



**INSTITUT FÜR INFORMATIONSSYSTEME
ABTEILUNG WISSENSBASIERTE SYSTEME**

**NONMONOTONIC PROBABILISTIC
REASONING UNDER VARIABLE-STRENGTH
INHERITANCE WITH OVERRIDING**

Thomas Lukasiewicz

INFSYS RESEARCH REPORT 1843-03-02
JANUARY 2003

Institut für Informationssysteme
Abtg. Wissensbasierte Systeme
Technische Universität Wien
Favoritenstraße 9-11
A-1040 Wien, Austria
Tel: +43-1-58801-18405
Fax: +43-1-58801-18493
sek@kr.tuwien.ac.at
www.kr.tuwien.ac.at



TECHNISCHE UNIVERSITÄT WIEN

INFSYS RESEARCH REPORT

INFSYS RESEARCH REPORT 1843-03-02, JANUARY 2003

NONMONOTONIC PROBABILISTIC REASONING UNDER VARIABLE-
STRENGTH INHERITANCE WITH OVERRIDING

20 JANUARY 2003

Thomas Lukasiewicz¹

Abstract. We present new probabilistic generalizations of Pearl's entailment in System Z [50, 32] and Lehmann's lexicographic entailment [39], called z_λ - and lex_λ -entailment, which are parameterized through a value $\lambda \in [0, 1]$ that describes the strength of the inheritance of purely probabilistic knowledge. In the special cases of $\lambda = 0$ and $\lambda = 1$, the notions of z_λ - and lex_λ -entailment coincide with the probabilistic notions of z - and lex -entailment recently introduced by the author in [45] and [46], respectively. We show that the notions of z_λ - and lex_λ -entailment have similar properties as their classical counterparts. In particular, they both satisfy the rationality postulates of System P and the property of Rational Monotonicity. Moreover, z_λ -entailment is weaker than lex_λ -entailment, and both z_λ - and lex_λ -entailment are proper generalizations of their classical counterparts.

¹Dipartimento di Informatica e Sistemistica, Università di Roma "La Sapienza", Via Salaria 113, I-00198 Rome, Italy; e-mail: lukasiewicz@dis.uniroma1.it. Alternate address: Institut für Informationssysteme, Technische Universität Wien, Favoritenstraße 9-11, A-1040 Vienna, Austria; e-mail: lukasiewicz@kr.tuwien.ac.at.

Acknowledgements: This research has been supported by a Marie Curie Individual Fellowship of the European Community programme "Human Potential" under contract number HPMF-CT-2001-001286 (Disclaimer: The author is solely responsible for information communicated and the European Commission is not responsible for any views or results expressed) and by the Austrian Science Fund under project N Z29-INF.

Copyright © 2003 by the authors

Contents

1	Introduction	1
2	Preliminaries	3
2.1	Probabilistic Knowledge Bases	3
2.2	Model-Theoretic Probabilistic Logic	4
2.3	Probabilistic Logic under Coherence	5
3	Nonmonotonic Probabilistic Logics	5
3.1	Key Ideas	5
3.2	λ -Consistency	6
3.3	System Z of Strength λ	7
3.4	Lexicographic Entailment of Strength λ	8
4	Properties	9
4.1	General Nonmonotonic Properties	9
4.2	Relationship between Probabilistic Formalisms	10
4.3	Relationship to Classical Formalisms	10
5	Special Cases	11
5.1	System Z and Lexicographic Entailment of Strength 0	11
5.2	System Z and Lexicographic Entailment of Strength 1	12
6	Summary and Outlook	12
A	Appendix: Proofs for Section 3	13
B	Appendix: Proofs for Section 4	13

1 Introduction

During the recent decades, there has been a significant amount of research in AI that focuses on probabilistic reasoning with interval restrictions for conditional probabilities, also called *conditional constraints* [42].

For example, suppose that we have the knowledge “ostriches are birds”, “birds have legs”, “birds fly with a probability of at least 0.95”, and “ostriches fly with a probability of at most 0.05”. What do we then conclude about the property of having legs of birds (resp., ostriches) and their ability to fly?

One important approach for handling conditional constraints is model-theoretic probabilistic logic, which can be traced back to Boole [10]. There is a wide spectrum of formal languages that have been explored in model-theoretic probabilistic logic, ranging from constraints for unconditional and conditional events [15, 19, 2, 18, 23, 33, 47, 41, 42, 44, 48] to rich languages that specify linear inequalities over events [21]. The main algorithmic tasks related to model-theoretic probabilistic logic are deciding satisfiability, deciding logical consequence, and computing tight logically entailed intervals.

In model-theoretic probabilistic logic, we conclude from the above knowledge that both birds and ostriches have legs, and that birds (resp., ostriches) fly with a probability of at least 0.95 (resp., at most 0.05).

Another important approach to probabilistic reasoning with conditional constraints is based on the coherence principle of de Finetti and generalizations of it [7, 11, 12, 13, 14, 27, 28, 29, 54], or on similar principles that have been adopted for lower and upper probabilities [51, 57]. The main tasks in this framework are checking the consistency of a probabilistic assessment, and the propagation of a given assessment to further conditional events.

In coherence-based probabilistic logic, we conclude from the above knowledge that birds (resp., ostriches) have (resp., do not have) legs, and that they fly with a probability of at least 0.95 (resp., at most 0.05).

The relationship between model-theoretic and coherence-based probabilistic logic has been recently explored in [9]. In particular, it turned out that probabilistic entailment under coherence is strictly weaker than model-theoretic probabilistic entailment, while the notion of consistency in probabilistic logic under coherence is strictly stronger than the notion of satisfiability in model-theoretic probabilistic logic. Furthermore, it has been shown that probabilistic entailment under coherence is a generalization of classical default entailment in System P , while model-theoretic probabilistic entailment is well-known to be a generalization of model-theoretic entailment in classical propositional logics.

Recently, the author has shown in [45, 46] that other formalisms for default reasoning from conditional knowledge bases can be extended to the probabilistic framework of conditional constraints, in order to overcome some serious drawbacks of model-theoretic and coherence-based probabilistic logic.

The literature contains several different proposals for default reasoning from conditional knowledge bases and extensive work on its desired properties. The core of these properties are the rationality postulates of System P proposed by Kraus et al. [34]. It turned out that these rationality postulates constitute a sound and complete axiom system for several classical model-theoretic entailment relations under uncertainty measures on worlds. In detail, they characterize classical model-theoretic entailment under preferential structures [55, 34], infinitesimal probabilities [1, 49], possibility measures [16], and world rankings [56, 31]. They also characterize an entailment relation based on conditional objects [17]. That these equivalences are not incidental is shown by Friedman and Halpern [22], who prove that many approaches are expressible as plausibility measures and thus they must, under some weak natural conditions, inevitably amount to the same notion of inference. A survey of the above relationships is given in [5, 24].

Mainly to solve problems with irrelevant information, the notion of rational closure as a more adventurous notion of entailment was introduced by Lehmann [38, 40]. It is equivalent to entailment in System Z by Pearl [50], to the least specific possibility entailment by Benferhat et al. [4], and to a conditional (modal)

logic-based entailment by Lamarre [37]. Finally, mainly to solve problems with property inheritance from classes to exceptional subclasses, the maximum entropy approach was proposed by Goldszmidt et al. [30]; lexicographic entailment was introduced by Lehmann [39] and Benferhat et al. [3]; conditional entailment was proposed by Geffner [25, 26]; and a belief function approach was suggested by Benferhat et al. [6].

The main ideas behind the probabilistic generalizations of sophisticated default reasoning formalisms in [45, 46] can be summarized as follows:

- The work [45] introduces probabilistic generalizations of Pearl’s entailment in System Z [50, 32] and Lehmann’s lexicographic entailment [39], which lie between model-theoretic probabilistic entailment and probabilistic entailment under coherence. That is, the new notions of entailment generalize their classical counterparts, they are stronger than entailment under coherence, and weaker than model-theoretic entailment. Roughly, the main difference between model-theoretic entailment and entailment under coherence is that the former realizes an inheritance of logical knowledge, while the latter does not. Intuitively, the new formalisms now add a strategy for resolving inconsistencies to model-theoretic probabilistic logic, and a restricted form of inheritance of logical knowledge to probabilistic logic under coherence. This is why the new notions of entailment are weaker than entailment in model-theoretic probabilistic logic and stronger than entailment in coherence-based probabilistic logic. The new formalisms can especially be used in place of model-theoretic probabilistic entailment when one wants to resolve inconsistencies related to conditioning on zero events.
- The companion paper [46] presents similar probabilistic generalizations of Pearl’s entailment in System Z [50, 32], Lehmann’s lexicographic entailment [39], and Geffner’s conditional entailment [25, 26]. The formalisms in [46], however, behave quite differently from the ones in [45]. Roughly, entailment in model-theoretic probabilistic logic realizes some inheritance of logical knowledge, but no inheritance of purely probabilistic knowledge. The new formalisms in [46] now add an inheritance of purely probabilistic knowledge and a strategy for resolving inconsistencies (due to the inheritance of logical and purely probabilistic knowledge) to entailment in model-theoretic probabilistic logic. This is why they are generally much stronger than entailment in model-theoretic probabilistic logic. Thus, they are especially useful where the notion of model-theoretic entailment is too weak, for example, in probabilistic logic programming [44, 43]. Other applications are deriving degrees of belief from statistical knowledge and degrees of belief, handling inconsistencies in probabilistic knowledge bases, and probabilistic belief revision.

In the present paper, we define a general approach to nonmonotonic probabilistic reasoning, which subsumes the two approaches in [45] and [46] as special cases. Roughly, the main idea behind this new approach is to add to the notion of logical (resp., g -coherent) entailment (i) some inheritance of purely probabilistic (resp., logical and purely probabilistic) knowledge, where the inheritance of purely probabilistic knowledge is controlled by a strength $\lambda \in [0, 1]$, and (ii) a mechanism for resolving inconsistencies due to the inheritance of logical and purely probabilistic knowledge.

The main contributions of this paper can be summarized as follows:

- We present new probabilistic generalizations of Pearl’s entailment in System Z [50, 32] and Lehmann’s lexicographic entailment [39], which are parameterized through a value $\lambda \in [0, 1]$ that describes the strength of the inheritance of purely probabilistic knowledge. In the special case of $\lambda = 0$ (resp., $\lambda = 1$), these new formalisms coincide with the formalisms presented in [45] (resp., [46]).
- We show that the new probabilistic formalisms of strength λ have similar properties as their classical counterparts. In particular, they both satisfy the rationality postulates of System P and the prop-

erty of Rational Monotonicity. Furthermore, entailment in System Z of strength λ is weaker than lexicographic entailment of strength λ .

- We show that the notions of entailment in System Z of strength λ and of lexicographic entailment of strength λ are proper generalizations of their classical counterparts. Moreover, they are weaker than some notion of logical entailment in model-theoretic probabilistic logic, and under certain conditions they coincide with this notion of entailment.

The rest of this paper is organized as follows. Section 2 gives some technical preliminaries, and recalls basic concepts from model-theoretic and coherence-based probabilistic logic. In Section 3, we introduce the new notions of entailment in System Z of strength λ , and of lexicographic entailment of strength λ . Section 4 explores some general properties of the new formalisms. In Section 5, we review the special cases of strength $\lambda = 0$ and $\lambda = 1$. Section 6 summarizes the main results and gives an outlook on future research. Note that detailed proofs of all results are given in Appendices A and B.

2 Preliminaries

In this section, we first recall probabilistic knowledge bases. We then recall the notions of satisfiability and of logical entailment from model-theoretic probabilistic logic, and the notions of g-coherence and of g-coherent entailment from probabilistic logic under coherence.

2.1 Probabilistic Knowledge Bases

We now recall the concept of a probabilistic knowledge base. We start by defining logical constraints and probabilistic formulas, which are interpreted by probability distributions over a set of possible worlds.

We assume a set of *basic events* $\Phi = \{p_1, \dots, p_n\}$ with $n \geq 1$. We use \perp and \top to denote *false* and *true*, respectively. We define *events* by induction as follows. Every element of $\Phi \cup \{\perp, \top\}$ is an event. If ϕ and ψ are events, then also $\neg\phi$ and $(\phi \wedge \psi)$. A *conditional event* is an expression of the form $\psi|\phi$ with events ψ and ϕ . A *conditional constraint* is an expression of the form $(\psi|\phi)[l, u]$ with events ψ, ϕ , and real numbers $l, u \in [0, 1]$. We define *probabilistic formulas* by induction as follows. Every conditional constraint is a probabilistic formula. If F and G are probabilistic formulas, then also $\neg F$ and $(F \wedge G)$. We use $(F \vee G)$ and $(F \Leftarrow G)$ to abbreviate $\neg(\neg F \wedge \neg G)$ and $\neg(\neg F \wedge G)$, respectively, where F and G are either two events or two probabilistic formulas, and adopt the usual conventions to eliminate parentheses. A *logical constraint* is an event of the form $\psi \Leftarrow \phi$.

A *world* I is a truth assignment to the basic events in Φ (that is, a mapping $I: \Phi \rightarrow \{\mathbf{true}, \mathbf{false}\}$), which is inductively extended to all events by $I(\perp) = \mathbf{false}$, $I(\top) = \mathbf{true}$, $I(\neg\phi) = \mathbf{true}$ iff $I(\phi) = \mathbf{false}$, and $I((\phi \wedge \psi)) = \mathbf{true}$ iff $I(\phi) = I(\psi) = \mathbf{true}$. We use \mathcal{I}_Φ to denote the set of all worlds for Φ . A world I *satisfies* an event ϕ , or I is a *model* of ϕ , denoted $I \models \phi$, iff $I(\phi) = \mathbf{true}$. We extend worlds I to conditional events $\psi|\phi$ by $I(\psi|\phi) = \mathbf{true}$ iff $I \models \psi \wedge \phi$, $I(\psi|\phi) = \mathbf{false}$ iff $I \models \neg\psi \wedge \phi$, and $I(\psi|\phi) = \mathbf{indeterminate}$ iff $I \models \neg\phi$. A *probabilistic interpretation* Pr is a probability function on \mathcal{I}_Φ (that is, a mapping $Pr: \mathcal{I}_\Phi \rightarrow [0, 1]$ such that all $Pr(I)$ with $I \in \mathcal{I}_\Phi$ sum up to 1). The *probability* of an event ϕ in the probabilistic interpretation Pr , denoted $Pr(\phi)$, is the sum of all $Pr(I)$ such that $I \in \mathcal{I}_\Phi$ and $I \models \phi$. For events ϕ and ψ with $Pr(\phi) > 0$, we write $Pr(\psi|\phi)$ to abbreviate $Pr(\psi \wedge \phi) / Pr(\phi)$. The *truth* of logical constraints and probabilistic formulas F in a probabilistic interpretation Pr , denoted $Pr \models F$, is defined as follows:

- $Pr \models \psi \Leftarrow \phi$ iff $Pr(\psi \wedge \phi) = Pr(\phi)$.

- $Pr \models (\psi|\phi)[l, u]$ iff $Pr(\phi) = 0$ or $Pr(\psi|\phi) \in [l, u]$.
- $Pr \models \neg F$ iff not $Pr \models F$.
- $Pr \models (F \wedge G)$ iff $Pr \models F$ and $Pr \models G$.

We say Pr satisfies F , or Pr is a model of F , iff $Pr \models F$. We say Pr satisfies a set of logical constraints and probabilistic formulas \mathcal{F} , or Pr is a model of \mathcal{F} , denoted $Pr \models \mathcal{F}$, iff Pr is a model of all $F \in \mathcal{F}$.

A probabilistic knowledge base $KB = (L, P)$ consists of a finite set of logical constraints L and a finite set of conditional constraints P such that (i) $l \leq u$ for all $(\psi|\phi)[l, u] \in P$, and (ii) $\psi_1|\phi_1 \neq \psi_2|\phi_2$ for any two distinct $(\psi_1|\phi_1)[l_1, u_1], (\psi_2|\phi_2)[l_2, u_2] \in P$. The following example illustrates the syntactic notion of a probabilistic knowledge base.

Example 2.1 The knowledge “ostriches are birds”, “birds have legs”, “birds fly with a probability of at least 0.95”, and “ostriches fly with a probability of at most 0.05” in our introductory example can be expressed by the following probabilistic knowledge base $KB = (L, P)$:

$$\begin{aligned} L &= \{bird \Leftarrow ostrich\}, \\ P &= \{(legs|bird)[1, 1], (fly|bird)[0.95, 1], (fly|ostrich)[0, 0.05]\}. \quad \square \end{aligned}$$

2.2 Model-Theoretic Probabilistic Logic

We now recall the model-theoretic notions of satisfiability and of logical entailment for probabilistic knowledge bases.

A set of logical constraints and probabilistic formulas \mathcal{F} is *satisfiable* iff a model of \mathcal{F} exists. A conditional constraint $(\psi|\phi)[l, u]$ is a *logical consequence* of \mathcal{F} , denoted $\mathcal{F} \models (\psi|\phi)[l, u]$, iff each model of \mathcal{F} is also a model of $(\psi|\phi)[l, u]$. It is a *tight logical consequence* of \mathcal{F} , denoted $\mathcal{F} \models_{tight} (\psi|\phi)[l, u]$, iff l (resp., u) is the infimum (resp., supremum) of $Pr(\psi|\phi)$ subject to all models Pr of \mathcal{F} with $Pr(\phi) > 0$. Here, we define $l = 1$ and $u = 0$, when $\mathcal{F} \models (\phi|\top)[0, 0]$.

A probabilistic knowledge base $KB = (L, P)$ is *satisfiable* iff $L \cup P$ is satisfiable. A conditional constraint $(\psi|\phi)[l, u]$ is a *logical consequence* of KB , denoted $KB \models (\psi|\phi)[l, u]$, iff $L \cup P \models (\psi|\phi)[l, u]$. It is a *tight logical consequence* of KB , denoted $KB \models_{tight} (\psi|\phi)[l, u]$, iff $L \cup P \models_{tight} (\psi|\phi)[l, u]$. We give an example to illustrate the above concepts.

Example 2.2 Consider again the probabilistic knowledge base $KB = (L, P)$ of Example 2.1. In model-theoretic probabilistic logic, KB represents the *logical knowledge* “all ostriches are birds” and “all birds have legs” (that is, in model-theoretic probabilistic logic, a logical constraint $\psi \Leftarrow \phi \in L$ has the same meaning as a conditional constraint $(\psi|\phi)[1, 1] \in P$), and the *probabilistic knowledge* “birds fly with a probability of at least 0.95” and “ostriches fly with a probability of at most 0.05”. It is not difficult to see that KB is satisfiable, and that some tight logical consequences of KB are given by:

$$\begin{aligned} KB \models_{tight} (legs|bird)[1, 1], \quad KB \models_{tight} (fly|bird)[0.95, 1], \\ KB \models_{tight} (legs|ostrich)[1, 1], \quad KB \models_{tight} (fly|ostrich)[0, 0.05]. \end{aligned}$$

Hence, under logical entailment, the logical property of having legs is inherited from the class of birds down to the subclass of ostriches. \square

2.3 Probabilistic Logic under Coherence

We now recall the notions of g-coherence and of g-coherent entailment. We define them by using some characterizations through concepts from default reasoning [9]. We first give some preparative definitions.

A probabilistic interpretation Pr verifies a conditional constraint $(\psi|\phi)[l, u]$ iff $Pr(\phi) > 0$ and $Pr \models (\psi|\phi)[l, u]$. We say Pr falsifies $(\psi|\phi)[l, u]$ iff $Pr(\phi) > 0$ and $Pr \not\models (\psi|\phi)[l, u]$. A set of conditional constraints P tolerates a conditional constraint $(\psi|\phi)[l, u]$ under a set of logical constraints L iff $L \cup P$ has a model that verifies $(\psi|\phi)[l, u]$. We say P is under L in conflict with $(\psi|\phi)[l, u]$ iff no model of $L \cup P$ verifies $(\psi|\phi)[l, u]$.

A conditional constraint ranking σ on a probabilistic knowledge base $KB = (L, P)$ maps each element of P to a nonnegative integer. It is admissible with KB iff every $P' \subseteq P$ that is under L in conflict with some $C \in P$ contains a conditional constraint C' such that $\sigma(C') < \sigma(C)$.

We are now ready to define the concept of g-coherence for KB . A probabilistic knowledge base KB is g-coherent iff there exists a conditional constraint ranking on KB that is admissible with KB .

We next define the notion of g-coherent entailment. Let $KB = (L, P)$ be a g-coherent probabilistic knowledge base, and let $(\psi|\phi)[l, u]$ be a conditional constraint. Then, $(\psi|\phi)[l, u]$ is a g-coherent consequence of KB , denoted $KB \Vdash^g (\psi|\phi)[l, u]$, iff $(L, P \cup \{(\psi|\phi)[p, p]\})$ is not g-coherent for all $p \in [0, l) \cup (u, 1]$. We say $(\psi|\phi)[l, u]$ is a tight g-coherent consequence of KB , denoted $KB \Vdash_{tight}^g (\psi|\phi)[l, u]$, iff l (resp., u) is the infimum (resp., supremum) of p subject to all g-coherent $(L, P \cup \{(\psi|\phi)[p, p]\})$. The following example illustrates the notions of g-coherence and g-coherent entailment.

Example 2.3 Consider again the probabilistic knowledge base $KB = (L, P)$ of Example 2.1. In probabilistic logic under coherence, KB represents the logical knowledge “all ostriches are birds”, the default logical knowledge “generally, birds have legs” (that is, in probabilistic logic under coherence, a logical constraint $\psi \Leftarrow \phi \in L$ does not have the same meaning as a conditional constraint $(\psi|\phi)[1, 1] \in P$), and the default probabilistic knowledge “generally, birds fly with a probability of at least 0.95” and “generally, ostriches fly with a probability of at most 0.05”. It is not difficult to see that KB is g-coherent, and that some tight g-coherent consequences of KB are given by:

$$\begin{aligned} KB \Vdash_{tight}^g (legs|bird)[1, 1], \quad KB \Vdash_{tight}^g (fly|bird)[0.95, 1], \\ KB \Vdash_{tight}^g (legs|ostrich)[0, 1], \quad KB \Vdash_{tight}^g (fly|ostrich)[0, 0.05]. \end{aligned}$$

Hence, under g-coherent entailment, the logical property of having legs is *not* inherited from the class of birds down to the subclass of ostriches. \square

3 Nonmonotonic Probabilistic Logics

In this section, we introduce new probabilistic generalizations of Pearl’s entailment in System Z and Lehmann’s lexicographic entailment. The new probabilistic formalisms are parameterized through a value $\lambda \in [0, 1]$ that describes the *strength* of the inheritance of purely probabilistic knowledge.

We first describe the main ideas behind the new formalisms, we then define the concept of λ -consistency for probabilistic knowledge bases, and we finally introduce the new notions of z_λ - and lex_λ -entailment.

3.1 Key Ideas

The property of *inheritance of knowledge* along subclass relationships can be divided into the properties of inheritance of logical knowledge and of inheritance of purely probabilistic knowledge. The *inheritance of*

logical knowledge (L-INH) is the following property (for all events ψ , ϕ , and ϕ^* , all probabilistic knowledge bases KB , and all $c \in \{0, 1\}$):

L-INH. If $KB \models (\psi|\phi)[c, c]$ and $\phi \Leftarrow \phi^*$ is valid, then $KB \models (\psi|\phi^*)[c, c]$.

The *inheritance of purely probabilistic knowledge (P-INH)* is defined as follows (for all events ψ , ϕ , and ϕ^* , all probabilistic knowledge bases KB , and all intervals $[l, u] \subseteq [0, 1]$ different from $[0, 0]$, $[1, 1]$, and $[1, 0]$):

P-INH. If $KB \models (\psi|\phi)[l, u]$ and $\phi \Leftarrow \phi^*$ is valid, then $KB \models (\psi|\phi^*)[l, u]$.

It is not difficult to verify that the notion of logical entailment satisfies (*L-INH*), but does not satisfy (*P-INH*), whereas the notion of g-coherent entailment satisfies neither (*L-INH*) nor (*P-INH*).

The basic idea behind the new probabilistic generalizations of Pearl's entailment in System Z and Lehmann's lexicographic entailment in this paper is that they add to the notion of logical (resp., g-coherent) entailment (i) some inheritance of purely probabilistic (resp., logical and purely probabilistic) knowledge, where the inheritance of purely probabilistic knowledge depends on a strength $\lambda \in [0, 1]$, and (ii) a mechanism for resolving inconsistencies due to the inheritance of logical and purely probabilistic knowledge.

The strength $\lambda \in [0, 1]$ determines to which extent purely probabilistic knowledge is inherited from classes down to subclasses. In the extreme cases of $\lambda = 0$ and $\lambda = 1$, purely probabilistic knowledge is not inherited at all [45] and completely inherited [46], respectively, while for $0 < \lambda < 1$, given the interval $[l, u]$ for the property of a class, some interval $[r, s] \supseteq [l, u]$ is inherited down to all subclasses, where the largeness of the interval $[r, s]$ depends on the strength λ (roughly, the smaller is λ , the larger is $[r, s]$).

3.2 λ -Consistency

We now introduce the notion of λ -consistency for probabilistic knowledge bases. We first give some preparative definitions.

A probabilistic interpretation Pr λ -verifies (resp., λ -falsifies) a conditional constraint $(\psi|\phi)[l, u]$ iff Pr verifies (resp., falsifies) $(\psi|\phi)[l, u]$ and $Pr(\phi) \geq \lambda$. A set of conditional constraints P λ -tolerates a conditional constraint C under a set of logical constraints L iff $L \cup P$ has a model that λ -verifies C . We say P is under L in λ -conflict with C iff no model of $L \cup P$ λ -verifies C . A conditional constraint ranking σ on a probabilistic knowledge base $KB = (L, P)$ is λ -admissible with KB iff every $P' \subseteq P$ that is under L in λ -conflict with some $C \in P$ contains some C' such that $\sigma(C') < \sigma(C)$.

We are now ready to define the notion of λ -consistency for probabilistic knowledge bases KB . We say KB is λ -consistent iff there exists a conditional constraint ranking σ on KB that is λ -admissible with KB .

The following theorem characterizes the λ -consistency of $KB = (L, P)$ through the existence of an ordered partition of P .

Theorem 3.1 *A probabilistic knowledge base $KB = (L, P)$ is λ -consistent iff there exists an ordered partition (P_0, \dots, P_k) of P such that every P_i , $0 \leq i \leq k$, is the set of all $C \in \bigcup_{j=i}^k P_j$ that are λ -tolerated under L by $\bigcup_{j=i}^k P_j$.*

We next give some other useful definitions. A *probability ranking* κ maps each probabilistic interpretation on \mathcal{I}_Φ to a member of $\{0, 1, \dots\} \cup \{\infty\}$ such that $\kappa(Pr) = 0$ for at least one interpretation Pr . We use the expression $\phi \succeq \lambda$ to abbreviate the probabilistic formula $\neg(\phi|\top)[0, 0] \wedge (\phi|\top)[\lambda, 1]$.

It is extended to all logical constraints and probabilistic formulas F as follows. If F is satisfiable, then $\kappa(F) = \min\{\kappa(Pr) \mid Pr \models F\}$; otherwise, $\kappa(F) = \infty$. A probability ranking κ is λ -admissible with a probabilistic knowledge base $KB = (L, P)$ iff $\kappa(\neg F) = \infty$ for all $F \in L$ and $\kappa(\phi \succeq \lambda) < \infty$ and $\kappa(\phi \succeq \lambda \wedge (\psi|\phi)[l, u]) < \kappa(\phi \succeq \lambda \wedge \neg(\psi|\phi)[l, u])$ for all $(\psi|\phi)[l, u] \in P$.

3.3 System Z of Strength λ

We now define a generalization of Pearl's entailment in System Z [50, 32] of strength $\lambda \in [0, 1]$ for λ -consistent probabilistic knowledge bases $KB = (L, P)$. The new notion of entailment is linked to an ordered partition of P , a conditional constraint ranking z_λ on KB , and a probability ranking κ^{z_λ} .

The z_λ -partition of KB is the unique ordered partition (P_0, \dots, P_k) of P such that every P_i , $i \in \{0, \dots, k\}$, is the set of all $C \in \bigcup_{j=i}^k P_j$ that are λ -tolerated under L by $\bigcup_{j=i}^k P_j$. Observe that by Theorem 3.1, every λ -consistent probabilistic knowledge base KB has a z_λ -partition. The following two examples show some z_λ -partitions.

Example 3.2 Let the probabilistic knowledge base $KB = (L, P)$ be given by:

$$\begin{aligned} L &= \{bird \Leftarrow eagle\}, \\ P &= \{(legs|bird)[1, 1], (fly|bird)[0.95, 1]\}. \end{aligned}$$

Then, for every $\lambda \in [0, 1]$, the z_λ -partition of KB is given as follows:

$$(P_0) = (\{(legs|bird)[1, 1], (fly|bird)[0.95, 1]\}). \quad \square$$

Example 3.3 For every $\lambda \in [0, 0.052632]$, the z_λ -partition of the probabilistic knowledge base $KB = (L, P)$ of Example 2.1 is given as follows:

$$(P_0) = (\{(legs|bird)[1, 1], (fly|bird)[0.95, 1], (fly|ostrich)[0, 0.05]\}).$$

For every other $\lambda \in [0, 1]$, the z_λ -partition is given as follows:

$$(P_0, P_1) = (\{(legs|bird)[1, 1], (fly|bird)[0.95, 1]\}, \{(fly|ostrich)[0, 0.05]\}). \quad \square$$

We next define z_λ and κ^{z_λ} . For every $j \in \{0, \dots, k\}$, each $C \in P_j$ is assigned the value j under z_λ . The probability ranking κ^{z_λ} on all probabilistic interpretations Pr is then defined as follows:

$$\kappa^{z_\lambda}(Pr) = \begin{cases} \infty & \text{if } Pr \not\models L \\ 0 & \text{if } Pr \models L \cup P \\ 1 + \max_{C \in P: Pr \not\models C} z_\lambda(C) & \text{otherwise.} \end{cases}$$

The probability ranking κ^{z_λ} defines a preference relation on probabilistic interpretations as follows. For probabilistic interpretations Pr and Pr' , we say Pr is z_λ -preferable to Pr' iff $\kappa^{z_\lambda}(Pr) < \kappa^{z_\lambda}(Pr')$. A model Pr of a set of logical constraints and probabilistic formulas \mathcal{F} is a z_λ -minimal model of \mathcal{F} iff no model of \mathcal{F} is z_λ -preferable to Pr .

We are now ready to define the notion of z_λ -entailment as follows. A conditional constraint $(\psi|\phi)[l, u]$ is a z_λ -consequence of KB , denoted $KB \Vdash^{z_\lambda} (\psi|\phi)[l, u]$, iff every z_λ -minimal model of $L \cup \{\phi \succeq \lambda\}$ satisfies $(\psi|\phi)[l, u]$. We say that $(\psi|\phi)[l, u]$ is a *tight* z_λ -consequence of KB , denoted $KB \Vdash_{tight}^{z_\lambda} (\psi|\phi)[l, u]$, iff l (resp., u) is the infimum (resp., supremum) of $Pr(\psi|\phi)$ subject to all z_λ -minimal models Pr of $L \cup \{\phi \succeq \lambda\}$.

We now illustrate the notion of z_λ -entailment through some examples. The following example shows that the notion of z_λ -entailment realizes an inheritance of logical properties from classes to non-exceptional subclasses, and an inheritance of purely probabilistic properties from classes to non-exceptional subclasses, where the latter inheritance depends on the strength λ .

Example 3.4 Consider the probabilistic knowledge base KB of Example 3.2. Some tight intervals under z_λ -entailment, where $\lambda \in \{0, 0.2, 0.4, 0.6, 0.8, 1\}$, are shown in Table 1. In particular, the logical knowledge of having legs is inherited from the class of birds down to the class of eagles, independently from λ , while the purely probabilistic knowledge of being able to fly with a probability of at least 0.95 is also inherited, but this is controlled by λ . \square

However, z_λ -entailment does not inherit properties from classes to subclasses that are exceptional relative to some other property (and thus, like its classical counterpart, shows the problem of *inheritance blocking*).

Example 3.5 Consider the probabilistic knowledge base KB of Example 2.1. Some tight intervals under z_λ -entailment, where $\lambda \in \{0, 0.2, 0.4, 0.6, 0.8, 1\}$, are shown in Table 2. In particular, for $\lambda \geq 0.2$, the logical property of having legs is not inherited from the class of birds to its exceptional subclass of ostriches. Note that in the case of $\lambda = 0$, this logical property is inherited, since there is no inheritance of purely probabilistic knowledge, and thus no conflict between the abilities to fly of birds and penguins. \square

Table 1: Tight intervals under z_λ - and lex_λ -entailment from KB in Example 3.2.

Conditional Event	$\lambda = 0$	$\lambda = 0.2$	$\lambda = 0.4$	$\lambda = 0.6$	$\lambda = 0.8$	$\lambda = 1$
$legs bird$	[1, 1]	[1, 1]	[1, 1]	[1, 1]	[1, 1]	[1, 1]
$legs eagle$	[1, 1]	[1, 1]	[1, 1]	[1, 1]	[1, 1]	[1, 1]
$fly bird$	[0.95, 1]	[0.95, 1]	[0.95, 1]	[0.95, 1]	[0.95, 1]	[0.95, 1]
$fly eagle$	[0, 1]	[0.75, 1]	[0.88, 1]	[0.92, 1]	[0.94, 1]	[0.95, 1]

Table 2: Tight intervals under z_λ -entailment from KB in Example 2.1.

Conditional Event	$\lambda = 0$	$\lambda = 0.2$	$\lambda = 0.4$	$\lambda = 0.6$	$\lambda = 0.8$	$\lambda = 1$
$legs bird$	[1, 1]	[1, 1]	[1, 1]	[1, 1]	[1, 1]	[1, 1]
$legs ostrich$	[1, 1]	[0, 1]	[0, 1]	[0, 1]	[0, 1]	[0, 1]
$fly bird$	[0.95, 1]	[0.95, 1]	[0.95, 1]	[0.95, 1]	[0.95, 1]	[0.95, 1]
$fly ostrich$	[0, 0.05]	[0, 0.05]	[0, 0.05]	[0, 0.05]	[0, 0.05]	[0, 0.05]

3.4 Lexicographic Entailment of Strength λ

We next define a generalization of Lehmann's lexicographic entailment [39] of strength $\lambda \in [0, 1]$ for λ -consistent probabilistic knowledge bases KB .

We use the z_λ -partition (P_0, \dots, P_k) of $KB = (L, P)$ to define a lexicographic preference relation on probabilistic interpretations as follows. For probabilistic interpretations Pr and Pr' , we say Pr is lex_λ -preferable to Pr' iff some $i \in \{0, \dots, k\}$ exists such that $|\{C \in P_i \mid Pr \models C\}| > |\{C \in P_i \mid Pr' \models C\}|$ and $|\{C \in P_j \mid Pr \models C\}| = |\{C \in P_j \mid Pr' \models C\}|$ for all $i < j \leq k$. A model Pr of a set of logical constraints and probabilistic formulas \mathcal{F} is a lex_λ -minimal model of \mathcal{F} iff no model of \mathcal{F} is lex_λ -preferable to Pr .

We now define the notion of lex_λ -entailment as follows. A conditional constraint $(\psi|\phi)[l, u]$ is a lex_λ -consequence of KB , denoted $KB \Vdash^{lex_\lambda} (\psi|\phi)[l, u]$, iff every lex_λ -minimal model of $L \cup \{\phi \succeq \lambda\}$ satisfies $(\psi|\phi)[l, u]$. We say $(\psi|\phi)[l, u]$ is a *tight* lex_λ -consequence of KB , denoted $KB \Vdash_{tight}^{lex_\lambda} (\psi|\phi)[l, u]$, iff l (resp., u) is the infimum (resp., supremum) of $Pr(\psi|\phi)$ subject to all lex_λ -minimal models Pr of $L \cup \{\phi \succeq \lambda\}$.

The following example shows that lex_λ -entailment realizes an inheritance of properties, without showing the problem of inheritance blocking.

Example 3.6 Some tight intervals under lex_λ -entailment from the probabilistic knowledge bases of Examples 3.2 and 2.1 are shown in Tables 1 and 3, respectively. As shown in Table 3, for all λ , having legs is inherited from the class of birds down to the exceptional subclass of ostriches. \square

Table 3: Tight intervals under lex_λ -entailment from KB in Example 2.1.

Conditional Event	$\lambda = 0$	$\lambda = 0.2$	$\lambda = 0.4$	$\lambda = 0.6$	$\lambda = 0.8$	$\lambda = 1$
$legs bird$	[1, 1]	[1, 1]	[1, 1]	[1, 1]	[1, 1]	[1, 1]
$legs ostrich$	[1, 1]	[1, 1]	[1, 1]	[1, 1]	[1, 1]	[1, 1]
$fly bird$	[0.95, 1]	[0.95, 1]	[0.95, 1]	[0.95, 1]	[0.95, 1]	[0.95, 1]
$fly ostrich$	[0, 0.05]	[0, 0.05]	[0, 0.05]	[0, 0.05]	[0, 0.05]	[0, 0.05]

4 Properties

In this section, we explore some properties of z_λ - and lex_λ -entailment. We first describe some general nonmonotonic properties. We then explore the relationship between the formalisms, and the one to their classical counterparts.

4.1 General Nonmonotonic Properties

We now analyze some general nonmonotonic properties of the new probabilistic entailment semantics introduced in this paper.

We first consider the postulates *Right Weakening (RW)*, *Reflexivity (Ref)*, *Left Logical Equivalence (LLE)*, *Cut*, *Cautious Monotonicity (CM)*, and *Or* proposed by Kraus et al. [34], which are commonly regarded as being particularly desirable for any reasonable notion of nonmonotonic entailment. The following result shows that z_λ - and lex_λ -entailment both satisfy (probabilistic versions of) these postulates. Here, $KB \Vdash^s (\phi|\varepsilon \vee \varepsilon')[l, u]$, where $s \in \{z_\lambda, lex_\lambda\}$, denotes that $Pr \models (\phi|\varepsilon)[l, u] \vee (\phi|\varepsilon')[l, u]$ for all s -minimal models Pr of $L \cup \{\varepsilon \succeq \lambda \vee \varepsilon' \succeq \lambda\}$.

Theorem 4.1 *Let $s \in \{z_\lambda, lex_\lambda\}$, let $KB = (L, P)$ be a λ -consistent probabilistic knowledge base, let $\varepsilon, \varepsilon', \phi, \psi$ be events, and let $l, l', u, u' \in [0, 1]$. Then,*

RW. *If $(\phi|\top)[l, u] \Rightarrow (\psi|\top)[l', u']$ is logically valid and $KB \Vdash^s (\phi|\varepsilon)[l, u]$, then $KB \Vdash^s (\psi|\varepsilon)[l', u']$.*

Ref. $KB \Vdash^s (\varepsilon|\varepsilon)[1, 1]$.

LLE. *If $\varepsilon \Leftrightarrow \varepsilon'$ is logically valid, then $KB \Vdash^s (\phi|\varepsilon)[l, u]$ iff $KB \Vdash^s (\phi|\varepsilon')[l, u]$.*

Cut. If $KB \models^s (\varepsilon | \varepsilon') [1, 1]$ and $KB \models^s (\phi | \varepsilon \wedge \varepsilon') [l, u]$, then $KB \models^s (\phi | \varepsilon') [l, u]$.

CM. If $KB \models^s (\varepsilon | \varepsilon') [1, 1]$ and $KB \models^s (\phi | \varepsilon') [l, u]$, then $KB \models^s (\phi | \varepsilon \wedge \varepsilon') [l, u]$.

Or. If $KB \models^s (\phi | \varepsilon) [l, u]$ and $KB \models^s (\phi | \varepsilon') [l, u]$, then $KB \models^s (\phi | \varepsilon \vee \varepsilon') [l, u]$.

Another desirable property is *Rational Monotonicity (RM)* [34], which describes a restricted form of monotony and allows to ignore certain kinds of irrelevant knowledge. The next theorem shows that z_λ - and lex_λ -entailment both satisfy (a weak form of) *RM*. Here, $KB \not\models^s \neg(\varepsilon' | \varepsilon) [1, 1]$, $s \in \{z_\lambda, lex_\lambda\}$, denotes that $Pr \models (\varepsilon' | \varepsilon) [1, 1]$ for some s -minimal models Pr of $L \cup \{\varepsilon \succeq \lambda\}$.

Theorem 4.2 *Let $s \in \{z_\lambda, lex_\lambda\}$, let $KB = (L, P)$ be a λ -consistent probabilistic knowledge base, and let $\varepsilon, \varepsilon', \psi$ be events. Then,*

RM. If $KB \models^s (\psi | \varepsilon) [1, 1]$ and $KB \not\models^s \neg(\varepsilon' | \varepsilon) [1, 1]$, then $KB \models^s (\psi | \varepsilon \wedge \varepsilon') [1, 1]$.

4.2 Relationship between Probabilistic Formalisms

We now explore the relationship between z_λ - and lex_λ -entailment. The following theorem shows that z_λ -entailment is weaker than lex_λ -entailment. Moreover, it shows that lex_λ -entailment of $(\psi | \phi) [l, u]$ from $KB = (L, P)$ is weaker than logical entailment of $(\psi | \phi) [l, u]$ from $L \cup P \cup \{\phi \succeq \lambda\}$.

Theorem 4.3 *Let $KB = (L, P)$ be a λ -consistent probabilistic knowledge base, and let $(\psi | \phi) [l, u]$ be a conditional constraint. Then,*

(a) $KB \models^{z_\lambda} (\psi | \phi) [l, u]$ implies $KB \models^{lex_\lambda} (\psi | \phi) [l, u]$.

(b) $KB \models^{lex_\lambda} (\psi | \phi) [l, u]$ implies $L \cup P \cup \{\phi \succeq \lambda\} \models (\psi | \phi) [l, u]$.

In general, the converse implications do not hold, as follows from Examples 3.5 and 3.6 for $\lambda = 1$. But, in the special case when $L \cup P \cup \{\phi \succeq \lambda\}$ is satisfiable, the notions of z_λ - and lex_λ -entailment of $(\psi | \phi) [l, u]$ from $KB = (L, P)$ both coincide with the notion of logical entailment of $(\psi | \phi) [l, u]$ from $L \cup P \cup \{\phi \succeq \lambda\}$. This result is expressed by the following theorem.

Theorem 4.4 *Let $KB = (L, P)$ be a λ -consistent probabilistic knowledge base, and let $(\psi | \phi) [l, u]$ be a conditional constraint such that $L \cup P \cup \{\phi \succeq \lambda\}$ is satisfiable. Then,*

(a) $KB \models^{z_\lambda} (\psi | \phi) [l, u]$ iff $KB \models^{lex_\lambda} (\psi | \phi) [l, u]$.

(b) $KB \models^{lex_\lambda} (\psi | \phi) [l, u]$ iff $L \cup P \cup \{\phi \succeq \lambda\} \models (\psi | \phi) [l, u]$.

4.3 Relationship to Classical Formalisms

We now analyze the relationship between the new formalisms of this paper and their classical counterparts. We first give some technical preparation.

The operator γ on conditional constraints, sets of conditional constraints, and probabilistic knowledge bases replaces each conditional constraint of the form $(\psi | \phi) [1, 1]$ by the classical default $\psi \leftarrow \phi$. We use the expressions \models^z and \models^{lex} to denote the classical notions of Pearl's entailment in System Z and Lehmann's lexicographic entailment, respectively.

The following theorem shows that the new notions of z_λ - and lex_λ -entailment for λ -consistent probabilistic knowledge bases generalize their classical counterparts for ε -consistent conditional knowledge bases.

Theorem 4.5 *Let $KB = (L, P)$ be a λ -consistent probabilistic knowledge base, where $P = \{(\psi_i|\phi_i)[1, 1] \mid i \in \{1, \dots, n\}\}$, and let $(\beta|\alpha)[1, 1]$ be a conditional constraint. Then,*

- (a) $KB \models^{z_\lambda} (\beta|\alpha)[1, 1]$ iff $\gamma(KB) \sim^z \beta \leftarrow \alpha$.
- (b) $KB \models^{lex_\lambda} (\beta|\alpha)[1, 1]$ iff $\gamma(KB) \sim^{lex} \beta \leftarrow \alpha$.

5 Special Cases

The notions of z_λ - and lex_λ -entailment of strength $\lambda = 0$ and $\lambda = 1$ are special cases explored in [45] and [46], respectively. In this section, we briefly review these special cases and some of their applications.

5.1 System Z and Lexicographic Entailment of Strength 0

Roughly, the notions of z_0 - and lex_0 -entailment add to logical entailment (resp., g-coherent entailment) a strategy for resolving inconsistencies due to the inheritance of logical knowledge (resp., a restricted form of inheritance of logical knowledge). This is why the notions of z_0 - and lex_0 -entailment are weaker than logical entailment and stronger than g-coherent entailment.

Hence, z_0 - and lex_0 -entailment are refinements of both logical and g-coherent entailment. They can be used in place of logical entailment, when we want to resolve probabilistic inconsistencies related to conditioning on zero events. Here, they are especially well-suited as they coincide with logical entailment as long as we condition on non-zero events. Furthermore, z_0 - and lex_0 -entailment can be used in place of g-coherent entailment, when we also want to have a restricted form of inheritance of logical knowledge.

The following example illustrates the use of lex_0 -entailment, instead of logical entailment, in order to resolve inconsistencies related to conditioning on zero events.

Example 5.1 Let the probabilistic knowledge base $KB = (L, P)$ be given as follows:

$$\begin{aligned} L &= \{bird \Leftarrow penguin\}, \\ P &= \{(legs|bird)[1, 1], (fly|bird)[1, 1], (fly|penguin)[0, 0.05]\}. \end{aligned}$$

It is not difficult to see that KB is satisfiable, and that some tight logical consequences of KB are as follows:

$$\begin{aligned} KB \models_{tight} (legs|bird)[1, 1], \quad KB \models_{tight} (fly|bird)[1, 1], \\ KB \models_{tight} (legs|penguin)[1, 0], \quad KB \models_{tight} (fly|penguin)[1, 0]. \end{aligned}$$

Here, the tight conclusions $KB \models_{tight} (legs|penguin)[1, 0]$ and $KB \models_{tight} (fly|penguin)[1, 0]$ are due to the fact that the logical property of being able to fly is under logical entailment inherited from birds to penguins, and is incompatible there with the knowledge that penguins are able to fly with a probability of at most 0.05. That is, our knowledge about penguins is inconsistent. This means that there does not exist any model Pr of $L \cup P$ such that $Pr(penguin) > 0$, and thus we are conditioning on the zero event $penguin$.

Hence, as far as the conditioning event $penguin$ is concerned, logical entailment does not provide the desired tight conclusions from KB , which are $(legs|penguin)[1, 1]$ and $(fly|penguin)[0, 0.05]$, rather than $(legs|penguin)[1, 0]$ and $(fly|penguin)[1, 0]$, respectively.

However, KB is 0-consistent, and the tight conclusions from KB under lex_0 -entailment are given by $(legs|penguin)[1, 1]$ and $(fly|penguin)[0, 0.05]$, respectively, which coincide with the desired ones.

Note that KB is also g-coherent, and that some tight g-coherent consequences of KB are as follows:

$$\begin{aligned} KB &\| \sim_{tight}^g (legs|bird)[1, 1], \quad KB \| \sim_{tight}^g (fly|bird)[1, 1], \\ KB &\| \sim_{tight}^g (legs|penguin)[0, 1], \quad KB \| \sim_{tight}^g (fly|penguin)[0, 0.05]. \end{aligned}$$

Hence, also g-coherent entailment does not provide the desired tight conclusions from KB . \square

5.2 System Z and Lexicographic Entailment of Strength 1

Roughly, the notions of z_1 - and lex_1 -entailment add to logical entailment (i) some inheritance of purely probabilistic knowledge, and (ii) a strategy for resolving inconsistencies due to the inheritance of logical and purely probabilistic knowledge. For this reason, the notions of z_1 - and lex_1 -entailment are generally much stronger than logical entailment.

The notions of z_1 - and lex_1 -entailment can especially be used where logical entailment is too weak, for example, in probabilistic logic programming [44, 43]. Other important applications are deriving degrees of belief from statistical knowledge and degrees of belief, handling inconsistencies in probabilistic knowledge bases, and probabilistic belief revision.

In particular, in reasoning from statistical knowledge and degrees of belief, z_1 - and lex_1 -entailment show a similar behavior as reference-class reasoning [53, 35, 36, 52] in a number of uncontroversial examples. They, however, also avoid many drawbacks of reference-class reasoning. In detail, they can handle complex scenarios and even purely probabilistic subjective knowledge as input. Moreover, conclusions are drawn in a global way from all the available knowledge as a whole. See [46] for further details on these issues.

The following example illustrates the use of lex_1 -entailment for reasoning from statistical knowledge and degrees of belief.

Example 5.2 Suppose that we have the statistical knowledge “all penguins are birds”, “between 90% and 95% of all birds fly”, “at most 5% of all penguins fly”, and “at least 95% of all yellow objects are easy to see”. Moreover, assume that we believe “Sam is a yellow penguin”. What do we then conclude about Sam’s property of being easy to see? Under reference-class reasoning, which is a machinery for dealing with statistical knowledge and degrees of belief, we conclude “Sam is easy to see with a probability of at least 0.95”. This is exactly what we obtain using the notion of lex_1 -entailment. More precisely, the above statistical knowledge can be represented by the following probabilistic knowledge base $KB = (L, P)$:

$$\begin{aligned} L &= \{bird \Leftarrow penguin\}, \\ P &= \{(fly|bird)[0.9, 0.95], (fly|penguin)[0, 0.05], (easy_to_see|yellow)[0.95, 1]\}. \end{aligned}$$

It is then not difficult to verify that KB is 1-consistent, and that $(easy_to_see|yellow \wedge penguin)[0.95, 1]$ is a tight conclusion from KB under lex_1 -entailment.

Note that KB is also satisfiable and g-coherent, and that $(easy_to_see|yellow \wedge penguin)[0, 1]$ is a tight conclusion from KB under logical and g-coherent entailment. That is, under logical and g-coherent entailment from KB , we do not conclude the desired tight interval $[0.95, 1]$ for $easy_to_see|yellow \wedge penguin$. \square

6 Summary and Outlook

We have presented the notions of z_λ - and lex_λ -entailment, which are probabilistic generalizations of Pearl’s entailment in System Z and Lehmann’s lexicographic entailment. They are parameterized through a value

$\lambda \in [0, 1]$ that describes the strength of the inheritance of purely probabilistic knowledge. In the special cases of $\lambda = 0$ and $\lambda = 1$, the new probabilistic formalisms coincide with the notions of z - and lex -entailment in [45] and [46], respectively. We have shown that z_λ - and lex_λ -entailment have similar properties as their classical counterparts. In particular, they both satisfy the rationality postulates of System P and the property of Rational Monotonicity. Moreover, z_λ -entailment is weaker than lex_λ -entailment, and both z_λ - and lex_λ -entailment have proper embeddings of their classical counterparts.

An interesting topic of future research is to develop algorithms for the new probabilistic formalisms and to analyze their computational complexity, which can be done along the lines of [45, 46].

Another exciting topic of future research is to develop and explore further nonmonotonic probabilistic logics. Besides extending other classical formalisms for default reasoning from conditional knowledge bases, one may also combine the new formalisms of this paper with some probability selection technique (as e.g. maximum entropy or center of mass).

A Appendix: Proofs for Section 3

Proof of Theorem 3.1. (\Leftarrow) Assume that there exists an ordered partition (P_0, \dots, P_k) of P such that every P_i , $0 \leq i \leq k$, is the set of all $C \in \bigcup_{j=i}^k P_j$ λ -tolerated under L by $\bigcup_{j=i}^k P_j$. We then define a conditional constraint ranking σ on KB as follows. For every $j \in \{0, \dots, k\}$, each $C \in P_j$ is assigned the value j under σ . We now show that σ is λ -admissible with KB . Towards a contradiction, assume the contrary. That is, some $P' \subseteq P$ is under L in λ -conflict with some $C \in P$, and $\sigma(C') \geq \sigma(C)$ for all $C' \in P'$. But this contradicts C being λ -tolerated under L by $\{C' \in P \mid \sigma(C') \geq \sigma(C)\}$. It thus follows that σ is λ -admissible with KB , and thus KB is λ -consistent.

(\Rightarrow) Assume that KB is λ -consistent. That is, there exists a conditional constraint ranking σ on KB that is λ -admissible with KB . We now show that a partition (P_0, \dots, P_k) exists as stated in the theorem. Towards a contradiction, assume the contrary. Hence, there exists some $P' \subseteq P$ that is under L in λ -conflict with every $C \in P'$. Let C' be a member of P' of smallest value under σ . Then, C' is λ -tolerated under L by P' , which is a contradiction. Thus, a partition (P_0, \dots, P_k) exists as stated in the theorem. \square

B Appendix: Proofs for Section 4

Proof of Theorem 4.1. Let s be any semantics among z_λ and lex_λ .

RW. Assume that $KB \Vdash^s (\phi|\varepsilon)[l, u]$. That is, $Pr \models (\phi|\varepsilon)[l, u]$ for all s -minimal models Pr of $L \cup \{\varepsilon \succeq \lambda\}$. Since $(\phi|\top)[l, u] \Rightarrow (\psi|\top)[l', u']$ is logically valid, it thus follows that $Pr \models (\psi|\varepsilon)[l', u']$ for all s -minimal models Pr of $L \cup \{\varepsilon \succeq \lambda\}$. That is, $KB \Vdash^s (\psi|\varepsilon)[l', u']$.

Ref. Since every Pr satisfies $(\varepsilon|\varepsilon)[1, 1]$, it trivially holds $KB \Vdash^s (\varepsilon|\varepsilon)[1, 1]$.

LLE. Assume that $KB \Vdash^s (\phi|\varepsilon)[l, u]$. That is, $Pr \models (\phi|\varepsilon)[l, u]$ for all s -minimal models Pr of $L \cup \{\varepsilon \succeq \lambda\}$. Since $\varepsilon \Leftrightarrow \varepsilon'$ is logically valid, it follows that $Pr \models (\phi|\varepsilon')[l, u]$ for all s -minimal models Pr of $L \cup \{\varepsilon' \succeq \lambda\}$. That is, $KB \Vdash^s (\phi|\varepsilon')[l, u]$.

Cut. Assume that $KB \Vdash^s (\varepsilon|\varepsilon')[1, 1]$ and $KB \Vdash^s (\phi|\varepsilon \wedge \varepsilon')[l, u]$. That is, $Pr \models (\varepsilon|\varepsilon')[1, 1]$ and $Pr \models (\phi|\varepsilon \wedge \varepsilon')[l, u]$ for all s -minimal models Pr of $L \cup \{\varepsilon' \succeq \lambda\}$ and $L \cup \{\varepsilon \wedge \varepsilon' \succeq \lambda\}$, respectively. Hence, $Pr \models (\phi|\varepsilon')[l, u]$ for all s -minimal models Pr of $L \cup \{\varepsilon' \succeq \lambda\}$. That is, $KB \Vdash^s (\phi|\varepsilon')[l, u]$.

CM. Assume that $KB \Vdash^s (\varepsilon|\varepsilon')[1, 1]$ and $KB \Vdash^s (\phi|\varepsilon')[l, u]$. That is, $Pr \models (\varepsilon|\varepsilon')[1, 1]$ and $Pr \models (\phi|\varepsilon')[l, u]$ for all s -minimal models Pr of $L \cup \{\varepsilon' \succeq \lambda\}$. It thus follows that $Pr \models (\phi|\varepsilon \wedge \varepsilon')[l, u]$ for all s -minimal models Pr of $L \cup \{\varepsilon \wedge \varepsilon' \succeq \lambda\}$. That is, $KB \Vdash^s (\phi|\varepsilon \wedge \varepsilon')[l, u]$.

Or. Suppose that $KB \Vdash^s (\phi|\varepsilon)[l, u]$ and $KB \Vdash^s (\phi|\varepsilon')[l, u]$. That is, $Pr \models (\phi|\varepsilon)[l, u]$ and $Pr \models (\phi|\varepsilon')[l, u]$ for all s -minimal models Pr of $L \cup \{\varepsilon \succeq \lambda\}$ and $L \cup \{\varepsilon' \succeq \lambda\}$, respectively. Thus, $Pr \models (\phi|\varepsilon)[l, u] \vee (\phi|\varepsilon')[l, u]$ for all s -minimal models Pr of $L \cup \{\varepsilon \succeq \lambda \vee \varepsilon' \succeq \lambda\}$. That is, $KB \Vdash^s (\phi|\varepsilon \underline{\vee} \varepsilon')[l, u]$. \square

Proof of Theorem 4.2. Let s be any semantics among z_λ and lex_λ . Assume that $KB \Vdash^s (\psi|\varepsilon)[1, 1]$ and $KB \Vdash^s \neg(\varepsilon'|\varepsilon)[1, 1]$. That is, $Pr \models (\psi|\varepsilon)[1, 1]$ for all s -minimal models Pr of $L \cup \{\varepsilon \succeq \lambda\}$, and $Pr \models (\varepsilon'|\varepsilon)[1, 1]$ for some s -minimal model Pr of $L \cup \{\varepsilon \succeq \lambda\}$. Hence, $Pr \models (\psi|\varepsilon \wedge \varepsilon')[1, 1]$ for all s -minimal models Pr of $L \cup \{\varepsilon \wedge \varepsilon' \succeq \lambda\}$. That is, $KB \Vdash^s (\psi|\varepsilon \wedge \varepsilon')[1, 1]$. \square

Proof of Theorem 4.3. (a) Suppose that $KB \Vdash^{z_\lambda} (\psi|\phi)[l, u]$. That is, every z -minimal model Pr of $L \cup \{\phi \succeq \lambda\}$ satisfies $(\psi|\phi)[l, u]$. Since every lex_λ -minimal model Pr of $L \cup \{\phi \succeq \lambda\}$ is also a z -minimal model of $L \cup \{\phi \succeq \lambda\}$, it follows that every lex_λ -minimal model Pr of $L \cup \{\phi \succeq \lambda\}$ satisfies $(\psi|\phi)[l, u]$. That is, $KB \Vdash^{lex_\lambda} (\psi|\phi)[l, u]$.

(b) Suppose that $KB \Vdash^{lex_\lambda} (\psi|\phi)[l, u]$. That is, every lex_λ -minimal model Pr of $L \cup \{\phi \succeq \lambda\}$ is also a model of $(\psi|\phi)[l, u]$. Assume first that $L \cup P \cup \{\phi \succeq \lambda\}$ is unsatisfiable. Then, $L \cup P \cup \{\phi \succeq \lambda\} \models (\psi|\phi)[l, u]$ trivially holds. Assume next that $L \cup P \cup \{\phi \succeq \lambda\}$ is satisfiable. It then follows that a probabilistic interpretation Pr is a lex_λ -minimal model of $L \cup \{\phi \succeq \lambda\}$ iff it is a model of $L \cup P \cup \{\phi \succeq \lambda\}$. Hence, every model of $L \cup P \cup \{\phi \succeq \lambda\}$ is also a model of $(\psi|\phi)[l, u]$. That is, $KB \models (\psi|\phi)[l, u]$. \square

Proof of Theorem 4.4. Immediate, as the existence of some model Pr of $L \cup P \cup \{\phi \succeq \lambda\}$ implies that a probabilistic interpretation Pr is a model of $L \cup P \cup \{\phi \succeq \lambda\}$ iff it is a lex_λ -minimal model of $L \cup \{\phi \succeq \lambda\}$ iff it is a z_λ -minimal model of $L \cup \{\phi \succeq \lambda\}$. \square

Proof of Theorem 4.5. Let s be any semantics among z_λ and lex_λ . Observe first that (P_0, \dots, P_k) is the z_λ -partition of KB iff $(\gamma(P_0), \dots, \gamma(P_k))$ is the classical z -partition of $\gamma(KB)$. The statement then follows from the observation that every s -minimal model Pr of $L \cup \{\alpha \succeq \lambda\}$ satisfies $(\beta|\alpha)[1, 1]$ iff every classical s -minimal model I of $L \cup \{\alpha\}$ satisfies β . \square

References

- [1] E. W. Adams. *The Logic of Conditionals*, volume 86 of *Synthese Library*. D. Reidel, Dordrecht, Netherlands, 1975.
- [2] S. Amarger, D. Dubois, and H. Prade. Constraint propagation with imprecise conditional probabilities. In *Proceedings UAI-91*, pages 26–34. Morgan Kaufmann, 1991.
- [3] S. Benferhat, C. Cayrol, D. Dubois, J. Lang, and H. Prade. Inconsistency management and prioritized syntax-based entailment. In *Proceedings IJCAI-93*, pages 640–645. Morgan Kaufmann, 1993.
- [4] S. Benferhat, D. Dubois, and H. Prade. Representing default rules in possibilistic logic. In *Proceedings KR-92*, pages 673–684. Morgan Kaufmann, 1992.
- [5] S. Benferhat, D. Dubois, and H. Prade. Nonmonotonic reasoning, conditional objects and possibility theory. *Artif. Intell.*, 92(1–2):259–276, 1997.

- [6] S. Benferhat, A. Saffiotti, and P. Smets. Belief functions and default reasoning. *Artif. Intell.*, 122(1–2):1–69, 2000.
- [7] V. Biazzo and A. Gilio. A generalization of the fundamental theorem of de Finetti for imprecise conditional probability assessments. *Int. J. Approx. Reasoning*, 24(2–3):251–272, 2000.
- [8] V. Biazzo, A. Gilio, T. Lukasiewicz, and G. Sanfilippo. Probabilistic logic under coherence: Complexity and algorithms. In *Proceedings ISIPTA-01*, pages 51–61, 2001. Extended Report INFSYS RR-1843-01-04, Institut für Informationssysteme, TU Wien, 2001.
- [9] V. Biazzo, A. Gilio, T. Lukasiewicz, and G. Sanfilippo. Probabilistic logic under coherence, model-theoretic probabilistic logic, and default reasoning in System *P*. *Journal of Applied Non-Classical Logics*, 12(2):189–213, 2002.
- [10] G. Boole. *An Investigation of the Laws of Thought, on which are Founded the Mathematical Theories of Logic and Probabilities*. Walton and Maberley, London, 1854. (reprint: Dover Publications, New York, 1958).
- [11] G. Coletti. Coherent numerical and ordinal probabilistic assessments. *IEEE Trans. Syst. Man Cybern.*, 24(12):1747–1754, 1994.
- [12] G. Coletti and R. Scozzafava. Characterization of coherent conditional probabilities as a tool for their assessment and extension. *Journal of Uncertainty, Fuzziness and Knowledge-based Systems*, 4(2):103–127, 1996.
- [13] G. Coletti and R. Scozzafava. Coherent upper and lower Bayesian updating. In *Proceedings ISIPTA-99*, pages 101–110, 1999.
- [14] G. Coletti and R. Scozzafava. Conditioning and inference in intelligent systems. *Soft Computing*, 3(3):118–130, 1999.
- [15] D. Dubois and H. Prade. On fuzzy syllogisms. *Computational Intelligence*, 4(2):171–179, 1988.
- [16] D. Dubois and H. Prade. Possibilistic logic, preferential models, non-monotonicity and related issues. In *Proceedings IJCAI-91*, pages 419–424. Morgan Kaufmann, 1991.
- [17] D. Dubois and H. Prade. Conditional objects as nonmonotonic consequence relationships. *IEEE Trans. Syst. Man Cybern.*, 24(12):1724–1740, 1994.
- [18] D. Dubois, H. Prade, L. Godo, and R. L. de Mántaras. Qualitative reasoning with imprecise probabilities. *Journal of Intelligent Information Systems*, 2:319–363, 1993.
- [19] D. Dubois, H. Prade, and J.-M. Touscas. Inference with imprecise numerical quantifiers. In Z. W. Ras and M. Zemankova, editors, *Intelligent Systems*, chapter 3, pages 53–72. Ellis Horwood, 1990.
- [20] T. Eiter and T. Lukasiewicz. Default reasoning from conditional knowledge bases: Complexity and tractable cases. *Artif. Intell.*, 124(2):169–241, 2000.
- [21] R. Fagin, J. Y. Halpern, and N. Megiddo. A logic for reasoning about probabilities. *Inf. Comput.*, 87:78–128, 1990.
- [22] N. Friedman and J. Y. Halpern. Plausibility measures and default reasoning. *J. ACM*, 48(4):648–685, 2001.

- [23] A. M. Frisch and P. Haddawy. Anytime deduction for probabilistic logic. *Artif. Intell.*, 69:93–122, 1994.
- [24] D. M. Gabbay and P. Smets, editors. *Handbook on Defeasible Reasoning and Uncertainty Management Systems*. Kluwer Academic, Dordrecht, Netherlands, 1998.
- [25] H. Geffner. *Default Reasoning: Causal and Conditional Theories*. MIT Press, 1992.
- [26] H. Geffner and J. Pearl. Conditional entailment: Bridging two approaches to default reasoning. *Artif. Intell.*, 53(2–3):209–244, 1992.
- [27] A. Gilio. Probabilistic consistency of conditional probability bounds. In *Advances in Intelligent Computing*, volume 945 of *LNCS*, pages 200–209. Springer, 1995.
- [28] A. Gilio. Probabilistic reasoning under coherence in System *P*. *Ann. Math. Artif. Intell.*, 34(1–3):5–34, 2002.
- [29] A. Gilio and R. Scozzafava. Conditional events in probability assessment and revision. *IEEE Trans. Syst. Man Cybern.*, 24(12):1741–1746, 1994.
- [30] M. Goldszmidt, P. Morris, and J. Pearl. A maximum entropy approach to nonmonotonic reasoning. *IEEE Trans. Pattern Anal. Mach. Intell.*, 15(3):220–232, 1993.
- [31] M. Goldszmidt and J. Pearl. Rank-based systems: A simple approach to belief revision, belief update and reasoning about evidence and actions. In *Proceedings KR-92*, pages 661–672. Morgan Kaufmann, 1992.
- [32] M. Goldszmidt and J. Pearl. Qualitative probabilities for default reasoning, belief revision, and causal modeling. *Artif. Intell.*, 84(1–2):57–112, 1996.
- [33] J. Heinsohn. Probabilistic description logics. In *Proceedings UAI-94*, pages 311–318. Morgan Kaufmann, 1994.
- [34] S. Kraus, D. Lehmann, and M. Magidor. Nonmonotonic reasoning, preferential models and cumulative logics. *Artif. Intell.*, 14(1):167–207, 1990.
- [35] H. E. Kyburg, Jr. *The Logical Foundations of Statistical Inference*. D. Reidel, Dordrecht, Netherlands, 1974.
- [36] H. E. Kyburg, Jr. The reference class. *Philos. Sci.*, 50:374–397, 1983.
- [37] P. Lamarre. A promenade from monotonicity to non-monotonicity following a theorem prover. In *Proceedings KR-92*, pages 572–580. Morgan Kaufmann, 1992.
- [38] D. Lehmann. What does a conditional knowledge base entail? In *Proceedings KR-89*, pages 212–222. Morgan Kaufmann, 1989.
- [39] D. Lehmann. Another perspective on default reasoning. *Ann. Math. Artif. Intell.*, 15(1):61–82, 1995.
- [40] D. Lehmann and M. Magidor. What does a conditional knowledge base entail? *Artif. Intell.*, 55(1):1–60, 1992.
- [41] T. Lukasiewicz. Local probabilistic deduction from taxonomic and probabilistic knowledge-bases over conjunctive events. *Int. J. Approx. Reasoning*, 21(1):23–61, 1999.

- [42] T. Lukasiewicz. Probabilistic deduction with conditional constraints over basic events. *J. Artif. Intell. Res.*, 10:199–241, 1999.
- [43] T. Lukasiewicz. Probabilistic logic programming under inheritance with overriding. In *Proceedings UAI-01*, pages 329–336. Morgan Kaufmann, 2001.
- [44] T. Lukasiewicz. Probabilistic logic programming with conditional constraints. *ACM Transactions on Computational Logic (TOCL)*, 2(3):289–339, 2001.
- [45] T. Lukasiewicz. Nonmonotonic probabilistic logics between model-theoretic probabilistic logic and probabilistic logic under coherence. Technical Report INFSYS RR-1843-02-02, Institut für Informationssysteme, TU Wien, 2002.
- [46] T. Lukasiewicz. Probabilistic default reasoning with conditional constraints. *Ann. Math. Artif. Intell.*, 34(1–3):35–88, 2002.
- [47] C. Luo, C. Yu, J. Lobo, G. Wang, and T. Pham. Computation of best bounds of probabilities from uncertain data. *Computational Intelligence*, 12(4):541–566, 1996.
- [48] N. J. Nilsson. Probabilistic logic. *Artif. Intell.*, 28:71–88, 1986.
- [49] J. Pearl. Probabilistic semantics for nonmonotonic reasoning: A survey. In *Proceedings KR-89*, pages 505–516. Morgan Kaufmann, 1989.
- [50] J. Pearl. System Z: A natural ordering of defaults with tractable applications to default reasoning. In *Proceedings TARK-90*, pages 121–135. Morgan Kaufmann, 1990.
- [51] R. Pelessoni and P. Vicig. A consistency problem for imprecise conditional probability assessments. In *Proceedings IPMU-98*, pages 1478–1485, 1998.
- [52] J. L. Pollock. *Nomic Probabilities and the Foundations of Induction*. Oxford University Press, Oxford, 1990.
- [53] H. Reichenbach. *Theory of Probability*. University of California Press, Berkeley, CA, 1949.
- [54] R. Scozzafava. Subjective conditional probability and coherence principles for handling partial information. *Mathware Soft Comput.*, 3(1):183–192, 1996.
- [55] Y. Shoham. A semantical approach to nonmonotonic logics. In *Proceedings LICS-87*, pages 275–279. IEEE Computer Society, 1987.
- [56] W. Spohn. Ordinal conditional functions: A dynamic theory of epistemic states. In W. Harper and B. Skyrms, editors, *Causation in Decision, Belief Change, and Statistics*, volume 2, pages 105–134. Reidel, Dordrecht, Netherlands, 1988.
- [57] P. Walley. *Statistical Reasoning with Imprecise Probabilities*. Chapman and Hall, 1991.