Practical Reasoning for Uncertain Agents

Nivea de C. Ferreira, Michael Fisher, and Wiebe van der Hoek

University of Liverpool, Department of Computer Science, UK
{niveacf, michael, wiebe}@csc.liv.ac.uk

Abstract. Logical formalisation of agent behaviour is desirable, not only in order to provide a clear semantics of agent-based systems, but also to provide the foundation for sophisticated reasoning techniques to be used on, and by, the agents themselves. The possible worlds semantics offered by modal logic has proved to be a successful framework in which to model mental attitudes of agents such as beliefs, desires and intentions. The most popular choices for modeling the informational attitudes involves annotating the agent with an *S5*-like logic for knowledge, or a *KD45*-like logic for belief. However, using these logics in their standard form, an agent cannot distinguish situations in which the evidence for a certain fact is 'equally distributed' over its alternatives, from situations in which there is only one, almost negligible, counterexample to the 'fact'. Probabilistic modal logics are a way to address this, but they easily end up being both computationally and conceptually complex, for example often lacking the property of compactness. In this paper, we propose a probabilistic modal logic $P_F KD45$, in which the probabilities of the possible worlds range over a finite domain of values, while still allowing the agent to reason about infinitely many options. In this way, the logic remains compact, implying that the agent still has to consider only finitely many possibilities for probability distributions during a reasoning task. We demonstrate a sound, compact and complete axiomatisation for $P_F KD45$ and show that it has several appealing features. Then, we discuss an implemented decision procedure for the logic, and provide a small example. Finally we show that, rather than specifying them beforehand, the finite set of possible probabilities can be obtained directly from the problem specification.

1 Introduction

In both reasoning *about* agents and in reasoning *within* agents, it is vital to choose tools that allow the representation of information at an appropriate level of abstraction, yet being simple enough to be mechanised. Logical formalisations of such informational aspects have been particularly successful, often using modal logics such as *S5* for knowledge, or *KD45* for belief. However, it is clear that, in realistic scenarios, such descriptions need to incorporate uncertainty. Without such descriptive flexibility, logical approaches cannot effectively represent real-world concerns and so cannot be used as the basis for practical reasoning in agents acting with uncertain information. While there have been some steps

J.J. Alferes and J. Leite (Eds.): JELIA 2004, LNAI 3229, pp. 82–94, 2004.

in developing logics of uncertainty or logics of probability (see Section 6) many of these (for example, probabilistic modal logics) are both computationally and conceptually complex. In particular, a significant drawback is that many such approaches lack the property of *compactness*[1].

In this paper we present a new probabilistic modal logic (called $P_F KD45$) that builds upon the natural framework of Kripke models (the basis of modal logics), while allowing reasoning about uncertainty. Importantly, in this logic, the probabilities of the possible worlds range over a finite domain of values, while still allowing the agent to reason about infinitely many options. In this way, the logic remains compact, ensuring that the agent only has to consider finitely many possibilities for probability distributions during a reasoning task.

The $P_F KD45$ logic extends, in some aspects, the system $P_F D$ previously introduced in [9], which in turn was inspired by the system from [3]. The basic modal operator $P^>$ allows us to write formulas such as $P^>_{0.5}\varphi$, meaning that the "agent believes φ with probability strictly greater than 0.5". The operators (which have self-explanatory meaning) $P^\geq$, $P^<$, $P^\leq$ and $P^=$ can then be defined in terms of the above basis. Since probabilities range from 0 to 1, $P^\geq_1$ corresponds to the modal operator $\Box$ or B. An important property of the logic is that it only allows probability measures (for each world) that are within a finite base set F. Although this semantically restricts probability assignments to a finite range, it is still possible to *express and reason about* arbitrary probabilities, since there is no restriction in the *language* that mirrors this semantic restriction. But again, in the *logic*, a particular axiom (Axiom $A7$; see later) ensures that arbitrary values collapse to values in the set F. The main motivation for using F is the restoration of compactness for the logic.

Logics that allow us to express that $Prob(\varphi) \sim x$ are, in general, not compact. Witness the set of premises Γ

$$\{Prob(q) > \alpha \mid \alpha \in Q \cap [0, 1)\} \tag{1}$$

Here, we have $\Gamma \models Prob(q) = 1$, and yet there is no finite subset of Γ that proves this conclusion. This has a computational counterpart: a mechanical device verifying whether a set of premises $\{Prob(\varphi) \sim x\}$ is satisfiable in $Q \cap [0, 1]$ in principle has to check an infinite number of assignments of probabilities to formulas φ. The advantage of the $P_F KD45$ logic is that the range of allowed probabilities is within a finite base set $F \subseteq [0, 1]$.

Although the use of the base set F causes a logical restriction, it is possible to highlight some interesting aspects (cf.[9]). For instance, it we take $F = \{0, 1\}$, we have classical modal logic. Alternatively, Driankov's linguistic estimates (as in [2]) *impossible, extremely unlikely, very low chance, small chance, it may, meaningful chance, most likely, extremely likely, certain* would be modelled by a 9-element F. In other words, the granularity of F can be chosen according to the intended application of the agent. However, since one of our main interests is to use the $P_F KD45$ logic for describing and implementing uncertain agents, then

[1] Compactness in the sense that inference in terms of infinite sets coincides with inference over finite sets.

having a mechanism for directly *calculating* the set F is very desirable. For this purpose, the basic idea is to have F determined by a set of arbitrary probability values which are directly extracted from the original agent specification. (Further discussion concerning this aspect will be provided in Section 6.)

In summary, contrary to many other logical approaches to probabilistic reasoning, our logic is both compact and conceptually simple. Thus, it represents a strong candidate for representing and reasoning about uncertainty within computational agents.

The paper is organised as follows. In Section 2 we present a description of the language and, in Section 3, we provide its semantics and establish its properties. Since the focus of $P_F D$ was not on a doxastic interpretation of modalities, we also include two additional properties in the $P_F KD45$ Logic (axioms $A8$ and $A9$; see later), in order to represent $KD45$-like belief. For instance, this allows us to have a probability distribution independent of worlds, and thus ensure that nested belief formulas are equivalent to formulas without nesting. Such issues are considered in more detail in Section 3. A decision procedure for the logic has been developed and implemented, and this is presented in Section 4. Due to space restrictions, only a small motivating example showing the versatility of the approach is provided in Section 5. Finally, related work and final remarks are presented in Section 6.

2 Language Description

The language L of $P_F KD45$ consists of a countable set of propositional symbols, the logical connectives $\neg$ and $\vee$ (with standard definitions for $\bot, \top, \wedge, \rightarrow, \leftrightarrow$), and parentheses. We also define a modal operator $P_x^>$, where x is a real number within the interval $[0,1]$.

Definition 1. *A set F is a* base *for a logic $P_F KD45$ if it satisfies:*

1. *F is finite;*
2. *$\{0,1\} \subseteq F \subseteq [0,1]$;*
3. *$x, y \in F$ and $(x+y \leq 1) \Rightarrow (x+y) \in F$;*
4. *$x \in F \Rightarrow (1-x) \in F$.*

The logic is defined relative to a fixed base set $F = \{x_0, x_1, ..., x_n\} \subseteq [0,1]$. It is assumed that $x_i < x_{i+1}$, if $i < n$ (implying $0 = x_0$ and $x_n = 1$). The basic operator is $P_x^>$, with intended meaning of $P_x^> \varphi$ being: "φ is believed to have a probability strictly greater than x".

The following abbreviations are used (from now on, x and y represent arbitrary values over $[0,1]$,and x_i, x_{i+1} are elements of the base set F.):

$D1$. $P_x^{\geq} \varphi \equiv \neg P_{1-x}^{>} \neg \varphi$
$D2$. $P_x^{<} \varphi \equiv P_{1-x}^{>} \neg \varphi$
$D3$. $P_x^{\leq} \varphi \equiv \neg P_{1-x}^{<} \neg \varphi$
$D4$. $P_x^{=} \varphi \equiv \neg P_x^{>} \varphi \wedge \neg P_x^{<} \varphi$

The inference rules ($R1$ and $R2$) and axioms ($A1$–$A9$) of $P_F KD45$ are:

$R1$ From φ and $\varphi \Rightarrow \psi$ infer ψ (modus ponens)
$R2$ From φ infer $P_1^{\geq}\varphi$ (necessitation rule)

$A1$ All propositional tautologies
$A2$ $P_1^{\geq}(\varphi \to \psi) \to [(P_x^{>}\varphi \to P_x^{>}\psi) \wedge (P_x^{>}\varphi \to P_x^{\geq}\psi) \wedge (P_x^{\geq}\varphi \to P_x^{\geq}\psi)]$
$A3$ $P_1^{\geq}(\varphi \to \psi) \to (P_x^{\geq}\varphi \to P_y^{>}\psi)$ (where $y < x$)
$A4$ $P_0^{\geq}\varphi$
$A5$ $P_{x+y}^{>}(\varphi \vee \psi) \to (P_x^{>}\varphi \vee P_y^{>}\psi)$ (where $x + y \in [0, 1]$)
$A6$ $P_1^{\geq}\neg(\varphi \wedge \psi) \to ((P_x^{>}\varphi \wedge P_y^{\geq}\psi) \to P_{x+y}^{>}(\varphi \vee \psi))$ (where $x + y \in [0, 1]$)
$A7$ $P_{x_i}^{>}\varphi \to P_{x_{i+1}}^{\geq}\varphi$
$A8$ $(P_0^{>} P_x^{\geq}\varphi \to P_x^{\geq}\varphi) \wedge (P_0^{>} P_x^{\leq}\varphi \to P_x^{\leq}\varphi)$
$A9$ $(P_x^{\geq}\varphi \to P_1^{\geq} P_x^{\geq}\varphi) \wedge (P_x^{\leq}\varphi \to P_1^{\geq} P_x^{\leq}\varphi)$

The axioms $A1$–$A6$ all reflect basic properties of probabilities. Axiom $A7$ reflects the peculiarity of having a base set F: it says that, if a probability is bigger than a certain value in F, it must be at least the next value. Axioms $A8$ and $A9$ are included to emphasize the relationship with the modal logic *KD45* and they make our agents doxastically introspective. The intuition behind these additional axioms is as follows. Axiom $A8$ denotes that, if the agent assigns a positive probability to some probabilistic judgement, then it incorporates this judgement. Axiom $A9$ states that the agent is absolutely sure about its own probabilistic beliefs (the focus of [9] was not on a doxastic interpretation of the modalities, and these introspective properties were not included).

Lemma 1. *The following theorems are derivable in* $P_F KD45$*:*
$\vdash P_1^{\leq}\varphi$ *and* $\vdash P_1^{\geq}\varphi \equiv P_1^{=}\varphi$

Remark 1. We can define a *belief* operator, 'B', using $B\varphi = P_1^{\geq}\varphi$, and can then infer the following.

a) 1. $\vdash \varphi \Rightarrow \vdash B\varphi$
   2. $\vdash B(\varphi \to \psi) \to (B\varphi \to B\psi)$
   3. $\vdash \neg B\bot$
   4. $B\varphi \to BB\varphi$
   5. $\neg B\varphi \to B\neg B\varphi$
b) We say that a formula in L is *modal* if it is built from atomic propositions, using only the logical connectives and the modal operator B. We claim that for all modal formulas, φ, $P_F KD45 \vdash \varphi$ iff $KD45 \vdash \varphi$.

   Proof. The $\Leftarrow$ part follows from **a** above; the $\Rightarrow$ part will be obvious from the semantics for $P_F KD45$ given later[2].

[2] Due to space limitations, full proofs are generally omitted, but can found in the associated technical report [1].

Below we present some further theorems of $P_F KD45$, though only give one example proof. We also utilise some additional notation:

- $(\varphi \bigtriangledown \psi)$ means $((\varphi \vee \psi) \wedge \neg(\varphi \wedge \psi))$, i.e. exclusive OR;
- $x{\uparrow} = min\{y \in F \mid y > x\}$ and $x{\downarrow} = max\{y \in F \mid y < x\}$.

Now, for all φ, ψ in the language and all $x \in [0, 1]$:

$T1.$ $(P_x^{\geq}\varphi \leftrightarrow (P_x^{>}\varphi \vee P_x^{=}\varphi)) \wedge (P_x^{\leq}\varphi(P_x^{<}\varphi \vee P_x^{=}\varphi))$
$T2.$ $P_x^{>}\varphi \bigtriangledown P_x^{=}\varphi \bigtriangledown P_x^{<}\varphi$
$T3.$ $\neg(P_x^{=}\varphi \wedge P_y^{=}\varphi)$ $(y \neq x)$
$T4.$ $(\neg P_x^{<}\varphi \leftrightarrow P_x^{\geq}\varphi) \wedge (\neg P_x^{>}\varphi \leftrightarrow P_x^{\leq}\varphi)$
$T5.$ $P_x^{=}\varphi \leftrightarrow (P_x^{\geq}\varphi \wedge P_x^{\leq}\varphi)$
$T6.$ $P_x^{>}\varphi \rightarrow P_y^{>}\varphi$ $y \leq x$
$T7.$ $P_x^{=}\varphi \leftrightarrow P_{1-x}^{=}\neg\varphi$

The following lemma shows the benefit of having a finite base F: it guarantees that we can express in the language that every formula has a probability.

Lemma 2. *For all $\varphi \in L$, the following is a $P_F KD45$-theorem:*
$P_{x_0}^{=}\varphi \bigtriangledown P_{x_1}^{=}\varphi \bigtriangledown ... \bigtriangledown P_{x_n}^{=}\varphi$ *(recall: $F = \{0 = x_0, x_1, ..., x_n = 1\}$)*

$T8.$ $P_x^{>}\varphi \rightarrow P_{x_i}^{=}\varphi \bigtriangledown P_{x_{i+1}}^{=}\varphi \bigtriangledown ... \bigtriangledown P_{x_n}^{=}\varphi,$ with $x_i = x{\uparrow}$;
$T9.$ $P_x^{<}\varphi \rightarrow P_{x_0}^{=}\varphi \bigtriangledown P_{x_1}^{=}\varphi \bigtriangledown ... \bigtriangledown P_{x_i}^{=}\varphi,$ with $x_i = x{\downarrow}$;
$T10.$ $(P_x^{>}\varphi \leftrightarrow P_{x\uparrow}^{\geq}\varphi) \wedge (P_y^{<}\varphi \leftrightarrow P_{y\downarrow}^{\leq}\varphi)$ $x \in [0, 1), y \in (0, 1]$
$T11.$ $[P_1^{\geq}\neg(\varphi \wedge \psi) \wedge P_x^{=}\varphi] \rightarrow [P_y^{=}\psi \leftrightarrow P_{x+y}^{=}(\varphi \vee \psi)]$ $x, y, x + y \in [0, 1]$

$T12$ $P_0^{>}P_x^{\sim}\varphi \rightarrow P_x^{\sim}\varphi,$ $\sim$ is one of $\{<, \leq, =, >\geq\}$
$T13$ $P_x^{\sim}\varphi \rightarrow P_1^{\geq}P_x^{\sim}\varphi,$ $\sim$ is one of $\{<, \leq, =, >\geq\}$

3 Semantics and Properties

Formulas $\varphi \in L$ are interpreted on Probabilistic Kripke Models over F.

Definition 2. *For each base set, F, $\mathcal{P}_F\mathcal{KD}45$ is the class of all models $M = \langle W, P_F, \pi\rangle$ for which:*

- *W is a non-empty set (of worlds);*
- *P_F is a function $P_F : W \rightarrow F$, satisfying $\sum_{w \in W} P_F(w) = 1$*
- *π is a valuation: $W \times L \rightarrow \{true, false\}$*
- *The truth definition for formulas is defined in a standard way, the modal clause reading:*

$$(M, w) \models P_x^{\geq}(\varphi) \text{ iff } \sum_{w' \text{ s.t. } (M,w') \models \varphi} P_F(w') \geq x$$

Note that the probability distribution is independent of the world. Let us call such a structure $\mathcal{P}_F\mathcal{KD}45$.

One can relate this semantics to a more standard Kripke semantics as follows. Given $M = \langle W, P_F, \pi\rangle$, first choose an arbitrary world w in the model M. Then, let W' be $\{w\} \cup \{w' \mid P_F(w') > 0\}$. Finally, define $R'(x, y)$ iff $P_F(y) > 0$, i.e., a world is accessible (from any world) if, and only if, its probability is positive. Let $M'_w = \langle W', R', \pi'\rangle$ be the model thus obtained, with π' being the restriction of π to W. The following gives a semantic motivation for coining our system *$P_F KD45$*:

Proposition 1. *Given a $P_F KD45$ model $M = \langle W, P_F, \pi\rangle$ and a world w, let $M'_w = \langle W', R', \pi'\rangle$ be obtained as described above. Moreover, let a purely modal formula from $P_F KD45$ be a formula in which all modal operators are $P_1^{\geq}$, or, equivalently, B. Then:*

1. *for every purely modal formula φ, we have $M, w \models \varphi$ iff $M'_w \models \varphi$;*
2. *the accessibility relation R' is serial, transitive and Euclidean.*

Lemma 3. *$P_F KD45$is sound with respect to $\mathcal{P}_F\mathcal{KD}45$, i. e., $P_F KD45 \vdash \varphi \Rightarrow \mathcal{P}_F\mathcal{KD}45 \models \varphi$.*

3.1 Completeness

Let φ be a consistent formula of *$P_F KD45$*. We will show how to construct a model that satisfies φ. Let Ψ be the set of sub-formulas of φ closed under a single negation and satisfying, for any $\sim$ within $\{<, >, \leq, \geq, =\}$, $(P_x^{\sim} \in \Psi \Rightarrow \{P_{x_i}^{=} | x_i \in F\} \subseteq \Psi)$. With Ψ being finite, say $|\Psi| = k$, we can define the Ψ-maximal consistent sets as $\Gamma_1, \Gamma_2, ..., \Gamma_n, n \leq 2^k$. Let γ_i be the conjunction of formulas in $\Gamma_i, i \leq n$. Then, we have:

i. $\vdash \neg(\gamma_i \wedge \gamma_j)$, where $i \neq j$;
ii. $\vdash (\gamma_1 \vee ... \vee \gamma_n)$
iii. $\vdash \psi \leftrightarrow \gamma_{\psi 1} \vee ... \vee \gamma_{\psi x}$, where $\gamma_{\psi 1} \vee ... \vee \gamma_{\psi x}$ are exactly those γ's which contain ψ as a conjunct, for each $\psi \in \Psi$.

Since φ is consistent and, by construction of the Γ's, there is at least one Γ_φ such that $\varphi \in \Gamma_\varphi$. Given this Γ_φ, we construct a set $\Phi \supseteq \Gamma_\varphi$ as follows. From Theorem $T8$, we know that for every consistent set Γ and formula ψ, at least one set of the sequence

$$\Gamma \cup \{P_0^{=}\psi\}, \Gamma \cup \{P_{x_1}^{=}\psi\}, \ldots, \Gamma \cup \{P_{x_{n-1}}^{=}\psi\}, \Gamma \cup \{P_1^{=}\psi\} \tag{2}$$

is also consistent. Now, we obtain Φ from Γ_φ as follows:

1. let $\Phi_0 = \Gamma_\varphi$ (this set is consistent);
2. for $i = 1$ to n, we know that there is some $x \in F$ such that $\Phi_{i-1} \cup \{P_x^{=}\gamma_i\}$ will be consistent, and we make the corresponding choice for Φ_i.

We let Φ be Φ_n; this is a consistent extension of Γ_φ, which contains a probability in F for every 'world' Γ_i $(i \leq n)$. We are now ready to define our canonical model $M^c = \langle W^c, P^c_F, \pi^c \rangle$ as follows:

1. $W^c = \{\Gamma_\varphi\} \cup \{\Gamma_i \mid \exists x > 0 P^{=}_{x} \gamma_i \in \Phi\}$.
2. $P^c_F(\Gamma_i) = x \Leftrightarrow P^{=}_{x} \gamma_i \in \Phi$
3. $\pi(\Gamma_i)(p) = true$ *iff* $p \in \Gamma_i$

Lemma 4 (Coincidence Lemma). *For all* $\psi \in \Psi$ *and* $\Gamma \in W^c$

$$M^c, \Gamma \models \psi \text{ iff } \psi \in \Gamma$$

Theorem 1 (Soundness and Completeness, Finite Models). *For any formula* φ*, we have* $\mathcal{P}_F\mathcal{KD}45 \models \varphi$ *iff* $P_F KD45 \vdash \varphi$*. Moreover, every consistent formula has a finite model.*

3.2 Nested Beliefs

Considering $P_F KD45$ as a language for representing properties within individual agents, we next show that *nested* belief formulas can be removed, i. e,. any nested belief formula is equivalent to some formula given without nesting.

Lemma 5 (Independence of Probability Distribution). *Let* $M = \langle W, P_F, \pi \rangle$ *be a* $\mathcal{P}_F\mathcal{KD}45$ *model. Then:*

$$\exists w \in W (M, w) \models P^{\geq}_{\gamma} \beta \Leftrightarrow \forall u \in W, (M, u) \models P^{\geq}_{\gamma} \beta\,.$$

We now demonstrate that nested beliefs are superfluous, in $P_F KD45$. This result is a generalisation of [10, Theorem 1.7.6.4], where it is proved for *S5*, which means that their result still goes through when weakening the logic to *KD45*, and even when having probabilistic operators.

Definition 3. *We say that a formula* ψ *is in* normal form *if it is a disjunction of conjunctions of the form* $\delta = \omega \wedge P^{\geq}_{\gamma_1}\beta_1 \wedge P^{\geq}_{\gamma_2}\beta_2 \wedge ... \wedge P^{\geq}_{\gamma_n}\beta_n \wedge P^{>}_{\kappa_1}\alpha_1 \wedge P^{>}_{\kappa_2}\alpha_2 \wedge \ldots \wedge P^{>}_{\kappa_k}\alpha_k$*, where* $\omega, \beta_i, \alpha_j$*,* $(i \leq n, j \leq k)$ *are all purely propositional formulas. The formula* δ *is called the canonical conjunction and the sub-formulas* $P^{\geq}_{\gamma_i}\beta_i$ *and* $P^{>}_{\kappa_j}\alpha_j$ *are called prenex formulas.*

Lemma 6. *If* ψ *is in normal form and contains a prenex formula* σ*, then* ψ *may be supposed to have the form* $\pi \vee (\lambda \wedge \sigma)$ *where* π, λ *and* σ *are in normal form.*

Proof. ψ is in normal form, so $\psi = \delta_1 \vee \delta_2 \vee ... \vee \delta_m$, where $\delta_{i's}$ are canonical conjunctions. Suppose σ occurs in δ_m. Then σ must be some conjunct $P^{\geq}_{\gamma}$, so that δ_m can be written as $(\lambda \wedge \sigma)$. Taking π to be $(\delta_1 \vee \delta_2 \vee ... \vee \delta_{m-1})$ gives the desired result $\psi = \pi \vee (\lambda \wedge \sigma)$.

This lemma guarantees that prenex formulas can always be moved to the outermost level.

Lemma 7 (Removal of Nested Beliefs). *We have, in $\mathcal{P}_F\mathcal{KD}45$:*

$$P_\alpha^{\geq}(\pi \vee (\lambda \wedge P_\gamma^{\geq}\beta)) \leftrightarrow (P_\alpha^{\geq}(\pi \vee \lambda) \wedge P_\gamma^{\geq}\beta) \vee (P_\alpha^{\geq}\pi \wedge \neg P_\gamma^{\geq}\beta) \tag{3}$$

$$P_\alpha^{\geq}(\pi \vee (\lambda \wedge P_\gamma^{>}\beta)) \leftrightarrow (P_\alpha^{\geq}(\pi \vee \lambda) \wedge P_\gamma^{>}\beta) \vee (P_\alpha^{\geq}\pi \wedge \neg P_\gamma^{>}\beta) \tag{4}$$

Proof. We sketch the proof of (3). As $(M, s) \models P_\gamma^{\geq}\beta \vee P_\gamma^{\geq}\beta$, there are two possible cases to consider.

First Case. Assuming $(M, s) \models P_\gamma^{\geq}\beta$ we aim to show that

$$P_\alpha^{\geq}(\pi \vee (\lambda \wedge P_\gamma^{\geq}\beta)) \leftrightarrow (P_\alpha^{\geq}(\pi \vee \lambda) \wedge P_\gamma^{\geq}\beta) \tag{5}$$

For $\rightarrow$, note that $(\pi \vee (\lambda \wedge P_\gamma^{\geq}\beta)) \rightarrow (\pi \vee \lambda)$ is a tautology. Hence, the truth of $P_\alpha^{\geq}(\pi \vee (\lambda \wedge P_\gamma^{\geq}\beta))$ in s implies that of $P_\alpha^{\geq}(\pi \vee \lambda)$ in s (using $A2$). This, together with $(M, s) \models P_\gamma^{\geq}\beta$ leads to

$$(M, s) \models P_\alpha^{\geq}(\pi \vee (\lambda \wedge P_\gamma^{\geq}\beta)) \rightarrow (P_\alpha^{\geq}(\pi \vee \lambda) \wedge P_\gamma^{\geq}\beta) \tag{6}$$

and this is valid for any state since $(M, s) \models P_\gamma^{\geq}\beta \iff \forall u \in S, (M, u) \models P_\gamma^{\geq}\beta$.

Concerning the converse, from $P_\alpha^{\geq}(\pi \vee \lambda) \wedge P_\gamma^{\geq}\beta$ we have that both $P_\alpha^{\geq}(\pi \vee \lambda)$ and $P_\gamma^{\geq}\beta$ are true in all $u \in S$. λ is true. $(\forall u)$ $(M, u) \models \lambda \iff \lambda \wedge P_\gamma^{\geq}\beta$ is also true. So,

$$(M, s) \models (P_\alpha^{\geq}(\pi \vee \lambda) \wedge P_\gamma^{\geq}\beta) \rightarrow P_\alpha^{\geq}(\pi \vee (\lambda \wedge P_\gamma^{\geq}\beta)) \tag{7}$$

Then, at this point we have:

$$(M, s) \models P_\gamma^{\geq}\beta \rightarrow (P_\alpha^{\geq}(\pi \vee (\lambda \wedge P_\gamma^{\geq}\beta)) \leftrightarrow (P_\alpha^{\geq}(\pi \vee \lambda) \wedge P_\gamma^{\geq}\beta)) \tag{8}$$

The second case is analogous, giving

$$(M, s) \models \neg P_\gamma^{\geq}\beta \rightarrow (P_\alpha^{\geq}(\pi \vee (\lambda \wedge P_\gamma^{\geq}\beta)) \leftrightarrow (P_\alpha^{\geq}\pi \wedge \neg P_\gamma^{\geq}\beta)) \tag{9}$$

After considering the two cases we can, finally, use the propositional tautology $[(p \rightarrow (q \leftrightarrow (p \wedge r))) \wedge (\neg p \rightarrow (q \leftrightarrow (\neg p \wedge s)))] \rightarrow [(q \leftrightarrow ((r \wedge p) \vee (s \wedge \neg p)))]$, together with (8) and (9) to conclude (3).

We can this way bring all the probabilistic operators to the outermost level:[3]

Theorem 2. *Every formula φ is equivalent to a formula ψ in normal form, i. e., a formula without nesting of probabilistic operators.*

[3] This result seems parallel to a result [7, Theorem 3.1.30] about a language with quantifiers, which proof is given with induction on φ.

4 Decision Procedure

As previously explained, the semantic definition for $P_F KD45$-formulas is based on Probabilistic Kripke Models. For each world w there is a set of worlds that w considers possible and each one of these possible worlds is specified according to the formulas it satisfies. For instance, if in the actual world w $P_1^{\geq} p$ holds, the probability values assigned to the possible worlds where p is true sum up to 1 (which, in this case, guarantees that all the worlds where p is false have probability zero).
In other words, by definition, the probability of a formula is given by the sum of the probability values assigned to the worlds that satisfy this formula, and satisfiability of a propositional formula is given by the assignment of truth-values to its symbols. So, by evaluating formulas, we identify the worlds where those formulas are satisfied. As a result, we can obtain the values that, once assigned to the set of possible worlds, can satisfy the modal formula present in the agent's specification.

The idea is to convert the set of formulas into constraint (in)equations. The inequation components represent all the possible truth valuations for the propositional symbols. A finite set of formulas is given, and a finite set of constraint (in)equations will be generated; each formula is converted into a set of (in)equational statements.

For instance, consider that the agent specification is expressed by the set of formulas: $\{P_{0.8}^{\geq} p, P_{0.7}^{\geq} q\}$ and $F = \{0, 0.1, 0.2, ..., 0.9, 1\}$. The four possible sets of worlds (characterised by the truth-assignments) are: $p1q1$ (where both p and q hold), $p1q0$ (in which p holds and the negation of q holds), $p0q1$ (in which the q and negation of p hold) and $p0q0$ (where both negations hold).

In the given example, the set of constraints generated is:

$$
\begin{aligned}
&p0q0 + p0q1 + p1q0 + p1q1 = 1 \\
&p1q0 + p1q1 \geq 0.8 \\
&p0q1 + p1q1 \geq 0.7
\end{aligned}
$$

The first equation expresses the fact that probability values have to sum up to 1. The two inequations represent constraints on the worlds in which p holds and worlds in which q holds, respectively.

Solving the constraint (in)equations determines which are the values in set F that obey the constraints imposed by the formulas and can be, consequently, applied to the set of worlds. Therefore, the decision procedure turns out to be a mechanism for finding all the possible probability assignments for the set of possible worlds that would satisfy the specified formulas, as long as this set of formulas is consistent. Otherwise, no possible assignment exists.

As mentioned above the decision procedure converts the set of formulas into a set of constraint (in)equations. Identifying the propositional symbols is essential for determining the inequation components, and the number of components grows exponentially in the number of propositional symbols. Each formula determines which components constitute each inequation. Finally, the inequations are produced and solved.

Theorem 3 (Decision Procedure). *A formula φ in $P_F KD45$ is satisfiable if, and only if, there is a solution for the set of (in)equations generated from φ within the domain F.*

5 Example

We present a simple example to show what an agent specification might look like in the $P_F KD45$ language. This is a variety of the the common "travel agent" scenario whereby once the travel agent believes you might be interested in a holiday, (s)he sends you information. The basic formulas are given as follows (the finiteness of the domain ensures that this example can indeed be represented in a propositional language).

A. $ask(you, x) \rightarrow P^{\geq}_{0.8}\ go(you, x)$, i. e., "if you ask for information about the destination x, then I believe that you wish to go to x with probability greater than, or equal to, 0.8"
B. $P^{=}_{1}\ [go(you, x) \rightarrow buy(you, holiday, x)]$, i. e., "I believe that, if you wish to go to x, then you will buy a holiday in x"
C. $P^{>}_{0.5} buy(you, holiday, x) \rightarrow sendinfo(you, x)$, i. e., "if I believe that you will buy a holiday for x with probability greater than 0.5, I send information about holidays at x"
D. $ask(you, x)$, i. e., "you ask for information on destination x"

From D and A and $R1$ we have: $P^{\geq}_{[0.8]}\ go(you, x)$ (Res1).
From Res1, $A3$ and item B: $P^{>}_{[0.7]} buy(you, holiday, x)$ (Res2).
From Res2 and $T6$: $P^{>}_{[0.5]} buy(you, holiday, x)$ (Res3).
From Res3 and item C: $sendinfo(you, x)$
Referring to the decision procedure execution, there are three formulas to be evaluated (the ones that express degrees of beliefs):

1. $P^{\geq}_{0.8}\ go(you, x)$ (from A)
2. $P^{=}_{1}\ [go(you, x) \rightarrow buy(you, holiday, x)]$ (from B)
3. $P^{>}_{0.5} buy(you, holiday, x)$ (from C)

We obtain 6 solutions when solving the first two rules. From this set, all solutions satisfy the antecedent of the third rule (as would be expected by the formal proof given above). Which means that, whatever solution is chosen as a possible value assignment, the antecedent of rule C is true. Or, independently of the assignment, $sendinfo(you, x)$ is a logical consequence of the knowledge theory, and six assignments can be considered as options when building a model for the agent specification.

In this case, the six assignments for $[B0G0, B0G1, B1G0, B1G1]$ are: $[0, 0, 0, 1]$, $[0, 0, .1, .9]$, $[0, 0, .2, .8]$, $[.1, 0, 0, .9]$, $[.1, 0, .1, .8]$ and $[.2, 0, 0, .8]$
(where "B" represents buy(...) and "G" go(...)).

5.1 Limiting F

In this section, we elaborate on ways to automatically *generate* an appropriate base set for a formula. In particular, we will look at sets F that are generated by some number $\frac{1}{d}$. For such an F, we will write $F = \frac{1}{d}$. In general we have that satisfiability is preserved when considering bigger sets F: if $F \subseteq F'$, then $\mathcal{P}_F\mathcal{KD}45$- satisfiability implies $\mathcal{P}_{F'}\mathcal{KD}45$- satisfiability. As a consequence, we have that a formula φ is $\mathcal{P}_F\mathcal{KD}45$-satisfiable for some F if, and only if, it is $\mathcal{P}_{F'}\mathcal{KD}45$-satisfiable for some *generated* F'. So, given a formula φ, can we generate a F which is sufficient for satisfiability of φ? If we succeed in this, the user of the specification language *$P_F KD45$* need not bother about a specific F, but instead can leave the system to generate it.

To get a feeling for how sensitive the matter of satisfiability is against the choice of F, suppose we have three atoms p, q and r, and let $L(p, q, r)$ be the set of conjunctions of literals over them: $L = \{(\neg)p \wedge (\neg)q \wedge (\neg)r\}$ and let our constraint φ be:

$$\bigwedge_{\psi \in L(p,q,r)} (P_0^{>}\psi \wedge P_{0.5}^{<}\psi) \tag{10}$$

If $F = \frac{1}{4}$, there is no model for (10), since every combination of atoms ψ would have a probability of $\frac{1}{4}$, giving the disjunction $\bigvee_{\psi} \in L(p, q, r)\psi$ a 'probability' of 2, which is, of course, not possible. One easily verifies that (10) is satisfiable for a set F generated by $\frac{1}{d}$ iff $d \geq 8$, giving enough 'space' for each of the ψ's.

A range F with few elements easily gives rise to unsatisfiability. Axiom $A7$ forces one to make 'big jumps' between constraints: if we have $P_{r_i}^{>}\varphi$ for a certain $r_i \in F$, we are forced to assign φ at least the next probability in F, viz., r_{i+1}.

We now sketch a way to construct an F from the formula φ in a most cautious way. Consider the formula φ. Rewrite all the occurrences of $P_x^{\sim}\psi$ in φ in such a way, that they all have a common denominator d: every $P_x^{\sim}\psi$ gets rewritten as a $P_{\frac{m}{d}}^{>}\psi$. Let $x_1, \ldots x_m$ be all the boolean combinations of atoms from φ $(m = 2^k)$. The formula φ gives rise to a number of v inequalities I:

$$I(d) = \begin{cases} \kappa_{1_1}x_1 + \kappa_{1_2}x_2 + \ldots + \kappa_{1_m} \sim_1 \frac{n_1}{d} \\ \ldots \qquad\qquad\qquad\qquad \sim \ \ldots \\ \kappa_{v_1}x_1 + \kappa_{v_2}x_2 + \ldots + \kappa_{v_m} \sim_k \frac{k_1}{d} \end{cases}$$

Since solutions of $I(d)$ are obtained by taking linear combinations of the inequalities: it is clear that they are (as linear combinations of the right hand sides) of the form $\frac{n}{d}$, for some n. Now, take the first x_i that is not yet determined, say the tightest constraint on x_i says that is is between $\frac{n_i}{d}$ and $\frac{n_i+t}{d}$ for certain n and t. Then we can safely add the constraint $x_i = \frac{n_i+t}{2d}$ to $I(d)$ and obtain a set of inequalities $I(2d)$. Doing this iteratively gives us the following:

Conjecture 1. Let φ be a formula in our language, with denominator d. Then, φ is satisfiable for some F iff φ if satisfiable for $F_\varphi = \frac{1}{2 \cdot d \cdot 2^k}$, where k is the number of atoms occurring in φ.

6 Related Work and Conclusion

Several methods have been developed to deal with uncertain information, often being split between numerical (or quantitative) or symbolic (or qualitative) ones [12]. $P_F KD45$ is a system that combines logic and probability. In this sense, it is related to other work that showed how this combination would be possible in different ways [6]. One of those possible approaches is the interpretation of the modal belief operator according to the concept of 'likelihood' (as in [8]). In this logic, instead of using numbers to express uncertainty one would have expressions like "p is likely to be a consistent hypothesis" (as a state is taken as a set of hypotheses "true for now"). That is, a qualitative notion of likelihood rather than explicit probabilities.

$P_F KD45$ was designed for reasoning with (exact) probabilities. Its Probabilistic Kripke Model semantics is similar to the one presented in [5,4]. In their formalism, a formula is typically a boolean combination of expressions of the form $a_1 w(\varphi_1) + ... + a_k w(\varphi_k) \geq c$, where $a_1, ..., a_k, c$ are integers, and each φ_i is propositional. The restriction of having φ's as purely propositional does not apply to $P_F KD45$. Besides, the system in [5,4] includes, as axioms, all the formulas of linear inequalities; consequently, their proofs of completeness rely on results in the area of linear programming. Our logic is conceptually simpler. Finally, $P_F KD45$ differs mainly from other systems for representing beliefs and probability by allowing only a finite range of probability values, an assumption that at the same time imposes restrictions about the values that can be assigned to the possible worlds and permits the restoration of compactness for the logic.

Maybe the work closest to ours is that of [11]. It considers languages for first order probabilities, and the compactness of $P_F KD45$ easily follows from [11, Theorem 11]. They also consider the case in which all the worlds are assigned the same probability function, but for a language that forbids iteration.

In this paper, we presented $P_F KD45$, a simple and compact logic combining modal logic with probability. Despite the inclusion of new axioms and slight changes in the semantics, it was shown how the logic preserves important results about soundness, completeness, finite model and decidability of the previous system $P_F D$ [9]. In addition, new results about nested beliefs have been presented, a decision procedure for the logic has been developed, and brief examples were given showing how the language can serve as an appropriate basis concerning the informational attitudes of an agent specification language. In summary, we proposed not only a complete axiomatization for the logic, but also a decision procedure that permits us to verify satisfiability of $P_F KD45$-formulas.

The use of a finite range F of probability values is a peculiar, and important, property of our logic. Although the use of a base F causes a logical restriction, it seems possible to chose its granularity according to the intended agent's application. Besides, as discussed earlier in Section 1, the compactness that it brings has significant benefits. Furthermore, a finite range of probability values reduces the computational effort required when building a model for the agent description.

Finally, this work on $P_F KD45$ represents one step towards our main goal: an agent programming language capable of specifying and implementing agents that deal with uncertain information, together with new mechanisms for handling such uncertainty in executable specifications. Future work will concentrate on developing an executable framework combining the probabilistic approach of $P_F KD45$ with the dynamic approach of and Temporal Logics. This will allow us to capture, in our simple an compact approach, the key aspects of uncertain agents working in an uncertain world.

Acknowledgements. The first author gratefully acknowledges support by the Brazilian Government under CAPES-scholarship. The authors thank the anonymous referees for their relevant comments.

References

1. N. de Carvalho Ferreira, M. Fisher and W. van der Hoek, *A Simple Logic for Reasoning about Uncertainty*, Technical Report, Department of Computer Science, University of Liverpool, (2004). Online version: `http://www.csc.liv.ac.uk/~niveacf/techreport/`
2. D. Driankov, 'Reasoning about Uncertainty: Towards a many-valued logic of belief', *IDA annual research report*, 113–120, Linköping University, (1987).
3. M. Fattarosi-Barnaba and G. Amati. 'Modal operators with probabilistic interpretations, I'. *Studia Logica*, **48**:383–393, 1989.
4. R. Fagin and J.Y. Halpern, 'Reasoning about knowledge and probability', *Journal of the ACM*, **41**(2), 340–367, (1994).
5. R. Fagin, J.Y. Halpern, and N. Megiddo, 'A logic for reasoning about probabilities', *Information and Computation*, **87**(1), 277–291, (1990).
6. T. Fernando, 'In conjunction with qualitative probability', *Annals of Pure and Applied Logic*, **92**(3), 217–234, (1998).
7. P. Hájek and T. Havránek, *Mechanizing Hypothesis Formation*, Springer, 1978.
8. J.Y. Halpern and M.O. Rabin, 'A logic to reason about likelihood', *Artificial Intelligence*, **32**(3), 379–405, (1987).
9. W. van der Hoek, 'Some considerations on the logic PFD', *Journal of Applied Non Classical Logics*, **7**(3), 287–307, (1997).
10. J.-J.Ch. Meyer and W. van der Hoek, *Epistemic Logic for AI and Computer Science*, Cambridge University Press, 1995.
11. Z. Ognjanovic and M. Rǎskovic, 'Some first-order probability logics', *Theoretical Computer Science*, 191–212, (2000).
12. S. Parsons and A. Hunter, 'A review of uncertainty handling formalisms', in *Applications of Uncertainty Formalisms*, A. Hunter and S. Parsons (eds), Springer (1998).