Learning about Probability

Jan van Eijck
CWI & ILLC, Amsterdam

INRIA talk, Rennes, April 16, 2013

Abstract

This talk proposes a 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 and an equivalence relation on lotteries over worlds, where a lottery over (finite) world set W is a function from W to the positive rational numbers.

Uncertainty about probability is modelled as equivalence of lotteries. The difference with the usual approach is that probability is linked to knowledge rather than belief, and that “agent a knows that φ” is equated with “agent a assigns probability 1 to φ.”

To motivate our approach, we formulate and prove a Certainty Theorem, stating that certainty in an epistemic probability model M corresponds to knowledge in the epistemic model that results when all lottery information gets erased from M. It follows immediately from this that the certainty operator in epistemic probability logic is an S5 operator.

If there is time, the talk will also introduce PRODEMO, a model checker for epistemic probability logic that can be used to keep track of information flow about aleatory acts among multiple agents.

Probability as a function of degree of information

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 [Lap14]

Relation between Probability and Knowledge

Agent a knows φ iff the probability a assigns to φ equals 1.

Let $P_a\varphi$ be the probability that agent a assigns to φ.

Certainty implies Truth

$$P_a\varphi = 1 \to \varphi.$$

Positive Introspection into Certainty

$$P_a\varphi = 1 \to P_a(P_a\varphi = 1) = 1.$$

Negative Introspection into Certainty

$$P_a\varphi < 1 \to P_a(P_a\varphi < 1) = 1.$$

Earlier proposals on combining knowledge and probability [FH94, Koo03b, Koo03a, BGK09, BS08, Gie09], and many more. These proposals do not equate knowledge with certainty.

Lotteries

A W-lottery l is a function from a set of worlds W to the set of positive rationals, i.e., $l : W \to \mathbb{Q}^+$. Two W-lotteries l, l' are equivalent if for some $q \in \mathbb{Q}^+$, $l' = (\lambda p \mapsto q * p) \cdot l$.

We say that two W-lotteries l, l' have the same scale if

$$\sum\{l(w) \mid w \in W\} = \sum\{l'(w) \mid w \in W\}.$$

A W-lottery l is normalized on $B \subseteq W$ if $\sum\{l(w) \mid w \in B\} = 1$.

If we have a lottery $l : W \to \mathbb{Q}^+$ and a block $B \subseteq W$ in a partition of W, then this determines a probability distribution P on B, by means of (we assume that $B \neq \emptyset$):

$$P(w) = \frac{l(w)}{\sum\{l(w') \mid w' \in B\}}.$$

Lotteries with Unknowns, or Lottery Functionals

To handle cases where it is given that no probability distribution for an event exists, we allow lotteries with unknown factors.

A W-lottery with unknowns $Q \subseteq \mathbf{P}$ (or: a W-lottery functional over Q) is a function from $(0..1)^Q$ to W-lotteries, where $(0..1)$ is the open unit interval $\subseteq \mathbb{Q}$.

Thus, the type of a W-lottery with unknowns Q is:

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

Constructing Lotteries from Lottery Functionals

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

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

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: Von Neumann's Trick

How to obtain fair results from a coin with unknown bias [vN51]:

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

Example: Von Neumann's Trick

How to obtain fair results from a coin with unknown bias [vN51]:

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

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

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

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.

Urn Example

Say there are two urns, U and V. U contains one black marble and two white marbles, V contains one black marble and one white marble. This is common knowledge among a, b and c. Now a selects one of the urns, without revealing which one to b, c. Then b picks a marble from it, without revealing the marble to a, c.

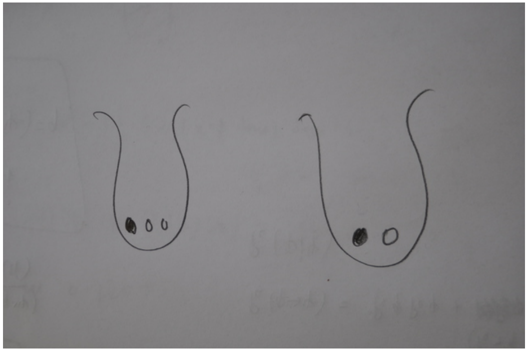

Representation

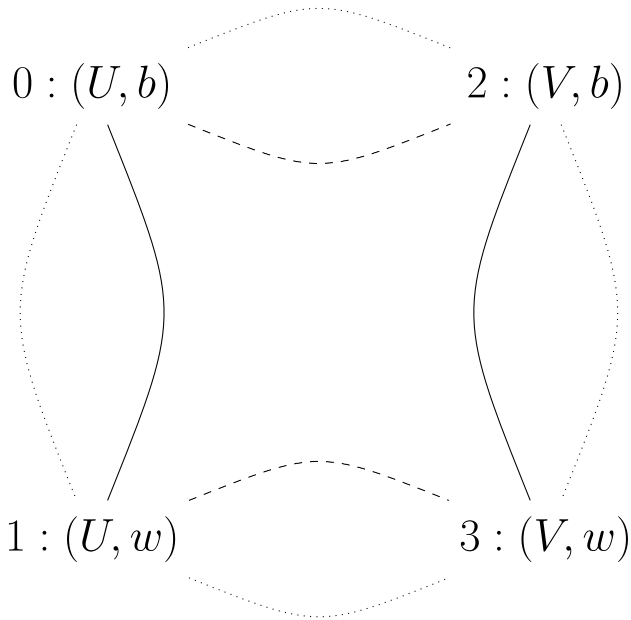

$$\{0 : \tfrac{1}{6}, 1 : \tfrac{1}{3}, 2 : \tfrac{1}{4}, 3 : \tfrac{1}{4}\}$$

Another Representation

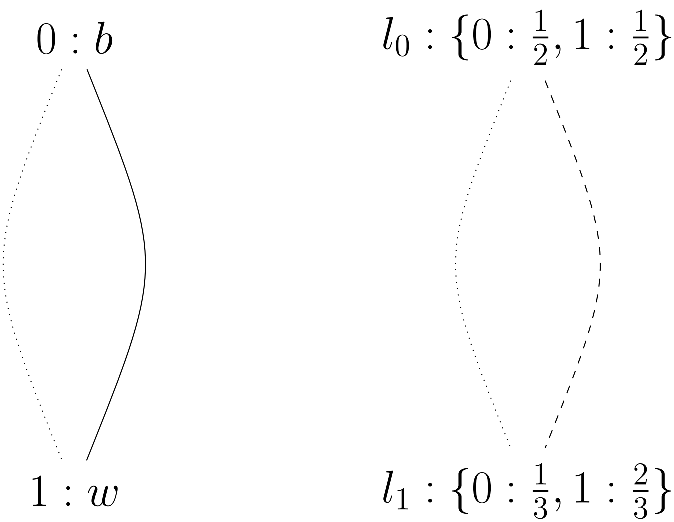

Lotteries over lotteries

Example from [Gne75]:

There are five urns with the following compositions: 2 urns with 2 white and 3 black balls each, 2 urns with 1 white and 4 black balls each, and one urn with 4 white balls and 1 black ball. A ball is chosen from one of the urns taken at random. It turns out to be white. What is the probability (after the experiment) that the ball was taken from the last urn?

Representation

$$0 : b$$

$$1 : w$$

$$l_0 : \{0 : \tfrac{3}{5}, 1 : \tfrac{2}{5}\}$$

$$l_1 : \{0 : \tfrac{4}{5}, 1 : \tfrac{1}{5}\}$$

$$l_1 : \{0 : \tfrac{1}{5}, 1 : \tfrac{4}{5}\}$$

$$L : \{l_0 : \tfrac{2}{5}, l_1 : \tfrac{2}{5}, l_2 : \tfrac{1}{5}\}$$

Another Representation

$$\begin{tikzcd}[arrows=dash]
 & l_0 : \{0 : \frac{3}{5}, 1 : \frac{2}{5}\} \arrow[d] \\
0 : b \arrow[dd] & l_1 : \{0 : \frac{3}{5}, 1 : \frac{2}{5}\} \arrow[d] \\
 & l_2 : \{0 : \frac{4}{5}, 1 : \frac{1}{5}\} \arrow[d] \\
1 : w & l_3 : \{0 : \frac{4}{5}, 1 : \frac{1}{5}\} \arrow[d] \\
 & l_4 : \{0 : \frac{1}{5}, 1 : \frac{4}{5}\}
\end{tikzcd}$$

Coin Tossing

Suppose Alice is tossing a coin while Bob is watching. Both know that the coin can either be fair or biased (say, with bias $\frac{2}{3}$ towards heads). Bob does not know which coin Alice is using, but Alice knows.

$$
\begin{array}{ll}
0 : H & l_0 : \{0 : \frac{1}{2}, 1 : \frac{1}{2}\} \\
 & \\
1 : T & l_1 : \{0 : \frac{2}{3}, 1 : \frac{1}{3}\}
\end{array}
$$

Standard Epistemic Models

A standard epistemic model 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 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.

$\mathcal{L}_0$ language of multi-agent epistemic logic:

$$\varphi ::= \top \mid p \mid \neg\varphi \mid \varphi \wedge \varphi \mid K_a\varphi$$

where p ranges over a set $\mathbf{P}$ of basic propositions and a ranges over a set of agents A.

Epistemic Probability Models

To change a standard epistemic model into an epistemic probability model, we assign to each agent an equivalence relation over a list of lotteries. This represents subjective probabilities.

An epistemic probability model is a tuple (W, V, R, L, I, E) where

- W, V, R are as above.
- L is a set of W-lotteries indexed by natural numbers (displayed as $\{l_0, l_1, \ldots\}$), with index set I, all lotteries having the same scale.
- E is a function that assigns an equivalence relation on I to each agent $a \in A$.

Epistemic Probability Language

$\mathcal{L}$ language of multi-agent epistemic probability logic:

$$\begin{aligned} \varphi &::= \top \mid p \mid l_j \mid \neg\varphi \mid \varphi \wedge \varphi \mid t_1 + \cdots + t_n \geq q \\ t &::= q \mid t \cdot P_a \varphi \end{aligned}$$

Abbreviations

- $\perp$, $\varphi_1 \vee \varphi_2$, $\varphi_1 \rightarrow \varphi_2$, $\varphi_1 \leftrightarrow \varphi_2$.
- $t < t'$ for $\neg t \geq t'$.
- $t > t'$ for $\neg t' \geq t$.
- $t \leq t'$ for $t' \geq t$.
- $t = t'$ for $t \geq t' \wedge t \leq t'$.
- $t \neq t'$ for $t > t' \vee t < t'$.
- $P_a(\varphi_1|\varphi_2) = t$ for $t \cdot P_a(\varphi_2) = P_a(\varphi_1 \wedge \varphi_2)$.

- l_j expresses that the current lottery index equals j.
- $P_a\varphi = q$ expresses that the probability of φ according to a equals q.
- $P_a(\varphi_1|\varphi_2) = q$ expresses that according to a, the probability of φ_1, conditional on φ_2, equals q.

Truth

Let $M = (W, V, R, L, I, E)$, let $w \in W$, let $i \in I$.

$$
\begin{array}{rll}
M, w, i \models \top & & \text{always} \\
M, w, i \models p & \text{iff} & p \in V(w) \\
M, w, i \models l_j & \text{iff} & i = j \\
M, w, i \models \neg\varphi & \text{iff} & \text{it is not the case that } M, w \models \varphi \\
M, w, i \models \varphi_1 \wedge \varphi_2 & \text{iff} & M, w, i \models \varphi_1 \text{ and } M, w, i \models \varphi_2 \\
M, w, i \models t_1 + \cdots + t_n \geq q & \text{iff} & [\![t_1]\!]_{w,i} + \cdots + [\![t_n]\!]_{w,i} \geq q
\end{array}
$$

Probability

$$\begin{aligned} [\![q]\!]_{w,i} &:= q \\ [\![t \cdot P_a\varphi]\!]_{w,i} &:= [\![t]\!]_{w,i} \times P_{a,w,i}(\varphi). \end{aligned}$$

P and D Functions

$$P_{a,w,i}(\varphi) = \frac{\sum\{D_{a,w,j}(\varphi) \mid iE_aj\}}{|\{j \mid iE_aj\}|}.$$

$P_{a,w,i}(\varphi)$ gives the average of the probabilities that a assigns to φ in w, for all lotteries that a confuses with l_i.

$$D_{a,w,i}(\varphi) = \frac{\sum\{l_i(u) \mid wR_au \text{ and } M, u, i \models \varphi\}}{\sum\{l_i(u) \mid wR_au\}}.$$

$D_{a,w,i}(\varphi)$ gives the probability that a assigns to φ in w, assuming that a knows l_i, i.e., assuming that a does not confuse l_i with any other lottery.

Note that $D_{a,w,i}(l_i) = 1$, and for all j with $j \neq i$, $D_{a,w,i}(l_j) = 0$.

Common Knowledge of Indifference Models

If $M = (W, V, R)$ is an epistemic model, then M^{indif} is the epistemic probability model (W, V, R', L, I, E) where

$$\begin{aligned}
W' &= W \\
V' &= V \\
R' &= R \\
L &= \{l_0\} \text{ where } l_0 = \lambda w \in W \mapsto 1 \\
I &= \{0\} \\
E &= \lambda a \in A \mapsto \{(0, 0)\}
\end{aligned}$$

Explanation: M^{indif} is the epistemic probability model that is the result of putting a uniform probability distribution on the worlds in M, and making this uniform probability distribution common knowledge.

Common Knowledge of Indifference about p, q

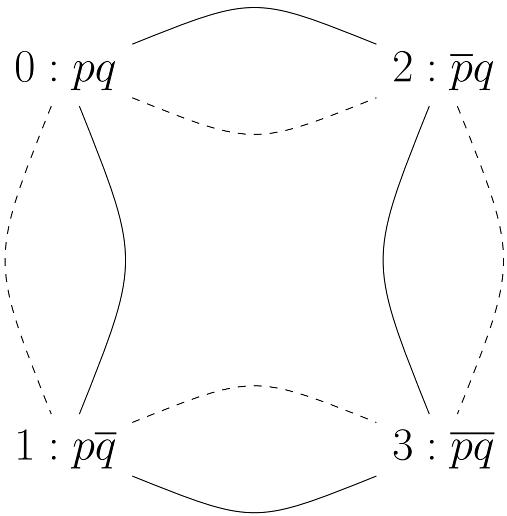

$$l_0 = \{0 : \tfrac{1}{4}, 1 : \tfrac{1}{4}, 2 : \tfrac{1}{4}, 3 : \tfrac{1}{4}\}$$

Erasing Probability Information From Models

If $M = (W, V, R, L, I, E)$ is an epistemic probability model, then we can map this to an epistemic model M° by putting $M^\circ = (W^\circ, V^\circ, R^\circ)$ with

- $W^\circ = \{(w, i) \mid w \in W, i \in I\}$
- $V^\circ = \lambda(w, i) \mapsto V(w)$
- R° is given by $(w, i)R_a^\circ(u, j)$ iff wR_au and iE_aj.

Note that if M is an epistemic probability model where the agents share a single lottery, then M° is the result of removing the lottery information.

Relation of Knowledge and Certainty

Define a translation $t : \mathcal{L}_0 \to \mathcal{L}$ from the language of multi-agent epistemic logic to the language of epistemic probability logic by means of:

$$\begin{aligned} t(p) &:= p \\ t(\neg\varphi) &:= \neg t(\varphi) \\ t(\varphi_1 \wedge \varphi_2) &:= t(\varphi_1) \wedge t(\varphi_2) \\ t(K_a\varphi) &:= P_a t(\varphi) = 1 \end{aligned}$$

This translates knowledge statements of $\mathcal{L}_0$ into certainty statements of $\mathcal{L}$, and allows us to prove the Certainty Theorem.

Certainty Theorem

Theorem 1 (Certainty) *For any epistemic probability model*

$$M = (W, V, R, L, I, E),$$

any world-index pair (w, i) *for* M, *any* $\varphi \in \mathcal{L}_0$*:*

$$M^\circ, (w, i) \models \varphi \textit{ iff } M, w, i \models t(\varphi).$$

Proof. Induction on the structure of φ. The only case you have to check is $K_a\varphi$. $\square$

This theorem motivates the following abbreviation for $\mathcal{L}$:

$$\text{Use } K_a\varphi \text{ for } P_a\varphi = 1.$$

This abbreviation reflects the equation of knowledge and certainty.

It follows immediately from the Certainty Theorem that the certainty operator $P_a\varphi = 1$ is an S5 operator.

Example: Uncertainty about q-bias

$$
\begin{tikzcd}
0 : pq \arrow[r, dash, dashed] \arrow[d, dash] & 2 : \overline{p}q \arrow[d, dash] \\
1 : p\overline{q} \arrow[r, dash, dashed] & 3 : \overline{pq}
\end{tikzcd}
\qquad
\begin{tikzcd}
l_0 = \{0 : \frac{1}{4}, 1 : \frac{1}{4}, 2 : \frac{1}{4}, 3 : \frac{1}{4}\} \arrow[d, dash] \\
l_1 = \{0 : \frac{1}{3}, 1 : \frac{1}{6}, 2 : \frac{1}{3}, 3 : \frac{1}{6}\}
\end{tikzcd}
$$

In this model, at world 0 and lottery l_0, the probability that a (represented by solid lines) assigns to p is 1, so $K_a p$ is true at $0, l_0$. $K_a q$ is false at $0, l_0$, for the probability that a assigns to q is less than 1. It equals the average of the probabilities that a assigns to q for the two lotteries.

Axioms

Propositional Logic Axioms

- All (instances of) tautologies of propositional logic are axioms.
- The Modus Ponens Rule: from $\vdash \varphi$ and $\vdash \varphi_1 \to \varphi_2$ conclude $\vdash \varphi_2$.

Probability Axioms

$$\begin{aligned}
(P_1) &\ \vdash\ P_a\top = 1 \\
(P_2) &\ \vdash\ P_a(\neg\varphi) = 1 - P_a\varphi \\
(P_3) &\ \vdash\ P_a(\varphi_1 \wedge \varphi_2) = P_a\varphi_1 * P_a\varphi_2
\end{aligned}$$

Derivable Principles

From (P_2), (P_3) we derive:

$$\vdash \ P_a(\varphi_1 \vee \varphi_2) = P_a\varphi_1 + P_a\varphi_2 - P_a\varphi_1 * P_a\varphi_2$$
$$\vdash \ P_a(\varphi_1 \to \varphi_2) = 1 + P_a\varphi_1 * P_a\varphi_2 - P_a\varphi_1$$

$$\begin{aligned} P_a(\varphi_1 \to \varphi_2) = 1 \ &\leftrightarrow \ 1 + P_a\varphi_1 * P_a\varphi_2 - P_a\varphi_1 = 1 \\ &\leftrightarrow \ P_a\varphi_1 * P_a\varphi_2 - P_a\varphi_1 = 0 \\ &\leftrightarrow \ P_a\varphi_1 * P_a\varphi_2 = P_a\varphi_1 \\ &\leftrightarrow \ P_a\varphi_1 = 0 \vee P_a\varphi_2 = 1 \end{aligned}$$

From this:

$$\vdash P_a\varphi_1 > 0 \wedge P_a(\varphi_1 \to \varphi_2) = 1 \to P_a\varphi_2 = 1 \qquad (*)$$

Formula $(^*)$ can be viewed as a probabilistic version of the K-axiom in epistemic logic.

Certainty Axioms

$$\begin{array}{lll}(C_1) & \vdash & P_a\varphi = 1 \rightarrow \varphi \\ (C_2) & \vdash & P_a\varphi \geq t \rightarrow P_a(P_a\varphi \geq t) = 1 \\ (C_3) & \vdash & P_a\varphi < t \rightarrow P_a(P_a\varphi < t) = 1\end{array}$$

Note that the following is derivable from (C_1) and (P_2):

$$\vdash P_a\varphi = 0 \rightarrow \neg\varphi.$$

Lottery Axiom

$$(L) \ \vdash \ P_a(l_{i1} \vee \cdots \vee l_{in}) = 1 \rightarrow P_a(p \wedge l_{i1}) + \cdots + P_a(p \wedge l_{in}) = n \cdot P_a p.$$

(L) expresses the definition of the probability of a basic fact as the average over its probabilities with respect to all accessible lotteries.

Probability Rule

(PR) If $\vdash \varphi_1 \rightarrow \varphi_2$ then $\vdash P_a\varphi_1 \leq P_a\varphi_2$.

From (PR) we derive:

- If $\vdash \varphi_1 \leftrightarrow \varphi_2$ then $\vdash P_a\varphi_1 = P_a\varphi_2$.

Also derivable is the necessitation rule for certainty:

- If $\vdash \varphi$ then $\vdash P_a\varphi = 1$.

Polynomial inequality axioms

- All instances of valid formulas about polynomial inequalities are axioms.

Cf. Halpern [Hal03, Section 7.7].

Commom Prior

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. If we want to impose this condition, we need a formula or set of formulas for it. In case both the number of agents and the number of atomic propositions are finite, we can express it in a single formula:

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

For infinite sets of agents or propositions we need an infinite number of formulas to express the fact that the agents have a common prior.

Completeness

This is a complete system for epistemic probability logic (I think).

Allowing Lotteries with Unknowns

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.

Model representing a coin with unknown bias:

$$\begin{tikzcd} 0 : p \arrow[d, dash] \\ 1 : \overline{p} \end{tikzcd}$$

$$l_0 : \{0 : B_p, 1 : 1 - B_p\}$$

A Paradox of John Maynard Keynes

Represent an urn with m white marbles and n black marbles as (m, n).

$$0 : (0, 2)$$

$$1 : (1, 1) \qquad\qquad l_0 = \{0 : \tfrac{1}{3}, 1 : \tfrac{1}{3}, 2 : \tfrac{1}{3}\}$$

$$2 : (2, 0)$$

Represent an urn as a stack of marbles:

$$\begin{tikzcd} 0 : WW \arrow[d, dash] \\ 1 : WB \arrow[d, dash] \\ 2 : BW \arrow[d, dash] \\ 3 : BB \end{tikzcd}$$

$$l_0 = \{0 : \tfrac{1}{4}, 1 : \tfrac{1}{4}, 2 : \tfrac{1}{4}, 3 : \tfrac{1}{4}\}$$

The ‘Paradox’

According to the first model, the probability that both marbles are black is $\frac{1}{3}$.

According to the second mode, the probability that both marbles are black is $\frac{1}{4}$.

How is this possible?

This is the question John Maynard Keynes posed in his book [Key63].

The Answer

The answer is that in the first model the principle of indifference was applied in the wrong way.

Compare also the puzzle about the probability that a family with two kids has two boys. This probability is $\frac{1}{4}$ rather than $\frac{1}{3}$, for the relevant cases are BB, BG, GB and GG. If we represent as $2B$, $B+G$ and $2G$, then we must bear in mind that the pattern $B+G$ is twice as likely as the two other patterns.

It is all a matter of representation. If we represent "consider an arbitrary urn with k marbles, either black or white" in terms of (m, n), then we have to add the lottery information that the pattern (m, n) gets lottery value $\binom{m+n}{m}$.

Updates: Public Announcement

If $M = (W, V, R, L, I, E)$ is an epistemic probability model, then $M^\varphi = (W', V', R', L', I', E')$ is the epistemic probability model given by

- $W' = \{w \in W \mid \text{ for some } j \in I \ : M, w, j \models \varphi\}$.
- V' is the restriction of V to W'.
- R' assigns to each agent a the relation R'_a that is the restriction of R_a to W'.
- L' is the set of lotteries from L with an index in $I' = \{j \in I \mid \text{ for some } w \in W \ : M, w, j \models \varphi\}$.
- E' assigns to each agent a the relation E'_a that is the restriction of E_a to I'.

A Language with Public Announcement

Let $\mathcal{L}^{\text{PA}}$ be the extension of $\mathcal{L}$ with public announcements, where $[\varphi_1]\varphi_2$ expresses that after the public announcement of φ_1, φ_2 holds.

More precisely:

$$M, w, i \models [\varphi_1]\varphi_2 \quad \text{iff} \quad M, w, i \models \varphi_1 \text{ implies } M^{\varphi_1}, w, i \models \varphi_2.$$

Calculus for epistemic probability logic with PA

Add the following principles for public announcement:

$$
\begin{array}{lll}
(PA_1) & \vdash & [\varphi]p \leftrightarrow (\varphi \to p) \\
(PA_2) & \vdash & [\varphi_1]\neg\varphi_2 \leftrightarrow \neg[\varphi_1]\varphi_2 \\
(PA_3) & \vdash & [\varphi_1](\varphi_2 \wedge \varphi_3) \leftrightarrow ([\varphi_1]\varphi_2 \wedge [\varphi_1]\varphi_3) \\
(PA_4) & \vdash & [\varphi_1](P_a\varphi_2 = q) \leftrightarrow (P_a([\varphi_1]\varphi_2 \mid \varphi_1) = q)
\end{array}
$$

Axiom (PA_4) was proposed by Johan van Benthem.

Plus the rule of announcement generalization:

From $\vdash \varphi_1$ derive $\vdash [\varphi_2]\varphi_1$.

Calculus for epistemic probability logic with PA

Add the following principles for public announcement:

$$(PA_1) \vdash [\varphi]p \leftrightarrow (\varphi \rightarrow p)$$
$$(PA_2) \vdash [\varphi_1]\neg\varphi_2 \leftrightarrow \neg[\varphi_1]\varphi_2$$
$$(PA_3) \vdash [\varphi_1](\varphi_2 \wedge \varphi_3) \leftrightarrow ([\varphi_1]\varphi_2 \wedge [\varphi_1]\varphi_3)$$
$$(PA_4) \vdash [\varphi_1](P_a\varphi_2 = q) \leftrightarrow (P_a([\varphi_1]\varphi_2 \mid \varphi_1) = q)$$

Axiom (PA_4) was proposed by Johan van Benthem.

Plus the rule of announcement generalization:

From $\vdash \varphi_1$ derive $\vdash [\varphi_2]\varphi_1$.

Completeness

If the above calculus of epistemic probability logic is complete, then adding the axioms and rule for public announcement gives a complete calculus for epistemic probability logic with public announcement.

Updates: Public Change

A $\mathbf{P}$ substitution is a finite list of bindings $p := \varphi$. This determines a substitution $\mathbf{P} \to \mathcal{L}$ in the usual way.

A public change is a substitution σ applied to the valuation at all worlds:

$$V^{M,\sigma,i}(w) = \{p \in \mathbf{P} \mid M, w, i \models \sigma(p)\}.$$

This depends on the lottery.

If we wish to avoid that dependence, then we have to restrict the formulas φ allowed in substitutions.

Updates: Non-determinate Change

In a probabilistic setting, a public change update can be non-determined, in the sense that different things might happen with certain probabilities.

To model this, we can represent an action as a lottery over substitutions.

Example:

$$\begin{tikzcd} & p := \top \\ \cdot \arrow[ur, "\frac{1}{2}"] \arrow[dr, "\frac{1}{2}"'] & \\ & p := \bot \end{tikzcd}$$

Use of Non-determinate change for creating coin flip models

$$\begin{tikzcd}
0 : p \arrow[d, dash, "l_0 : \{0 : \frac{1}{2}, 1 : \frac{1}{2}\}"] \\
1 : \overline{p}
\end{tikzcd}$$

A Puzzle of Lewis Carroll

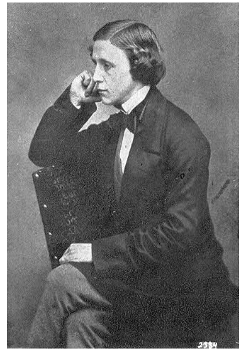

An urn contains a single marble, either white or black. Mr A puts another marble in the urn, a white one. The urn now contains two marbles. Next, Mrs B draws one of the two marbles from the urn. It turns out to be white. What is the probability that the other marble is also white? (Gardner [Gar81])

Solution with PRODEMO

Call the first white marble p and the second one q. Mrs B does not know whether she is drawing from $\neg p + q$ or from $p + q$.

Let's start with a model of complete ignorance about p, for two agents a, b:

```
m1 :: Pem Prp
m1 = initPM [a,b] [P 0]
```

```
*PRODEMO> m1
MO [a,b] [0,1] [(0,[p]),(1,[])]
[(a,[[0,1]]),(b,[[0,1]])] [0,1]
[[(0,1 % 2),(1,1 % 2)]] [(a,[[0]]),(b,[[0]])] [0]
```

First update

Result of telling a the value of p, while b does not learn this fact.

```
m2 :: Pem Prp
m2 = upd [P 0] m1 um1
```

This gives:

```
*PRODEMO> m2
MO [a,b] [0,1]
[(0,[p]),(1,[])]
[(a,[[0],[1]]),(b,[[0,1]])]
[0,1]
[[(0,1 % 2),(1,1 % 2)]]
[(a,[[0]]),(b,[[0]])]
[0]
```

Putting a second white marble in the urn.

A public change that makes q true:

```
m3 :: Pem Prp
m3 = upd_pc [P 0,Q 0] m2 [(Q 0,Top)]
```

The result:

```
*PRODEMO> m3
MO [a,b] [0,1] [(0,[p,q]),(1,[q])]
[(a,[[0],[1]]),(b,[[0,1]])]
[0,1]
[[(0,1 % 2),(1,1 % 2)]]
[(a,[[0]]),(b,[[0]])]
[0]
```

Removing either p or q from the bag

Nobody knows which of these two takes place. Note that removing p from the bag has as precondition that p is true, and similarly for q.

Result of updating with this:

```
m4 :: Pem Prp
m4 = upd [P 0,Q 0] m3 um2
```

Here is what this model looks like:

```
*PRODEMO> m4
MO [a,b] [0,1,2]
[(0,[q]),(1,[p]),(2,[])]
[(a,[[0,1],[2]]),(b,[[0,1,2]])] [0,1,2]
[[(0,1 % 3),(1,1 % 3),(2,1 % 3)]]
[(a,[[0]]),(b,[[0]])] [0]
```

What is the probability that the other marble is also white?

In our setting: what is the probability of $p \vee q$? It is different for a and b.

```
*PRODEMO> prob m4 a 0 0 p_or_q
1 % 1
*PRODEMO> prob m4 a 1 0 p_or_q
1 % 1
*PRODEMO> prob m4 a 2 0 p_or_q
0 % 1
*PRODEMO> prob m4 b 0 0 p_or_q
2 % 3
*PRODEMO> prob m4 b 1 0 p_or_q
2 % 3
*PRODEMO> prob m4 b 2 0 p_or_q
2 % 3
```

Work in Progress

- Implementation of epistemic probability model checking with PRODEMO.
- Completeness proofs for various axiom systems (e.g., for extensions of the language with (relativized) common knowledge operators).
- Investigation of ‘lottery bisimulations’.
- Extensions: Handling of Lottery Functionals
- Design and analysis of probabilistic protocol languages for epistemic probability updating.
- Connections with Bayesian learning, treating Baysian learning as increase of knowledge rather than change of belief.

- See `homepages.cwi.nl/~jve/software/prodemo/`

References

[BGK09] J. van Benthem, J. Gerbrandy, and B. Kooi. Dynamic update with probabilities. Studia Logica, 93:67–96, 2009.

[BS08] Alexandru Baltag and Sonja Smets. Probabilistic dynamic belief revision. Synthese, 165(2):179–202, 2008.

[FH94] R. Fagin and J.Y. Halpern. Reasoning about knowledge and probability. Journal of the ACM, pages 340–367, 1994.

[Gar81] Martin Gartner. Mathematical Circus. Vintage, 1981.

[Gie09] Gierasimszuk. Bridging learning theory and dynamic epistemic logic. Synthese, 169:371–374, 2009.

[Gne75] B.V Gnedenko. The Theory of Probability. Mir Publishers, Moscow, 1975.

[Hal03] J. Halpern. Reasoning About Uncertainty. MIT Press, 2003.

[Key63] J.M. Keynes. A Treatise on Probability. Macmillan, London, 1963.

[Koo03a] Barteld P. Kooi. Knowledge, Chance, and Change. PhD thesis, Groningen University, 2003.

[Koo03b] Barteld P. Kooi. Probabilistic dynamic epistemic logic. Journal of Logic, Language and Information, 12(4):381–408, 2003.

[Lap14] M. le Comte Laplace. Essai Philosophique sur les Probabilités. Courcier, Paris, 1814.

[vN51] John von Neumann. Various techniques used in connection with random digits. Technical Report 12: 36, National Bureau of Standards Applied Math Series, 1951.