Epistemic Probability Logic Simplified

Jan van Eijck [1]

CWI and ILLC
Science Park 123
1098 XG Amsterdam

François Schwarzentruber [2]

ENS Rennes
Campus de Ker Lann
35170 BRUZ

Abstract

We propose a simplified logic for reasoning about (multi-agent) epistemic probability models, and for epistemic probabilistic model checking. Epistemic probability models are multi-agent Kripke models that assign to each agent an equivalence relation on worlds, together with a function from worlds to positive rationals (a lottery). The difference with the usual approach is that probability is linked to knowledge rather than belief, and that knowledge is equated with certainty.

A first contribution of the paper is a comparison of a semantics for epistemic probability in terms of models with multiple lotteries and models with a single lottery. We give a proof that multiple lottery models can always be replaced by single lottery models. As multiple lotteries represent multiple subjective probabilities, our result connects subjective and intersubjective probability.

Next, we define an appropriate notion of bisimulation, and use it to prove an adaptation of the Hennessy-Milner Theorem and to prove that some finite multiple lottery models only have infinite single lottery counterparts. We then prove completeness, and state results about model checking complexity. In particular, we show the PSPACE-completeness of the model checking in the dynamic version with action models.

The logic is designed with model checking for epistemic probability logic in mind; a prototype model checker for it exists. This program can be used to keep track of information flow about aleatory acts among multiple agents.

Keywords: Probability. Epistemic modal logic. Lottery. Hennessy Milner Theorem. Dynamic epistemic logic. Complexity theory.

[1] jve@cwi.nl

[2] francois.schwarzentruber@ens-rennes.fr

1 Probability as a function of degree of information

A classical view of probability theory is that probability measures degree of information. Here is a characteristic quote from [17]:

> *Dans les choses qui ne sont que vraisemblables, la différence des données que chaque homme a sur elles, est une des causes principales de la diversité des opinions que l'on voit régner sur les mêmes objects.* (Laplace)

We present a multi-agent logic of probability and knowledge, with a very natural product update, yielding a simplification of the logic proposed in [7], which is in turn based on [15] and [6]. We show how probability measures on Kripke models can be defined in a straightforward way from lotteries. We propose a complete logic for lottery models, define an appropriate notion of bisimulation (different from the notion in [14,15]), and prove a Hennessey-Milner result for this notion. We prove that every model with lotteries is equivalent to a single-lottery model, where all agents share the same lottery. Finally, we investigate the model checking complexity of the logic.

This paper presents a logic of probability and knowledge where the two are related as follows:

Agent a knows ϕ if and only if the probability a assigns to ϕ equals 1.

Our proposal has obvious relations to earlier proposals on combining knowledge and probability [10,15,14,7,5,13] and many more. A key difference is that these proposals do not equate knowledge with certainty. An exception to this is [1].

A possible reason for not equating knowledge with certainty is the well-known difference between impossible in practice and impossible in theory which arises when measuring probabilities in uncountable spaces, where one equates "the probability of ϕ equals 1" with "ϕ is almost certain". An infinite process of fair coin throwing that results in an infinite sequence of 1s is practically imposssible (its probability is 0), but the sequence *is* in the sample space. Since we are careful to work with countable models and with lotteries that are bounded (Definition 2.2), this difficulty does not arise for us.

In real applications, knowledge and certainty are strongly related. We present our simplified framework of epistemic probability logic in Section 2. In particular, we will present models with a single lottery and in Section 3 we prove that the semantics with a single lottery and the semantics with several lotteries are equivalent, by constructing single lottery models from multiple lottery models. This throws light on the relation between subjective probability (modeled by multiple lotteries) and intersubjective probability (modeled by single lotteries). In Section 4 we define the appropriate notion of bisimulation, and use it to prove a Hennessy-Milner Theorem for epistemic probability logic. Section 5 gives an axiomatization for epistemic probability logic based on [10] and proves that the $S5$ axioms for certainty can be derived. In Section 6, we deal with the model checking procedure, and show that it runs in polynomial time. Section 7 explains how to add action model update in DEL style (but simplified), and gives a PSPACE-completeness proof for the model checking problem that results from adding a dynamic operator to the language.

2 Epistemic probabilistic logic

We present our epistemic lottery models (with the variant with a single lottery even for the multi-agent case). We then present the language of our version of epistemic probabilistic logic and its semantics and finally we show how to embed standard epistemic logic in our framework.

2.1 Epistemic lottery models

We start out from the definition of standard epistemic models.

Definition 2.1 A standard epistemic model $\mathcal{M}$ for a set $\mathbf{P}$ of propositions and a set A of agents is a tuple (W, V, R) where

- W is a non-empty, finite or countable set of worlds,
- V is a valuation function that assigns to every $w \in W$ a subset of $\mathbf{P}$.
- R is a function that assigns to every agent $a \in A$ an equivalence relation R_a on W.

To turn a standard epistemic model into an epistemic probability model, we assign to each agent a lottery, representing the subjective probabilities the agent assigns.

Definition 2.2 A W**-lottery** L is a function from a (finite or countable) set W to the set of positive (non-zero) rationals, i.e., $L : W \to \mathbb{Q}^+$. A W-lottery L is **bounded** on $V \subseteq W$ if $\sum_{v \in V} L(v) < \infty$.

Definition 2.3 An **epistemic multiple lottery model** $\mathcal{M}$ is a tuple $(W, V, R, \mathbb{L})$ where W, V, R are as in Definition 2.1 and $\mathbb{L}$ is a function that assigns to every agent $a \in A$ a W-lottery that is bounded on every R_a equivalence class.

We say that an epistemic lottery model is *normalized* if $\mathbb{L}_a$ restricted to E is a probability measure for all agents a and for all R_a-equivalence classes E. By the boundedness condition, all epistemic lottery models can be normalized.

Now, we define an epistemic lottery model where the lotteries are the same for each agent, that is $\mathbb{L}_a = \mathbb{L}_b$ for all agents a, b. We will write L instead of $\mathbb{L}_a$ for a given agent a. Models where there is a single lottery seem easier to manipulate. Formally:

Definition 2.4 An **epistemic single lottery model** $\mathcal{M}$ is a tuple (W, V, R, L) where W, V, R are as in Definition 2.1 and L is a W-lottery that is bounded on every R_a equivalence class, for every agent a.

2.2 Epistemic probability logic language

The language $\mathcal{L}$ of multi-agent epistemic probability logic is defined as follows.

Definition 2.5 Let p range over $\mathbf{P}$, a over A, q over $\mathbb{Q}$. Then $\mathcal{L}$ is given by:

$$\begin{array}{rcl} \phi & ::= & \top \mid p \mid \neg\phi \mid \phi \wedge \phi \mid t_a \geq 0 \mid t_a = 0 \\ t_a & ::= & q \mid q \cdot P_a\phi \mid t_a + t_a \end{array}$$

The intention in $t_a + t_a$ is that both indices are the same.

Some useful abbreviations:

- $\bot$, $\phi_1 \vee \phi_2$, $\phi_1 \to \phi_2$, $\phi_1 \leftrightarrow \phi_2$.

- $t \geq t'$ for $t + (-1)t' \geq 0$.
- $t < t'$ for $\neg t \geq t'$.
- $t > t'$ for $\neg t' \geq t$.
- $t \leq t'$ for $t' \geq t$.
- $t \neq t'$ for $\neg t = t'$.
- $P_a(\phi_1|\phi_2) = q$ for $P_a(\phi_2) > 0 \wedge q \cdot P_a(\phi_2) = P_a(\phi_1 \wedge \phi_2)$.

t_a generates linear expressions dealing with subjective probabilities of agent a. A formula of the form $t_a = 0$ or $t_a > 0$ is called an *a-probability formula*.

Given these, we have:

- $P_a\phi = q$ expresses that the probability of ϕ according to a equals q.
- $P_a(\phi_1|\phi_2) = q$ expresses that according to a, the probability of ϕ_1, conditional on ϕ_2, equals q.

The truth definition for $\mathcal{L}$ is given below.

Definition 2.6 Let $\mathcal{M} = (W, V, R, \mathbb{L})$ be an epistemic lottery model and let $w \in W$.

$$\begin{aligned}
\mathcal{M}, w &\models \top \quad \text{always} \\
\mathcal{M}, w &\models p \text{ iff } p \in V(w) \\
\mathcal{M}, w &\models \neg\phi \text{ iff it is not the case that } \mathcal{M}, w \models \phi \\
\mathcal{M}, w &\models \phi_1 \wedge \phi_2 \text{ iff } \mathcal{M}, w \models \phi_1 \text{ and } \mathcal{M}, w \models \phi_2 \\
\mathcal{M}, w &\models t_a \geq 0 \text{ iff } [\![t_a]\!]_w^{\mathcal{M}} \geq 0 \\
\mathcal{M}, w &\models t_a = 0 \text{ iff } [\![t_a]\!]_w^{\mathcal{M}} = 0.
\end{aligned}$$

$$\begin{aligned}
[\![q]\!]_w^{\mathcal{M}} &:= q \\
[\![q \cdot P_a\phi]\!]_w^{\mathcal{M}} &:= q \times P_{a,w}^{\mathcal{M}}(\phi) \\
[\![t_a + t'_a]\!]_w^{\mathcal{M}} &:= [\![t_a]\!]_w^{\mathcal{M}} + [\![t'_a]\!]_w^{\mathcal{M}}
\end{aligned}$$

$$P_{a,w}^{\mathcal{M}}(\phi) = \frac{\sum\{\mathbb{L}_a(u) \mid wR_au \text{ and } \mathcal{M}, u \models \phi\}}{\sum\{\mathbb{L}_a(u) \mid wR_au\}}.$$

Notice that $\mathbb{L}_a(u) > 0$ for all $u \in W$ so that there is no division by zero. Also, $\sum\{\mathbb{L}_a(u) \mid wR_au\} < \infty$, by the boundedness condition on $\mathbb{L}_a$. So $P_{a,w}$ is well-defined. The interpretation of formulas in epistemic single lottery models is similar except that we directly use L instead of $\mathbb{L}_a$ for a given agent a.

2.3 Relating Knowledge to Certainty

We use $K_a(\phi)$ as an abbreviation for $P_a(\phi) = 1$. This interprets knowledge as certainty and makes K_a behave as an $S5$-operator. [3]

Example 2.7 [Agents with different priors]

[3] If you have still qualms about this, then please read: "This interprets knowledge as almost-certainty."

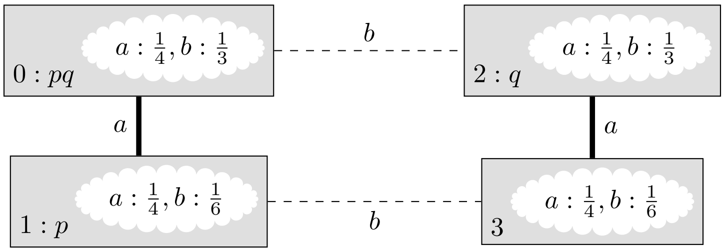

In the model of this example, loterries are $\mathbb{L}_a = \{0 : \frac{1}{4}, 1 : \frac{1}{4}, 2 : \frac{1}{4}, 3 : \frac{1}{4}\}$ and $\mathbb{L}_b = \{0 : \frac{1}{3}, 1 : \frac{1}{6}, 2 : \frac{1}{3}, 3 : \frac{1}{6}\}$. At world 0, the probability that a (represented by solid lines) assigns to p is 1, so $K_a p$ is true at 0. $K_a q$ is false at 0, for the probability that a assigns to q is less than 1. In fact, we have:

Proposition 2.8 *Let ϕ be a formula of standard epistemic logic. The following statements are equivalent:*

(i) *ϕ is satisfiable in a standard epistemic model;*

(ii) *$tr(\phi)$ is satisfiable in an epistemic single lottery model*

(iii) *$tr(\phi)$ is satisfiable in an epistemic lottery model*

where tr is defined by $tr(K_a\phi) = P_a tr(\phi) = 1$.

Proof. $(iii) \Rightarrow (i)$. If $tr(\phi)$ is satisfiable in an epistemic lottery model, we extract a standard epistemic model by dropping the lotteries and we prove that ϕ is true.

$(i) \Rightarrow (ii)$. Suppose that ϕ is satisfiable in a standard epistemic model. As $S5_n$ has the finite model property [8], there is a *finite* standard epistemic model for ϕ. We transform this standard epistemic model into an epistemic single lottery model $\mathcal{M} = (W, V, R, L)$ by adding a fake single lottery L that assigns 1 to all worlds. As the model is finite, it is guaranteed that the W-lottery L is bounded on every R_a equivalence class. We prove that $tr(\phi)$ is true in $\mathcal{M}$.

$(ii) \Rightarrow (iii)$. Because an epistemic single lottery model is an epistemic lottery model. $\square$

Proposition 2.8 can be generalized to set of formulas.

Proposition 2.9 *Let Σ be a formula of standard epistemic logic. The following statements are equivalent:*

(i) *Σ is satisfiable in a standard epistemic model;*

(ii) *$\{tr(\phi) \mid \phi \in \Sigma\}$ is satisfiable in an epistemic single lottery model*

(iii) *$\{tr(\phi) \mid \phi \in \Sigma\}$ is satisfiable in an epistemic lottery model*

where tr is defined by $tr(K_a\phi) = P_a tr(\phi) = 1$.

Proof. The proof is essentially the same one than for proposition 2.8 except for $(i) \Rightarrow (ii)$. Suppose that Σ is satisfiable in a standard epistemic model. The finite model property argument does not work anymore. Nevertheless, we suppose that the standard epistemic model has at most a countable number of worlds. We transform this standard epistemic model into an epistemic lottery model $\mathcal{M} = (W, V, R, \mathbb{L})$

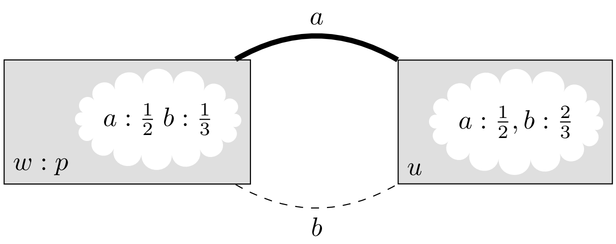

Fig. 1. Example of an epistemic probability model $\mathcal{M}$

by adding a fake lottery $\mathbb{L}$ as follows: we consider $\{w_0, w_1, \dots\}$ a (possibly finite) enumeration of worlds in W. We define $\mathbb{L}_a(w_k) = \frac{1}{k^2}$. As $\sum_{k=0}^{+\infty} \frac{1}{k^2} < +\infty$, it is guaranteed that the W-lottery $\mathbb{L}_a$ is bounded on every R_a equivalence class. We prove that $\{tr(\phi) \mid \phi \in \Sigma\}$ is true in $\mathcal{M}$. $\square$

Remark 2.10 Notice that if we define a belief operator $B_a\phi$ by $P_a(\phi) > \alpha$ for some $\alpha \in (\frac{1}{2}, 1]$, the formula $B_a p \wedge B_a q \wedge \neg B_a(p \wedge q)$ is satisfiable. That is, B_a behaves as a non-normal operator and not as a $KD45$ operator. This provides a way out of the so-called *lottery paradox* [16].

3 Single lottery versus multiple lottery models

In this section, we prove that the semantics given in terms of epistemic multiple lottery models (definition 2.3) and the semantics given in terms of epistemic *single* lottery models (definition 2.4) are equivalent, in the sense that for each multiple lottery model there is an equivalent single lottery model.

Philosophically, this suggests that objective probability, or at least intersubjective probability, can be defined from subjective probabilities. In any case, epistemic single lottery models are easier to handle because we attribute the same value to a world for each agent.

Proposition 3.1 *Given an epistemic lottery model $\mathcal{M} = (W, V, R, \mathbb{L})$, given a world w, there exists an epistemic single lottery model $\mathcal{M}' = (W', V', R', L)$ and a world $w' \in W'$ such that for all formulas ϕ, $\mathcal{M}, w \models \phi$ iff $\mathcal{M}', w' \models \phi$.*

Before starting the proof, let us consider an example. We start with the model $\mathcal{M}$ depicted in Figure 1 consisting in two worlds w and u. p is true in w and false in u. The lottery for agent a assigns probability $\frac{1}{2}$ to w and $\frac{1}{2}$ to u. The lottery for agent b assigns probability $\frac{1}{3}$ to w and $\frac{2}{3}$ to u.

In order to get a model with a single lottery, we unravel the model $\mathcal{M}$ and we obtain the infinite epistemic lottery model of Figure 2. The proof formalizes this transformation.

Proof. The construction goes as follows. The set of worlds W' is the set of all sequences of the form:

$$w_0 a_0 w_1 a_1 \dots w_{n-1} a_{n-1} w_n$$

such that, $n \geq 0$, $w_0 = w$, w_i are worlds, a_i are agents, $(w_i, w_{i+1}) \in R_{a_i}$ and $a_i \neq a_{i+1}$. For any sequence $\boldsymbol{s}$, we write $end(\boldsymbol{s})$ for the last world in the sequence

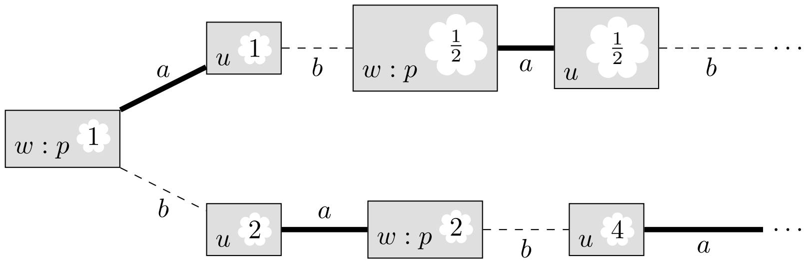

Fig. 2. Epistemic lottery model $\mathcal{M}'$ obtained by unraveling the model $\mathcal{M}$

that is, $end(w_0a_0w_1a_1 \dots w_{n-1}a_{n-1}w_n) = w_n$. The valuation V' is defined by $V'(\boldsymbol{s}) = V(end(\boldsymbol{s}))$. The relation R'_a is defined as follows:

$$\begin{aligned} R'_a = \; & \{(\boldsymbol{s}, \boldsymbol{s}) \mid \boldsymbol{s} \in W'\} \cup \\ & \{(\boldsymbol{s}, \boldsymbol{s}au), (\boldsymbol{s}au, \boldsymbol{s}) \mid \boldsymbol{s}, \boldsymbol{s}au \in W'\} \cup \\ & \{(\boldsymbol{s}au_1, \boldsymbol{s}au_2) \mid s, \boldsymbol{s}au_1, \boldsymbol{s}au_2 \in W'\} \end{aligned}$$

The lottery is defined by induction on sequences as follows:

- $L(w) = 1$;
- $L(sau) = \frac{\mathbb{L}_a(u)L(s)}{\mathbb{L}_a(end(s))}$.

It can now be proved by induction on ϕ that for all sequences s, $\mathcal{M}, end(s) \models \phi$ iff $\mathcal{M}', s \models \phi$. □

Notice that our construction produces models with infinite sets of worlds. We will prove in the next Section (Proposition 4.3) that this is unavoidable. What this means is that the logic, when interpreted over the class of single lottery models, does not have the finite model property. Also, it suggests that the logic is not expressive enough to characterize models that are built from a single lottery. What is needed to make such characterization possible? We leave this question for future work.

4 Bisimulation

In this section, we pick up a yarn in the story about bisimulation from [14], we modify (simplify) the definition so that it suits our logic, and we prove a Hennessy-Milner result for our new version. Next, we use our notion of bisimulation to prove that some finite multiple lottery models cannot have finite single lottery counterparts.

If X is a set of worlds with $X \subseteq R_a(w)$ then we use $\mathbb{L}_{a,w}(X)$ for $\frac{\sum_{x \in X} \mathbb{L}_a(x)}{\sum_{v \in R_a(w)} \mathbb{L}_a(v)}$.

Given two epistemic lottery models $\mathcal{M} = (W, V, R, \mathbb{L}), \mathcal{M}' = (W', V', R', \mathbb{L}')$, we say that a relation B is a bisimulation over $W \times W'$ if wBw' implies:

(i) w and w' satisfy the same atomic propositions;

(ii) for every set $E \subseteq R_a(w)$ there exists a set $E' \subseteq R'_a(w')$ such that:
- for all $u' \in E'$, there exists $u \in E$ such that uBu';
- and $\mathbb{L}_{a,w}(E) \leq \mathbb{L}_{a,w'}(E')$.

(iii) for every set $E' \subseteq R'_a(w')$ there exists a set $E \subseteq R_a(w)$ such that:

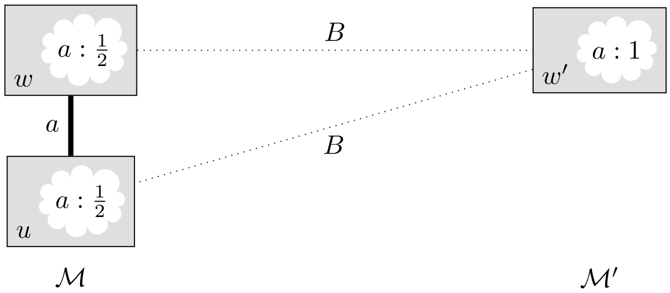

Fig. 3. Two models $\mathcal{M}$, $\mathcal{M}'$ and a bisimulation relation B

- for all $u \in E$, there exists $u' \in E'$ such that uBu';
- and $\mathbb{L}_{a,w'}(E') \leq \mathbb{L}_{a,w}(E)$.

If there exists a bisimulation B such that wBw' we say that w and w' are bisimilar, notation $w \leftrightarroweq w'$ or $\mathcal{M}, w \leftrightarroweq \mathcal{M}', w'$ if there is danger of ambiguity.

Figure 3 shows two models $\mathcal{M}$, $\mathcal{M}'$ and a bisimulation B. We see that condition (ii) and condition (iii) require inequalities and not equalities.

Proposition 4.1 *Let $\mathcal{M}, \mathcal{M}'$ be two models and w and w' be worlds of $\mathcal{M}$ and $\mathcal{M}'$, respectively. If $\mathcal{M}, w \leftrightarroweq \mathcal{M}', w'$ then $\mathcal{M}, w$ and $\mathcal{M}', w'$ satisfy the same formulas.*

Proof. We will prove by induction on ϕ and t_a:

(a) for all w, w' with $\mathcal{M}, w \leftrightarroweq \mathcal{M}', w'$: $\mathcal{M}, w \models \phi$ iff $\mathcal{M}', w' \models \phi$,

(b) for all w, w' with $\mathcal{M}, w \leftrightarroweq \mathcal{M}', w'$: $[\![t_a]\!]^{\mathcal{M}}_w = [\![t_a]\!]^{\mathcal{M}'}_{w'}$.

Let C be a bisimulation that witnesses $\mathcal{M}, w \leftrightarroweq \mathcal{M}', w'$.

For $\top$ (a) holds trivially. For atoms p, use property (i) of C. The cases of $\neg\phi$ and $\phi_1 \wedge \phi_2$ are straightforward, using the induction hypothesis for (a).

For the cases of $t_a \geq 0$ and $t_a = 0$, we assume that $[\![t_a]\!]^{\mathcal{M}}_w = [\![t_a]\!]^{\mathcal{M}'}_{w'}$, from which the statements $\mathcal{M}, w \models t_a \geq 0$ iff $\mathcal{M}', w' \models t_a \geq 0$ and $\mathcal{M}, w \models t_a = 0$ iff $\mathcal{M}', w' \models t_a = 0$ follow.

Next, we show (b). The key issue here is to show that $P^{\mathcal{M}}_{a,w}(\phi) = P^{\mathcal{M}'}_{a,w'}(\phi)$. From this (b) easily follows.

Let

$$E = \{v \in R^{\mathcal{M}}_a(w) \mid \mathcal{M}, v \models \phi\}.$$

By property (ii) of C there is a set $E' \subseteq R^{\mathcal{M}'}_a(w')$ such that:

(i) for all $u' \in E'$ there exists $u \in E$ with uCu';

(ii) $\mathbb{L}_{a,w}(E) \leq \mathbb{L}'_{a,w'}(E')$.

From (i) we get that for all $u' \in E'$: $\mathcal{M}', u' \models \phi$. Use this, plus (ii) and the fact that $[\![P_a\phi]\!]^{\mathcal{M}}_w = \mathbb{L}_{a,w}(E)$ to get:

$$[\![P_a\phi]\!]^{\mathcal{M}}_w \leq \mathbb{L}'_{a,w'}(E') \leq [\![P_a\phi]\!]^{\mathcal{M}'}_{w'}.$$

Let

$$E' = \{v' \in R^{\mathcal{M}'}_a(v) \mid \mathcal{M}', v' \models \phi\}.$$

By property (iii) of C there is a set $E \subseteq R_a^{\mathcal{M}}(w)$ such that:

(i) for all $u \in E$ there exists $u' \in E'$ with uCu';

(ii) $\mathbb{L}'_{a,w'}(E') \leq \mathbb{L}_{w,a}(E)$.

From (i) we get that for all $u \in E$: $\mathcal{M}, u \models \phi$. Use this, plus (ii) and the fact that $[\![P_i\phi]\!]_{w'}^{\mathcal{M}'} = \mathbb{L}'_{a,w'}(E')$ to get:

$$[\![P_i\phi]\!]_{w'}^{\mathcal{M}'} \leq \mathbb{L}_{a,w}(E) \leq [\![P_i\phi]\!]_{w}^{\mathcal{M}}.$$

Together this gives $[\![P_i\phi]\!]_{w}^{\mathcal{M}} = [\![P_i\phi]\!]_{w'}^{\mathcal{M}'}$. □

Now we adapt the proof of the Hennessy-Milner Theorem [8, p. 69] to our epistemic lottery logic. We say that a model $\mathcal{M}$ is image-finite iff for all worlds w in $\mathcal{M}$, and for all agents a, $R_a(w)$ is finite.

Proposition 4.2 *Let $\mathcal{M}, \mathcal{M}'$ be two image-finite models and w and w' be respectively two worlds of $\mathcal{M}$ and $\mathcal{M}'$. $\mathcal{M}, w$ and $\mathcal{M}', w'$ are bisimilar if, and only if $\mathcal{M}, w$ and $\mathcal{M}', w'$ satisfy the same formulas.*

Proof. We show the right to left direction and for that, we prove that the relation $\leftrightsquigarrow$ of modal equivalence on the two models is itself a bisimulation.

Condition (i) is immediate: if w and w' satisfy the same formulas, they satisfy the same atomic propositions. Assume that $w \leftrightsquigarrow w'$ and let E be a subset of $R_a(w)$. We will prove condition (ii) by arriving at a contradiction by assuming that there is no $E' \subseteq R'_a(w')$ such that

- for all $u' \in E'$, there exists $u \in E$ such that $u \leftrightsquigarrow u'$;
- and $\mathbb{L}_a(E) \leq \mathbb{L}_a(E')$.

That is, we assume that for every set $E' \subseteq R'_a(w')$ such that $\mathbb{L}_a(E) \leq \mathbb{L}_a(E')$ there exists $u' \in E'$ such that for all $u \in E$ we have $u \not\leftrightsquigarrow u'$.

Let $\mathcal{S}' = \{E'_1, \ldots, E'_n\}$ be an enumeration of sets $E' \subseteq R'_a(w')$ such that $\mathbb{L}_a(E) \leq \mathbb{L}_a(E')$. For all $i \in \{1, \ldots, n\}$, there exists $u' \in E'_i$ and a (finite) collection of formulas $(\psi_{i,u})_{u \in E}$ such that for all $u \in E$, $\mathcal{M}, u \models \psi_{i,u}$ and $\mathcal{M}', u' \not\models \psi_{i,u}$.

Let $\phi = \bigwedge_{i=1..n} \bigvee_{u \in E} \psi_{i,u}$. On the one hand, we have that for all $u \in E$, $\mathcal{M}, u \models \phi$. Thus, if we pose $\alpha = \mathbb{L}_a(E)$, we have $\mathcal{M}, w \models P_a(\phi) \geq \alpha$.

On the other hand, for all $i \in \{1, \ldots, n\}$, there exists a world $u' \in E'_i$ such that $\mathcal{M}', u' \models \bigwedge_{u \in E} \neg\psi_{i,u}$. That is $\mathcal{M}', u' \models \bigvee_{i=1..n} \bigwedge_{u \in E} \neg\psi_{i,u}$, that is to say $\mathcal{M}', u' \not\models \phi$. In particular, the set $\{u' \in R_a(w') \mid \mathcal{M}', u' \models \phi\}$ is not in $\mathcal{S}'$ and is therefore of probability strictly lower that α. So, $\mathcal{M}, w' \not\models P_a(\phi) \geq \alpha$.

So w and w' do not satisfy the same formulas and there is a contradiction hence condition (ii) holds. Condition (iii) is symmetrical and may be checked in a similar way. □

Proposition 4.3 *There is no finite epistemic single lottery model $\mathcal{M}'$ that is bisimilar to the model $\mathcal{M}$ from Figure 1.*

Proof. Suppose there is a finite epistemic single lottery model $\mathcal{M}' = (W', V', R', L)$ that is bisimilar to the model $\mathcal{M}$ from Figure 1.

Let $R'_a(w_1), \dots R'_a(w_\ell)$ be an enumeration of a-equivalence classes in $\mathcal{M}'$. Let $R'_b(w_1), \dots R'_b(w_n)$ be an enumeration of b-equivalence classes in $\mathcal{M}'$.

As $\mathcal{M}$ and $\mathcal{M}'$ are bisimilar, we have:

- $\sum_{u \in R'_a(w_i) | p \in V(u)} L(u) = \sum_{u \in R'_a(w_i) | p \notin V(u)} L(u)$;
- $2 \sum_{u \in R'_b(w_i) | p \in V(u)} L(u) = \sum_{u \in R'_b(w_i) | p \notin V(u)} L(u)$.

Now:

- On the one hand,
$$\begin{aligned} \textstyle\sum_{u \in W' | p \notin V(u)} L(u) &= \textstyle\sum_{i=1}^{\ell} \sum_{u \in R'_a(w_i) | p \notin V(u)} L(u) \\ &= \textstyle\sum_{i=1}^{\ell} \sum_{u \in R'_a(w_i) | p \in V(u)} L(u) \\ &= \textstyle\sum_{u \in W' | p \in V(u)} . \end{aligned}$$
- On the other hand,
$$\begin{aligned} \textstyle\sum_{u \in W' | p \notin V(u)} &= \textstyle\sum_{i=1}^{n} \sum_{u \in R'_b(w_i) | p \notin V(u)} L(u) \\ &= 2 \textstyle\sum_{i=1}^{n} \sum_{u \in R'_b(w_i) | p \in V(u)} L(u) \\ &= 2 \textstyle\sum_{u \in W' | p \in V(u)} . \end{aligned}$$

Thus, $\sum_{u \in W' | p \notin V(u)} L(u) = 0$ which contradicts the definition of model $\mathcal{M}'$. We have proved by contradiction that there is no finite epistemic single lottery model $\mathcal{M}'$ that is bisimilar to $\mathcal{M}$.

□

5 Axiomatization

Figure 4 shows a complete axiomatization of epistemic probabilistic logic. We show in subsection 5.1 that the principles of standard epistemic logic $S5$ (where $K_a\phi$ is replaced by $P_a\phi = 1$) are derivable from the axiomatization. In subsection 5.2, we adapt the proof of completeness of [10] to our simplified logic.

5.1 Principles of $S5$ are derivable

Principles of $S5$ are the following:

- the necessitation rule: if $\vdash_{S5} \phi$ then $\vdash_{S5} K_a\phi$;
- the K-principle: $K_a\phi_1 \wedge K_a(\phi_1 \to \phi_2) \to K_a\phi_2$;
- the T-axiom or truth axiom : $K_a\phi \to \phi$;
- the 4-axiom (positive introspection): $K_a\phi \to K_aK_a\phi$;
- the 5-axiom (negative introspection): $\neg K_a\phi \to K_a\neg K_a\phi$.

First, remark that **ProbaT** corresponds to the T-axiom. Now, we prove that all other $S5$-principles are derivable.

Proposition 5.1 *The necessitation rule for certainty is derivable:*

If $\vdash \phi$ *then* $\vdash P_a\phi = 1$.

Proof. From $\vdash \phi$ derive $\vdash \phi \leftrightarrow \top$. From this and **ProbaRule**, $\vdash P_a(\phi) = P_a(\top)$ and with **ProbaTrue**, gives $\vdash P_a(\phi) = 1$. □

Propositional Logic Axioms

All instances of tautologies of propositional logic are axioms (**CPL**)
From $\vdash \phi$ and $\vdash \phi_1 \to \phi_2$ conclude $\vdash \phi_2$. (**ModusPonens**)

Probability Rule

If $\vdash \phi_1 \leftrightarrow \phi_2$ then $\vdash P_a\phi_1 = P_a\phi_2$. (**ProbaRule**)

Probability Axioms

$\vdash \quad P_a\phi \geq 0$ (**ProbaNonNeg**)
$\vdash \quad P_a\top = 1$ (**ProbaTrue**)
$\vdash \quad P_a(\phi_1 \wedge \phi_2) + P_a(\phi_1 \wedge \neg\phi_2) = P_a\phi_1$ (**ProbaAdditivity**)

$\vdash \quad \psi \to P_a\psi = 1$ for all a-probability formulas ψ (**ProbaProba**)

Certainty Axioms

$\vdash \quad P_a\phi = 1 \to \phi$ (**ProbaT**)

Linear (in)equality axioms

All instances of valid formulas about linear inequalities (**Linear**)

Fig. 4. Axiomatization

Proposition 5.2 *The K-principle for certainty is derivable:*

$$\vdash P_a\phi_1 = 1 \wedge P_a(\phi_1 \to \phi_2) = 1 \to P_a\phi_2 = 1 \qquad (*)$$

Proof. From **ProbaRule**, **ProbaAdditivity** and **Linear**, we have: $\vdash P_a(\top \wedge \phi) + P_a(\top \wedge \neg\phi) = 1$. With **ProbaRule** this gives: $\vdash P_a\phi + P_a(\neg\phi) = 1$. And therefore: $\vdash P_a(\neg\phi) = 1 - P_a\phi$. Using this, we derive $P_a(\phi_1 \to \phi_2) = P_a(\neg(\phi_1 \wedge \neg\phi_2)) = 1 - P_a(\phi_1 \wedge \neg\phi_2)$. From this: $P_a(\phi_1 \to \phi_2) = 1 \leftrightarrow P_a(\phi_1 \wedge \neg\phi_2) = 0$. On the other hand, using **ProbaAdditivity**: $P_a(\neg\phi_2) = P_a(\phi_1 \wedge \neg\phi_2) + P_a(\neg\phi_1 \wedge \neg\phi_2)$. Therefore: $P_a\phi_1 = 1 \wedge P_a(\phi_1 \to \phi_2) = 1 \to P_a(\neg\phi_2) = P_a(\neg\phi_1 \wedge \neg\phi_2) = 0$ By Propositional logic axioms, this proves $(*)$. $\square$

Formula $(*)$ is a theorem of epistemic probability logic; it can be viewed as a probabilistic version of the K-axiom in epistemic logic.

Note that the following is derivable from **ProbaT** and **ProbaAdditivity**:

$$\vdash P_a\phi = 0 \rightarrow \neg\phi \quad \textbf{(ProbaTfalse)}$$
$$\vdash \phi \rightarrow P_a\phi > 0 \quad \textbf{(ProbaGeq0)}$$

Note that the following formulas are theorem because they are instantiation of **ProbaProba**:

$$\vdash P_a\phi = 1 \rightarrow P_a(P_a\phi = 1) = 1 \quad \textbf{(4)}$$
$$\vdash P_a\phi > 0 \rightarrow P_a(P_a\phi > 0) = 1 \quad \textbf{(5)}$$

They corresponds respectively to axiom 4 (positive introspection) and axiom 5 (negative introspection) in standard epistemic logic $S5$.

5.2 Soundness and Completeness

Theorem 5.3 *The EPL calculus is sound.*

Proof. All axioms are valid in all EPL models. All rules preserve validity. □

Theorem 5.4 *The EPL calculus is complete.*

The proof is given in the appendix.

6 Model checking

Here is the algorithm for model checking, where again it is assumed that the input model $\mathcal{M}$ is a normalized epistemic lottery model.

```
function mc(M = (W, V, R, 𝕃), φ)
    if T[φ] is defined then
        return T[φ];
    endIf
    match φ do
        case ⊤:
            T[φ] := W;
            return T[φ];
        case p:
            T[φ] := {w ∈ W | p ∈ V(w)};
            return T[φ];
        case ¬φ:
            T[φ] := W \ mc(φ);
            return T[φ];
        case φ1 ∧ φ2:
            T[φ] := mc(M, φ1) ∩ mc(M, φ2);
            return T[φ];
        case t_a ≥ q:
            T[φ] := {w ∈ W | get(M, t_a, w, i) ≥ q};
            return T[φ];
    endMatch
endFunction
```

```
function get(M, t, w, i)
    match t do
        case q:
            return q;
        case q · P_a(φ):
            Σ := mc(M, φ);
            v := Σ_{u∈Σ|wR_a u} 𝕃_a(u)
            return q × v;
        case t_1 + t_2:
            return get(M, t_1, w) + get(M, t_2, w);
    endMatch
endFunction
```

Theorem 6.1 *A call to $mc(\mathcal{M}, \phi)$ returns the set $\{w \in W \mid \mathcal{M}, w \models \phi\}$.*

Proof. By induction on ϕ. □

Theorem 6.2 *A call to $mc(\mathcal{M}, \phi)$ requires $O(|\phi|^2 \times |W|^3)$ elementary operations.*

Proof. A call to $mc(\mathcal{M}, \phi)$ calls $mc(\mathcal{M}, \psi)$ where ψ is a subformula of ϕ. As the algorithm $mc(\mathcal{M}, \phi)$ is based on memoization: for a given ψ, the call $mc(\mathcal{M}, \psi)$ is called at most once. So it is sufficient to compute an upper bound of the number of elementary operations performed in one call $mc(\mathcal{M}, \psi)$. Then we multiply this upper bound by an upper bound of the number of calls, that is the number of subformulas of ϕ which is $O(|\phi|)$.

We may represent subsets of W by an array of Booleans $W \rightarrow \{0, 1\}$. The case $\phi_1 \times \phi_2$ uses the intersection operation that requires $O(|W|)$ operations. Let us study the case $t_a \geq q$. We first browse all the $O(|W|)$ worlds w and we check whether $\sum_k get(\mathcal{M}, t_k, w) \geq q$ holds or not. We make n calls to get and $n = O(\phi)$. Each call to get requires at most $O(|W|)$ for browsing successors of w by R_a. We count the call to $mc(\mathcal{M}, \phi)$ as $O(1)$ since it is done once with the memoïzation and its effective computation is counted apart.

Conclusion: each call to $mc(\mathcal{M}, \psi)$ costs at most $O(|\phi| \times |W|^2)$. There are at most $|\phi|$ such calls so the global complexity is bounded by $O(|\phi|^2 \times |W|^3)$. □

7 Updates

7.1 Example

Consider the following story. An urn contains a single marble, either white or black. Mr A and Mrs B know this, and they also know that both possibilities are equally likely. Next, Mr A looks in the urn, while Mrs B is watching. Mr A puts another marble in the urn, a white one, and Mrs B sees this. The urn now contains two marbles. Next, Mrs B draws one of the two marbles from the urn. It turns out to be white. What is the probability, according to Mr A, that the other marble is also white? What is the probability, according to Mrs B, that the other marble is also white? (This is a multi-agent variation on a puzzle by Lewis Carroll, see [12].)

Call the first white marble p and the second one q. We start with a situation where there is nothing in the urn, and both agents know this. Update this with the action of tossing a fair coin and making p true in case the coin shows heads. It is assumed that the two agents a and b see that the action happens, but do not see what the outcome is. The action model for this (solid arrows for a, dashed arrows for b):

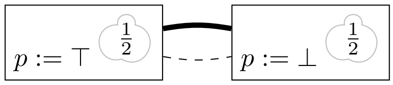

The update of the initial model looks like this:

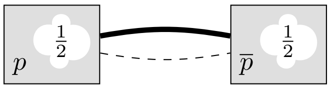

The action where a takes a look, while b sees this but does not observe what a sees:

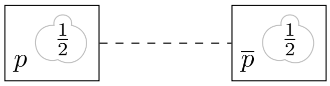

The situation after a has taken a look:

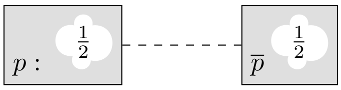

The action of putting another white marble (represented by q) in the urn:

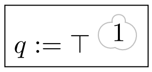

The result of updating with this:

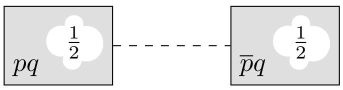

Extracting a white marble from the box is represented as either the act of removing p or the act of removing q, with neither a nor b seeing the difference. The act of removing p (making p false) has as precondition that p is true, the act of removing q has as precondition that q is true.

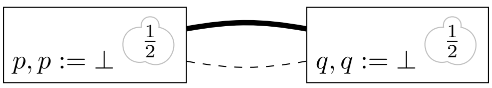

The result of updating with this is the following model:

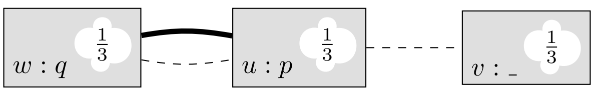

We see that in w it holds that $P_a(p \vee q) = 1$, $P_b(p \vee q) = \frac{2}{3}$. Same values in u, while in v it holds that $P_a(p \vee q) = 0$, $P_b(p \vee q) = \frac{2}{3}$.

7.2 Definitions

Formally, an action model $\mathcal{E}$ for epistemic probability logic is the result of replacing the valuation function in an epistemic lottery model by a pair of functions PRE and

POST that assign to every world (or: event) a precondition and a postcondition, where the precondition ϕ is a formula of the epistemic probability language, and the postcondition is a finite set of bindings $p := \phi$, with p in the set of basic proposition letters of the epistemic probability language, and ϕ a formula of the language.

Update is defined as the product construction of [4], with the extra proviso that $L_a(w, e) = L_a(w) \times L'_a(e)$, where L_a is the a-lottery of the input model and L'_a is the a-lottery of the update model. Let $\mathcal{M} \times \mathcal{E}$ denote the product of $\mathcal{M}$ and $\mathcal{E}$. If the initial epistemic lottery model and the update model are both normalized, then the product defines an epistemic lottery model (not necessarily normalized, for some (w, e) pairs may drop out by the update restriction). Our update definition is a considerable simplification of the update defined in [7].

We consider a probabilistic version of the language extended with the dynamic operator of [4] $[\pi]\psi$ where π defined by $\pi ::= \mathcal{E}, e \mid \pi \cup \pi$. This allows updates with pointed action models and choice between such updates, by means of the union operator $\cup$. Call the new language DEPL. The truth conditions are defined as follows:

- $\mathcal{M}, w \models [\mathcal{E}, e]\psi$ iff $\mathcal{M}, w \models PRE(e)$ implies $\mathcal{M} \otimes \mathcal{E}, (w, e) \models \psi$;
- $\mathcal{M}, w \models [\pi \cup \pi']\psi$ iff $\mathcal{M}, w \models [\pi]\psi$ and $\mathcal{M}, w \models [\pi']\psi$.

7.3 Model checking with updates

To study model checking in DEPL, we adapt the model checking procedure written in Section 6. Now, array T is replaced by $T_{\mathcal{M}}$ where $\mathcal{M}$ is the current epistemic lottery model. The implemented version works as follows:

case $[\mathcal{E}, e]\psi$:

$T_{\mathcal{M}}(PRE(e)) = mc(\mathcal{M}, PRE(e))$;

$T_{\mathcal{M}}([\mathcal{E}, e]\psi) := \left\{ w \in W \mid \begin{array}{l} w \notin T_{\mathcal{M}}(PRE(e)) \\ (w, e) \in mc(\mathcal{M} \times \mathcal{E}, \psi) \end{array} \right\}$;

return $T_{\mathcal{M}}([\mathcal{E}, e]\psi)$;

This leads to an algorithm which is running in exponential time and that uses an exponential amount of memory. We may write an algorithm that only use a polynomial amount of memory in the size of the initial model and the size of the formula, that is inspired by the algorithm provided in [2]: we browse the product models *on the fly*. Thus our model checking in the dynamic case is in PSPACE. Nevertheless, the PSPACE-hardness bound for DEL without probability, with $\cup$ operator (and where preconditions are all $\top$) shown in [2] does not provide a lower bound because we can not reduce DEL without probability on models without constraints to the model checking problem in DEPL. Nevertheless the idea of the proof of [2] can be adapted and it provides the following lemma.

Lemma 7.1 *The model checking problem when the initial models and action models are $S5$-models, when we have the $\cup$-operator in the language and when there are at least two agents is PSPACE-hard. The result holds even if all preconditions of the action models are propositional formulas*

Proof. Without loss of generality, we only consider in this proof quantified Boolean formulas of the form $\forall p_1 \exists p_2 \forall p_3 \ldots \forall p_{2k-1} \exists p_{2k} \psi(p_1, \ldots p_{2k})$, where $\psi(p_1, \ldots, p_{2k})$ is a Boolean formula over the atomic propositions $p_1, \ldots, p_{2k}$.

The *quantified Boolean formula satisfiability problem* takes as an input a natural number k and a quantified Boolean formula $\phi \triangleq \forall p_1 \exists p_2 \forall p_3 \dots \forall p_{2k-1} \exists p_{2k} \psi(p_1, \dots, p_{2k})$. It returns yes iff ϕ is true in quantified Boolean logic.

Let ϕ be such a quantified Boolean formula. We define a pointed epistemic model $\mathcal{M}, w^0$, $2k$ pointed event models $\mathcal{E}_1, e_1, \dots, \mathcal{E}_{2k}, e_{2k}$, a pointed event model $\mathcal{E}_{\circlearrowleft}, e_{\circlearrowleft}$ and an epistemic formula ψ' that are computable in polynomial time in the size of ϕ such that:

$$\begin{gathered}\phi \text{ is satisfiable in quantified Boolean logic}\\ \text{iff}\\ \mathcal{M}, w^0 \models [\mathcal{E}_1, e_1 \cup \mathcal{E}_{\circlearrowleft}, e_{\circlearrowleft}]\langle \mathcal{E}_2, e_2 \cup \mathcal{E}_{\circlearrowleft}, e_{\circlearrowleft}\rangle \dots\\ [\mathcal{E}_{2k-1}, e_{2k-1} \cup \mathcal{E}_{\circlearrowleft}, e_{\circlearrowleft}]\,\langle \mathcal{E}_{2k}, e_{2k} \cup \mathcal{E}_{\circlearrowleft}, e_{\circlearrowleft}\rangle \psi'.\end{gathered}$$

where

- $\mathcal{M}$ is depicted below:

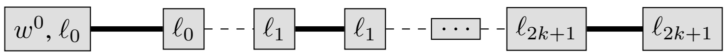

  where $\ell_0, \ell_1, \dots, \ell_{2k+1}$ are distinct propositional letters.
- For all $i \in \{1, \dots, 2k\}$, $\mathcal{E}_i, e_i$ is the action model depicted below:

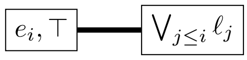

- $\mathcal{E}_{\circlearrowleft}, e_{\circlearrowleft}$ is the action model made up of a single event $e_{\circlearrowleft}$ with precondition $\top$;
- ψ' is the formula ψ where all p_i occurrences are substituted by $(\hat{K}_a \hat{K}_b)^i K_a K_b \ell_i$.

p_i is true is interpreted as the existence of a branch that stop at ℓ_i world. Making the product with $\mathcal{E}_i, e_i$ will add such a branch in the model whereas making the product with $\mathcal{E}_{\circlearrowleft}, e_{\circlearrowleft}$ will leave the epistemic model as it is. The universal and existential choices of values for the p_i's are simulated by the dynamic epistemic operators. □

Proposition 7.2 *Model checking for DEPL with at least two agents and with the* $\cup$ *operator is PSPACE-complete.*

Proof. Membership in PSPACE comes from the remark above. We polynomially reduce the model checking of $S5$-models that is PSPACE-hard (Lemma 7.1) in the model checking of DEPL. To do so, we add to the models 'artificial' lotteries and we recursively replace all subformulas $K_a \phi$ by $P_a(\phi) = 1$. Thus, we obtain the lower-bound. □

8 Connections, Further Work

The assumption that agents have a common prior, widely used in epistemic game theory, is not built into our concept of an epistemic probability model. It follows from Proposition 3.1 that "having a common prior" does *not* coincide with "having the same lottery."

If we want to impose common prior conditions, say for proposition p, then a natural way to express this would be by means of:

$$\bigwedge_{a,b\in A} P_a p = P_b p.$$

Currently, this is not in our language, but if we allow such expressions, then this formula rules out models like the following:

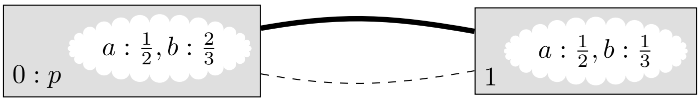

This model describe a situation where a and b 'agree to disagree' on the probability of p. If they are both willing to take bets on the truth of p, they are not rational, for then they make themselves vulnerable to a pair of bets that forms a Dutch book [3]. In finite models with a single lottery Dutch books cannot occur.

Question 1 *Can we strengthen the language to allow for an axiom that forces lotteries to be single?*

If we want to allow lotteries with unknowns in our models, then the language should be extended with expressions B_p with meaning: the (unknown) probability of p, and lotteries should allow for factors B_p. To handle cases where it is *given* that no probability distribution for an event exists, we can allow lotteries with *unknown factors*. A W-lottery with unknowns $X \subseteq \mathbf{P}$ (or: a W-lottery functional over X) is a function from $(0..1)^X$ to W-lotteries, where $(0..1)$ is the open unit interval $\subseteq \mathbb{Q}$. Thus, the type of a W-lottery with unknowns X is:

$$(X \to (0..1)) \to W \to \mathbb{Q}^+$$

Let B be a function that assigns probabilities to the members of X, i.e., $B : X \to (0..1)$. Let l be a W-lottery with values summing up to 1 over W, and let V be a valuation for W. Then $L_{l,V,B}$ is the W-lottery given by:

$$\begin{aligned} L_{l,V,B}(w) &= l(w) \\ &\times \prod\{B(p) \mid p \in Q, p \in V(w)\} \\ &\times \prod\{1 - B(p) \mid p \in \mathbf{P}, p \notin V(w)\}. \end{aligned}$$

Then for all $w \in W$, $L_{l,V,B}(w) \in (0..1) \subseteq \mathbb{Q}$, so $L_{l,V,B}$ is a W-lottery. The function $B \mapsto L_{l,V,B}$ is a lottery functional.

Example 8.1 [Von Neumann's Trick] How to obtain fair results from a coin with unknown bias [18]:

> Toss the coin twice. If the results match, forget both results and start over. If the results differ, use the first result.

Here is the explanation. Represent the coin as a lottery functional for the set $\{h\}$. Let B assign a probability to h. That is, $B_h = b$ is the coin bias. Then the probabilities of the four possible outcomes of Von Neumann's procedure are represented by the following lottery:

$$\{hh : b^2, ht : b - b^2, th : b - b^2, tt : (1 - b)^2\}.$$

This shows that the cases ht and th are equally likely, so interpreting the first as h and the second as t gives indeed a model of a fair coin.

Example 8.2 [Model representing a coin with unknown bias]

Model checking and model update for the epistemic probability logic of this paper is implemented in [9]. This allows to solve urn problems in a multi-agent setting by means of epistemic model checking. This extension generates lots of further logical questions. Also, it can serve as a solid basis for the design and analysis of probabilistic protocol languages for epistemic probability updating. Hooking up to more sophisticated model checkers like NuSMV (`nusmv.fbk.eu`) is future work. Finally, we would like to further explore the obvious connections with Bayesian learning.

Acknowledgement

Thanks to Alexandru Baltag, Bryan Renne and Joshua Sack for illuminating conversations on the topic of this paper. We also thank the reviewers for their support and comments.

References

[1] Achimescu, A. C., "Games and Logics for Informational Cascades," Master's thesis, ILLC, Amsterdam (2014).

[2] Aucher, G. and F. Schwarzentruber, *On the complexity of dynamic epistemic logic*, in: *Proceedings of TARK 2013*, 2013.

[3] Aumann, R., *Agreeing to disagree*, Annals of Statistics **4(6)** (1976), pp. 1236–1239.

[4] Baltag, A., L. Moss and S. Solecki, *The logic of public announcements, common knowledge, and private suspicions*, in: I. Bilboa, editor, *Proceedings of TARK'98*, 1998, pp. 43–56.

[5] Baltag, A. and S. Smets, *Probabilistic dynamic belief revision*, Synthese **165** (2008), pp. 179–202.

[6] Benthem, J. v., *Conditional probability meets update logic*, Journal of Logic, Language and Information **12** (2003), pp. 409–421.

[7] Benthem, J. v., J. Gerbrandy and B. Kooi, *Dynamic update with probabilities*, Studia Logica **93** (2009), pp. 67–96.

[8] Blackburn, P., M. de Rijke and Y. Venema, "Modal Logic," Cambridge Tracts in Theoretical Computer Science, Cambridge University Press, 2001.

[9] Eijck, J. v., *Learning about probability* (2013), available from `homepages.cwi.nl:~/jve/software/prodemo`.

[10] Fagin, R. and J. Halpern, *Reasoning about knowledge and probability*, Journal of the ACM (1994), pp. 340–367.

[11] Fagin, R., J. Y. Halpern and N. Megiddo, *A logic for reasoning about probabilities*, Information and computation **87** (1990), pp. 78–128.

[12] Gartner, M., "Mathematical Circus," Vintage, 1981.

[13] Gierasimszuk, *Bridging learning theory and dynamic epistemic logic*, Synthese **169** (2009), pp. 371–374.

[14] Kooi, B. P., "Knowledge, Chance, and Change," Ph.D. thesis, Groningen University (2003).
[15] Kooi, B. P., *Probabilistic dynamic epistemic logic*, Journal of Logic, Language and Information **12** (2003), pp. 381–408.
[16] Kyburg, H., "Probability and the Logic of Rational Belief," Wesleyan University Press, Middletown, CT, 1961.
[17] Laplace, M. l. C., "Essai Philosophique sur les Probabilités," Courcier, Paris, 1814.
[18] von Neumann, J., *Various techniques used in connection with random digits*, Technical Report 12: 36, National Bureau of Standards Applied Math Series (1951).

Appendix: Completeness

The completeness works as follows. We prove that:

Proposition 8.3 *If ϕ is consistent, then ϕ is satisfiable.*

We adapt the proof from [10]. First we construct a canonical epistemic probability model. Contrary to the proof in [10], the epistemic relations are inferred from probabilities.

Let $SF(\phi)$ be the set of all subformulas of ϕ augmented with the negations of subformulas. Let us define the canonical model $\mathcal{M} = (W, V, R, \mathbb{L})$. W is the set of all maximal consistent subsets of $SF(\phi)$. W is not empty because ϕ is supposed to be consistent. Valuations are defined as follows: $V(w) = \mathbf{P} \cap w$.

Let $sat(w) = \{\psi \mid w \vdash \psi\}$, that is, $sat(w)$ is the set of formulas that are provable from w. Relations are defined as follows: $wR_a u$ iff $sat(w)$ and $sat(u)$ contain the same a-probability formulas.

Now it remains to define $\mathbb{L}$. Let us consider an agent a and an equivalence class $R_a(w)$ in the canonical model $\mathcal{M}$. All worlds u of $R_a(w)$ contain the same a-probability formulas. In the sequel, we are transforming all the a-probability formulas in a system of linear inequations that is consistent.

For all $u \in W$, we write ϕ_u the conjunction of all formulas in u. We have:

- $\vdash \phi_u \rightarrow \neg\phi_v$ if $u \neq v$ by **CPL**.

Given any formulas ψ of $SF(\phi)$, we have

- $\vdash \psi \leftrightarrow \bigvee_{u \in W \mid \psi \in u} \phi_u$ by **CPL**.

Let ψ be any formula of $SF(\phi)$. By axioms **ProbaRule** and **ProbaAdditivity**, we have:

- $\vdash P_a(\psi) = \sum_{u \in W \mid \psi \in u} P_a(\phi_u)$.

Thus, if we take any a-probability formula ψ, and we replace any term $P_a(\chi)$ by $\sum_{u \in W \mid \psi \in u} P_a(\phi_u)$, we obtain

$$\sum_{u \in W} c_u P_a(\phi_u) \geq b$$

where $c_u, b \in \mathbb{Q}$. Now, when we evaluate the value of $P_a(\phi_u)$ in w, we should obtain non-zero if, and only if, $u \in R_a(w)$. Let us prove it.

- If $u \in R_a(w)$, we have:

(i) $\vdash \phi_u \rightarrow P_a(\phi_u) > 0$ by **ProbaGeq0**;
(ii) $P_a(\phi_u) > 0 \in sat(u)$;

(iii) $P_a(\phi_u) > 0 \in sat(w)$ because $u \in R_a(w)$.
There, $P_a(\phi_u) > 0$ should be also true in w.

- If $u \notin R_a(w)$, u and w differ by at least one a-probability formula $\psi \in SF(\phi)$ such that $\psi \in w$ and $\psi \notin u$ without loss of generality. We have:

(i) $\vdash \phi_w \to \psi$ by **CPL**;
(ii) $\vdash \psi \to \neg\phi_u$ by **CPL**;
(iii) $\vdash \psi \to P_a(\psi) = 1$ axiom **ProbaProba**;
(iv) $\vdash \phi_w \to P_a(\psi) = 1$ by **CPL** and **ModusPonens**;
(v) $\vdash \phi_w \to P_a(\neg\phi_u) = 1$ by 2. and 4.
(vi) $\vdash \phi_w \to P_a(\phi_u) = 0$.
Therefore $P_a(\phi_u) = 0$ should be true in w.

Thus, ψ should be equivalent to

$$\sum_{u \in R_a(w)} c_u P_a(\phi_u) \geq b$$

where $c_u, b \in \mathbb{Q}$. This yields a system of linear inequations made up of inequations $\sum_{u \in R_a(w)} c_u x_u \geq b$ when $\psi \in w$ or $\sum_{u \in R_a(w)} c_u x_u < b$ when $\psi \notin w$, plus $\sum_{u \in R_a(w)} x_u = 1$ and $x_u > 0$ for all $u \in R_a(w)$. The set $sat(w)$ is consistent so the above system, which is a rephrasing of some inequations that are in $sat(w)$, is also consistent and therefore satisfiable [11, Theorem 2.2]. Let $(x_u*)_{u \in R_w(a)}$ be a solution. We define $\mathbb{L}_a(u) = x_u*$.

Lemma 8.4 (truth lemma) *For all formulas $\psi \in SF(\phi)$, we have $\mathcal{M}, w \models \psi$ iff $\psi \in w$.*

Proof. By induction on ψ. □