Knowledge, Belief, Probabilities, Updates, Model Checking

Jan van Eijck
CWI & ILLC, Amsterdam

LoLaCo Guest Lecture, Nov 17, 2014

Abstract

The talk will present epistemic probability models with probabilistic updates, and will discuss an implementation that allows model checking the results of updates in a multi-agent setting.

I will also try to convey the attractions of functional programming.

How are Logic and Probability Theory Related?

- Logic = Reasoning about Certainty
- Probability Theory = Reasoning about Uncertainty
- Wikipedia defines Epistemic or Bayesian probability as "... an extension of propositional logic that enables reasoning with hypotheses, i.e., the propositions whose truth or falsity is uncertain."
- But logic has something to say, too, about reasoning under uncertainty: epistemic logic, doxastic logic, default logic, ...
- How are all these enterprises related?

The Usefulness of Probability Theory

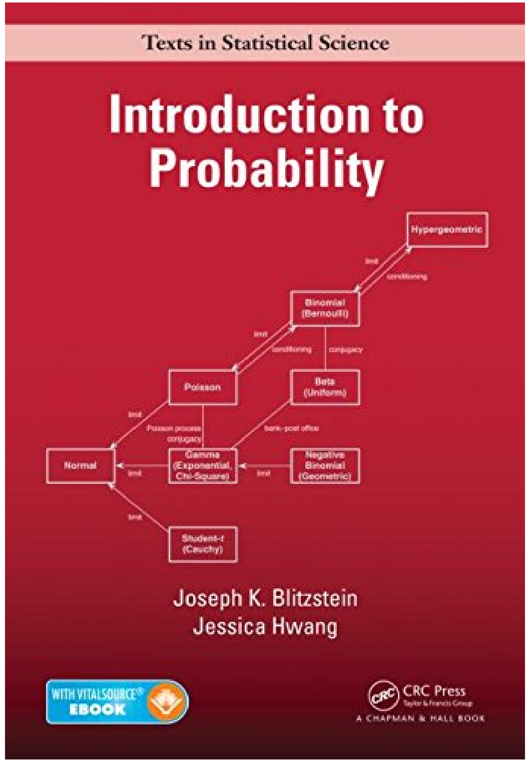

Quote from [BH14]:

> "Life: Life is uncertain, and probability is the logic of uncertainty. While it isn't practical to carry out a formal probability calculation for every decision made in life, thinking hard about probability can help us avert some common fallacies, shed light on coincidences, and make better predictions."

Prime example: probability of having disease D after positive test

You are from a population with a statistical chance of 1 in 100 of having disease D. The initial screening test for this has a false positive rate of 0.2 and a false negative rate of 0.1. You tested positive (T).

Should you believe you have disease D?

- You reason: if I test positive then, given that the test is quite reliable, the probability that I have D is quite high. So I believe that I have D.

- You use pen and paper and calculate:

$$P(D|T) = \frac{P(T|D)P(D)}{P(T)} = \frac{P(T|D)P(D)}{P(T|D)P(D) + P(T|\neg D)P(\neg D)}$$

  Filling in $P(T|D) = 0.9$, $P(D) = 0.01$, $P(\neg D) = 0.99$, $P(T|\neg D) = 0.2$ gives $P(D|T) = \frac{1}{23}$. You don't believe you have D but you agree to undergo further testing.

Analysis

“Now the discrepancy between 4% and 80 or 90% is no small matter, particularly if the consequence of an error involves either unnecessary surgery or (in the reverse case) leaving a cancer untreated. But decisions similar to these are constantly being made based upon ”intuitive feel” – i.e., without the benefit of paper and pen, let alone Bayesian networks (which are simpler to use than paper and pen!).” [KN11]

Amos Tversky and Daniel Kahneman [TK74].

Successes and Failures of Probabilistic Analysis

The German tank problem Given a list of serial numbers of tanks that were captured or destroyed, estimate the total number of tanks. Find an estimate of the number of tanks produced each month. The probabilistic analysis of this turned out to be vastly more reliable than the intelligence estimates.

Successes and Failures of Probabilistic Analysis

The German tank problem Given a list of serial numbers of tanks that were captured or destroyed, estimate the total number of tanks. Find an estimate of the number of tanks produced each month. The probabilistic analysis of this turned out to be vastly more reliable than the intelligence estimates.

Statistical estimate for tanks produced in August 1942: 327. Intelligence estimate: 1550. German records: 342.

People v. Collins Testimony of bystanders on a robbery committed in Los Angeles in 1968. Robbery was committed by a black male, with a beard and moustache, and a caucasian female with blonde hair tied in a ponytail. They had escaped in a yellow motor car.

The prosecutor invited the jury to calculate the probability that a pair who fitted the description were not the robbers, by multiplication. Black man with beard: 1 in 10, white woman with ponytail, 1 in 10, and so on.

The jury accepted this and the pair was convicted. (Fortunately, the conviction was overruled after appeal.)

See Wikipedia for a description, and [KN11] for a Bayesian analysis of the fallacy.

Lucia de B(erk) Probabilistic analysis of unexpected deaths in the Juliana Children's Hospital in The Hague. Same pattern of reasoning as in People v. Collins …

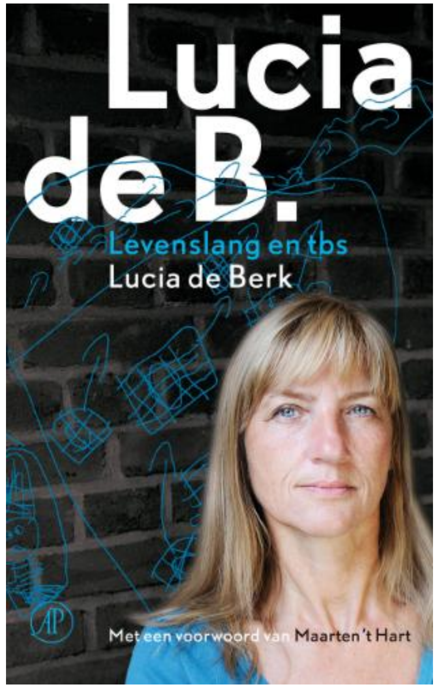

The Drake Equation Probabilistic argument for estimating the number of active, radio-communicative extraterrestial civilizations in the Milky Way galaxy (Frank Drake, 1961).

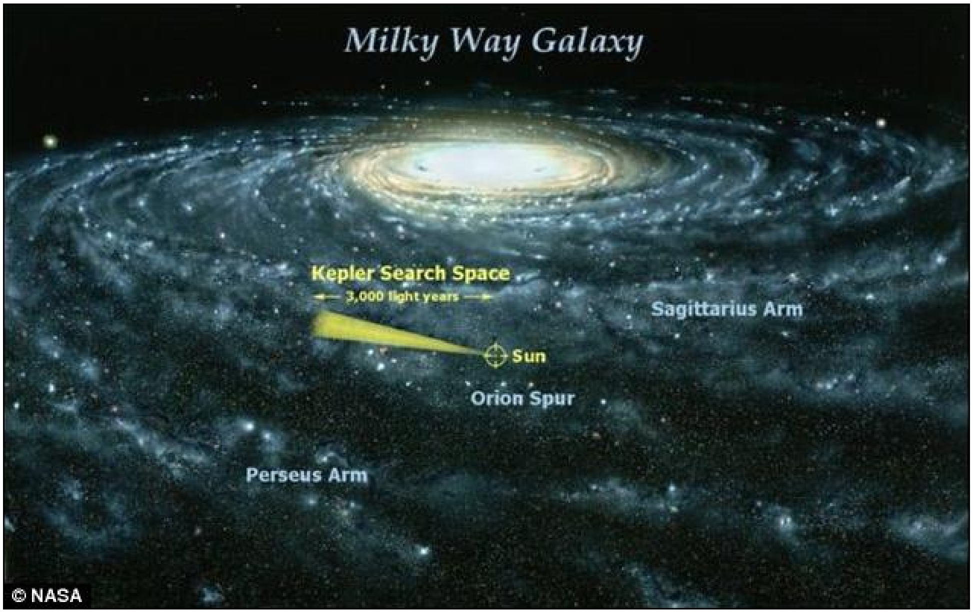

The Drake Equation Probabilistic argument for estimating the number of active, radio-communicative extraterrestial civilizations in the Milky Way galaxy (Frank Drake, 1961).

Also see: the Fermi question.

The Drake Equation Probabilistic argument for estimating the number of active, radio-communicative extraterrestial civilizations in the Milky Way galaxy (Frank Drake, 1961).

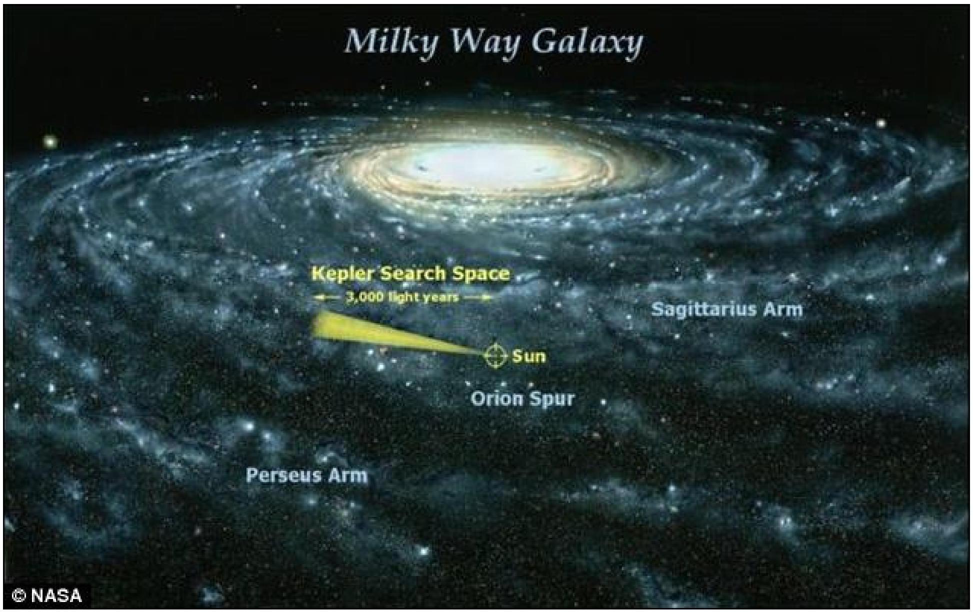

Also see: the Fermi question.

"Where is everybody?"

Belief and Probability

In the perspective of epistemic logic, our body of knowledge consists of true facts that we are certain about. But in the practice of everyday life and in the pursuit of science such absolute certainty is very rare.

- Can I safely cross this road?
- Should I bring my umbrella?
- Can I trust this bank?
- Is it safe to order from this cheap website?
- Can I trust this estimate of the mass of the planet Saturn?[1]

[1] Pierre Simon Laplace made a famous calculation of this, including an estimate for the uncertainty, using the astronomical data that were available to him in the early Nineteenth Century.

Decision Making under Uncertainty

An agent faces a choice between a finite number of possible courses of action $a_1, \ldots, a_n$. The agent is uncertain about the state of the world: she considers states $s_1, \ldots, s_m$ possible. There is a table of consequences c, with $c(s_i, a_j)$ giving the consequences of performing action a_j in state s_i.

Suppose there is a preference ordering R on the consequences, with cRc' expressing that either the agent is indifferent between c and c', or the agent strictly prefers c to c'. Assume R is transitive and reflexive. Then define cPc' as $cRc' \wedge \neg c'Rc$, so that cPc' expresses that the agent strictly prefers c to c'. The relation P is transitive and irreflexive.

A utility function $u : C \to \mathbb{R}$ represents R if u satisfies $u(c) \geq u(c')$ iff cRc'.

How can the agent pick the best available action?

The Von Neuman and Morgenstern Decision Tool

Von Neumann and Morgenstern [NM44] showed how to turn this into a tool for decision making if one adds a probability measure P on the state set. So assume $P(s_i) \geq 0$ and $\sum_{i=1}^{n} P(s_i) = 1$. Then a utility function u on the consequences induces a utility function U on the actions, by means of

$$U(a_j) = \sum_{i=1}^{n} P(s_i)u(s_i, a_j).$$

A rational agent who disposes of a utility function u representing her preferences and a probability measure on what she thinks is possible will perform the action a_j that maximizes $U(a_j)$ ...

This is the reason why expositions of probability theory often make strong claims about the applicability of their subject.

Varieties of Belief

- Betting belief (or: Bayesian belief) in φ: $P(\varphi) > P(\neg\varphi)$. Compare [ER14].
- Threshold belief in φ: $P(\varphi) > t$, for some specific t with $\frac{1}{2} \leq t < 1$. Also known as Lockean belief.
- Stable belief in φ: For all consistent ψ: $P(\varphi|\psi) > P(\neg\varphi|\psi)$ (Hannes Leitgreb [Lei10]).
- Strong belief in φ. Defined for plausibility models, e.g., locally connected well-preorders. An agent strongly believes in φ if φ is true in all most plausible accessible worlds. This yields a KD45 notion of belief (reflexive, euclidean, and serial).
- Subjective certainty belief in φ: $P(\varphi) = 1$. This is a notion used in epistemic game theory [Aum99].

The Lottery Puzzle

If Alice believes of each of the tickets 000001 through 111111 that they are not winning, then this situation is described by the following formula:

$$\bigwedge_{t=000001}^{111111} B_a \neg t.$$

If her beliefs are closed under conjunction, then this follows:

$$B_a \bigwedge_{t=000001}^{111111} \neg t.$$

But actually, she believes, of course, that one of the tickets is winning:

$$B_a \bigvee_{t=000001}^{111111} t.$$

This is a contradiction. The difficulty arises if we assume belief is closed under conjunction.

So it seems we need an operator B_i that does not satisfy (Dist).

$$B_i(\varphi \to \psi) \to B_i\varphi \to B_i\psi \qquad \text{(Dist-B)}$$

This means: B_i is not a normal modal operator.

Epistemic Neighbourhood Models

An **Epistemic Neighbourhood Model** $\mathcal{M}$ is a tuple

$$(W, R, N, V)$$

where

- W is a non-empty set of worlds.
- R is a function that assigns to every agent $i \in Ag$ an equivalence relation $\sim_i$ on W. We use $[w]_i$ for the $\sim_i$ class of w, i.e., for the set $\{v \in W \mid w \sim_i v\}$.
- N is a function that assigns to every agent $i \in Ag$ and world $w \in W$ a collection $N_i(w)$ of sets of worlds—each such set called a neighbourhood of w—subject to a set of conditions.
- V is a valuation function that assigns to every $w \in W$ a subset of *Prop*.

Conditions

(c) $\forall X \in N_i(w) : X \subseteq [w]_i$. This ensures that agent i does not believe any propositions $X \subseteq W$ that she knows to be false. If X contains a world in $w' \in W - [w]_i$ that the agent knows is not possible with respect to the actual world w, then she knows that X cannot be the case and hence she does not believe X.

(f) $\emptyset \notin N_i(w)$. This ensures that no logical falsehood is believed.

(n) $[w]_i \in N_i(w)$. This ensures that what is known is also believed.

(a) $\forall v \in [w]_i : N_i(v) = N_i(w)$. This ensures that if X is believed, then it is known that X is believed.

(m) $\forall X \subseteq Y \subseteq [w]_i$: if $X \in N_i(w)$, then $Y \in N_i(w)$. This says that belief is monotonic: if an agent believes X, then she believes all propositions $Y \supseteq X$ that follow from X.

(d) If $X \in N_i(w)$ then $[w]_i - X \notin N_i(w)$. This says that if i believes a proposition X then i does not believe the negation of that proposition.

Language

$$\varphi ::= \top \mid p \mid \neg\varphi \mid (\varphi \wedge \varphi) \mid K_i\varphi \mid B_i\varphi.$$

Semantics:

$$\mathcal{M}, w \models K_i\varphi \text{ iff } \quad \text{for all } v \in [w]_i : \mathcal{M}, v \models \varphi.$$

$$\mathcal{M}, w \models B_i\varphi \text{ iff } \quad \text{for some } X \in N_i(w), \text{for all } v \in X : \mathcal{M}, v \models \varphi.$$

Example

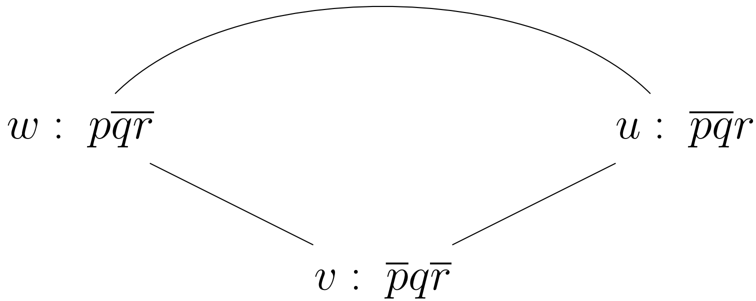

$$N(w) = N(v) = N(u) = \{\{w, v\}, \{v, u\}, \{w, u\}, \{w, v, u\}\}$$

Example

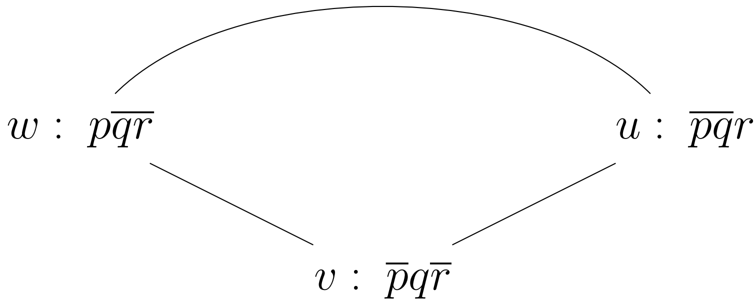

$$N(w) = N(v) = N(u) = \{\{w, v\}, \{v, u\}, \{w, u\}, \{w, v, u\}\}$$

In all worlds, $K(p \vee q \vee r)$ is true.

Example

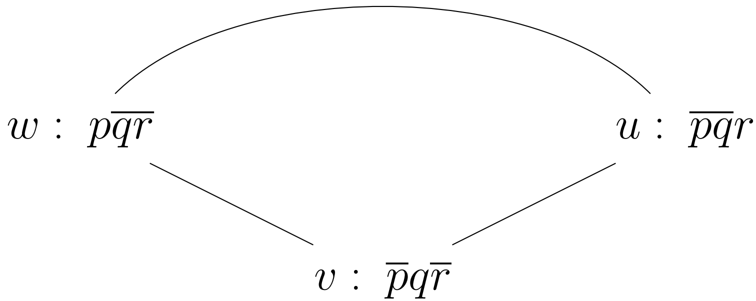

$$N(w) = N(v) = N(u) = \{\{w, v\}, \{v, u\}, \{w, u\}, \{w, v, u\}\}$$

In all worlds, $K(p \vee q \vee r)$ is true.

In all worlds $B\neg p$, $B\neg q$, $B\neg r$ are true.

Example

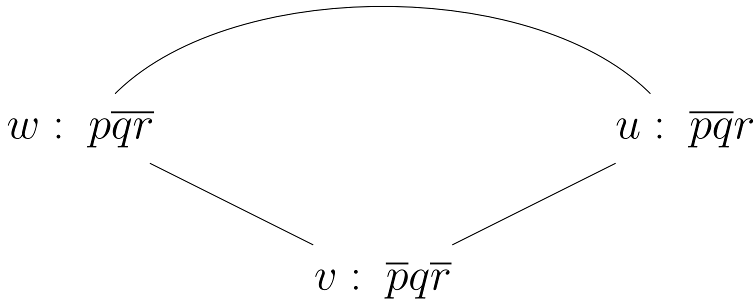

$$N(w) = N(v) = N(u) = \{\{w, v\}, \{v, u\}, \{w, u\}, \{w, v, u\}\}$$

In all worlds, $K(p \vee q \vee r)$ is true.

In all worlds $B\neg p$, $B\neg q$, $B\neg r$ are true.

In all worlds $B(\neg p \wedge \neg q)$, $B(\neg p \wedge \neg r)$, $B(\neg q \wedge \neg r)$ are false.

AXIOMS

(Taut)	All instances of propositional tautologies
(Dist-K)	$K_i(\varphi \to \psi) \to K_i\varphi \to K_i\psi$
(T)	$K_i\varphi \to \varphi$
(PI-K)	$K_i\varphi \to K_iK_i\varphi$
(NI-K)	$\neg K_i\varphi \to K_i\neg K_i\varphi$
(F)	$\neg B_i\bot$.
(PI-KB)	$B_i\varphi \to K_iB_i\varphi$
(NI-KB)	$\neg B_i\varphi \to K_i\neg B_i\varphi$
(KB)	$K_i\varphi \to B_i\varphi$
(M)	$K_i(\varphi \to \psi) \to B_i\varphi \to B_i\psi$
(D)	$B_i\varphi \to \neg B_i\neg\varphi$.

RULES

$$\frac{\varphi \to \psi \quad \varphi}{\psi}\ \text{(MP)} \qquad \frac{\varphi}{K_i \varphi}\ \text{(Nec-K)}$$

Further details: see [ER14] and [BvBvES14].

Knowledge, Certainty, Belief

One way to make the connection between epistemic logic and probability theory is by interpreting $K_i\varphi$ as "agent i assigns φ probability 1", or, "agent i is certain that φ is true."

Interpret $B_i\varphi$ as "agent i assigns φ higher probability than $\neg\varphi$", or, "agent i assigns φ probability greater than $\frac{1}{2}$."

As it turns out, the only thing we have to do is remove the neighbourhood function and add a weight function to an epistemic model.

If W is the set of worlds of an epistemic model, a weight function L assigns to every agent i a function $L_i : W \to \mathbb{Q}^+$, subject to the constraint that the sum of the L_i values over each epistemic partition cell of i is bounded.

Knowledge, Certainty, Belief

One way to make the connection between epistemic logic and probability theory is by interpreting $K_i\varphi$ as "agent i assigns φ probability 1", or, "agent i is certain that φ is true."

Interpret $B_i\varphi$ as "agent i assigns φ higher probability than $\neg\varphi$", or, "agent i assigns φ probability greater than $\frac{1}{2}$."

As it turns out, the only thing we have to do is remove the neighbourhood function and add a weight function to an epistemic model.

If W is the set of worlds of an epistemic model, a weight function L assigns to every agent i a function $L_i : W \to \mathbb{Q}^+$, subject to the constraint that the sum of the L_i values over each epistemic partition cell of i is bounded.

If $X \subseteq W$ then $L_i(X)$ is shorthand for $\sum_{x \in X} L_i(x)$.

Boundedness

The boundedness condition excludes cases where $[w]_i$ is infinite and each v in $[w]_i$ gets the same positive value c. It does not exclude infinite epistemic partition cells, however.

Example 1 *Let* $[w]_i = \mathbb{N}$*, and let* $L_i(n) = \frac{1}{2^n}$*. Then:*

$$L_i([w]_i) = \sum_{n \in \mathbb{N}} \frac{1}{2^n} = 2 < \infty.$$

Epistemic Weight Models

An **Epistemic Weight Model** $\mathcal{M}$ is a tuple (W, R, V, L), where

- W is a non-empty set of worlds.
- R is a function that assigns to every agent $i \in Ag$ an equivalence relation $\sim_i$ on W.
- V is a valuation function that assigns to every $w \in W$ a subset of *Prop*.
- L is a function that assigns to every agent $i \in Ag$ a weight L_i, where L_i is a function from W to $\mathbb{Q}^+$, the set of positive rationals, with the constraint that for each $w \in W$,

$$L_i([w]_i) < \infty.$$

Single Weight Models

An epistemic weight model $\mathcal{M} = (W, R, V, L)$ is single (or: a single weight model) if for all $i, j \in Ag$ it holds that $L_i = L_j$.

Example 2 *Take any epistemic model* $\mathcal{M} = (W, R, V)$ *with* W *finite. Let* L *be the function that maps* i *to the weight* $L_i = \lambda w.1$. *Then* (W, R, V, L) *is an epistemic single weight model.*

Example 3 *Two agents* i, j *consider betting on a horse race. Three horses take part in the race, and there are three possible outcomes:* a *for "*a *wins the race",* b *for "* b *wins the race", and* c *for "*c *wins the race." Neither agent knows which horse will win;* i *takes the winning chances to be* $3 : 2 : 1$, j *takes them to be* $1 : 2 : 1$. *In a picture:*

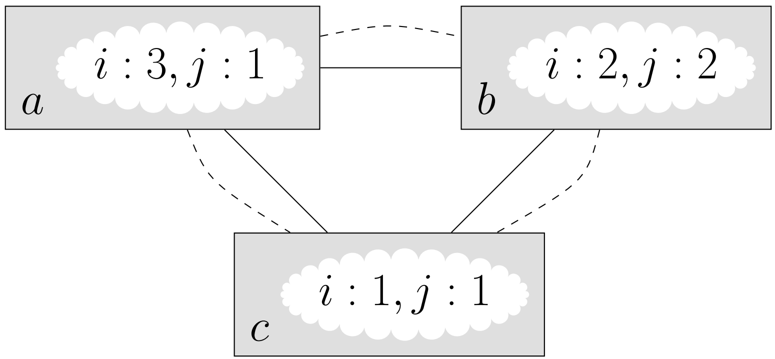

In all worlds, i *assigns probability* $\frac{1}{2}$ *to* a, $\frac{1}{3}$ *to* b *and* $\frac{1}{6}$ *to* c, *while* j *assigns probability* $\frac{1}{4}$ *to* a *and to* c, *and probability* $\frac{1}{2}$ *to* b.

Example 4 *Same situation as in example 3, but now agent j (dashed lines) considers c impossible.*

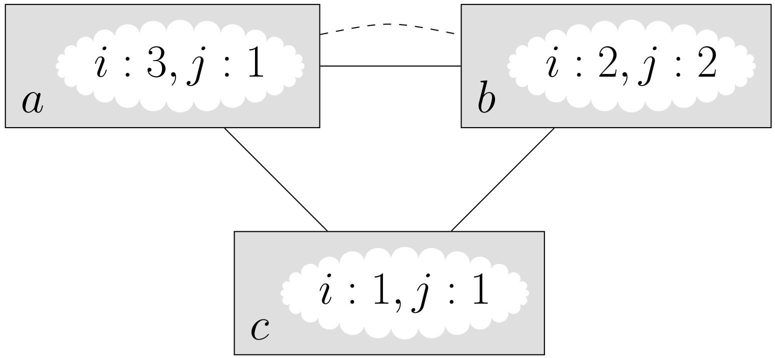

The probabilities assigned by i remain as before. The probabilities assigned by j have changed, as follows. In worlds a and b, j assigns probability $\frac{1}{3}$ to a and $\frac{2}{3}$ to b. In world c, j is sure of c.

Example 5 *Two agents i (solid lines) and j (dashed lines) are uncertain about the toss of a coin. i holds it for possible that the coin is fair f and that it is biased $\overline{f}$, with a bias $\frac{2}{3}$ for heads h. j can distinguish f from $\overline{f}$. The two agents share the same weight (so this is a single weight model), and the weight values are indicated as numbers in the picture.*

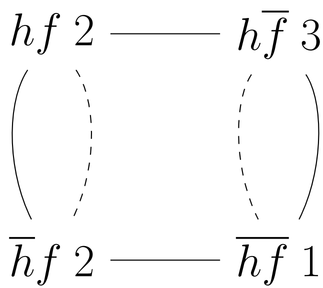

In world hf, i assigns probability $\frac{5}{8}$ to h and probability $\frac{1}{2}$ to f. In world hf, j assigns probability $\frac{1}{2}$ to h and probability 1 to f. In other words, j is certain that the coin is fair.

Epistemic Probability Language

Let i range over *Ag*, p over *Prop*, and q over $\mathbb{Q}$. Then the language of epistemic probability logic is given by:

$$\varphi ::= \top \mid p \mid \neg\varphi \mid (\varphi \wedge \varphi) \mid t_i \geq 0 \mid t_i = 0$$

$$t_i ::= q \mid q \cdot P_i\varphi \mid t_i + t_i \text{ where all indices } i \text{ are the same.}$$

Truth for Epistemic Probability Logic

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

$$\begin{array}{rcl}
\mathcal{M}, w \models \top & & \text{always} \\
\mathcal{M}, w \models p & \text{iff} & p \in V(w) \\
\mathcal{M}, w \models \neg\varphi & \text{iff} & \text{it is not the case that } \mathcal{M}, w \models \varphi \\
\mathcal{M}, w \models \varphi_1 \wedge \varphi_2 & \text{iff} & \mathcal{M}, w \models \varphi_1 \text{ and } \mathcal{M}, w \models \varphi_2 \\
\mathcal{M}, w \models t_i \geq 0 & \text{iff} & [\![t_i]\!]_w^{\mathcal{M}} \geq 0 \\
\mathcal{M}, w \models t_i = 0 & \text{iff} & [\![t_i]\!]_w^{\mathcal{M}} = 0.
\end{array}$$

$$\begin{array}{rcl}
[\![q]\!]_w^{\mathcal{M}} & := & q \\
[\![q \cdot P_i\varphi]\!]_w^{\mathcal{M}} & := & q \times P_{i,w}^{\mathcal{M}}(\varphi) \\
[\![t_i + t_i']\!]_w^{\mathcal{M}} & := & [\![t_i]\!]_w^{\mathcal{M}} + [\![t_i']\!]_w^{\mathcal{M}}
\end{array}$$

$$P_{i,w}^{\mathcal{M}}(\varphi) = \frac{L_i(\{u \in [w]_i \mid \mathcal{M}, u \models \varphi\})}{L_i([w]_i)}.$$

Example 6 *A normalized model for the horse racing situation from Example 3 is given in the picture:*

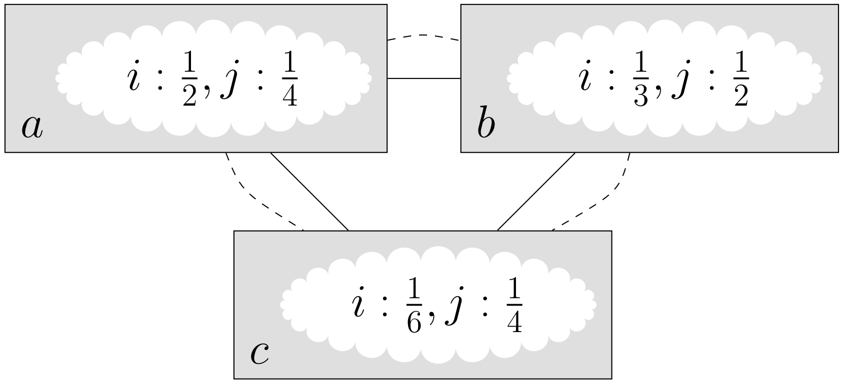

Example 7 *[Continued from Example 5] The model from Example 5 is an epistemic weight model where the two agents share the same weight. It is also possible to give each agent its own weight, and to normalize the weight functions using the epistemic accessibilities.*

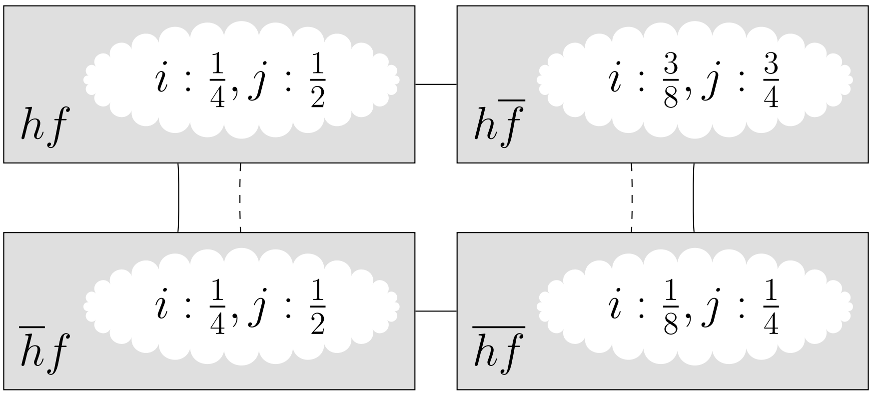

Fact 1 *Formulas of epistemic probability logic are invariant for (the appropriate notion of) bisimulation [ES14].*

Fact 2 *On epistemic weight models with finite epistemic partition cells for every agent, invariance for formulas of epistemic probability logic implies bisimilarity [ES14].*

Fact 3 *A sound and complete for the language of epistemic probability logic, interpreted in epistemic probability models, is given in [ES14].*

AXIOMS

(Taut)	All instances of propositional tautologies
(Linear)	All instances of valid formulas about linear inequalities
(ProbNonNeg)	$P_i\varphi \geq 0$
(ProbTrue)	$P_i\top = 1$
(ProbAdd)	$P_i(\varphi_1 \wedge \varphi_2) + P_i(\varphi_1 \wedge \neg\varphi_2) = P_i\varphi_1$
(ProbProbGeq)	$t_i \geq 0 \rightarrow P_i(t_i \geq 0) = 1$
(ProbProbEq)	$t_i = 0 \rightarrow P_i(t_i = 0) = 1$
(ProbT)	$P_i\varphi = 1 \rightarrow \varphi$

RULES

$$\frac{\varphi \rightarrow \psi \quad \varphi}{\psi}\ \text{(MP)} \qquad \frac{\varphi_1 \leftrightarrow \varphi_2}{P_i\varphi_1 = P_i\varphi_2}\ \text{(ProbRule)}$$

From Weight Models to Neighbourhood Models

If $\mathcal{M} = (W, R, V, L)$ is an epistemic weight model, then $\mathcal{M}^\bullet$ is the tuple (W, R, V, N) given by replacing the weight function by a function N, where N is defined as follows, for $i \in Ag$, $w \in W$.

$$N_i(w) = \{X \subseteq [w]_i \mid L_i(X) > L_i([w]_i - X)\}.$$

Fact 4 *For any epistemic weight model $\mathcal{M}$ it holds that $\mathcal{M}^\bullet$ is a neighbourhood model.*

Fact 5 *The calculus of epistemic-doxastic neighbourhood logic is sound for interpretation in epistemic probability models. Probabilistic beliefs are neighbourhoods.*

Translating Knowledge and Belief

If φ is a formula of the language of epistemic/doxastic logic, then $\varphi^\bullet$ is the formula of the language of epistemic probability logic given by the following instructions:

$$\begin{aligned}
\top^\bullet &= \top \\
p^\bullet &= p \\
(\neg\varphi)^\bullet &= \neg\varphi^\bullet \\
(\varphi_1 \wedge \varphi_2)^\bullet &= \varphi_1^\bullet \wedge \varphi_2^\bullet \\
(K_i\varphi)^\bullet &= P_i(\varphi^\bullet) = 1 \\
(B_i\varphi)^\bullet &= P_i(\varphi^\bullet) > P_i(\neg\varphi^\bullet).
\end{aligned}$$

Theorem 6 *For all formulas of epistemic/doxastic logic φ, for all epistemic weight models $\mathcal{M}$, for all worlds w of $\mathcal{M}$:*

$$\mathcal{M}^\bullet, w \models \varphi \textit{ iff } \mathcal{M}, w \models \varphi^\bullet.$$

Theorem 7 *Let* $\vdash$ *denote derivability in the calculus of EDNL. Let* $\vdash'$ *denote derivability in the calculus of EPL. Then* $\vdash \varphi$ *implies* $\vdash' \varphi^\bullet$.

Implementation

Building epistemic models from partitions . . .

```
type Erel a = [[a]]
```

```
data Agent = Ag Int

a,b,c,d,e :: Agent
a = Ag 0; b = Ag 1; c = Ag 2; d = Ag 3; e = Ag 4

data Prp = P Int | Q Int | R Int | S Int
```

Epistemic models

```
data EpistM state = Mo
                    [state]
                    [Agent]
                    [(state,[Prp])]
                    [(Agent,Erel state)]
                    [state]
```

```haskell
example1 :: EpistM Int
example1 = Mo
  [0..3]
  [a,b,c]
  []
  [(a,[[0],[1],[2],[3]]),
   (b,[[0],[1],[2],[3]]),(c,[[0..3]])]
  [1]
```

Epistemic Formulas

```
data Frm a = Tp
           | Info a
           | Prp Prp
           | N (Frm a)
           | C [Frm a]
           | D [Frm a]
           | Kn Agent (Frm a)
```

Truth Definition

...

```
isTrueAt :: Ord state =>
         EpistM state -> state -> Frm state -> Bool
isTrueAt m w Tp = True
isTrueAt m w (Info x) = w == x
isTrueAt
  m@(Mo worlds agents val acc points) w (Prp p) =
  let props = apply val w
  in elem p props
isTrueAt m w (N f) = not (isTrueAt m w f)
isTrueAt m w (C fs) = and (map (isTrueAt m w) fs)
isTrueAt m w (D fs) = or  (map (isTrueAt m w) fs)
isTrueAt
 m@(Mo worlds agents val acc points) w (Kn ag f) =
  let
    r = rel ag m
    b = bl r w
  in
    and (map (flip (isTrueAt m) f) b)
```

Public Announcement

```
upd_pa :: Ord state =>
       EpistM state -> Frm state -> EpistM state
upd_pa m@(Mo states agents val rels actual) f =
 (Mo sts' agents val' rels' actual') where
   sts'    = [ s | s <- states, isTrueAt m s f ]
   val'    = [ (s, ps) | (s,ps) <- val,
                             s `elem` sts']
   rels'   = [(ag,restrict sts' r) |
                                  (ag,r) <- rels ]
   actual'= [ s | s <- actual, s `elem` sts' ]

upds_pa ::  Ord state =>
         EpistM state -> [Frm state] -> EpistM state
upds_pa = foldl upd_pa
```

Example: Sum and Product (Hans Freudenthal)

A says to S and P: I have chosen two integers x, y such that $1 < x < y$ and $x + y \leq 100$. In a moment, I will inform S only of $s = x + y$, and P only of $p = xy$. These announcements remain private. You are required to determine the pair (x, y). He acts as said. The following conversation now takes place:

1. P says: "I do not know the pair."
2. S says: "I knew you didn't."
3. P says: "I now know it."
4. S says: "I now also know it."

Determine the pair (x, y).

A model checking solution with DEMO [vE05, vE07] (based on a DEMO program written by Ji Ruan) was presented in [DRV05]. An optimized version of that solution is in [vE13].

The list of candidate pairs:

```
pairs :: [(Int,Int)]
pairs  = [ (x,y) |  x <- [2..100], y <- [2..100],
                    x < y, x+y <= 100 ]
```

The solution:

```
solution = upds_pa msnp
     [k_a_statement_1e,statement_2e,statement_3e]
```

This is checked in a matter of seconds:

```
*DEMO_S5> solution
Mo [(4,13)] [a,b] [(a,[[(4,13)]]),(b,[[(4,13)]])]
   [(4,13)]
```

Extending This With Weights

```
data EpistWM state = WMo
                     [state]
                     [Agent]
                     [(state,[Prp])]
                     [(Agent,Erel state)]
                     [(Agent,[(state,Rational)])]
                     [state]
```

- Representation of probabibility information by means of weight functions was designed with implementation of model checking in mind.
- Just extend the epistemic models with a weight table for each agent.
- Implementations of model checkers for these logics can be found in [Eij13] and in [San14] ...
- The implementations can deal with Monty Hall style puzzles, urn puzzles, Bayesian updating by drawing from urns or tossing (possibly biased) coins, and 'paradoxes' such as the puzzle of the three prisoners (below).
- Efficiency was not a goal, but these implementation can be made very efficient with a little effort.

Aside: The Puzzle of the Three Prisoners

Alice, Bob and Carol are in prison. It is known that two of them will be shot, the other freed. The warden knows what is going to happen, so Alice asks him to reveal the name of one other than herself who will be shot, explaining to him that since there must be at least one, this will not reveal any new information. The warden agrees and says that Bob will be shot. Alice is cheered up a little by this, for she concludes that her chance of surviving has now improved from $\frac{1}{3}$ to $\frac{1}{2}$. Is this correct? How does this agree with the intuition that the warden has not revealed new information?

Many sources, e.g. [Jef04].

How to Move on From Here

- Combine EPL with network information for the agents, where the network is given by a relation, and where links starting from an agent can be added ("start following") and deleted ("stop following, unfollow"). Interpret announcements as group messages to all followers. See [RT11] and current work by Jerry Seligman and Thomas Agotnes. But: this can all be done with epistemic PDL with a binary follow relation F added.

- Further analysis of the connection between neighbourhood logics and probabilistic logics [ER14]. This is also connected to work of Wes Holliday and Thomas Icard.

- Add bias variables X for the representation of unknown biases. Collaboration in progress with Joshua Sack.

- Work with the epistemic PDL version of the probabilistic logic, as an extension of LCC from [BvEK06]. This gives us common knowledge, and a nice axiomatisation by means of epistemic program transformation [Ach14].
- Achieve better efficiency, by using methods proposed by Kaile Su.
- Towards analysis of real-life protocols. Compare the use of epistemic model checking by Malvin Gattinger [Gat13, Gat14b, Gat14a].
- Consider weak weight models, where the weight functions assign pairs of values (x, y), with x giving the lower probability L and $x + y$ the upper probability U. Belief of i in φ is now modelled as $L_i(\varphi) > H_i(\neg\varphi)$. This connects up to weak Bayesianism and imprecise probability theory [Wal91].
- Consolidate what we know about the topic in a state-of-the-art

textbook [BvBvES14].

References

[Ach14] Andreea Christina Achimescu. Games and logics for informational cascades. Master's thesis, ILLC, Amsterdam, February 2014.

[Aum99] R.J. Aumann. Interactive epistemology I: Knowledge. International Journal of Game Theory, 28(3):263–300, 1999.

[BH14] Joseph K. Blitzstein and Jessica Hwang. Introduction to Probability. CRC Press, 2014.

[BvBvES14] Alexandru Baltag, Johan van Benthem, Jan van Eijck, and Sonja Smets. Reasoning about communication and action. Book Manuscript, ILLC, 2014.

[BvEK06] J. van Benthem, J. van Eijck, and B. Kooi. Logics of communication and change. Information and Computation, 204(11):1620–1662, 2006.

[DRV05] Hans van Ditmarsch, Ji Ruan, and Rineke Verbrugge. Model checking sum and product. In Shichao Zhang and Ray Jarvis, editors, AI 2005: Advances in Artificial Intelligence: 18th Australian Joint Conference on Artificial Intelligence, volume 3809 of Lecture Notes in Computer Science, pages 790–795. Springer-Verlag GmbH, 2005.

[Eij13] Jan van Eijck. Learning about probability. available from homepages.cwi.nl:~/jve/software/prodemo, 2013.

[ER14] Jan van Eijck and Bryan Renne. Belief as willingness to bet. Manuscript, ILLC, Amsterdam, 2014.

[ES14] Jan van Eijck and François Schwarzentruber. Epistemic probability logic simplified. In Advances in Modal Logic, 2014.

[Gat13] Malvin Gattinger. Epistemic crypto logic — functional programming and model checking of cryptographic protocols. Technical report, ILLC, Amsterdam, 2013. Exam paper for the course 'Functional Specification of Algorithms'.

[Gat14a] Malvin Gattinger. Dynamic epistemic logic for guessing games and cryptographic protocols. Master's thesis, ILLC, University of Amsterdam, June 2014.

[Gat14b] Malvin Gattinger. More epistemic crypto logic. Technical report, ILLC, Amsterdam, February 2014.

[Jef04] Richard Jeffrey. Subjective Probability — The Real Thing. Cambridge University Press, 2004.

[KN11] Kevin B. Korb and Ann E. Nicholson. Bayesian Artificial Intelligence — Second Edition. Chapman and Hall/CRC, 2011.

[Lei10] Hannes Leitgreb. Reducing belief simpliciter to degrees of belief. Technical report, Ludwig Maximilian University of Munich, 2010.

[NM44] John von Neumann and Oscar Morgenstern. Theory of Games and Economic Behavior. Princeton University Press, 1944.

[RT11] Ji Ruan and Michael Thielscher. A logic for knowledge flow in social networks. In D. Wang and M. Reynolds, editors, AI 2011, number 7106 in LNAI, pages 511–520. Springer, 2011.

[San14] Thomas Santoli. Haskell project epistemic logic. Technical report, ILLC, Summer 2014.

[TK74] A. Tversky and D. Kahneman. Judgment under uncertainty: Heuristics and biases. Science, 185:1124–1131, 1974.

[vE05] Jan van Eijck. DEMO program and documentation, 2005. Available from `http://www.cwi.nl/~jve/demo/`.

[vE07] Jan van Eijck. DEMO — a demo of epistemic modelling. In Johan van Benthem, Dov Gabbay, and Benedikt Löwe, editors, Interactive Logic — Proceedings of the 7th Augustus de Morgan Workshop, number 1 in Texts in Logic and Games, pages 305–363. Amsterdam University Press, 2007.

[vE13] Jan van Eijck. DEMO-S5. Technical report, CWI, Amsterdam, 2013.

[Wal91] Peter Walley. Statistical Reasoning with Imprecise Probabilities. Chapman and Hall, London, 1991.