Update, Probability, Knowledge and Belief

Jan van Eijck
CWI & ILLC, Amsterdam

December 1, 2015

Abstract

The paper considers two kinds of models for logics of knowledge and belief, neighbourhood models and epistemic weight models, and traces connections. Epistemic weight models combine knowledge and probability by using epistemic accessibility relations and weights to define subjective probabilities. We present a new Probability Comparison Calculus that is sound and complete for epistemic weight models. This is a further simplification of the calculus for probabilistic epistemic weight models that was presented in AIML 2014.

The paper gives a definition of generic epistemic probabilistic update, again a simplification of earlier proposals, with examples of how this is used in epistemic probabilistic model checking. This application illustrates, among other things, that update by public announcement and update by Bayesian conditioning are two sides of the same coin. We end with a proposal for capturing the distinction between risk and uncertainty in epistemic weight models, and we put the question of the axiomatisation of the logic of imprecise weight models on the agenda.

Probability and 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.
> Tr: When concerned with things that are only likely true, the difference in how informed every man is about them is one of the principal causes of the diversity of opinions about the same objects.
>
> [Lap14]

A Bayesian learner is an agent who uses new information to update a subjective probability distribution that somehow captures what she knows or believes about the world. In a multi-agent setting, various learners could receive differents pieces of information, and, if we are to believe the Count of Laplace, these differences explain to a large extent the differences in viewpoints on the world that people have.

If multi-agent logics of knowledge of belief are extended with update procedures that implement the processing of new information, then this also gives a perspective on learning in a multi-agent setting. So it is natural to combine probability theory and the update perspective on knowledge and belief from dynamic modal logic.

Indeed, there exists already a modest tradition in combining DEL (Dynamic Epistemic Logic) and probability theory. The first combinations are in Kooi's thesis [Koo03], in Van Benthem's [Ben03], and in the combined effort of Van Benthem cs in [BGK09]. Inspiration for this goes back to work of Fagin and Halpern in the 1990s [FHM90a]. A simplified combined system of DEL and probability theory is presented in [ES14]. Further simplication of the base logic is provided in [DR15]. A full blown probabilistic logic of communication and change, in the spirit of [BvEK06], is given in [Ach14].

A natural notion of belief that turns up in a setting of probabilistic dynamic epistemic logic is betting belief (or: Bayesian belief) in φ: $P(\varphi) > P(\neg\varphi)$. Van Eijck & Renne [ER14] give a logical calculus for this that we will briefly review below. A variation on this is threshold belief in φ: $P(\varphi) > t$, for some specific t with $\frac{1}{2} \leq t < 1$. This is also known as Lockean belief. John Locke suggests in his work that a person's belief in a proposition φ is somehow connected to that person's confidence in φ. This confidence should then be connected in turn to the evidence that the person has for φ. If it makes sense to talk about degree of belief at all, then subjective probability is one way of making this precise [Fol09].

A logic with KD45 belief and an explicit belief comparison operator is presented in [JG13]. See [Nar07] for an overview of the extensive literature on belief comparison operators. Related to neighbourhood models for belief are the evidence models proposed in [BFDP14].

Next, there is stable belief in φ: For all consistent ψ: $P(\varphi|\psi) > P(\neg\varphi|\psi)$ (Leitgeb [Lei10]). This is still defined in terms of probability. The notion of strong belief is not; strong belief in φ is 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). Baltag & Smets [BS06, BS08] give details.

Finally, there is subjective certainty belief in φ: $P(\varphi) = 1$. This is used in epistemic game theory (Aumann [Aum99]).

To briefly compare these notions of belief, consider the well known lottery puzzle or lottery paradox, first considered in [Kyb61].

The Lottery Puzzle

If Alice believes of each of the tickets 000001 through 100000 that they are not winning, then this situation is described by the following formula:

$$\bigwedge_{t=000001}^{100000} B_a \neg t.$$

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

$$B_a \bigwedge_{t=000001}^{100000} \neg t.$$

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

$$B_a \bigvee_{t=000001}^{100000} t.$$

This is a contradiction.

Given that this is an argument with two premisses and a conclusion, there are precisely three ways to get around it.

1. Deny that Alice believes that her ticket is not winning.
2. Block the inference from $\bigwedge_{t=000001}^{100000} B_a \neg t$ to $B_a \bigwedge_{t=000001}^{100000} \neg t$.
3. Deny that Alice believes that there is a winning ticket.

It all depends on what we mean by "Alice believes that φ". Betting belief would accept that Alice believes that her ticket is not winning, but would block the inference from a conjunction of beliefs to belief in a conjunction. Lockean belief would maybe reject that Alice believes that her ticket is not winning, depending on the treshold, but would also reject the inference from a a conjunction of beliefs to belief in a conjunction. Stable belief would reject the claim that Alice believes that her ticket is not winning, for it is conceivable that Alice has to update her belief with the information that she has won the prize. Under all of these notions of belief, and indeed under most reasonable notions that one can think up, it is accepted that Alice believes that there is a winning ticket.

Advantage of (1) is that there is no need to sacrifice closure of belief under conjunction. A disadvantage of (1) is that it imposes a rather severe restriction of what counts as belief. Proponents of (1): many philosophers, easily recognizable from the fact that they call the lottery puzzle *the lottery paradox*.

Advantage of (2): sacrifice closure of belief under conjunction is maybe not so bad after all. Lots of nice logical properties remain (see below). Another advantage of (2) is that there is no need to artificially restrict what counts as belief. Proponents of (2) are subjective probabilists like Jeffrey [Jef04] and decision theorists like Kyburg [Kyb61].

In this paper, we will pursue the second way out. How can we drop the closure of belief under conjunction? For that 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. A semantics in terms of standard Kripke models will not work. What we need is neighbourhood semantics [Che80, Ch. 8] See also [Zve10], [HKP09], and [BFDP14].

Epistemic Neighbourhood Models

Definition 1 (Epistemic Neighbourhood Models) *An* **Epistemic Neighbourhood Model** $\mathcal{M}$ *is a tuple*

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

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* $\{v \in W \mid w \sim_i v\}$.
- V *is a valuation function that assigns to every* $w \in W$ *a subset of Prop.*
- 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.

Conditions are:

(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.

(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.

Extra conditions (corresponding to additional properties of epistemic neighbourhood models) are:

(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. This ensures that beliefs are consistent. If one wants to model Lockean belief with a threshold below one half, this should be blocked.

(sc) $\forall X, Y \subseteq [w]_a$: if $[w]_a - X \notin N_a(w)$ and $X \subsetneq Y$, then $Y \in N_a(w)$. If the agent does not believe the complement $[w]_a - X$, then she must believe any strictly weaker Y implied by X. Property (sc) is a form of "strong commitment": if the agent does not believe the complement $\overline{X}$, then she must believe any strictly weaker Y implied by X. This is a principle of strong commitment that should be blocked if one want to allow the possibility that some pieces of information has zero weight.

The language of epistemic doxastic logic (ED) has operators for knowledge and belief.

Definition 2 (ED Language) *Let p range over a set of basic propositions P and i over a finite set of agents A.*

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

Definition 3 (Truth in Neighbourhood Models) *Key clauses are*

$$\begin{array}{lll} \mathcal{M}, w \models K_i\varphi & \textit{iff} & \textit{for all } v \in [w]_i : \mathcal{M}, v \models \varphi. \\ \mathcal{M}, w \models B_i\varphi & \textit{iff} & \textit{for some } X \in N_i(w), \textit{for all } v \in X : \mathcal{M}, v \models \varphi. \end{array}$$

Our first example illustrates that neighbourhood belief is not closed under conjunction.

Example 1

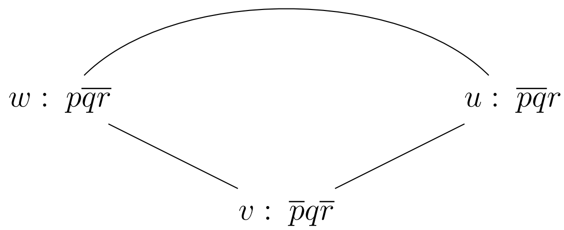

$$N(w) = N(v) = N(u) = \{\{w, v\}, \{v, u\}, \{w, u\}, \{w, v, u\}\}$$

- *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.*

Example 1 illustrates that the lottery puzzle is "solved" in neighbourhood models for belief by non-closure of belief under conjunction.

The following calculus for epistemic doxastic neighbourhood logic is complete for epistemic doxastic neighbourhood models (see [ER14] and [BvBvES14]). The axioms (M), (D), (SC) correspond to the conditions (m), (d), and (sc).

Definition 4 (ED Calculus) 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)	$B_i\varphi \to \neg B_i\neg\varphi$.
(SC)	$\check{B}_a\varphi \wedge \check{K}_a(\neg\varphi \wedge \psi) \to B_a(\varphi \vee \psi)$

RULES

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

Use $\vdash_{ED}$ for derivability in this calculus.

Epistemic Weight Models

An alternative semantics for the language of epistemic doxastic logic can be given with respect to *Epistemic Weight Models*. As it turns out, the calculus given above is *incomplete* for this alternative semantics.

Definition 5 (Epistemic Weight Models) *An* **epistemic weight model** *for agents I and basic propositions P is a tuple $\mathcal{M} = (W, R, V, L)$ where*

- *W is a non-empty countable set of worlds,*
- *R assigns to every agent $i \in I$ an equivalence relation $\sim_i$ on W,*
- *V assigns to every $w \in W$ a subset of P,*
- *L assigns to every $i \in I$ a function $\mathbb{L}_i$ from W to $\mathbb{Q}^+$ (the positive rationals), subject to the following boundedness condition (*).*

$$\forall i \in I \forall w \in W \sum_{u \in [w]_i} \mathbb{L}_i(w) < \infty. \tag{*}$$

where $[w]_i$ is the cell of w in the partition induced by $\sim_i$.

Definition 6 (Truth in Epistemic Weight Models) *Key clauses are:*

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

$$\mathcal{M}, w \models B_i\varphi \textit{ iff}$$

$$\sum\{\mathbb{L}_i(v) \mid v \in [w]_i, \mathcal{M}, v \models \varphi\} > \sum\{\mathbb{L}_i(v) \mid v \in [w]_i, \mathcal{M}, v \models \neg\varphi\}.$$

Definition 7 (Agreement) *Let* $\mathcal{M} = (W, R, V, N)$ *be a neighbourhood model and let* L *be a weigth function for* $\mathcal{M}$*. Then* L agrees with $\mathcal{M}$ *if it holds for all agents* i *and all* $w \in W$ *that*

$$X \in N_i(w) \textit{ iff } \mathbb{L}_i(X) > \mathbb{L}_i([w]_i - X).$$

The following theorem shows that the ED calculus is incomplete for epistemic weight models.

Theorem 8 *[ER14] There exists an epistemic neighbourhood model* $\mathcal{M}$ *that has no agreeing weight function.*

Proof. Adaptation of example 2 from [WF79, pp. 344-345]

Let *Prop* $:= \{a, b, c, d, e, f, g\}$. Assume a single agent 0. Define:

$$\mathcal{X} := \{efg, abg, adf, bde, ace, cdg, bcf\}.$$

$$\mathcal{X}' := \{abcd, cdef, bceg, acfg, bdfg, abef, adeg\}.$$

Notation: xyz for $\{x, y, z\}$.

$$\mathcal{Y} := \{Y \mid \exists X \in \mathcal{X} : X \leq Y \leq W\}.$$

Let $\mathcal{M} := (W, R, V, N)$ be defined by $W :=$ ***Prop***, $R_0 = W \times W$, $V(w) = \{w\}$, and for all $w \in W$, $N_0(w) = \mathcal{Y}$.

Check that $\mathcal{X}' \cap \mathcal{Y} = \emptyset$. So $\mathcal{M}$ is a neighbourhood model.

Toward a contradiction, suppose there exists a weight function L that agrees with $\mathcal{M}$. Since each letter $p \in W$ occurs in exactly three of the seven members of $\mathcal{X}$, we have:

$$\sum_{X \in \mathcal{X}} \mathbb{L}_0(X) = \sum_{p \in W} 3 \cdot \mathbb{L}_0(\{p\}).$$

Since each letter $p \in W$ occurs in exactly four of the seven members of $\mathcal{X}'$, we have:

$$\sum_{X \in \mathcal{X}'} \mathbb{L}_0(X) = \sum_{p \in W} 4 \cdot \mathbb{L}_0(\{p\}).$$

On the other hand, from the fact that $\mathbb{L}_0(X) > L_o(W - X)$ for all members X of $\mathcal{X}$ we get:

$$\sum_{X \in \mathcal{X}} \mathbb{L}_0(X) > \sum_{X \in \mathcal{X}} \mathbb{L}_0(W - X) = \sum_{X \in \mathcal{X}'} \mathbb{L}_0(X).$$

Contradiction. So no such $\mathbb{L}_0$ exists. $\square$

Strengthening the Axiom System

It is possible to strengthen the axiom system to get completeness. For this we need the so-called Scott axioms [Sco64]. Intuitively, what the Scott axioms say is this: If agent a knows the number of true φ_i is less than or equal to the number of true ψ_i, agent a believes φ_1, and the remaining φ_i are each consistent with her beliefs, then agent a believes one of the ψ_i.

It turns out that this is expressible in the KB language [Seg71]. In Segerberg notation:

$(\varphi_1, \ldots, \varphi_m \mathbb{I}_a \psi_1, \ldots, \psi_m)$

abbreviates a KB formula expressing that agent a knows that the number of true φ_i's is less than or equal to the number of true ψ_i's.

Put another way, $(\varphi_i \mathbb{I}_a \psi_i)_{i=1}^m$ is true if and only if every one of a's epistemically accessible worlds satisfies at least as many ψ_i as φ_i.

The Scott axioms can now be phrased as follows.

(Scott) $\quad [(\varphi_i \mathbb{I}_a \psi_i)_{i=1}^m \wedge B_a\varphi_1 \wedge \bigwedge_{i=2}^m \check{B}_a \varphi_i] \to \bigvee_{i=1}^m B_a \psi_i$

Fact 9 *Adding the Scott axioms to the KB calculus yields a system that is sound and complete for epistemic weight models [ER14].*

What does this mean? At least that qualitative and quantitative belief are different.

Since any epistemic weight model determines a neighbourhood model, one may interpret this as saying that there are subtleties about belief that are not captured by probability.

What are natural examples of situations that are correctly described by a neighbourhood model that cannot be extended to a weight model? Maybe one has to think of propositions that have their weight determined by context?

We will now present the system of simplified probabilistic epistemic logic of [ES14], but with alternative syntax, following [DR15].

Definition 10 (Probabilistic Epistemic Logic Simplified: Language)

$$\begin{array}{rcl} \varphi & ::= & \top \mid p \mid \neg\varphi \mid \varphi \wedge \varphi \mid \Phi \leq_i \Phi \\ \Phi & ::= & \varphi \mid \varphi \oplus \Phi \end{array}$$

Abbreviations:

As usual for $\perp, \vee, \rightarrow, \leftrightarrow$.

$\Phi <_i \Psi$ for $\Phi \leq_i \Psi \wedge \neg\Psi \leq_i \Phi$.

$\Phi =_i \Psi$ for $\Phi \leq_i \Psi \wedge \Psi \leq_i \Phi$.

$B_i\varphi$ for $(\neg\varphi) <_i \varphi$, $\check{B}_i\varphi$ for $(\neg\varphi) \leq_i \varphi$. "Belief as willingness to bet"

$K_i\varphi$ for $\top \leq_i \varphi$, $\check{K}_i\varphi$ for $\perp <_i \varphi$. "Knowledge as certainty"

Definition 11 (Probabilistic Epistemic Logic Simplified: Semantics) *Let $\mathcal{M} = (W, R, V, L)$ be an epistemic weight model, let $w \in W$.*

$$\begin{array}{rcl} [\![\varphi]\!]_{\mathcal{M}} & := & \{w \in W \mid \mathcal{M}, w \models \varphi\} \\ [\![\varphi]\!]_{\mathcal{M}}^{w,i} & := & [\![\varphi]\!]_{\mathcal{M}} \cap [w]_i \\ \mathbb{L}_{w,i}\varphi & := & \displaystyle\sum_{u \in [\![\varphi]\!]_{\mathcal{M}}^{w,i}} \mathbb{L}_i(u) \end{array}$$

$$\begin{array}{rcl} \mathcal{M}, w \models \top & & \textit{always} \\ \mathcal{M}, w \models p & \textit{iff} & p \in V(w) \\ \mathcal{M}, w \models \neg\varphi & \textit{iff} & \textit{not } \mathcal{M}, w \models \varphi \\ \mathcal{M}, w \models \varphi_1 \wedge \varphi_2 & \textit{iff} & \mathcal{M}, w \models \varphi_1 \textit{ and } \mathcal{M}, w \models \varphi_2 \\ \mathcal{M}, w \models \Phi \leq_i \Psi & \textit{iff} & \displaystyle\sum_{\varphi \in \Phi} \mathbb{L}_{w,i}\varphi \leq \sum_{\psi \in \Psi} \mathbb{L}_{w,i}\psi \end{array}$$

Weight function and epistemic accessibility relation together determine probability:

$$P_{w,i}^{\mathcal{M}}\varphi := \frac{\mathbb{L}_{w,i}\varphi}{\mathbb{L}_{w,i}\top} \left(= \frac{\sum_{u \in [\![\varphi]\!]_{\mathcal{M}} \cap [w]_i} \mathbb{L}_i(u)}{\sum_{u \in [w]_i} \mathbb{L}_i(u)} \right)$$

In a slogan: "Probabilities are weights normalized for epistemic partition cells."

Example 2 *Two bankers i, j consider buying stocks in three firms a, b, c that are involved in a takeover bid. There are three possible outcomes: a for "a wins", b for " b wins", and c for "c wins." i takes the winning chances to be $3:2:1$, j takes them to be $1:2:1$.*

i: solid lines, j: dashed lines.

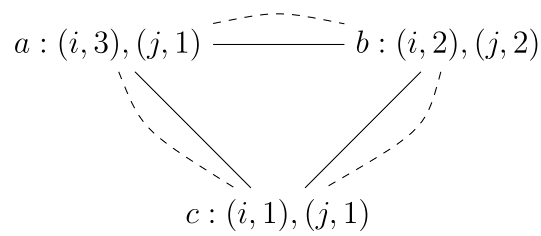

We see that i is willing to bet $1 : 1$ on a, while j is willing to bet $3 : 1$ against a.

It follows that in this model i and j have an opportunity to gamble, for, to put it in Bayesian jargon, they do not have a common prior.

Example 3 *Suppose j has foreknowledge about what firm c will do.*

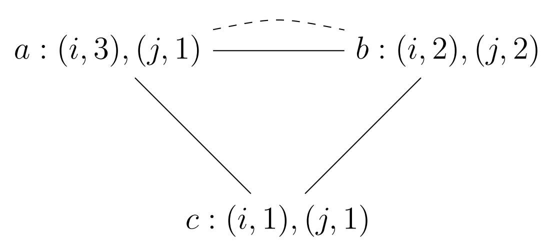

The probabilities assigned by i remain as before. The probabilities assigned by j have changed, as follows. In worlds a and b, j assigns probability $\frac{1}{3}$ to a and $\frac{2}{3}$ to b. In world c, j is sure of c.

- We may suppose that this new model results from j being informed about the truth value of c, while i is aware that j received this information, but without i getting the information herself.
- So i is aware that j's subjective probabilities have changed, and it would be unwise for i to put her beliefs to the betting test. For although i cannot distinguish the three situations, she knows that j can distinguish the c situation from the other two.
- Willingness of j to bet against a at any odds can be interpreted by i as an indication that c is true, thus forging an intimate link between action and information update.

A model $\mathcal{M} = (W, R, V, L)$ is **single weight** if $\forall i, j \in L \forall w \in W : \mathbb{L}_i(w) = \mathbb{L}_j(w)$.

Theorem 12 *[ES14]: Every epistemic weight model has an equivalent single weight model.*

Theorem 13 *[ES14]: There are finite epistemic weight models that only have infinite single weight counterparts.*

These theorems are proved with an appropriate notion of bisimulation for our language. If $X \subseteq W$ then we use $\mathbb{L}_i(X)$ for $\sum_{x \in X} \mathbb{L}_i(x)$.

Definition 14 (Bisimulation) *Let $\mathcal{M} = (W, R, V, L)$ and $\mathcal{M}' = (W', R', V', L')$ be two epistemic weight models, and let B be a relation on $W \times W'$. Then B is a bisimulation if wBw' implies:*

Invar *w and w' satisfy the same atomic formulas.*

Zig *For every i, every set $E \subseteq [w]_i$ there exists a set $E' \subseteq [w']_i$ such that*

- *for all $u' \in E'$ there exists $u \in E$ with uBu',*
- $\mathbb{L}_i(E) \leq \mathbb{L}'_i(E')$.

Zag *Similarly in the other direction.*

Example 4 *Two models $\mathcal{M}$, $\mathcal{M}'$ and a bisimulation relation B*

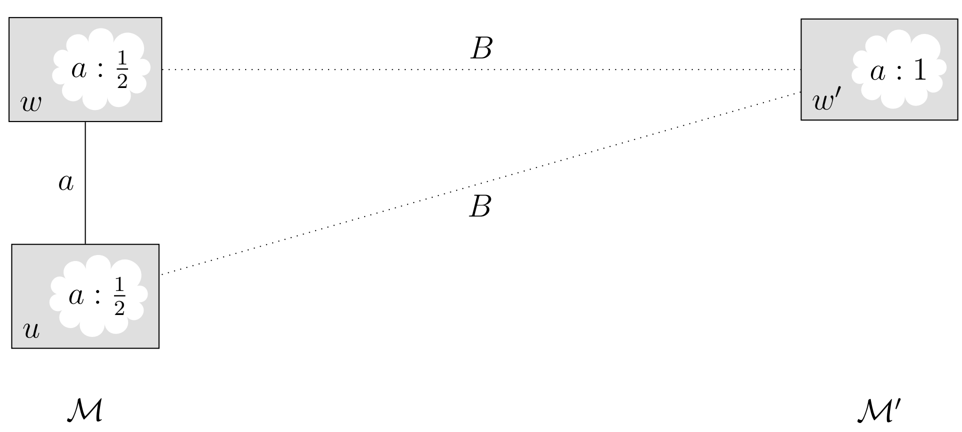

This notion of bisimulation is well-behaved:

- We can prove a Hennessy-Milner theorem
- Bisimulations are closed under composition and union.

Example 5 (Fair or Biased?) *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.*

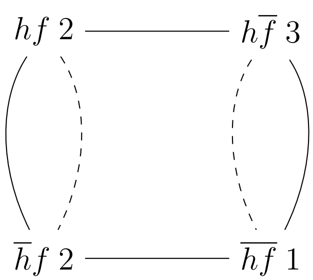

In world hf, i assigns probability $\frac{5}{8}$ to h and probability $\frac{1}{2}$ to f, and j assigns probability $\frac{1}{2}$ to h and probability 1 to f.

Example 6 (Fair or Biased; Normalized Version) *Give each agent its own weight, and normalize the weight functions using the epistemic accessibilities.*

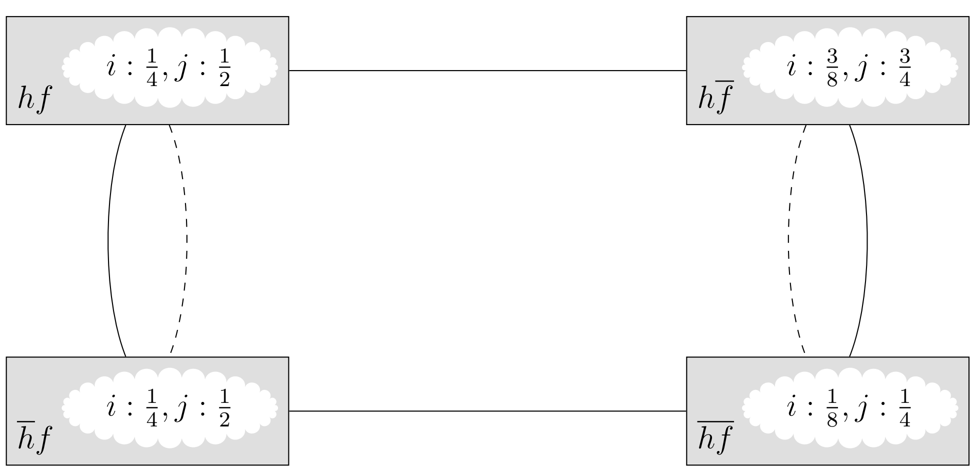

Definition 15 (SC Calculus)

Taut	*instances of propositional tautologies*
ProbT	$(\top \leq_i \varphi) \to \varphi$
ProbImpl	$\top \leq_i (\varphi \to \psi) \to (\varphi \leq_i \psi)$
PropPos	$(\Phi \leq_i \Psi) \to \top \leq_i (\Phi \leq_i \Psi)$
PropNeg	$(\Phi >_i \Psi) \to \top \leq_i (\Phi >_i \Psi)$
PropAdd	$(\varphi \wedge \psi) \oplus (\varphi \wedge \neg\psi) =_i \varphi$
Tran	$(\Phi \leq_i \Psi) \wedge (\Psi \leq_i \Xi) \to (\Phi \leq_i \Xi)$
Tot	$(\Phi \leq_i \Psi) \vee (\Psi \leq_i \Phi)$
ComL	$(\Phi_1 \oplus \Phi_2 \leq_i \Psi) \leftrightarrow (\Phi_2 \oplus \Phi_1 \leq_i \Psi)$
ComR	$(\Phi \leq_i \Psi_1 \oplus \Psi_2) \leftrightarrow (\Phi \leq_i \Psi_2 \oplus \Psi_1)$
Add	$(\Phi_1 \leq_i \Psi_1) \wedge (\Phi_2 \leq_i \Psi_2) \to (\Phi_1 \oplus \Phi_2 \leq_i \Psi_1 \oplus \Psi_2)$
Succ	$(\Phi \oplus \top \leq_i \Psi \oplus \top) \to (\Phi \leq_i \Psi)$
MP	*From* $\vdash \varphi$ *and* $\vdash \varphi \to \psi$ *derive* $\vdash \psi$
PR	*From* $\vdash \varphi \to \psi$ *derive* $\vdash \varphi \leq_i \psi$

Definition 16 (Derivability) $\Gamma \vdash \varphi$ *holds if either* $\varphi \in \Gamma$*, or* φ *is an axiom, or* φ *follows by means of the rules of the calculus from axioms or members of* Γ*, while taking care that application of PR only is allowed when the set of premisses* Γ *is empty.*

The deduction theorem holds for this calculus:

Theorem 17 $\Gamma \cup \{\varphi\} \vdash \psi$ *iff* $\Gamma \vdash \varphi \to \psi$.

Proof. The proof for this is folklore, but see [HN12] for an explanation of why the restriction on the use of PR (counterpart to the rule of necessitation in standard modal logic) is crucial for this. □

The deduction theorem will help us to state a number of useful derivable principles.

1. From ProbImpl, with propositional reasoning:

$$\vdash \top \leq_i (\varphi \leftrightarrow \psi) \to (\varphi =_i \psi).$$

2. Consider the following instance of ProbT: $\vdash (\top \leq_i \bot) \to \bot$. By propositional logic, $\vdash \neg(\top \leq_i \bot)$, i.e., $\vdash \bot <_i \top$.

3. $\vdash \Phi =_i \Phi$ by Tot.

4. From PropAdd, $\vdash (\varphi \wedge \bot) \oplus (\varphi \wedge \top) =_i \varphi$. By propositional reasoning, $\vdash \bot \oplus \varphi =_i \varphi$.

5. $\vdash \bot \oplus \Phi =_i \Phi$ by an easy induction using the previous two items.

6. Plugging in $\top$ as a special case in (3), we get $\vdash \bot \oplus \top =_i \top$.

7. Plugging in $\bot$ as a special case in (3), we get $\vdash \bot \oplus \bot =_i \bot$.

8. Assume $\vdash \varphi \leftrightarrow \psi$. Then also $\vdash \varphi \to \psi$ and $\vdash \psi \to \varphi$. From the former, with PR, $\vdash \varphi \leq_i \psi$, and from these, we get $\vdash \varphi =_i \psi$. This gives the derived inference rule:

   From $\vdash \varphi \leftrightarrow \psi$, derive $\vdash \varphi =_i \psi$.

9. From PropAdd, $\vdash (\top \wedge \varphi) \oplus (\top \wedge \neg\varphi) =_i \top$, so by propositional reasoning, $\vdash \varphi \oplus \neg\varphi =_i \top$.

10. Assume $\vdash \varphi =_i \top$. Then $\vdash \neg\varphi \oplus \varphi =_i \neg\varphi \oplus \top$, and hence $\vdash \neg\varphi \oplus \top =_i \top$, by the previous item. Since $\vdash \bot \oplus \top =_i \top$ we get $\vdash \neg\varphi \oplus \top =_i \bot \oplus \top$, and therefore, by Succ, $\vdash \neg\varphi =_i \bot$. By the deduction theorem, we have derived $\vdash \varphi =_i \top \to \neg\varphi =_i \bot$.

11. In a similar way we can derive: $\vdash \varphi =_i \bot \to \neg\varphi =_i \top$.

12. Assume $\vdash \top \leq_i \varphi$ and $\vdash \top \leq_i \varphi \to \psi$. From this: $\vdash \varphi =_i \top$ and $\vdash \varphi \to \psi =_i \top$. From $\vdash \varphi =_i \top$, $\vdash \neg\varphi \wedge \neg\psi =_i \bot$ and from $\vdash \varphi \to \psi =_i \top$ we get that $\vdash \varphi \wedge \neg\psi =_i \bot$. Since $\vdash \neg\psi =_i (\varphi \wedge \neg\psi) \oplus (\neg\varphi \wedge \neg\psi)$, by PropAdd, we derive that $\vdash \neg\psi =_i \bot$, and it follows that $\vdash \psi =_i \top$, hence $\vdash \top \leq_i \psi$. By the deduction theorem, we have derived the distribution principle for certainty: $\vdash (\top \leq_i \varphi \wedge \top \leq_i (\varphi \to \psi) \to \top \leq_i \psi$.

13. From Tot, by the definitions of $<_i$, $>_i$ and $=_i$:

$$\vdash \Phi <_i \Psi \vee \Phi =_i \Psi \vee \Phi >_i \Psi.$$

14. By Tran and the definition of $<_i$:

$$\vdash \Phi <_i \Psi \wedge \Psi \leq_i \Xi \rightarrow \Phi <_i \Xi,$$

$$\vdash \Phi <_i \Psi \wedge \Psi \leq_i \Xi \rightarrow \Phi <_i \Xi,$$

$$\vdash \Phi <_i \Psi \wedge \Psi <_i \Xi \rightarrow \Phi <_i \Xi.$$

15. By Add, and the definition of $=_i$:

$$\vdash \Phi_1 =_i \Phi_2 \wedge \Psi_1 =_i \Psi_2 \rightarrow \Phi_1 \oplus \Psi_1 =_i \Phi_2 \oplus \Psi_2.$$

Definition 18 (Translating Knowledge and Belief) *If φ is an ED formula, then $\varphi^\bullet$ is the formula of the language of epistemic comparison logic given by the following instructions:*

$$\begin{array}{rcl}
\top^\bullet & = & \top \\
p^\bullet & = & p \\
(\neg\varphi)^\bullet & = & \neg\varphi^\bullet \\
(\varphi_1 \wedge \varphi_2)^\bullet & = & \varphi_1^\bullet \wedge \varphi_2^\bullet \\
(K_i\varphi)^\bullet & = & \top \leq_i (\varphi^\bullet) \\
(B_i\varphi)^\bullet & = & \neg\varphi^\bullet <_i \varphi^\bullet.
\end{array}$$

If Γ is a set of formulas then $\Gamma^\bullet$ is the corresponding set of $^\bullet$-transations.

Theorem 19 *For all φ in the ED language: if $\Gamma \vdash_{ED} \varphi$ then $\Gamma^\bullet \vdash_{SC} \varphi^\bullet$.*

Proof. We show that the (translations of the) axioms and rules of the ED calculus are all derivable in the SCD calculus.

K_i necessitation.

Assume $\vdash_{SC} \varphi$. Then $\vdash_{SC} \top \rightarrow \varphi$, hence by PR $\vdash_{SC} \top \leq_i \varphi$. By the translation convention, $\vdash_{SC} K_i\varphi$.

Dist-K $K_i(\varphi \rightarrow \psi) \rightarrow K_i\varphi \rightarrow K_i\psi$

Assume $\vdash \top \leq_i (\varphi \rightarrow \psi)$ and $\vdash \top \leq_i \varphi$. Then by propositional reasoning, $\vdash \top =_i (\varphi \rightarrow \psi)$, and by PropAdd and propositional reasoning, $\vdash \bot =_i \varphi \wedge \neg\psi$. From $\vdash \top \leq_i \varphi$, by propositional reasoning, $\vdash \top =_i \varphi$, and substitution yields $\vdash \bot =_i \top \wedge \neg\psi$, and therefore $\vdash \bot =_i \neg\psi$. It follows by PropAdd that $\vdash \top =_i \varphi$.

T: Follows from PropT.

PI-K: Follows from ProbPos

NI-K: Follows from ProbNeg

F: From ProbT, $\vdash (\top \leq_i \bot) \rightarrow \bot$, hence by propositional reasoning, $\vdash \bot <_i \top$, and therefore by propositional reasoning $\vdash \bot \leq_i \top$.

PI-KB: For this assume $\vdash \top \leq_i \varphi$. From this: $\vdash \varphi =_i \top$. Earlier we derived $\vdash \neg\varphi \oplus \varphi =_i \top$ and $\top =_i \bot \oplus \top$. From this, with propositional reasoning: $\vdash \neg\varphi \oplus \top =_i \bot \oplus \top$, and hence by Succ, $\vdash \neg\varphi =_i \bot$. From this, with $\vdash_i \bot <_i \top$ (derived earlier), $\vdash_i \neg\varphi <_i \varphi$, i.e., $\vdash B_i\varphi$.

M: Assume $\vdash \top \leq_i (\varphi \rightarrow \psi)$ and $\vdash \neg\varphi <_i \varphi$. We have to show that $\vdash \neg\psi <_i \psi$. From $\vdash \top \leq_i (\varphi \rightarrow \psi)$ we can derive $\vdash \varphi \leq_i \psi$ (todo: show derivation). From $\vdash \varphi \leq_i \psi$ derive $\vdash \neg\psi \leq_i \neg\varphi$ (todo: show derivation). From $\vdash \neg\varphi <_i \varphi$ and $\vdash \varphi \leq_i \psi$ and by Trans $\vdash \neq \varphi < \psi$. From this and $\vdash \neg\psi \leq_i \neg\varphi$ by Trans $\vdash \neq \psi < \psi$.

D: From $\vdash \neg\varphi <_i \varphi$ by Tot $\vdash \neg\varphi \leq_i \varphi$.

SC: Assume $\vdash (\neg\varphi <_i \varphi)$ and $\vdash \bot <_i (\neg\varphi \wedge \psi)$. We have to show $\vdash (\neg\varphi \wedge \neg\psi) <_i (\varphi \vee \psi)$. For this we can use the equivalence of $\varphi \vee \psi$ and $(\varphi \wedge \neg\psi) \vee (\varphi \wedge \psi) \vee (\neg\varphi \wedge \psi)$. □

Theorem 20 *Every consistent formula φ determines a canonical epistemic weight model $\mathcal{M}_\varphi$.*

Proof. Suppose φ is consistent, i.e., $\nvdash \neg\varphi$. We construct a canonical epistemic weight model for φ.

Let Φ be the set of all subformulas of φ, closed under single negations. The subformulas of $\Psi \leq_i \Xi$ are all subformulas of formulas ψ that occur as $\oplus$ terms in Ψ or Ξ, plus the results $\Psi' \leq_i \Xi'$ of leaving out $\oplus$ terms in Ψ or Ξ while taking care that Ψ' and Ξ' are not empty.

We define the canonical model $\mathcal{M}_\varphi = (W, R, V, L)$. W is the set of all maximal consistent subsets of Φ. W is non-empty because φ is supposed to be consistent.

Valuations are defined as follows: $V(\mathbf{w}) = \mathit{Prop} \cap \mathbf{w}$.

Let sat($\mathbf{w}$) $= \{\psi \in \Phi \mid \mathbf{w} \vdash \psi\}$, that is, sat($\mathbf{w}$) is the set of Φ-formulas that are provable from $\mathbf{w}$.

Notice that it follows by the soundness of the probability comparison calculus that all members of sat($\mathbf{w}$) are true in $\mathbf{w}$.

Relations are defined as follows: $\mathbf{w}R_i\mathbf{u}$ iff sat($\mathbf{w}$) and sat($\mathbf{u}$) contain the same

i-comparison formulas. Clearly, all R_i are equivalence relations.

Now it remains to define L. Consider an agent i and an equivalence class $R_i(\mathbf{w})$ in the canonical model $\mathcal{M}_\varphi$. All worlds $\mathbf{u}$ of $R_i(\mathbf{w})$ contain the same i-comparison formulas.

We show how to transform all these i-comparison formulas in a system of linear inequalities that is consistent.

For all $\mathbf{u} \in W$, we write $\varphi_\mathbf{u}$ for the conjunction of all formulas in $\mathbf{u}$. We have:

- $\vdash \varphi_\mathbf{u} \rightarrow \neg\varphi_\mathbf{v}$ if $\mathbf{u} \neq \mathbf{v}$ by propositional logic.

Given any formula ψ of Φ, we have

- $\vdash \psi \leftrightarrow \bigvee_{\{\mathbf{u} \in W \mid \psi \in \mathbf{u}\}} \varphi_\mathbf{u}$ by propositional logic.

Since the $\varphi_\mathbf{u}$ are all mutually inconsistent we can prove in the calculus:

- $\vdash \psi =_i \bigoplus \{\varphi_\mathbf{u} \mid \mathbf{u} \in W \text{ and } \psi \in \mathbf{u}\}$.

Now, when we i-compare $\varphi_\mathbf{u}$ to $\bot$ in $\mathbf{w}$, we should obtain $>_i$ iff $\mathbf{u} \in R_i(\mathbf{w})$. Let us prove this fact.

- If $\mathbf{u} \in R_i(\mathbf{w})$, we have:
  1. $\vdash \varphi_\mathbf{u} \rightarrow \bot <_i \varphi_\mathbf{u}$ by (ProbT) and (PropAdd).
  2. $\bot <_i \varphi_\mathbf{u} \in \mathrm{sat}(\mathbf{u})$;
  3. $\bot <_i \varphi_\mathbf{u} \in \mathrm{sat}(\mathbf{w})$ because $\mathbf{u} \in R_i(\mathbf{w})$.

  Therefore, $\bot <_i \varphi_\mathbf{u}$ follows from $\mathbf{w}$.
- Suppose $\mathbf{u} \notin R_i(\mathbf{w})$. Then $\mathbf{u}$ and $\mathbf{w}$ differ by at least one i-comparison formula $\Psi \leq_i \Xi \in \Phi$. Without loss of generality, assume $\Psi \leq_i \Xi \in \mathbf{w}$ and $\Psi \leq_i \Xi \notin \mathbf{u}$. Then we have:
  1. $\vdash \varphi_\mathbf{w} \rightarrow \Psi \leq_i \Xi$ by propositional logic;
  2. $\vdash \Psi \leq_i \Xi \rightarrow \neg\varphi_\mathbf{u}$ by propositional logic;
  3. $\vdash \Psi \leq_i \Xi \rightarrow \top \leq_i (\Psi \leq_i \Xi)$ by axiom (PropT);
  4. $\vdash \varphi_\mathbf{w} \rightarrow \top \leq_i (\Psi \leq_i \Xi)$ by propositional logic;

5. $\vdash \varphi_{\mathbf{w}} \to \top \leq_i (\neg\varphi_{\mathbf{u}})$ by 2. and 4.
6. $\vdash \varphi_{\mathbf{w}} \to (\varphi_{\mathbf{u}} \leq_i \bot)$ from the above by PropAdd and propositional reasoning.

Therefore $\varphi_{\mathbf{u}} \leq_i \bot$ follows from $\mathbf{w}$.

Thus, ψ has the same i-weight as $\{\varphi_{\mathbf{u}} \mid \mathbf{u} \in R_i(\mathbf{w}) \text{ and } \psi \in \mathbf{u}\}$. We can prove in the calculus:

- $\vdash \psi =_i \bigoplus\{\varphi_{\mathbf{u}} \mid \mathbf{u} \in R_i(\mathbf{w}) \text{ and } \psi \in \mathbf{u}\}$.

Now let $\Psi \leq_i \Xi$ be any i-comparison formula of $\mathbf{w}$. Then we can replace any $\oplus$ term ψ occurring in either Ψ or Ξ by a list of terms $\bigoplus\{\varphi_{\mathbf{u}} \mid \mathbf{u} \in R_i(\mathbf{w}) \text{ and } \psi \in \mathbf{u}\}$ with the same i-weight. Let the result of this be $\Psi' \leq_i \Xi'$. Regrouping the $\oplus$ terms in Ψ' and Ξ', using the abbreviation $n\chi$ for $\underbrace{\chi \oplus \cdots \oplus \chi}_{n \text{ times}}$, 0 for $\bot \oplus \cdots \oplus \bot$, m for $\underbrace{\top \oplus \cdots \oplus \top}_{m \text{ times}}$, and replacing $\oplus$ by $+$ and $\leq_i$ by $\leq$ gives a linear inequality

$$a_1\varphi_{\mathbf{u}_1} + \cdots + a_n\varphi_{\mathbf{u}_n} + k \leq b_1\varphi_{\mathbf{v}_1} + \cdots + b_m\varphi_{\mathbf{v}_m} + l$$

where a_i, b_j, k, l are non-negative integers, and the $\varphi_{\mathbf{u}}$ and $\varphi_{\mathbf{v}}$ figure as variables. Applying this recipe to each i-comparison formula in $\mathbf{w}$, we get a system of linear inequalities made up of i-inequalities in $\mathbf{w}$.

The set sat($\mathbf{w}$) is consistent so the above system, which is a rephrasing of inequations that are in sat($\mathbf{w}$), is also consistent and therefore satisfiable [FHM90b, Theorem 2.2]. Let $(x_u^*)_{\mathbf{u} \in R_i(\mathbf{w})}$ be a solution, and define $\mathbb{L}_i(\mathbf{u}) = x_u^*$. □

Lemma 21 (Truth Lemma) *For all formulas $\psi \in \Phi$, we have $\mathcal{M}_\varphi, \mathbf{w} \models \psi$ iff $\psi \in \mathbf{w}$.*

Proof. Induction on ψ. □

Theorem 22 (Completeness of Epistemic Comparison Logic) *The calculus of epistemic comparison logic is complete for epistemic weight models:*

$$\textit{If } \models \varphi \textit{ then } \vdash_{SC} \varphi.$$

Proof. Let $\not\vdash \varphi$. Then $\neg\varphi$ is consistent, and one can find a maximal consistent set $\mathbf{w}$ in the closure of $\neg\varphi$ with $\neg\varphi \in \mathbf{w}$. By the Truth Lemma, $\mathcal{M}_{\neg\varphi}, \mathbf{w} \models \neg\varphi$, i.e., $\mathcal{M}_{\neg\varphi}, \mathbf{w} \not\models \varphi$. Therefore, $\not\models \varphi$. □

From Epistemic Probability Models to Epistemic 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 $i \in Ag$, $w \in W$.

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

Fact 23 *For any epistemic weight model $\mathcal{M}$ it holds that $\mathcal{M}^\bullet$ is a neighbourhood model.*

Fact 24 *The calculus of epistemic-doxastic neighbourhood logic is sound for interpretation in epistemic probability models. Probabilistic beliefs are neighbourhoods.*

Theorem 25 *For all ED formulas φ, for all epistemic probability models $\mathcal{M}$, for all worlds w of $\mathcal{M}$:*

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

Theorem 26 *Let $\vdash_{ED}$ denote derivability in the neighbourhood calculus for ED. Let $\vdash_{SC}$ denote derivability in the probability comparison calculus. Then $\vdash_{ED} \varphi$ implies $\vdash_{SC} \varphi^\bullet$.*

Updates

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? We can model this with public announcement update.

Example 7 (Disease and Test)

$$\begin{tikzcd}
dt\ 0.9 \arrow[r, dash] \arrow[d, dash] & d\bar{t}\ 0.1 \arrow[d, dash] & & dt\ 0.9 \arrow[d, dash] \\
\bar{d}t\ 0.2 * 99 \arrow[r, dash] & \overline{dt}\ 0.8 * 99 \arrow[r, phantom, "\Rightarrow !t \Rightarrow"] & & \bar{d}t\ 0.2 * 99
\end{tikzcd}$$

The model shows that after the update with t the probability of d equals $\frac{0.9}{0.9+0.2*99} = \frac{9}{207} = \frac{1}{23}$.

Example 8 (Applying Bayes' Rule) *Compare this with applying Bayes' Rule:*

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

Filling in $P(T|D) = 0.9, P(D) = 0.01, P(\neg D) = 0.99, P(T|\neg D) = 0.2$ *gives* $P(D|T) = \frac{1}{23}$.

Public announcement update of an epistemic weight model and application of Bayes' rule give the same result.

In [BGK09], update models for probabilistic epistemic logic are built from sets of formulas that are mutually exclusive. We will stay a bit closer to the original update model from [BMS98, BM04]. A weighted update model is like a weighted epistemic model, but with the valuation function replaced by a function that assigns preconditions and actions (substitutions) to events.

Definition 27 (Binding) *A binding is a map from proposition letters to formulas, represented by a finite set of links*

$$\{p_1 \mapsto \varphi_1, \ldots, p_n \mapsto \varphi_n\}$$

where the p_k *are all different, and where no* φ_k *is equal to* p_k*. It is assumed that each* p *that does not occur in a left-hand side of a binding is mapped to itself.*

We use $\mathbf{S}$ *for the set of bindings.*

Definition 28 (Update Model) *An update model for probabilistic comparison logic is a tuple*

$$(E, \mathbf{pre}, S, T, K)$$

where

- *E is a non-empty set of events*
- $\mathbf{pre}$ *is a function $E \to \mathcal{L}$ that assigns a $\mathcal{L}$ formula φ (the precondition) to each event $e \in E$.*
- *S is a function $E \to \mathbf{S}$ that assigns a binding to each event $e \in E$.*
- *T is a function that maps each agent i to an equivalence $\sim_i$ on E.*
- *K is a function from agents to $\mathbb{L}$-functions for events.*

Update is a product operation, as in [BMS98, BM04]. The new i-weight for (w, e) is computed as the product the weights of w and of e.

Definition 29 (Update Product) *The update execution of static model $\mathcal{M} = (W, R, V, L)$ with action model $\mathbf{A} = (E, \mathbf{pre}, S, T, K)$ is a tuple*

$$\mathcal{M} \circledast \mathbf{A} = (W' R', V', L')$$

where:

- $W' = \{(w, e) \mid \mathcal{M}, w \models \mathbf{pre}(e)\}$.
- *R'_i is given by*
$$\{((w_1, e_1), (w_2, e_2)) \mid w_1 \sim_i w_2 \textit{ and } e_1 \sim_i e_2\}.$$
- $V'(p) = \{(w, e) \in W' \mid \mathcal{M}, w \models S(e)(p)\}$.
- $\mathbb{L}_i(w, e) = \mathbb{L}_i(w) \times \mathbb{L}_i(e)$.

We will illustrate this definition with a model of an example from 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 [Gar81]? As a slight modification, let's assume that while Mr a puts his white marble in the urn, he takes a look at what is already in there, and Mrs b notices this.

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 $p + q$.

To model the complete sequence of events, we start out from a model of complete ignorance about p, for two agents a (solid) and b (dashed):

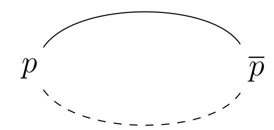

Update model for a learning whether p is the case, while b observes that a learns, but without learning herself.

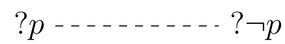

Result of the update with this:

The action of putting a second white marble into the urn:

$$q := \top$$

Result of the update with this:

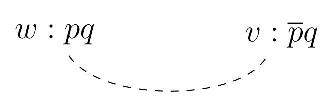

The final action: an update model for removing either p or q from the urn. Neither Mr a nor Mrs b knows which of these two takes place. Note that removing p from the urn has as precondition that p is true, and similarly for q.

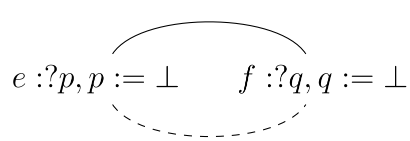

The result of updating with this:

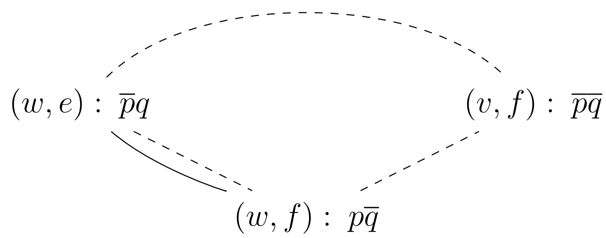

The relevant property to check now is $p \vee q$. Since all possibilities are equally likely for both agents, in world (w, e), the probability for Mr a that $p \vee q$ is 1, and for Mrs b it is $\frac{2}{3}$.

Further Work: Risk Versus Uncertainty

Recall the famous distinction made in [Kni21] between **risk** and **uncertainty**: a risky choice is a choice involving *known* probabilities, an uncertain choice is a choice involving *unknown* probabilities. This distinction is used by Keynes in [Key21] to explain how the insurance business works, where underwriters distinguish routinely between risks which are properly insurable and risks which cannot be calculated or made into a "book" that covers all possibilities "…and which cannot form the basis of a regular business of insurance" [Key21, p. 21]

If we want to incorporate the useful distinction between risk and uncertainty in our weight models, we should allow for propositions that agents are not willing to bet on, because they do not trust their own estimations of the probability.

A traditional way to tackle this is by distinguishing between *lower* and *upper* probabilities. If these are wide apart, the agent is not willing to bet at all [Sha76, Hal03]. Or think of it as buying and selling. The lower probability is like the price for which I can sell, the higher probability is like the price for which I can buy.

In an epistemic weight model, we can simply replace weights by *weights with margin of error*. Instead of a single value, we assign a pair of values (x, y), and we say that the lower value is x, and the higher value $x + y$. Thus, y gives the spread. Old style weights are a special case, with (x, y) such that $y = 0$.

Let $\mathbb{L}_i$ give the lower value, and $\mathbb{L}_i^\epsilon$ the spread. Add a new comparison relation $\Phi <_i \Psi$ to the language. Interpretation instructions for $\Phi \leq_i \Psi$ and $\Phi <_i \Psi$ now get modified as follows:

$$\mathcal{M}, w \models \Phi \leq_i \Psi \text{ iff } \sum_{\varphi \in \Phi} (\mathbb{L}_{w,i}\varphi + \mathbb{L}^{\epsilon}_{w,i}\varphi) \leq \sum_{\psi \in \Psi} \mathbb{L}_{w,i}\psi.$$

$$\mathcal{M}, w \models \Phi <_i \Psi \text{ iff } \sum_{\varphi \in \Phi} (\mathbb{L}_{w,i}\varphi + \mathbb{L}^{\epsilon}_{w,i}\varphi) < \sum_{\psi \in \Psi} \mathbb{L}_{w,i}\psi.$$

Note that $\Phi \leq_i \Psi$ and $\Phi <_i \Psi$ are no longer interdefinable. Also, totality does no longer hold, for it may be that neither $\Phi \leq_i \Psi$ nor $\Psi \leq_i \Phi$.

If we use $\Phi =_i \Psi$ for the conjunction of $\Phi \leq_i \Psi$ and $\Psi \leq_i \Phi$ we see that $\Phi =_i \Psi$ holds iff the lower weights of Φ and Ψ are the same, and the spreads of both Φ and Ψ equal 0.

A bet on φ is safe iff the $\mathbb{L}$-value of φ is greater than one half (or, equivalently, the $\mathbb{L}$ plus $\mathbb{L}^{\epsilon}$ value of $\neg\varphi$ is less than one half).

This perspective is connected to what is sometimes called imprecise probability theory (See Walley [Wal91]). The distinction between lower and spread probabilities makes the probabilistic belief logic weaker. It is not quite clear yet what that means. Can we now get a precise match between imprecise weight models and epistemic doxastic neighbourhood models? Also on the agenda for further work is the question of the axiomatisation of imprecise epistemic probability logic. Finally, what is the appropriate notion of generic imprecise update?

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.

[Ben03] J. van Benthem. Conditional probability meets update logic. *Journal of Logic, Language and Information*, 12(4):409–421, 2003.

[BFDP14] Johan van Benthem, David Fernández-Duque, and Eric Pacuit. Evidence logic: A new look at neighborhood structures. *Annals of Pure and Applied Logic*, 165(1):106–133, 2014.

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

[BM04] A. Baltag and L.S. Moss. Logics for epistemic programs. *Synthese*, 139(2):165–224, 2004.

[BMS98] A. Baltag, L.S. Moss, and S. Solecki. The logic of public announcements, common knowledge, and private suspicions. In I. Bilboa, editor, *Proceedings of TARK'98*, pages 43–56, 1998.

[BS06] A. Baltag and S. Smets. Conditional doxastic models: A qualitative approach to dynamic belief revision. *Electronic Notes in Theoretical Computer Science (ENTCS)*, 165:5–21, 2006.

[BS08] A. Baltag and S. Smets. A qualitative theory of dynamic interactive belief revision. In G. Bonanno, W. van der Hoek, and M. Wooldridge, editors, *Logic and the Foundations of Game and Decision Theory (LOFT 7)*, Texts in Logic and Games, pages 11–58. Amsterdam University Press, 2008.

[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.

[Che80] B.F. Chellas. *Modal Logic: An Introduction*. Cambridge University Press, 1980.

[DR15] James P. Delgrande and Bryan Renne. The logic of qualitative probability. In *Proceedings of the Twenty-Fourth International Conference on Artificial Intelligence (IJCAI 2015)*, pages 2904–2910, Buenos Aires, 2015.

[ER14] Jan van Eijck and Bryan Renne. Belief as willingness to bet. E-print, arXiv.org, December 2014. arXiv:1412.5090v1 [cs.LO].

[ES14] Jan van Eijck and François Schwarzentruber. Epistemic probability logic simplified. In Rajeev Goré, Barteld Kooi, and Agi Kurucz, editors, *Advances in Modal Logic, Volume 10*, pages 158–177, 2014.

[FHM90a] Ronald Fagin, Joseph Y Halpern, and Nimrod Megiddo. A logic for reasoning about probabilities. *Information and computation*, 87(1):78–128, 1990.

[FHM90b] Ronald Fagin, Joseph Y Halpern, and Nimrod Megiddo. A logic for reasoning about probabilities. *Information and computation*, 87(1):78–128, 1990.

[Fol09] Richard Foley. Beliefs, degrees of belief, and the Lockean thesis. In F. Huber and C. Schmidt-Petri, editors, *Degrees of Belief*, volume 342 of *Synthese Library*. Springer, 2009.

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

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

[HKP09] H. H. Hansen, C. Kupke, and E. Pacuit. Neighbourhood structures: Bisimilarity and basic model theory. *Logical Methods in Computer Science*, 5:1–38, 2009.

[HN12] Raul Hakli and Sara Negri. Does the deduction theorem fail for modal logic? *Synthese*, 187(3):849–867, 2012.

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

[JG13] Dick de Jongh and Sujata Ghosh. Comparing strengths of belief explicitly. *Logic Journal of the IGPL*, 21:488–514, 2013.

[Key21] J.M. Keynes. *A Treatise on Probability*. Macmillan, London, 1921. Republished by Dover.

[Kni21] F. H. Knight. *Risk, Uncertainty, and Profit*. Hart, Schaffner & Marx; Houghton Mifflin Company, Boston, MA, 1921.

[Koo03] Barteld P. Kooi. *Knowledge, Chance, and Change*. PhD thesis, Groningen University, 2003.

[Kyb61] H.E. Kyburg. *Probability and the Logic of Rational Belief*. Wesleyan University Press, Middletown, CT, 1961.

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

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

[Nar07] Louis Narens. *Theories of Probability*. World Scientific, 2007.

[Sco64] Dana Scott. Measurement structures and linear inequalities. *Journal of Mathematical Psychology*, 1:233–247, 1964.

[Seg71] K. Segerberg. Qualitative probability in a modal setting. In J. Fenstad, editor, *Proceedings of the 2nd Scandinavian Logic Symposium*, pages 341–352, Amsterdam, 1971. North Holland.

[Sha76] Glenn Shafer. *A Mathematical Theory of Evidence*. Princeton University Press, Princeton, 1976.

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

[WF79] P. Walley and T.L. Fine. Varieties of modal (classificatory) and comparative probability. *Synthese*, 41:321–374, 1979.

[Zve10] Jonathan A. Zvesper. *Playing with Information*. PhD thesis, ILLC, University of Amsterdam, 2010.