{W}_{EM}^I(\mathcal{I} \cup (f, af'))$  are considered in the change of the annotation function, Vacuity and Weak Success follow trivially. Further, Lemma 2 shows that Inclusion, Consistency Preservation, and Relevance are satisfied while Lemma 3 shows that Uniformity 1 is satisfied.

(Only-If) Suppose BWOC that an operator  $\bullet$  satisfies all postulates and  $\bullet \notin \mathbf{AFO}$ . Then, one of four conditions must hold: (i) it does not satisfy Lemma 2 or (ii) it does not satisfy Lemma 3. However, by those previous arguments, if it satisfies all postulates, these arguments must be true as well – hence a contradiction.  $\square$

## 6 Conclusions

We have proposed an extension of the PreDeLP language that allows sentences to be annotated with probabilistic events; such events are connected to a probabilistic model, allowing a clear separation of interests between certain and uncertain knowledge. After presenting the language, we focused on characterizing belief revision operations over P-PreDeLP KBs. We presented a set of postulates inspired in the ones presented for non-prioritized revision of classical belief bases, and then proceeded to study a construction based on these postulates and prove that the two characterizations are equivalent.

As future work, we plan to study other kinds of operators, such as more general ones that allow the modification of the EM, as well as others that operate at different levels of granularity. Finally, we are studying the application of P-PreDeLP to real-world problems in cyber security and cyber warfare domains.

**Acknowledgments.** The authors are partially supported by UK EPSRC grant EP/J008346/1 (“PrOQAW”), ERC grant 246858 (“DIADEM”), ARO project 2GDATXR042, DARPA project R.0004972.001, Consejo Nacional de Investigaciones Científicas y Técnicas (CONICET) and Universidad Nacional del Sur (Argentina).

The opinions in this paper are those of the authors and do not necessarily reflect the opinions of the funders, the U.S. Military Academy, or the U.S. Army.

## References

1. Martínez, M.V., García, A.J., Simari, G.R.: On the use of presumptions in structured defeasible reasoning. In: Proc. of COMMA. (2012) 185–196
2. Hansson, S.: Semi-revision. J. of App. Non-Classical Logics **7**(1-2) (1997) 151–175
3. Alchourrón, C.E., Gärdenfors, P., Makinson, D.: On the logic of theory change: Partial meet contraction and revision functions. J. Sym. Log. **50**(2) (1985) 510–530

4. Gärdenfors, P.: Knowledge in flux: modeling the dynamics of epistemic states. MIT Press, Cambridge, Mass. (1988)
5. Hansson, S.O.: Kernel contraction. *J. Symb. Log.* **59**(3) (1994) 845–859
6. Doyle, J.: A truth maintenance system. *Artif. Intell.* **12**(3) (1979) 231–272
7. Falappa, M.A., Kern-Isberner, G., Simari, G.R.: Explanations, belief revision and defeasible reasoning. *Artif. Intell.* **141**(1/2) (2002) 1–28
8. Li, H., Oren, N., Norman, T.J.: Probabilistic argumentation frameworks. In: Proc. of TFAA. (2011) 1–16
9. Thimm, M.: A probabilistic semantics for abstract argumentation. In: Proc. of ECAI 2012. (2012) 750–755
10. Hunter, A.: Some foundations for probabilistic abstract argumentation. In: Proc. of COMMA 2012. (2012) 117–128
11. Fazzinga, B., Flesca, S., Parisi, F.: On the complexity of probabilistic abstract argumentation. In: Proc. of IJCAI 2013. (2013)
12. Haenni, R., Kohlas, J., Lehmann, N.: Probabilistic argumentation systems. Springer (1999)
13. Chesñevar, C.I., Simari, G.R., Alsinet, T., Godo, L.: A logic programming framework for possibilistic argumentation with vague knowledge. In: Proc. of UAI 2004. (2004) 76–84
14. Hunter, A.: A probabilistic approach to modelling uncertain logical arguments. *Int. J. Approx. Reasoning* **54**(1) (2013) 47–81
15. Gottlob, G., Lukasiewicz, T., Martínez, M.V., Simari, G.I.: Query answering under probabilistic uncertainty in Datalog+/- ontologies. AMAI (2013)
16. Nilsson, N.J.: Probabilistic logic. *Artif. Intell.* **28**(1) (1986) 71–87
17. Khuller, S., Martínez, M.V., Nau, D.S., Sliva, A., Simari, G.I., Subrahmanian, V.S.: Computing most probable worlds of action probabilistic logic programs: scalable estimation for  $10^{30,000}$  worlds. AMAI **51**(2-4) (2007) 295–331
18. Simari, G.I., Martínez, M.V., Sliva, A., Subrahmanian, V.S.: Focused most probable world computations in probabilistic logic programs. AMAI **64**(2-3) (2012) 113–143
19. Rahwan, I., Simari, G.R.: Argumentation in Artificial Intelligence. Springer (2009)
20. García, A.J., Simari, G.R.: Defeasible logic programming: An argumentative approach. TPLP **4**(1-2) (2004) 95–138
21. Lloyd, J.W.: Foundations of Logic Programming, 2nd Edition. Springer (1987)
22. Simari, G.R., Loui, R.P.: A mathematical treatment of defeasible reasoning and its implementation. *Artif. Intell.* **53**(2-3) (1992) 125–157
23. Stolzenburg, F., García, A., Chesñevar, C.I., Simari, G.R.: Computing Generalized Specificity. *Journal of Non-Classical Logics* **13**(1) (2003) 87–113
24. Dung, P.M.: On the acceptability of arguments and its fundamental role in non-monotonic reasoning, logic programming and  $n$ -person games. *Artif. Intell.* **77** (1995) pp. 321–357
25. Falappa, M.A., Kern-Isberner, G., Reis, M., Simari, G.R.: Prioritized and non-prioritized multiple change on belief bases. *J. Philosophical Logic* **41**(1) (2012) 77–113