

Specifying and Reasoning about Uncertain Agents

Nivea de Carvalho Ferreira, Michael Fisher & Wiebe van der Hoek

Department of Computer Science, University of Liverpool, Liverpool L69 3BX, U.K.

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 a ‘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_FKD45 , 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 axiomatization for P_FKD45 and show that it has several appealing features. Then, we discuss an implemented decision procedure for the logic, and provide a small example.

Key words:

Agents, Knowledge and Belief, Probabilistic Modal Logic, Decision Procedure

Email address: {niveacf,michael,wiebe}@csc.liv.ac.uk (Nivea de Carvalho Ferreira, Michael Fisher & Wiebe van der Hoek).

URL: <http://www.csc.liv.ac.uk/~{niveacf,michael,wiebe}> (Nivea de Carvalho Ferreira, Michael Fisher & Wiebe van der Hoek).

¹ The first author gratefully acknowledges the support given by the Brazilian Government through a CAPES-scholarship.

1 Introduction

Agent technology is increasingly used in contemporary systems. The overall idea is that an agent aims at maximizing its performance, based on environmental evidence and its knowledge, or beliefs. In this context, the representation of beliefs plays an important role in the agent description. This is the reason why, when considering the agent’s representation, the chosen formalism often characterises the agent’s state of “mind”. And as a consequence, reasoning with beliefs (within that representation) becomes a crucial aspect for successful agent design.

One possible approach to an agent’s representation of knowledge (or belief) is the use of a formal language, whose syntax and semantics are precisely defined. In this way, a logical agent description and its associated semantics are consequently strongly linked. As information about the world may be vague, imperfect, uncertain, or ambiguous, agents should be able to represent and reason under uncertainty in order to operate in such an environment. By considering interaction with the “real” world, we require agent descriptions to incorporate some elements of uncertainty. Note that, in this paper, we will not consider multi-agent settings, in which agents have uncertain predictions about other agents’ uncertainty.

Here we use a possible worlds semantics (also known as Kripke structures, or models [18]) as the semantical basis that characterises modal logics of knowledge and belief, and as a means of expressing uncertainty with respect to the true state of the world. Given a situation of a system, one could draw a map of states considered possible and, consequently, be able to determine what is believed in that situation. In this context, a set of possible worlds would represent the doxastic possibilities. In other words, by having worlds that are named “possible”, an agent expresses its “doubts” about which is the “real” situation, i.e., its uncertainty about the true state of the world. The more worlds an agent considers possible, the more uncertain it is, and the less it believes. This is what makes possible worlds a qualitative measure of an agent’s uncertainty [14].

The most popular choices for modelling informational attitudes such as beliefs, involves annotating the agent with a *KD45*-like logic [8,20]. However, when using logics such as *KD45* in its 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 logics (cf. [22]) and probabilistic modal logics ([12]) are a way to address this. In particular, probabilistic logics of knowledge and belief ([13,6,25]) aim at removing the limitations implied by classical epistemic and doxastic logic. In epistemic logic the formalization is restricted to sentences such as “agent knows φ ” or “agent does not know φ ”, in which no quantification of the agent’s certainty is possible.

We present a logic that builds upon the natural framework of Kripke models, while allowing us to reason about uncertainty. For us it is both important and interesting to capture, and express, the notion of degrees of uncertainty within the agent itself. This means, intuitively, that we want to express statements like: “agent i believes that the probability of statement b being true is greater than x ”. In this sense, the agent can have more (or less) confidence in certain facts. More specifically, we introduce the P_FKD45 Logic which extends, in some aspects, the system P_FD given in [25] (which in turn was inspired by the logic from [9]). We propose not only a complete axiomatization for the logic, but also a decision procedure that permits us to verify the satisfiability of P_FKD45 -formulae. We claim that P_FKD45 represents a good candidate for representing and reasoning about uncertainty within computational agents, especially because, contrary to many logical approaches to probabilistic reasoning, it is conceptually simple and logically *compact* — we come back to this in the concluding section.

This paper is organised as follows. In Section 2 we present a description of the language, including basic definitions and a complete axiomatization. In the subsequent section we provide the semantics and establish some meta-logical properties such as soundness, completeness and the finite model property. An implemented decision procedure for the logic is presented in Section 4. Finally, we consider related work and provide concluding remarks in Section 5.

2 Logic: Language and Axioms

In this paper we introduce the logic P_FKD45 , an extension of the P_FD system, introduced by van der Hoek in [25], which in turn was inspired by the work of Fattorosi and Amati [9]. The P_FKD45 basic modal operator $P^{>}$ allows us to write formulae such as $P_{0.5}^{>}\varphi$, meaning that the “agent believes that the probability of φ being true is strictly greater than 0.5”. The other operators (with self-explanatory meaning) P^{\geq} , $P^{<}$, P^{\leq} and $P^{=}$ can each be defined in terms of the basic one. Since probabilities range from 0 to 1, and the probability of a formula is given by the sum of values associated to the worlds in which this formula holds, P_1^{\geq} is identified with the classical modal operator B for belief.

A peculiar property of the semantics is that it only allows probability measures (for each world) that are in some finite base set F . The motivation for this, as will be explained in Section 5, is the restoration of compactness for the logic. In Section 2.1 it will also become clear that, although this restricts probability assignments to a finite range, it is still possible to express and reason about arbitrary probabilities.

2.1 Language Description

The language L of P_FKD45 (as for P_FD originally described in [25]) consists of a countable set of propositional symbols, the logical connectives \neg and \vee (with standard definitions for $\perp, \top, \wedge, \rightarrow, \leftrightarrow$), and parentheses. We also define a modal operator $P_x^>$, where x is a rational number within the closed interval $[0, 1]$.

The logic is defined relative to a finite fixed base set F with $\{0, 1\} \subseteq F = \{x_0, x_1, \dots, x_n\} \subseteq [0, 1]$. It is assumed that $x_i < x_{i+1}$, if $i < n$ (implying $0 = x_0$ and $x_n = 1$). This is no restriction on the *language* of the logic: here, the basic operator is $P_x^>$, where $x \in [0, 1]$, with the intended meaning of $P_x^>\varphi$ being: “it is believed that the probability of φ being true is (strictly) greater than x ”. From the definition of the basic modal operator $P^>$ we derive the definitions of the other modal operators (x representing an arbitrary value over $[0, 1]$):

- 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.** $P_x^=\varphi \equiv \neg P_x^>\varphi \wedge \neg P_x^<\varphi$

In Section 2.2 we establish that these operators have the expected properties. For instance, we show that $P_x^{\leq}\varphi$ is equivalent to $(P_x^<\varphi \vee P_x^=\varphi)$.

The probability values used for assignments are taken from a finite set of rational numbers, which we call set F . This set includes the extreme values 0 and 1 in order to reason about absolute certainty, and is moreover closed under restricted (if the sum does not exceed 1) addition— to represent reasoning about mutually exclusive disjunctions— and complement with respect to 1 — to deal with negation. The formal definition of this set is as follows.

Definition 1 *A set F is a base for the logic P_FKD45 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) $x \in F \Rightarrow (1 - x) \in F$.

Let d be such that $0 < d \leq 1$. We say that D is *generated* by d , notation $D = \vec{d}$, if $D = \{x \in [0, 1] \mid \exists k \in \mathbb{N} x = k \cdot d\}$. Not only is every generated set a base set, but the converse also holds:

Observation 2 *F is a base set iff $F = \vec{d}$ for some $d \in (0, 1]$.*

2.2 Axioms and some Theorems

We do not restrict the *reasoning* about probabilities to numbers from F , only semantically the values for probabilities come from this set. In other words, even though F is a finite set of specific values, one can still represent facts about numbers not in F . So, let x and y be arbitrary values over $[0, 1]$ and let x_i, x_{i+1} be elements of F ($0 \leq i < n$), the inference rules (R1 and R2) and axioms (A1-A9) of P_FKD45 are defined in the following way.

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

We say that $P_FKD45 \vdash \varphi$, if there is a proof for φ , using the axioms and inference rules of P_FKD45 , and in such a case we say that φ is *derivable*. According to the first axiom, all the propositional tautologies are part of the system. The axioms A2-A6 all reflect basic properties of probabilities. Axioms A2 and A3 assume a certain implication: it has probability 1. Given that, A2 expresses that any probability for the consequent is at least the probability of the antecedent, and that ‘greater than’ implies ‘at least’. Axiom A3 says that the probability of the consequent is greater than any number that is smaller than the lower bound for the antecedent. Axiom A4 guarantees that no probability is negative, and A5 guarantees that the probability of a disjunction is obtained from the probability of its disjuncts. By Axiom A6 we have that the probability of a disjunction of mutual exclusive assertions is at least the sum of the individual probabilities of those assertions.

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. In other words, it enforces that arbitrary values collapse to the values present in F . Axioms A8 and A9 are new w.r.t. [25] and emphasize the relation with the modal logic $KD45$, making our agents doxastically introspective. (This will be made precise later in Theorem 5.) Axiom A8 denotes that, if the agent assigns a positive probability to some probabilistic judgment, then it incorporates this judgment. Axiom A9 states that the agent is absolutely sure about its own probabilistic beliefs.

To highlight the use of P_FKD45 , we present a simple planning example proposed in Kushmerick *et al.* [19] and show how P_FKD45 can be used to specify and reason about its properties.

Example 3 (Robot With a Bomb) *The example is described as follows. A robot is given two packages, and told that exactly one of them contains a bomb. It needs to defuse the bomb, and the only way to do so is to ‘dunk’ the package containing the bomb in the toilet. Placing a package in the toilet might (as we consider here with probability 0.1) ‘clog’ the toilet, and that is to be avoided.*

Suppose both goals are desired, i.e., we want to defuse bomb and have an unclogged toilet. Furthermore, assume that we want this to happen with probability at least 0.8.

Using the predicates² $in(Package, Bomb)$, $dunk(Package)$, and $clogged(Toilet)$, a possible specification would be:

A. $in(package1, bomb) \leftrightarrow \neg in(package2, bomb)$

B. $P_{0.5}^= in(package1, bomb) \wedge P_{0.5}^= in(package2, bomb)$

C. *One valuation chosen among the following:*

(1) $\neg dunk(package1) \wedge \neg dunk(package2)$

(2) $\neg dunk(package1) \wedge dunk(package2)$

(3) $dunk(package1) \wedge \neg dunk(package2)$

(4) $dunk(package1) \wedge dunk(package2)$

D1. $\neg dunk(package1) \wedge \neg dunk(package2) \rightarrow$

$P_1^= (\neg defused(bomb) \wedge P_1^= \neg clogged(toilet))$

D2. $\neg dunk(package1) \wedge dunk(package2) \rightarrow$

$P_{0.5}^= (\neg defused(bomb) \wedge P_{0.9}^= \neg clogged(toilet))$

D3. $dunk(package1) \wedge \neg dunk(package2) \rightarrow$

$P_{0.5}^= (\neg defused(bomb) \wedge P_{0.9}^= \neg clogged(toilet))$

D4. $dunk(package1) \wedge dunk(package2) \rightarrow$

$P_1^= (defused(bomb) \wedge P_{0.8}^= \neg clogged(toilet))$

E. $P_{0.8}^{\geq} (defused(bomb) \wedge \neg clogged(toilet))$

We now demonstrate some properties that follow from definitions and axioms presented above. All proofs are to be found in Appendix A, unless stated otherwise.

Lemma 4 *The following theorems are derivable from P_FKD45 . A proof is to be found in [25]:*

L1. $P_1^{\leq} \varphi$

L2. $P_1^{\geq} \varphi \leftrightarrow P_1^= \varphi$

² Although we use predicate symbols for brevity, the example is purely propositional.

Theorem 5 Define the belief operator ‘ B ’ using $B\varphi = P_1^{\geq}\varphi$. Based on this, we can infer the following.

- a) All $KD45$ -properties of B are derivable in P_FKD45 :
- (1) $\vdash \varphi \Rightarrow \vdash B\varphi$
 - (2) $\vdash B(\varphi \rightarrow \psi) \rightarrow (B\varphi \rightarrow B\psi)$
 - (3) $\vdash B\varphi \rightarrow BB\varphi$ (from A9)
 - (4) $\vdash \neg B\varphi \rightarrow B\neg B\varphi$
- 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 formulae, φ , $P_FKD45 \vdash \varphi$ iff $KD45 \vdash \varphi$.

Below we present some further theorems of P_FKD45 (Proofs are to be found in [25]) since P_FKD45 extends P_FD , the proof is similar to that in [25]. In our schemes, x, y represent arbitrary values over $[0, 1]$. We also define:

$$\nabla(\varphi_1, \varphi_2, \dots, \varphi_k) \equiv \left(\bigvee_i \varphi_i \right) \wedge \left(\bigwedge_{1 \leq i < j \leq k} \neg(\varphi_i \wedge \varphi_j) \right)$$

Theorem 6 For all φ, ψ in the language and all $x \in [0, 1]$:

- T1.** $P_x^{\geq}\varphi \leftrightarrow (P_x^>\varphi \vee P_x^=\varphi) \wedge P_x^{\leq} \leftrightarrow (P_x^<\varphi \vee P_x^=\varphi)$
- T2.** $\nabla(P_x^>\varphi, P_x^=\varphi, P_x^<\varphi)$
- 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.** $P_x^=\varphi \leftrightarrow P_{1-x}^=\neg\varphi$
- T8.** $(P_1^{\geq}\neg(\varphi \wedge \psi) \wedge P_x^=\varphi) \leftrightarrow (P_y^=\psi \rightarrow P_{x+y}^=(\varphi \vee \psi))$

Recall that our logic is based on a finite set of probability values F . Although the use of a base set may seem restrictive, we have seen that Axiom A7 ensures the possibility of arbitrary values being used. In addition, the following lemma shows a benefit of having a finite base F : we can express, in the language, that every formula has a probability. The proof is again similar to one given in [25].

Lemma 7 For all $\varphi \in L$, the following is a P_FKD45 -theorem:

$$\nabla(P_{x_0}^=\varphi, P_{x_1}^=\varphi, \dots, P_{x_n}^=\varphi) \quad \text{where } F = \{0 = x_0, x_1, \dots, x_n = 1\}$$

3 Semantics and Properties

Formulae are interpreted on *Probabilistic Kripke Models over F* (or $\mathcal{P}_F KD45$ models). $P_x^{\geq}\varphi$ is true at a world w if the probability values assigned to the possible

worlds that verify φ sum up to at least x .

3.1 Semantics

The classical Kripke model semantics refers to a collection of possible worlds. Different worlds may have different interpretations for sentences. A probabilistic Kripke model adds the concept of a probability distribution to the picture of possible worlds. That is, there is an assignment of probability values to the set of possible worlds in accordance with the formulae specified. Naturally, once those worlds model sentences of the language, we have an assignment of probabilities to the sentences themselves. In our case, assignments are provided *over* F , meaning that probability values assigned are in the assumed base set.

Definition 8 For each base set, F , $\mathcal{P}_F KD45$ 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: $W \times L \rightarrow \{\mathbf{true}, \mathbf{false}\}$

The truth definition for formulae is inductively defined as:

- $(M, w) \models p$ iff $\pi(w, p) = \mathbf{true}$, for atomic sentences p
- $(M, w) \models \neg\varphi$ iff not $M \models \varphi$
- $(M, w) \models \varphi \wedge \psi$ iff $M \models \varphi$ and $M \models \psi$
- $(M, w) \models P_x^>\varphi$ iff $(\sum_{\{w' | (M, w') \models \varphi\}} P_F(w')) > x$

Here, $M \models \varphi$ is short for $\forall w \in W, (M, w) \models \varphi$, and $P_F KD45 \models \varphi$ abbreviates that for each $\mathcal{P}_F KD45$ model M , we have $M \models \varphi$. In that case, we say that φ is valid.

In contrast to the definition of P_F found in [25] which allowed different sets of worlds to have different probability distributions, the probability distribution we define here is independent of the world. This topic is discussed in detail later, in Section 3.3. Figure 1 shows an example of a probabilistic model.

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 set W of the model M . Then, let W' be $\{w\} \cup \{w' \mid P_F(w') > 0\}$. Finally, define $R'(x, y)$ if, and only if, $P_F(y) > 0$, 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 (we omit the proof, but it follows directly from Theorem 5) gives a semantic motivation for coining our system $P_F KD45$:




Fig. 1. An example of a probabilistic belief model: $p_1 + p_2 + p_3 = 1$.

Proposition 9 Given a P_FKD45 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_FKD45 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 $M'_w, w \models \varphi$;
- (2) the accessibility relation R' is serial, transitive and Euclidean.

Note that these are precisely the properties of models for $KD45$ modal logics [20].

3.2 Logical Properties

In Section 2.2 the axiomatic system of P_FKD45 was presented. This system is *sound* with respect to the semantics of Definition 8 if, for all φ , we have $P_FKD45 \vdash \varphi \Rightarrow \mathcal{P}_F KD45 \models \varphi$. In other words, what is derivable is valid. Conversely, P_FKD45 is *complete* if the implication holds in the other direction. In other words, what is valid is derivable.

Lemma 10 (P_FKD45 Soundness) For all φ : $P_FKD45 \vdash \varphi \Rightarrow \mathcal{P}_F KD45 \models \varphi$.

The completeness property refers to the opposite direction. Having a complete system ensures that valid sentences are derivable from the theory. The completeness of P_FKD45 is what we show next. Let φ be a consistent formula of P_FKD45 i.e., $P_FKD45 \not\vdash \neg\varphi$. Let Ψ be the set of sub-formulae of φ closed under single negation and satisfying, for any \sim within $\{<, >, \leq, \geq, =\}$, $(P_x^\sim \psi \in \Psi) \Rightarrow \{P_{x_i}^\sim \psi \mid x_i \in F\} \subseteq \Psi$. We will now show how to construct a model that satisfies φ . Given a set of formulas Δ , a set Σ is Δ -maximal consistent if (i) $\Sigma \subseteq \Delta$, (ii) Σ is consistent, and, finally (iii) for no $\delta \in \Delta \setminus \Sigma$ is $\Sigma \cup \{\delta\}$ consistent. In words: Σ is a maximal consistent part of Δ . With Ψ being finite, say $|\Psi| = k$, we can define the Ψ -maximal consistent sets as $\Gamma_1, \Gamma_2, \dots, \Gamma_n, n \leq 2^k$. Let γ_i be the conjunction of formulae in $\Gamma_i, i \leq n$. Let $\Psi' = \Psi \cup \{P_x^\sim \gamma_i \mid x \in F, i \leq n\}$. The following is standard, and is not specific for our logic; we omit its proof.

Proposition 11

- i. $\vdash \neg(\gamma_i \wedge \gamma_j)$, where $i \neq j$;
- ii. $\vdash (\gamma_1 \vee \dots \vee \gamma_n)$
- iii. $\vdash \psi \leftrightarrow \gamma_{\psi_1} \vee \dots \vee \gamma_{\psi_r}$, where $\gamma_{\psi_1} \vee \dots \vee \gamma_{\psi_r}$ are exactly those γ 's which contain ψ as a conjunct, for each $\psi \in \Psi$.

Since φ is consistent and hence there is at least one Γ_i such that $\varphi \in \Gamma_i$, as $\{\varphi\} = \Sigma_0$ can be extended to a Ψ -maximal consistent set in a finite variant of the Lindenbaum construction, [1, p. 197] (enumerate the members of Ψ as ψ_1, \dots, ψ_k , let Σ_{j+1} be $\Sigma_j \cup \{\psi\}$ if the latter set is consistent, else put $\Sigma_{j+1} = \Sigma_j$; finally, put $\Gamma_i = \Gamma_\varphi = \Gamma_k$). Given this Γ_φ , we construct a set $\Phi \supseteq \Gamma_\varphi$ as follows. From Lemma 7, we know that for every consistent set Γ and formula ψ , at least one set of the sequence (1) is also consistent:

$$\Gamma \cup \{P_0^= \psi\}, \Gamma \cup \{P_{r_1}^= \psi\}, \dots, \Gamma \cup \{P_{r_{n-1}}^= \psi\}, \Gamma \cup \{P_1^= \psi\} \quad (1)$$

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 $r \in F$ such that $\Phi_{i-1} \cup \{P_r^= \gamma_i\}$ 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 $M^c = \langle W^c, P_F^c, \pi^c \rangle$ as follows:

- (1) $W^c = \{\Gamma_\varphi\} \cup \{\Gamma_i \mid \exists r > 0 P_r^= \gamma_i \in \Phi\}$.
- (2) $(P_F^c(\Gamma_i) = r) \Leftrightarrow (P_r^= \gamma_i \in \Phi)$
- (3) $\pi(\Gamma_i, p) = (p \in \Gamma_i)$

Proposition 12 *The model M^c is indeed a $\mathcal{P}_F\mathcal{KD}45$ model.*

This all leads us to the following coincidence lemma.

Lemma 13 (Coincidence) *For all $\psi \in \Psi$ and $\Gamma \in W^c$, $M^c, \Gamma \models \psi$ iff $\psi \in \Gamma$*

Now completeness follows immediately, we only sketch the proof: if $\mathcal{P}_F\mathcal{KD}45 \not\models \varphi$, then $\neg\varphi$ is consistent and, by the Coincidence Lemma, $\neg\varphi$ is satisfied in M^c, Γ for some set Γ , and hence $\mathcal{P}_F\mathcal{KD}45 \not\models \varphi$. Also note that M^c is a finite model, which follows from the fact that it is constructed for a given φ , which has finitely many propositional atoms and, moreover, the set F is finite, and hence the probabilities assigned by P_F^c have a finite range.

Theorem 14 (Soundness and Completeness, Finite Models) *For any formula φ , we have $\mathcal{P}_F\mathcal{KD}45 \models \varphi$ iff $\mathcal{P}_F\mathcal{KD}45 \vdash \varphi$. Moreover, every consistent formula has*

a finite model.

Observation 15 *If n is the number of atoms occurring in φ , and f the cardinality of F , it is easy to see that there are at most $2^n \cdot f^{2^n - 1}$ pairs $(M, s) = (\langle W, P_F, \pi \rangle, s)$ that we need to consider for φ 's satisfiability: there are 2^n different valuations as candidate for the current state s , and we can add all those different valuations to W . The model is then determined if we have fixed a probability $P_F(w)$ for every valuation in w , but the probability of the last of those is determined once the probability of the others is fixed. For each $w \in W$ except the last one, we can make f different choices for $P_F(w)$. The model M can possibly be smaller by deleting all states w from W for which $P_F(w) = 0$.*

3.3 Nested Beliefs

Considering P_FKD45 as a language for representing properties within individual agents, we next show that *nested* belief formulae can be removed, i.e., any nested belief formula is equivalent to some formula without nesting. First we show that the truth of probabilistic formulas is independent of the world of evaluation.

Lemma 16 *Let $M = \langle W, P_F, \pi \rangle$ be a $\mathcal{P}_F\mathcal{K}D45$ model. Then, $[\exists w \in W(M, w) \models P_\gamma^\geq \beta] \Leftrightarrow [\forall u \in W(M, u) \models P_\gamma^\geq \beta]$.*

We are now going to show that nested beliefs are superfluous, in P_FKD45 . This result is a generalisation of [20, THEOREM 1.7.6.4], where it is proved for $S5$. This means that that result still goes through when weakening the logic to $KD45$, and even when having probabilistic operators.

Definition 17 *We say that a formula ψ is in normal form if it is a disjunction of conjunctions each of the form*

$$\delta = \omega \wedge P_{\gamma_1}^\geq \beta_1 \wedge P_{\gamma_2}^\geq \beta_2 \wedge \dots \wedge P_{\gamma_n}^\geq \beta_n \wedge P_{\kappa_1}^> \alpha_1 \wedge P_{\kappa_2}^> \alpha_2 \wedge \dots \wedge P_{\kappa_k}^> \alpha_k,$$

where $\omega, \beta_i, \alpha_j, (i \leq n, j \leq k)$ are all purely propositional formulae. The formula δ is called the *canonical conjunction* and the sub-formulae $P_{\gamma_i}^\geq \beta_i$ and $P_{\kappa_j}^> \alpha_j$ are called *prenex formulae*.

Lemma 18 *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.*

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

Lemma 19 (Removal of Nested Beliefs) *We have the following two equivalences in $\mathcal{P}_F\mathcal{K}D45$:*

$$\models P_\alpha^\geq (\pi \vee (\lambda \wedge P_\gamma^\geq \beta)) \leftrightarrow (P_\alpha^\geq (\pi \vee \lambda) \wedge P_\gamma^\geq \beta) \vee (P_\alpha^\geq \pi \wedge \neg P_\gamma^\geq \beta) \quad (2)$$

$$\models P_\alpha^\geq(\pi \vee (\lambda \wedge P_\gamma^\gt\beta)) \leftrightarrow (P_\alpha^\geq(\pi \vee \lambda) \wedge P_\gamma^\gt\beta) \vee (P_\alpha^\geq\pi \wedge \neg P_\gamma^\gt\beta) \quad (3)$$

From this result it immediately follows (proof is omitted) that we can bring all the probabilistic operators to the outermost level, giving us:

Theorem 20 *Every formula φ is equivalent to a formula, ψ , in normal form, i.e., a formula without nesting of probabilistic operators.*

Intuitively, having nested beliefs being reduced to a non-nested formula corresponds to the ideal rationality feature of the agents we aim to design. That is, an agent that is certain about its own uncertainties.

For instance, suppose $P_{0.8}^-r$. Then $P_1^-P_{0.8}^-r$, and hence $P_1^-(P_{0.8}^-r \vee P_{0.2}^-q)$. If we assume that $Prob(r) \neq 0.8$ and $Prob(q) \neq 0.2$, then $Prob(r) = R$ and $Prob(q) = Q$ for some R (different from 0.8) and some $Q \neq 0.2$. Consequently, it is everywhere the case that $P_1^-(P_R^-r \vee P_Q^-q)$. Moreover, $P_0^-P_{0.8}^-r$ and $P_0^-P_{0.2}^-q$ would hold, and hence the probability of their disjunction is also 0.

Hence we do not model situations in which, for example, an agent has a fair coin with probability 0.5, and an unfair one (e.g., $Prob(head) = 0.8$) with probability 0.5, and it is then able to reason about the probability of the possible coin toss outcome h . In short, $P_{0.5}^-(P_{0.5}^-h) \wedge P_{0.5}^-(P_{0.8}^-h)$. In fact, a formula such as $P_{0.5}^-(P_{0.8}^-head)$ is always false in our semantics. This is so because the outermost probability in this case must always be either 0 or 1, once the probabilities are assigned by one and the same *perfect reasoner* agent. Therefore, what we *can* model is the uncertainty of an agent who knows that the coin toss situation mentioned above reduces to $Prob(h) = 0.5(0.5 + 0.8) = 0.65$.

4 Decision Procedure

We now describe a practical decision procedure for P_FKD45 . This procedure aims to find a finite model for the agent's specification. That is, given the agent's description (as a set of probabilistic formulae in the language), the objective of the decision procedure is to determine a set of probability values that can be assigned to the set of possible worlds in order to satisfy the given formula, if such a set of values exists.

According to the P_FKD45 semantics, we know that $P_x^\gt\varphi$ holds if the sum of all the accessible worlds that satisfy φ is (strictly) greater than x . In other words, each probabilistic operator imposes a restriction on the values that can be assigned to the worlds that satisfy the formulae it describes. As a consequence, having a set of probabilistic formulae is like having a set of constraints over the values to be

assigned to the possible worlds. This gives us a hint of how a simple decision procedure might be produced — by translating formulae into a set of numerical constraints on the possible valuations for propositions and then invoking an appropriate constraint solver.

Thus, given a finite set of P_FKD45 formulae, we generate a finite set of constraint (in)equalities. The components of the inequalities represent all the possible truth valuations of the propositional symbols, and it is to those components that probability values (that express the graded beliefs of the agent) are assigned. Solving the (in)equalities produces, as a result, the set of values that can be assigned to the set of possible worlds in order to satisfy the formulae presented. Consequently, once we have generated all the constraints for the formulae in the specification, we send this set of constraints to an appropriate solver. If the solver succeeds in finding a solution, this gives a set of probability assignments to the set of possible worlds; if the solver fails, no such assignment exists.

The approach can be summarized via the diagram in Figure 2, while the algorithm comprising the decision procedure is given in Figure 3.

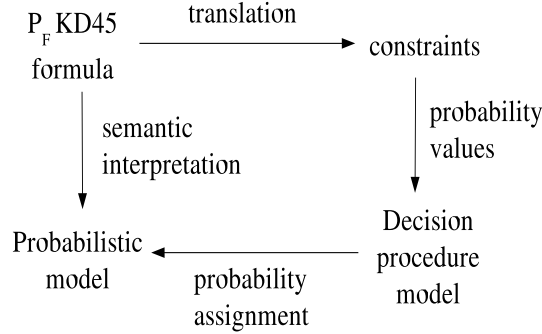


Fig. 2. Relation between P_FKD45 and decision procedure.

How are the in-equalities generated from a given formula, φ , of which the satisfiability is to be tested? First of all, we know from Theorem 20 that we may assume that φ has an equivalent normal form $\varphi = \delta_1 \vee \dots \vee \delta_m$, where every formula δ is of the form $\omega \wedge P_{\gamma_1}^{\geq} \beta_1 \wedge P_{\gamma_2}^{\geq} \beta_2 \wedge \dots \wedge P_{\gamma_n}^{\geq} \beta_n \wedge P_{\kappa_1}^{>} \alpha_1 \wedge P_{\kappa_2}^{>} \alpha_2 \wedge \dots \wedge P_{\kappa_k}^{>} \alpha_k$. All the formulas $\omega, \beta_i, \alpha_j$ ($i \leq n, j \leq k$) are propositional. If ω has a propositional model, this is a candidate for the state s (this is line 11 of the procedure) for the pair (M, s) that we are after. Moreover, each δ gives rise to a set of constraints $Con(At(\varphi), F, \{P_{\gamma_1}^{\geq} \beta_1, \dots, P_{\gamma_n}^{\geq} \beta_n, P_{\kappa_1}^{>} \alpha_1, \dots, P_{\kappa_k}^{>} \alpha_k\})$ as follows. $At(\varphi)$ is the set of all atoms p, q, r, \dots occurring in φ . Put every β_i and α_j into Disjunctive Normal Form $DNF(\beta_i)$ and $DNF(\alpha_j)$, respectively, where each atom from $At(\varphi)$ is used. For instance, if $\beta_i = p \vee q$, and $At(\varphi) = \{p, q, r\}$, then $DNF(\beta_i) = (\neg p \wedge q \wedge \neg r) \vee (\neg p \wedge q \wedge r) \vee (p \wedge \neg q \wedge \neg r) \vee (p \wedge \neg q \wedge r) \vee (p \wedge q \wedge \neg r) \vee (p \wedge q \wedge r)$.

Now, let us take this example with $At(\varphi) = \{p, q, r\}$ and $\beta_i = p \vee q$ a little further. $Con(At(\varphi), F, \{P_{\gamma_1}^{\geq} \beta_1, \dots, P_{\gamma_n}^{\geq} \beta_n, P_{\kappa_1}^{>} \alpha_1, \dots, P_{\kappa_k}^{>} \alpha_k\})$ consists of the following inequalities, where (*) is an overall-constraint, $f_0 \dots f_7$ are F -specific constraints,

and $C(P_{\gamma_i}^{\geq} \beta_i)$ is generated by $P_{\gamma_i}^{\geq} \beta_i$; similarly for $C(P_{\kappa_j}^> \alpha_j)$.

$$\begin{array}{lll}
(*) & p0q1r0 + p0q1r1 + p1q0r0 + p1q0r1 + p1q1r0 + p1q1r1 = & 1 \\
f_0 & p0q0r0 & \in F \\
\dots & \dots & \dots \\
f_7 & p1q1r1 & \in F \\
\dots & \dots & \dots \\
(C(P_{\gamma_i}^{\geq} \beta_i)) & p1q0r0 + p1q0r1 + p1q1r0 + p1q1r1 + p0q1r0 + p0q1r1 \geq & \gamma_i \\
\dots & \dots & \dots \\
(C(P_{\kappa_j}^> \alpha_j)) & \dots & > \kappa_j \\
\dots & \dots & \dots
\end{array}$$

Think of an expression such as $p0q0r0 + p0q0r1 > \rho$ as the constraint that the probability of those worlds satisfying $\neg p \wedge \neg q$ must be greater than ρ . The constraint (*) expresses the fact that a tautology should have probability 1, but also that there are no other worlds than those satisfying one of the combinations appearing as a term within it.

We now establish correctness of the decision procedure. We begin by noting that, if a P_FKD45 formula is satisfiable, then any model produced must satisfy one of the constraint sets C_i generated by the procedure from δ_i outlined above.

Lemma 21 *Given a P_FKD45 formula φ , and a translation of the formula into a class \mathcal{C} of sets of constraints C_i , one for each δ_i in φ 's normal form (as described above), then: M is a model for φ if, and only if, M is a solution for a solvable set of constraints $C \in \mathcal{C}(\varphi)$.*

Given the particular form of linear inequalities generated, and the fact that constraint solvers exist for such constraints, we also have the following result.

Observation 22 *A solution for any set of constraints C_i will be found if one exists.*

Note that it is guaranteed that the solver finds a solution, if there is one. That is, solving a finite set of constraint (in)equalities over a finite set F of probabilities is a decidable problem.

Given these results, the correctness of the decision procedure is straightforward.

Theorem 23 (Decision Procedure) *A formula φ in P_FKD45 , over the base set F , is satisfiable if, and only if, our decision procedure generates a model for φ and,*

P_FKD45 decision procedure:

Input: P_FKD45 formula φ .

Output: “yes, (M, s) ” for some (M, s) such that $M, s \models \varphi$ if it exists
“no, \emptyset ” if φ is unsatisfiable.

```

1. function  $dec(\varphi)$  returns “yes,  $(M, s)$ ” or “no,  $\emptyset$ ”
2.   let  $At(\varphi)$  be the set of atoms in  $\varphi$ 
3.   let  $\psi = \delta_1 \vee \dots \vee \delta_m$  be the normal form of  $\varphi$ 
4.   // ** we find  $\psi$  using (the proof of) Theorem 20
5.    $i := 1$ ;
6.   found := false
7.   // ** found = model for  $\varphi$  has been found
8.   while (( $i \leq m$ ) and (not found)) do
9.     let  $\delta_i = \omega \wedge P_{\gamma_1}^{\geq} \beta_1 \wedge P_{\gamma_2}^{\geq} \beta_2 \wedge \dots \wedge P_{\gamma_n}^{\geq} \beta_n \wedge P_{\kappa_1}^{>} \alpha_1 \wedge P_{\kappa_2}^{>} \alpha_2 \wedge \dots \wedge P_{\kappa_k}^{>} \alpha_k$ 
10.    if  $\omega$  has a propositional model  $\pi$  then
11.      let  $s := \pi$ ;
12.      let  $C_i = Con(At(\varphi), F, \{P_{\gamma_1}^{\geq} \beta_1, \dots, P_{\gamma_n}^{\geq} \beta_n, P_{\kappa_1}^{>} \alpha_1, \dots, P_{\kappa_k}^{>} \alpha_k\})$ 
13.      // ** cf. page 13
14.      apply the constraint solver to  $C_i$ 
15.      if this has a solution then
16.        read off  $M$  from it
17.        found := true
18.      end-if
19.    end-if
20.     $i := i+1$ 
21.  end-while
22.  if found then
23.    return “yes,  $(M, s)$ ”
24.  else return “no,  $\emptyset$ ”
25.  end-if
26. end-function

```

Fig. 3. A decision procedure for P_FKD45 .

conversely, every solvable set C_i in \mathcal{C} corresponds to a model for φ .

Example 3 (Continued) Depending on the action taken (valuation taken for both $dunk(package1)$ and $dunk(package2)$), the last sentence causes (or not — when D4) the set of probabilistic beliefs to be unsatisfiable. In other words, what we can show is that there are 4 possible cases: 3 in which the set of beliefs is unsatisfiable

$$\begin{aligned} &\neg dunk(package1) \wedge \neg dunk(package2), \\ &\neg dunk(package1) \wedge dunk(package2), \\ &dunk(package1) \wedge \neg dunk(package2) \end{aligned}$$

and one that shows a consistent set

$$\text{dunk}(\text{package1}) \wedge \text{dunk}(\text{package2}).$$

The implemented decision procedure automatically finds such a consistent set, indicating how the action of the robot can be planned.

Of course, the number of models that we mention in Observation 15 is huge, and indeed, finding a model for a formula (if it exists) is NP complete. This follows from an observation in [11] about *Integer Programming*, as which our satisfiability problem can be phrased.

Definition 24 [11, p. 245] *The following problem is called Integer Programming. INSTANCE. Given are the following. A finite set X of pairs (\bar{x}, b) , where \bar{x} is an m -tuple of integers and b is an integer, an m -tuple \bar{c} of integers, and an integer B . QUESTION: Is there an m -tuple \bar{y} of integers such that $\bar{x} \cdot \bar{y} \leq b$ for all $(\bar{x}, b) \in X$ and such that $\bar{c} \cdot \bar{y} \geq B$ (where the dot-product $\bar{u} \cdot \bar{v}$ of two m -tuples $\bar{u} = (u_1, u_2, \dots, u_m)$ and $\bar{v} = (v_1, v_2, \dots, v_m)$ is given by $\sum_{i=1}^m u_i v_i$?)*

Theorem 25 [11, p. 245] *The problem of Integer Programming as defined in Definition 24 is NP-complete.*

Recall how we assume our formula φ to be in normal form $\psi = \delta_1 \vee \dots \vee \delta_m$, where each δ_i is of the form $\delta_i = \omega \wedge P_{\gamma_1}^{\geq} \beta_1 \wedge P_{\gamma_2}^{\geq} \beta_2 \wedge \dots \wedge P_{\gamma_n}^{\geq} \beta_n \wedge P_{\kappa_1}^> \alpha_1 \wedge P_{\kappa_2}^> \alpha_2 \wedge \dots \wedge P_{\kappa_k}^> \alpha_k$. We now argue how finding a solution for the inequalities induced by δ_i (that is, steps 12–14 in our procedure) is an Integer Programming problem.

First of all, the dimension m of Definition 24 is going to be 2^n , where n is the number of propositional atoms in φ , say these atoms are p^1, \dots, p^n . We know from Observation 2 that F is generated by some d , that is, $F = \{0, \frac{1}{d}, \frac{2}{d}, \dots, \frac{d}{d}\}$. Now, every entry y_i in \bar{y} corresponds to $d \cdot P_i$, where P_0 is the variable $p_0^1 p_0^2 \dots p_0^{n-1} p_0^n$, $P_1 = p_0^1 p_0^2 \dots p_0^{n-1} p_1^n$, $P_2 = p_0^1 p_0^2 \dots p_1^{n-1} p_0^n$, etc. Now consider a subformula $P_{\gamma}^{\geq} \beta$ in Δ_i . We know γ is a rational number, say $\gamma = \frac{s}{t}$. The constraint that $P_{\gamma}^{\geq} \beta$ gives rise to looks like

$$k_0 \cdot p_0^1 \dots p_0^n + \dots + k_{2^n} \cdot p_1^1 \dots p_1^n \geq \frac{s}{t} \quad (4)$$

Where each k_i is either 1 (if the corresponding disjunct occurs in $DNF(\beta)$) or 0 (otherwise). Equation (4) is equivalent to

$$t \cdot d \cdot k_0 \cdot p_0^1 \dots p_0^n + \dots + t \cdot d \cdot k_{2^n} \cdot p_1^1 \dots p_1^n \geq d \cdot s \quad (5)$$

and hence, for each $P_{\gamma}^{\geq} \beta$, we add (\bar{x}, b) to X , with $\bar{x} = (t \cdot d \cdot \kappa_0, \dots, t \cdot d \cdot \kappa_{2^n})$ and $b = d \cdot s$.

Next, for every subformula $P_{\kappa}^>\alpha$, the constraint that we would obtain looks like

$$k_0 \cdot p_0^1 \dots p_0^n + \dots + k_{2^n} \cdot p_1^1 \dots p_1^n > \alpha \quad (6)$$

where, as before, k_i is 1 if the corresponding disjunct occurs in α , and 0 otherwise. But since we know that each value for $p_{z_1}^1 p_{z_2}^2 \dots p_{z_n}^n$ is a multiple of d , the left hand side of (6) must also be a multiple of d . Hence, let $\alpha_F^>$ be the smallest element in F for which $\alpha_F^>$. It is easy to see that, given that each variable $p_{z_1}^1 p_{z_2}^2 \dots p_{z_n}^n$ must be in F , equation (6) is equivalent to (7), and we can proceed to add a member (\bar{x}, b) to X as we did above.

$$k_0 \cdot p_0^1 \dots p_0^n + \dots + k_{2^n} \cdot p_1^1 \dots p_1^n \geq \alpha_F^> \quad (7)$$

Finally, we address the overall constraint (*):

$$(*) \quad d \cdot p_0^1 p_0^2 \dots p_0^{n-1} p_0^n + d \cdot p_0^1 p_0^2 \dots p_1^{n-1} p_0^n + \dots + d \cdot p_1^1 p_1^2 \dots p_1^{n-1} p_1^n = d$$

For this constraint, we add $(\bar{x}, b) = ((1, 1, \dots, 1), d)$ to X , and we choose also $\bar{c} = (1, 1, \dots, 1)$ and $B = d$.

It is easy to see that this shows that the question whether $\delta_i = \omega \wedge P_{\gamma_1}^{\geq} \beta_1 \wedge P_{\gamma_2}^{\geq} \beta_2 \wedge \dots \wedge P_{\gamma_n}^{\geq} \beta_n \wedge P_{\kappa_1}^> \alpha_1 \wedge P_{\kappa_2}^> \alpha_2 \wedge \dots \wedge P_{\kappa_k}^> \alpha_k$ can be made true by assigning probabilities in F to each combination of propositional atoms, is equivalent to the translated Integer Programming problem.

Admittedly, Theorem 25 together with the observations above provide a negative result for our decision procedure: even solving the constraints for the probabilistic part is NP-complete, and, obviously, to check whether the propositional part of the given formula (step 10 in our procedure) is also NP-complete. We have not, looked into whether there are ways to cope with this NP-completeness, but what we *did* do is implement a toy system that finds a model for a formula, if it has one (see [3]). The development of our P_FKD45 decision procedure involved Prolog programming, making use of Sicstus Prolog v3.12.2 and its inbuilt module `c1p(FD)`, a constraint logic programming solver over finite (integer) domains [2]. In short, we have an executable module that receives, as input, a set of probabilistic belief formulae and outputs all the possible probability assignments that satisfy the restrictions that those formulae impose. Furthermore, we also provide as input the set F that should be considered for that particular set of formulae.

Further development of the system incorporates also temporal operators allowing us to reason also about (discrete) time information ([3,4]). However, this is a story to be told in detail at another occasion.

5 Conclusion

We have shown how the P_FKD45 logic preserves important results about soundness, completeness, and decidability of its predecessor P_FD [25]. We have also presented new results about nested beliefs, and a description and implementation of a decision procedure for the logic. A brief example was used to show how the language can serve as an appropriate agent specification language.

Our logic is conceptually simple and *compact*. Compactness refers to the fact that a set of sentences is satisfiable if, and only if, every subset of it is satisfiable. Logics that allow us to express that $Prob(\varphi) \sim r$ are, in general, not compact, witness the set of premises Γ

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

Then, obviously, we have $\Gamma \models Prob(q) = 1$, but 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 $\{Prob(\varphi) \sim r\}$ is satisfiable in $\mathbb{Q} \cap [0, 1]$, in principle has to check an infinite number of assignments of probabilities to formulae φ . For these reasons, we assume that the range of allowed probabilities is a finite set $F \subseteq [0, 1]$. Thus, the compactness of our approach has benefits, especially with respect to computational tractability aspects. This is a decisive feature considering our aim of having P_FKD45 as part of our executable framework for agents. A logic similar to ours (but also allowing predicates) and with the same motivation was presented in [24]. Its semantics is based on measure spaces, and all formulas are interpreted on a set (of what we call states) – even the propositional ones. Decidability is reduced “to an easy problem of linear programming, which can be easily solved” ([24, p. 3]). We are more explicit about the completeness proof and the algorithm to find a model.

Although the use of the set base F causes logical restrictions, it is possible to highlight some interesting aspects (cf.[25]). For instance, if we take $F = \{0, 1\}$, we have the classical modal logic. Having Driankov’s linguistic estimates (as in [5]) *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 . And the same analogy can be used for any finite range of values to be assumed. In other words, the granularity of F can be chosen according to the intended agent’s application.

P_FKD45 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 alternative ways [17,10]. One of those possible approaches is the interpretation of the modal belief operator according to the concept of “likelihood” as in [16]. In this logic, in-

stead of using numbers to express uncertainty one would have expressions like “*p is likely to be a consistent hypothesis*” (since a state is taken as a set of hypotheses “true for now”). That is, a qualitative notion of likelihood of events rather than explicit probabilities. A more detailed comparison between the notions of likelihood and probability can be found in [15].

The P_FKD45 logic was designed for reasoning with (exact) probabilities, and its Probabilistic Kripke Model semantics is similar to the one presented in [7,6]. In those formalisms, a formula is typically a Boolean combination of expressions of the form $a_1w(\varphi_1) + \dots + a_nw(\varphi_n) \geq c$, where a_1, \dots, a_n, c are integers. The system in [7,6] includes, as axioms, all the formulae of linear inequalities, and consequently, their proofs of completeness to rely on results in the area of linear programming. Our logic is conceptually simpler. Furthermore, in P_FKD45 we can express statements such as “*p is true, although its probability is less than 0.1*”, something that their formalism is unable to represent, once they do not allow “*w-free formulae*”. The paper [13] is close to ours: it also considers a fixed probability for the agent that does not depend on the current state, and identifies a probability of 1 with belief.

Another approach can be found in [21]. The probabilistic epistemic logic used there is a special case of the one presented in [6]. The additional feature in Milch and Koller’s work is the fact of having an algorithm for finding the probability of the defined formula model using Bayesian networks. This allows them to model the formulae without constructing the probabilistic epistemic model explicitly. That is, there is no explicit representation of an agent’s probability distribution, or enumeration of states/worlds. This is possible once it is assumed that agents have a common prior probability distribution over outcomes and their beliefs differs only by different observations.

Finally, P_FKD45 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. In [23], Ognjanović and Răsković present a probabilistic logic suitable for describing events in a discrete sample space. Informally, their basic operator expresses that the probability of a certain event is in a particular set, a notion intuitively similar to ours.

In summary, we have a language that is conceptually simple and compact. Besides, by having an implemented module for deciding the probability attribution to the possible worlds, we have shown how that the language is a powerful tool not only for theoretical reasoning, but also for effective implementation of computational agents that deal with uncertainty.

Acknowledgements

We thank two reviewers for their helpful comments, and the editor of this issue for his support. The first author wishes to gratefully acknowledge the support given by the Brazilian Government through a CAPES-scholarship while doing this research.

References

- [1] P. Blackburn, M. de Rijke, Y. Venema, *Modal Logic*, Cambridge University Press, 2001.
- [2] M. Carlsson, G. Ottosson, B. Carlson, An open-ended finite domain constraint solver, in: *Proceedings of the Ninth International Symposium on Programming Languages: Implementations, Logics, and Programs (PLILP'97)*, 1997.
- [3] N. de Carvalho Ferreira, *Programming Uncertain Agents*, Ph.D. thesis, Department of Computer Science, University of Liverpool, UK (2006).
- [4] N. de Carvalho Ferreira, M. Fisher, W. van der Hoek, A Logical Implementation of Uncertain Agents, in: A. C. C. Bento, G. Dias (eds.), *Progress in Artificial Intelligence*, vol. 3808, 2005.
- [5] D. Driankov, Reasoning about Uncertainty: Towards a many-valued logic of belief, in: *IDA annual research report*, Linköping University, 1987.
- [6] R. Fagin, J. Y. Halpern, Reasoning About Knowledge and Probability, *Journal of the ACM* 41 (2) (1994) 340–367.
- [7] R. Fagin, J. Y. Halpern, N. Megiddo, A Logic for Reasoning about Probabilities, *Information and Computation* 87 (1/2) (1990) 78–128.
- [8] R. Fagin, J. Y. Halpern, Y. Moses, M. Y. Vardi, *Reasoning about Knowledge*, The MIT Press, Cambridge, Massachusetts, 1995.
- [9] M. Fattorosi-Barnaba, G. Amati, Modal Operators with Probabilistic Interpretations, *Studia Logica* 46 (1987) 383–393.
- [10] T. Fernando, In Conjunction with Qualitative Probability, *Annals of Pure and Applied Logic* 92 (3) (1998) 217–234.
- [11] M. R. Garey, D. S. Johnson, *Computers and Intractability – A Guide to the Theory of NP-Completeness*, Freeman and Company, New York, USA, 1979.
- [12] J. Halpern, *Reasoning about Uncertainty*, MIT Press, 2003.
- [13] J. Y. Halpern, The Relationship between Knowledge, Belief, and Certainty, *Annals of Mathematics and Artificial Intelligence* 4 (1991) 301–322.
- [14] J. Y. Halpern, *Reasoning About Uncertainty*, MIT Press, 2003.

- [15] J. Y. Halpern, D. A. McAllester, Likelihood, Probability, and Knowledge, *Computational Intelligence* 5 (1989) 151–160.
- [16] J. Y. Halpern, M. O. Rabin, A Logic to Reason about Likelihood, *Artificial Intelligence* 32 (3) (1987) 379–405.
- [17] C. Howson, Logic and Probability, *British Journal for the Philosophy of Science* 48 (1997) 517–531.
- [18] S. Kripke, Semantical Considerations on Modal Logic, *Acta Philosophica Fennica* 16 (1963) 83–94.
- [19] N. Kushmerick, S. Hanks, D. Weld, An algorithm for Probabilistic Planning, *Artificial Intelligence* 76 (1995) 239–286.
- [20] J. Meyer, W. van der Hoek, *Epistemic Logic for AI and Computer Science*, Cambridge University Press, 1995.
- [21] B. Milch, D. Koller, Probabilistic Models for Agent’s Beliefs and Decisions, in: C. Boutilier, M. Goldszmidt (eds.), *Proceedings of the 16th Conference in Uncertainty in Artificial Intelligence (UAI)*, Morgan Kaufmann, 2000.
- [22] N. Nilsson, Probabilistic Logic, *Artificial Intelligence* 28 (1) (1986) 71–87.
- [23] Z. Ognjanović, M. Răsković, Some Probability Logics with New Types of Probability Operators, *Journal of Logic and Computation* 9 (2) (1999) 181–195.
- [24] M. Rašković, Classical Logic with Some Probability Operators, *Publications de l’institute mathématique* 53 (1993) 1–3.
- [25] W. van der Hoek, Some Considerations on the Logic $PFD\sim$, *Journal of Applied Non Classical Logics* 7 (3) (1997) 287–307.

A Proofs

Proof of Lemma 4:

$$\begin{aligned}
\mathbf{L1.} \quad & \neg P_1^{\leq} \varphi \Rightarrow_{D3} P_0^{\leq} \neg \varphi \Rightarrow_{D2} P_1^{\geq} \varphi \\
& \Rightarrow_{D1} \neg P_0^{\geq} \neg \varphi \Rightarrow_{A4} \perp \\
\mathbf{L2.} \quad & P_1^{\geq} \varphi \Rightarrow_{L1} P_1^{\geq} \varphi \wedge P_1^{\leq} \varphi \Rightarrow_{D1, D3} \neg P_0^{\geq} \neg \varphi \wedge \neg P_0^{\leq} \neg \varphi \\
& \Rightarrow_{D2} \neg P_1^{\leq} \varphi \wedge \neg P_1^{\geq} \varphi \Rightarrow_{D4} P_1^{\equiv} \varphi \\
& P_1^{\equiv} \varphi \Rightarrow_{D4} \neg P_1^{\geq} \varphi \wedge \neg P_1^{\leq} \varphi \Rightarrow_{A1} \neg P_1^{\leq} \varphi \\
& \Rightarrow_{D2} \neg P_0^{\geq} \neg \varphi \Rightarrow_{D1} P_1^{\geq} \varphi
\end{aligned}$$

□

Proof of Observation 2: From right to left is obvious, so suppose F is a base set. Let $d = \min\{y - x \mid x, y \in F, y - x > 0\}$. That is, d is the smallest positive difference between any two members of F . Let x and y be elements of F for which $d = y - x$. Since F is a base, $(1 - x) \in F$ and, since $0 < d \leq 1$, we have $1 - d \in [0, 1)$, and hence $y + (1 - x) = 1 - d$ is also in F . It follows that $d \in F$, and hence all $k \cdot d \leq 1$ are such that $k \cdot d \in F$. But also, these must make up *all* elements in F , since if for some $f \in F$, f would not be a multiple of d , we would either find a k such that $k \cdot d < f < (k + 1) \cdot d$ or a biggest k with $k \cdot d < f < 1$. In both cases, $f - k \cdot d$ would be smaller than d , which is in contradiction with how d is chosen, i.e., as the minimal distance between two members of F . □

Proof of Theorem 5:

a) All $KD45$ -properties of B are derivable in P_FKD45 :

(1) $\vdash \varphi \Rightarrow \vdash B\varphi$ (directly from R2)

(2) $\vdash B(\varphi \rightarrow \psi) \rightarrow (B\varphi \rightarrow B\psi)$ (from A2)

(3) $\vdash \neg B\perp$:

By R2, we have that $\vdash P_1^{\geq}(\perp \rightarrow \perp)$ (†)

$P_1^{\geq}\perp \Rightarrow_{(\dagger), A1} P_1^{\geq}\perp \wedge P_1^{\geq}(\perp \rightarrow \perp)$

$\Rightarrow_{A3} P_0^{\geq}\perp \qquad \Rightarrow_{D2} P_1^{<}\top$

$\Rightarrow_{A0} \neg(\neg P_1^{<}\top \wedge \neg P_1^{\geq}\top) \qquad \Rightarrow_{D4} \neg P_1^{\neq}\top$

$\Rightarrow_{R2, L2} \perp$

(4) $\vdash B\varphi \rightarrow BB\varphi$ (from A9)

(5) $\vdash \neg B\varphi \rightarrow B\neg B\varphi$

$\neg B\varphi \equiv \neg P_1^{\geq}\varphi \qquad \equiv P_0^{\geq}\neg\varphi$

$\Rightarrow_{A7} P_{x_i}^{\geq}\neg\varphi \qquad \Rightarrow_{A8} P_1^{\geq}P_{x_1}^{\geq}\neg\varphi$

$\Rightarrow_{R2} P_1^{\geq}P_{x_1}^{\geq}\neg\varphi \wedge P_1^{\geq}(\neg\varphi \rightarrow \neg\varphi)$

$\Rightarrow_{A3} P_1^{\geq}P_{x_1}^{\geq}\neg\varphi \wedge P_1^{\geq}(P_{x_1}^{\geq}\neg\varphi \rightarrow P_0^{\geq}\neg\varphi)$

$\Rightarrow_{A2} P_1^{\geq}P_0^{\geq}\neg\varphi \qquad \equiv B\neg B\varphi$

b) The ' \Leftarrow ' part follows from item a) above; the ' \Rightarrow ' part will be obvious from the semantics for P_FKD45 given later, together with the soundness of P_FKD45 and the completeness of $KD45$. □

Proof of Lemma 10: It is straightforward to check that all the axioms are valid on

Proof of Proposition 12: We have to show that $P_F^c(\{\Gamma \mid \Gamma \in W^c\}) = 1$. We will show that the following variant of Axiom A6 is derivable:

$$\mathbf{A6}' \quad P_1^{\geq} \neg(\varphi \wedge \psi) \rightarrow ((P_x^{\leq} \varphi \wedge P_y^{\leq} \psi) \rightarrow P_{x+y}^{\leq}(\varphi \vee \psi)) \quad (\text{where } x + y \in [0, 1])$$

Once $\mathbf{A6}'$ has been established, we reason as follows: Let r_i be such that $P_{r_i}^{\leq}(\gamma_i) \in \Phi$, i.e., $P_F^c(\Gamma_i) = r_i$. By Lemma 7 and T3, we know that such an r_i exists, that it is unique, and a member of F . All γ_i 's are mutually exclusive and their disjunction is derivable (Proposition 11). But then, using $\mathbf{A6}'$, we conclude, from $(P_{r_1}^{\leq} \gamma_1 \wedge P_{r_2}^{\leq} \gamma_2) \in \Phi$, that $P_{r_1+r_2}^{\leq}(\gamma_1 \vee \gamma_2) \in \Phi$. And since also $(\gamma_1 \vee \gamma_2)$ and γ_3 are mutually exclusive, we find $P_{r_1+r_2+r_3}^{\leq}(\gamma_1 \vee \gamma_2 \vee \gamma_3) \in \Phi$. Eventually, we find $P_r^{\leq} \gamma \in \Phi$, where $r = \sum_{i \leq n} r_i$ and $\gamma = \bigvee_{i \leq n} \gamma_i$. Since (use Proposition 11) $P_1^{\geq} \gamma \in \Phi$, we find $\sum_{i \leq n} r_i = 1$, i.e., $P_F^c(\{\Gamma \mid \Gamma \in W^c\}) = 1$. We finally prove derivability of $\mathbf{A6}'$.

Assume $P_1^{\geq} \neg(\varphi \wedge \psi) \wedge ((P_x^{\leq} \varphi \wedge P_y^{\leq} \psi)$. We know (Lemma 7 and T3) that $x, y \in F$. By Observation 2, F is generated by some d , so let $x = k \cdot d$ and $y = m \cdot d$. If $k = m = x = y = 0$, we have to prove that $P_0^{\leq}(\varphi \vee \psi)$, which is easy: by Axiom A4, the only other possibility for the disjunction would be $P_0^{\geq}(\varphi \vee \psi)$, and then, by Axiom A5, we would have $P_0^{\geq} \varphi \vee P_0^{\geq} \psi$, which, (use T2) contradicts the assumption $P_0^{\leq} \varphi \wedge P_0^{\leq} \psi$. Without lack of generality, assume that $x = k \cdot d \neq 0$. Let z be a number $(k-1) \cdot d < z < k \cdot d$. Since we have $P_x^{\leq} \varphi$, using A3, we have $P_z^{\geq} \varphi$. Using A6, we then conclude $P_t^{\geq}(\varphi \vee \psi)$, where $t = z + y$, i.e., $(m+k-1) \cdot d < t < (m+k) \cdot d$, with $(m+k) \cdot d$ the smallest number in F greater than t . Using A7, we find $P_{(m+k) \cdot d}^{\geq}(\varphi \vee \psi)$, or $P_{x+y}^{\geq}(\varphi \vee \psi)$. This means we have either $P_{x+y}^{\geq}(\varphi \vee \psi)$ or $P_{x+y}^{\leq}(\varphi \vee \psi)$. The first of these options cannot be, since with A5 it would yield $P_x^{\geq} \varphi \vee P_y^{\geq} \psi$, which is impossible given $P_x^{\leq} \varphi \wedge P_y^{\leq} \psi$. \square

Proof of Lemma 13: We consider the modal case. Suppose $P_s^{\geq} \delta \in \Gamma$, and let $\delta \equiv (\gamma_1 \vee \dots \vee \gamma_v)$. Either $\Gamma = \Gamma_\varphi$, and we immediately obtain $\Phi \vdash P_s^{\geq} \delta$, or else let γ be the characteristic formula for Γ , then, since $\Gamma_\varphi \neq \Gamma \in W^c$, we have $\Phi \vdash P_0^{\geq} \gamma$, and, by A2, we have $\Phi \vdash P_0^{\geq} P_s^{\geq} \delta$. By A8, we conclude $\Phi \vdash P_s^{\geq} \delta$.

Then, $\Phi \vdash P_s^{\geq}(\gamma_1 \vee \dots \vee \gamma_v)$, and, by T8, with $s \uparrow$ being the first member of F greater than s , $\Phi \vdash \bigvee_{r_i \geq s \uparrow} P_{r_i}^{\leq}(\gamma_1 \vee \dots \vee \gamma_v)$. Let t_1, \dots, t_v be such that $\{P_{t_1}^{\leq} \gamma_1, \dots, P_{t_v}^{\leq} \gamma_v\} \subseteq \Phi$. Let $t = t_1 + \dots + t_v$. By definition of P_F^c , we have $P_F^c(\Gamma_i) = t_i$ ($i \leq v$), and, by induction, $M^c, \Gamma_i \models \gamma_i$. Since every two different γ_i and γ_j logically exclude each other (Proposition 11), we have $P_F^c(\{\Gamma_i \mid M^c, \Gamma_i \models \gamma_i\}) = t$, and hence $M^c, \Gamma \models P_t^{\leq} \delta$. Now, obviously $t > s$, since we have that Φ is consistent, $\Phi \vdash P_s^{\geq} \delta$ and $\Phi \vdash P_t^{\leq} \delta$. We conclude that $M^c, \Gamma \models P_s^{\geq} \delta$.

Conversely, suppose $M^c, \Gamma \models P_s^{\geq} \delta$. Then, for some $r \in F$, both $r > s$ and $P_F^c(\{\Delta \mid \delta \in \Delta\}) = r$. Again, assuming $\delta \equiv (\gamma_1 \vee \dots \vee \gamma_v)$, there are r_1, \dots, r_v , such that $r_1 + \dots + r_v = r$, $\gamma_i \in \Gamma_i$ and $P_F^c(\Gamma_i) = r_i$ ($i \leq v$). By definition of M^c ,

we have $P_{r_x}^{\leq} \gamma_{r_x} \in \Phi$. If $v = 1$, we have $\gamma_1 = \delta$ and $r_1 = r$, and hence $P_r^{\leq} \delta$. For $v \geq 2$, we show by induction on v that if $\vdash \neg(\gamma_i \wedge \gamma_j) (i \neq j \leq v)$ then $\vdash \bigwedge_{i \leq v} P_{r_i}^{\leq} \gamma_i \rightarrow P_{r_1 + \dots + r_v}^{\leq} (\gamma_1 \vee \dots \vee \gamma_v)$. For $v = 2$ this is immediate from T8, Theorem 6. Suppose it holds for v , and consider $\vdash \neg(\gamma_i \wedge \gamma_j) (i \neq j \leq v + 1)$. It follows that $\vdash P_1^{\geq} ((\gamma_1 \vee \dots \vee \gamma_v) \wedge \gamma_{v+1})$. Now assume $P_{r_i}^{\leq} \gamma_i (i \leq v + 1)$. By induction, we have $P_{r_1 + \dots + r_v}^{\leq} (\gamma_1 \vee \dots \vee \gamma_v)$. Using T8 again yields $P_{r_1 + \dots + r_{v+1}}^{\leq} (\gamma_1 \vee \dots \vee \gamma_{v+1})$. This proves $\Phi \vdash P_r^{\leq} \delta$, and hence, by T1, $\Phi \vdash P_r^{\geq} \delta$ and by A3, $\Phi \vdash P_s^{\geq} \delta$ (\dagger). Now, to arrive at a contradiction, suppose $P_s^{\geq} \delta \notin \Gamma$. By construction of our set of formulae Ψ then, we know that for some $t \leq s$, $P_t^{\leq} \delta \in \Gamma$. Since $P_F^c(\Gamma) > 0$, we have $\Phi \vdash P_0^{\geq} P_t^{\leq} \delta$, and hence, by A8, $\Phi \vdash P_t^{\leq} \delta$, which is in contradiction with (\dagger). \square

Proof of Lemma 16: The \Leftarrow direction is trivial, since $W \neq \emptyset$. For \Rightarrow , observe that $(M, w) \models P_x^{\geq} \varphi$ iff $P_F(\{w' : (M, w') \models \varphi\}) \geq x$ iff $(M, u) \models P_x^{\geq} \varphi$ \square

Proof of Lemma 18: ψ is in normal form, so $\psi = \delta_1 \vee \delta_2 \vee \dots \vee \delta_m$, where $\delta_{i's}$ are canonical conjunctions. Suppose σ occurs in δ_m . Then σ must be some conjunct P_γ^{\geq} , so that δ_m can be written as $(\lambda \wedge \sigma)$. Taking π to be $(\delta_1 \vee \delta_2 \vee \dots \vee \delta_{m-1})$ gives the desired result $\psi = \pi \vee (\lambda \wedge \sigma)$. \square

Proof of Lemma 19: We sketch the proof of (2). As $(M, s) \models P_\gamma^{\geq} \beta \vee \neg P_\gamma^{\geq} \beta$, there are two possible cases to consider.

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

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

For the ' \rightarrow ' direction, 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

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

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 $u \in S$. ($\forall u$) $(M, u) \models \lambda$ iff $\lambda \wedge P_\gamma^{\geq} \beta$ is also true. So,

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

$$(M, s) \models P_\gamma^{\geq} \beta \rightarrow (P_\alpha^{\geq} (\pi \vee (\lambda \wedge P_\gamma^{\geq} \beta)) \leftrightarrow (P_\alpha^{\geq} (\pi \vee \lambda) \wedge P_\gamma^{\geq} \beta)) \quad (\text{A.1})$$

Second Case. Assuming that $(M, s) \models \neg P_\gamma^{\geq} \beta$, we will show that

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

For the ‘ \rightarrow ’ direction, suppose that $(M, s) \models P_\alpha^\geq(\pi \vee (\lambda \wedge P_\gamma^\geq\beta))$. If this holds for s , it holds for all u . So,

$$\forall u, (M, u) \models P_\alpha^\geq(\pi \vee (\lambda \wedge P_\gamma^\geq\beta)) \quad (\text{A.2})$$

By a similar argument, $\forall u (M, u) \models \neg P_\gamma^\geq\beta$, if $(M, s) \models \neg P_\gamma^\geq\beta$, and hence

$$(M, u) \models P_1^\geq\neg P_\gamma^\geq\beta \quad (\text{A.3})$$

Combining A.2 and A.3, we get $M, u \models P_\alpha^\geq\pi$. Hence:

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

For the converse, $\pi \rightarrow (\pi \vee \sigma)$ is a tautology. So, we can say that

$$(M, s) \models (P_\alpha^\geq\pi \wedge \neg P_\gamma^\geq\beta) \rightarrow P_\alpha^\geq(\pi \vee (\lambda \wedge P_\gamma^\geq\beta)), \text{ and, consequently,}$$

$$(M, s) \models \neg P_\gamma^\geq\beta \rightarrow (P_\alpha^\geq(\pi \vee (\lambda \wedge P_\gamma^\geq\beta)) \leftrightarrow (P_\alpha^\geq\pi \wedge \neg P_\gamma^\geq\beta)) \quad (\text{A.4})$$

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 (A.1) and (A.4) to conclude (2). \square

Proof of Lemma 21: If M is a model for φ , it must be a model for one δ_i ’s in φ ’s normal form. It is clear that this M then satisfies the constraints $C_i \in \mathcal{C}(\varphi)$, where C_i is generated by δ_i . Conversely, every probabilistic model M_i that satisfies the constraints of $C_i \in \mathcal{C}(\varphi)$, is a model for δ_i , and hence for φ . \square

Proof of Theorem 23: We split the proof into two cases.

- φ is satisfiable.
The constraints generated will be passed on to the constraint solver. By Lemma 21, a solution must exist if a model exists, and by Observation 22, the solver will find a solution if one exists.
- φ is unsatisfiable.
By Lemma 21, the constraints generated should have no solution since there is no model. By Observation 22, the solver will, indeed, fail to find a solution.

\square