Update, Probability, Knowledge and Belief

Jan van Eijck
CWI & ILLC, Amsterdam

LogiCIC Workshop 2015, November 27, 2015

Abstract

The talk considers two kinds of models for logics of knowledge and belief, neighbourhood models and epistemic weight models, and traces connections.

We present a new Probability Comparison Calculus that is sound and complete for epistemic weight models. Epistemic weight models combine knowledge and probability by using epistemic accessibility relations and weights to define subjective probabilities. This is a further simplification of the calculus for probabilistic epistemic weight models that was presented in AIML 2014 [ES14].

At the end of the talk we turn to generic update, and present some examples of how this is handled in PRODEMO [Eij13], our prototype model checker for probabilistic epistemic logic.

Probability and 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]

Probability and 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]

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.

State of the Art: Existing Combinations of DEL and Probability Theory

- Kooi's thesis [Koo03], Van Benthem [Ben03] , Van Benthem CS [BGK09]
- Inspiration for this: work of Fagin and Halpern in the 1990s [FHM90]
- Simplified version: [ES14]. Epistemic model checker for this: [Eij13]. Further simplication presented here based on: [DR15].
- Belief as willingness to bet: [ER14]. Logic with explicit belief comparison operator: [JG13], or [Nar07] for an overview of the literature. Related: evidence models [BFDP14].
- Probabilistic Logic of Communication and Change: [Ach14].

Some Varieties of Belief

Some Varieties of Belief

- Betting belief (or: Bayesian belief) in φ: $P(\varphi) > P(\neg\varphi)$. Van Eijck & Renne [ER14].

Some Varieties of Belief

- Betting belief (or: Bayesian belief) in φ: $P(\varphi) > P(\neg\varphi)$. Van Eijck & Renne [ER14].
- Threshold belief in φ: $P(\varphi) > t$, for some specific t with $\frac{1}{2} \leq t < 1$. Also known as Lockean belief.

Some Varieties of Belief

- Betting belief (or: Bayesian belief) in φ: $P(\varphi) > P(\neg\varphi)$. Van Eijck & Renne [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: $P(\varphi|\psi) > P(\neg\varphi|\psi)$ (Leitgeb [Lei10]).

Some Varieties of Belief

- Betting belief (or: Bayesian belief) in φ: $P(\varphi) > P(\neg\varphi)$. Van Eijck & Renne [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: $P(\varphi|\psi) > P(\neg\varphi|\psi)$ (Leitgeb [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). Baltag & Smets [BS06, BS08]

Some Varieties of Belief

- Betting belief (or: Bayesian belief) in φ: $P(\varphi) > P(\neg\varphi)$. Van Eijck & Renne [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: $P(\varphi|\psi) > P(\neg\varphi|\psi)$ (Leitgeb [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). Baltag & Smets [BS06, BS08]
- Subjective certainty belief in φ: $P(\varphi) = 1$. Used in epistemic game theory (Aumann [Aum99]).

The Lottery Puzzle

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:

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

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:

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

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.

Three Possible Reactions to the Lottery Puzzle

Three Possible Reactions to the Lottery Puzzle

1. Deny that Alice believes that her ticket is not winning.

Three Possible Reactions to the Lottery Puzzle

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

Three Possible Reactions to the Lottery Puzzle

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.

Three Possible Reactions to the Lottery Puzzle

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. $\ldots$?? $\ldots$Deny that Alice believes that there is a winning ticket.

Pursuing the second strategy: How can we drop the closure of belief under conjunction?

We need an operator B_i that does not satisfy (Dist).

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

This means: B_i is not a normal modal operator.

What we need is neighbourhood semantics [Che80, Ch. 8]. See also [Zve10], [HKP09], and [BFDP14].

Epistemic Neighbourhood Models

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

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

where

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.

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 $\{v \in W \mid w \sim_i v\}$.

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*.

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

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.

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.

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

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.

(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.

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.

(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.

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.

(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

Extra conditions

(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.

Extra conditions

(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.

(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.

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

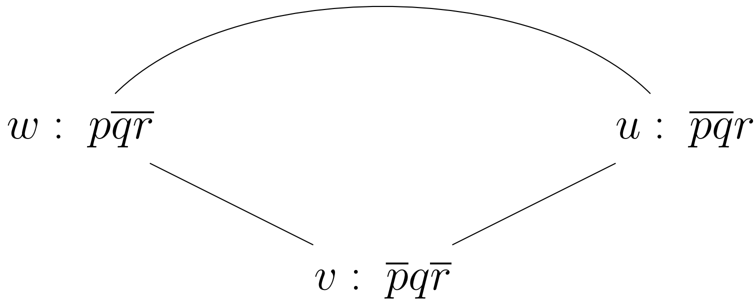

$$N(w) = N(v) = N(u) = \{\{w, v\}, \{v, u\}, \{w, u\}, \{w, v, u\}\}$$

Example

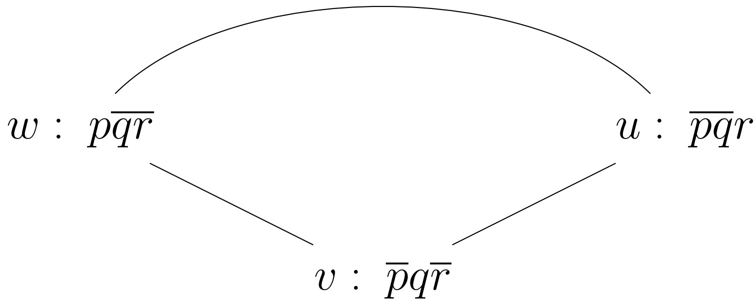

$$N(w) = N(v) = N(u) = \{\{w, v\}, \{v, u\}, \{w, u\}, \{w, v, u\}\}$$

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

Example

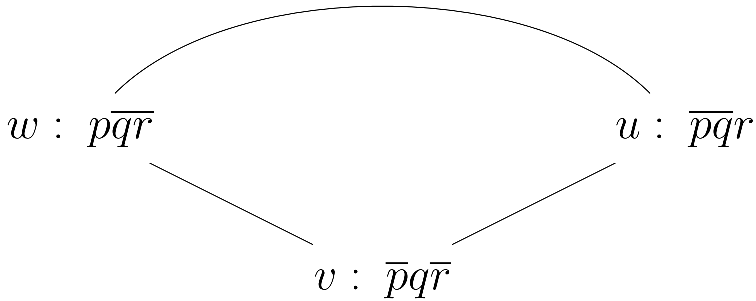

$$N(w) = N(v) = N(u) = \{\{w, v\}, \{v, u\}, \{w, u\}, \{w, v, u\}\}$$

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

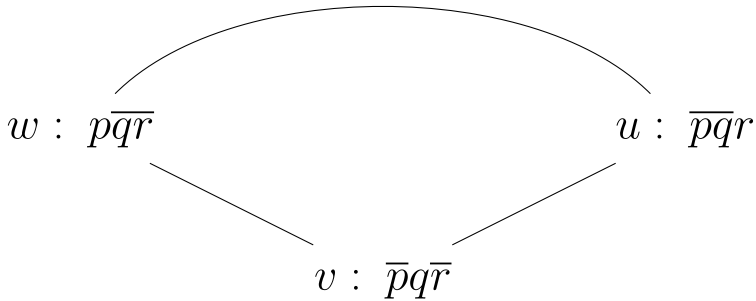

$$N(w) = N(v) = N(u) = \{\{w, v\}, \{v, u\}, \{w, u\}, \{w, v, u\}\}$$

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.

Example

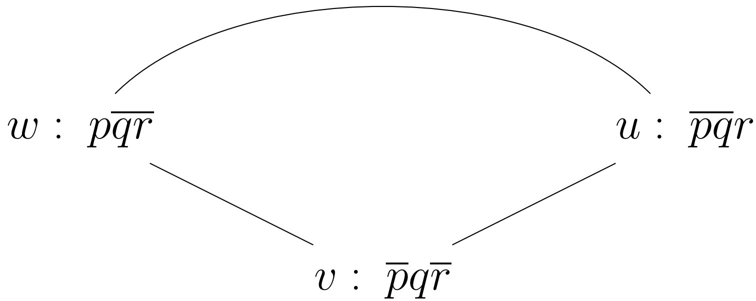

$$N(w) = N(v) = N(u) = \{\{w, v\}, \{v, u\}, \{w, u\}, \{w, v, u\}\}$$

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.

The lottery puzzle is “solved” in neighbourhood models for belief by non-closure of belief under conjunction.

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

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

Completeness for Epistemic Neighbourhood Models

See [ER14] and [BvBvES14].

Completeness for Epistemic Neighbourhood Models

See [ER14] and [BvBvES14].

Alternative Semantics

An alternative semantics can be given with respect to Epistemic Weight Models.

Completeness for Epistemic Neighbourhood Models

See [ER14] and [BvBvES14].

Alternative Semantics

An alternative semantics can be given with respect to Epistemic Weight Models.

It turns out that the calculus is incomplete for this alternative semantics.

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 boundedness condition (*) below.

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

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

Interpretation of KB language in Epistemic Weight Models

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

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

Interpretation of KB language in Epistemic Weight Models

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

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

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

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

Incompleteness of KB Calculus for Probability Models

There exists an epistemic neighbourhood model $\mathcal{M}$ that has no agreeing weight function.

Adaptation of example 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 $\mathcal{M}$.

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:

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

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:

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

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:

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

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.

Strengthening the Axiom System

Scott Axioms, intuitively:

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.

Strengthening the Axiom System

Scott Axioms, intuitively:

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 this is expressible in the KB language.

Segerberg notation [Seg71]:

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

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.

Scott Axioms

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

Scott Axioms

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

Fact 1 *Adding the Scott axioms to the KB calculus yields a system that is sound and complete for epistemic weight models [ER14].*

Scott Axioms

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

Fact 1 *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.

Scott Axioms

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

Fact 1 *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.

Scott Axioms

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

Fact 1 *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?

Scott Axioms

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

Fact 1 *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?

Probabilistic Epistemic Logic Simplified: [ES14], but with Alternative Syntax [DR15]

EPISTEMIC COMPARISON LOGIC: LANGUAGE

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

Probabilistic Epistemic Logic Simplified: [ES14], but with Alternative Syntax [DR15]

Epistemic Comparison Logic: Language

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

Abbreviations:

As usual for $\bot, \vee, \to, \leftrightarrow$.

$\Phi <_i \Psi$ 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 $\bot <_i \varphi$. “Knowledge as certainty”

Semantics for this language

Let $\mathcal{M} = (W, R, V, L)$, let $w \in W$.

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

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

Weight + Access = Prob

Weight function and epistemic accessibility relation together determine probability:

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

Slogan

"Probabilities are weights normalized for epistemic partition cells."

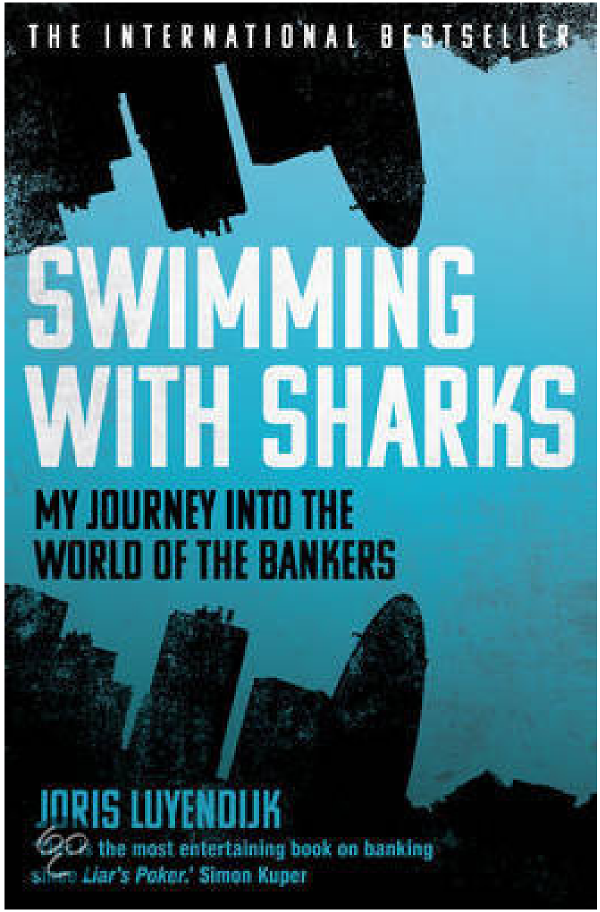
THE INTERNATIONAL BESTSELLER
SWIMMING WITH SHARKS
MY JOURNEY INTO THE WORLD OF THE BANKERS
JORIS LUYENDIJK
the most entertaining book on banking
Liar's Poker.' Simon Kuper

Example: Willingness to Bet in Investment Banking

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.

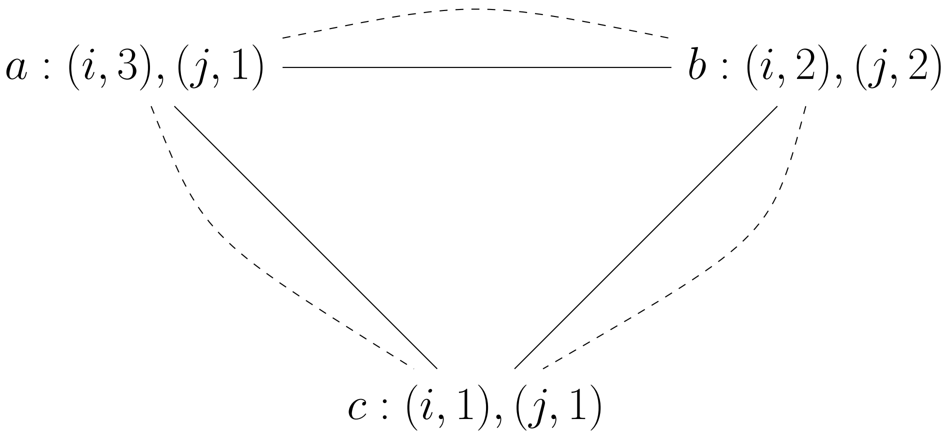

Belief as Willingness to Bet

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.

Foreknowledge in Investment Banking

Suppose j has foreknowledge about what firm c will do.

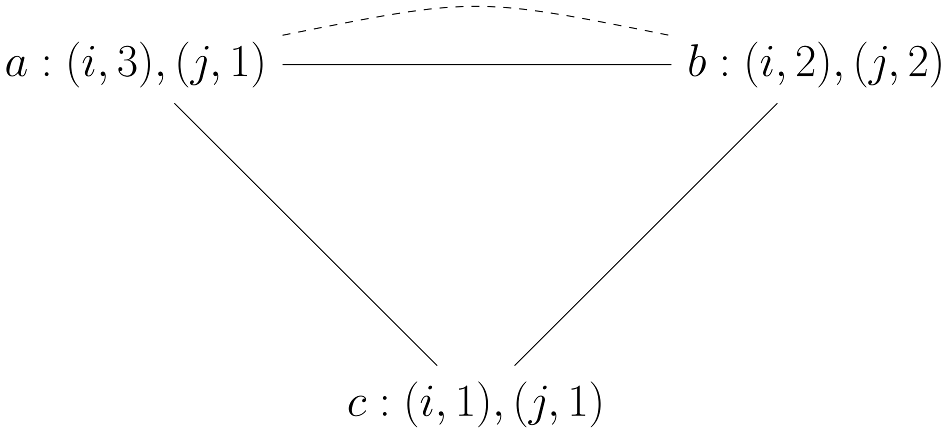

The probabilities assigned by i remain as before. The probabilities assigned by j have changed, as follows. In worlds a and b, j assigns probability $\frac{1}{3}$ to a and $\frac{2}{3}$ to b. In world c, j is sure of c.

- 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.

- 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.

- 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.

Multiple Versus Single Weight Models

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

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

To prove this we need an appropriate notion of **bisimulation**.

Bisimulation

If $X \subseteq W$ then we use $\mathbb{L}_i(X)$ for $\sum_{x \in X} \mathbb{L}_i(x)$.

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:

Bisimulation

If $X \subseteq W$ then we use $\mathbb{L}_i(X)$ for $\sum_{x \in X} \mathbb{L}_i(x)$.

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.

Bisimulation

If $X \subseteq W$ then we use $\mathbb{L}_i(X)$ for $\sum_{x \in X} \mathbb{L}_i(x)$.

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',
- $\mathbb{L}_i(E) \leq \mathbb{L}'_i(E')$.

Bisimulation

If $X \subseteq W$ then we use $\mathbb{L}_i(X)$ for $\sum_{x \in X} \mathbb{L}_i(x)$.

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 Bisimulation

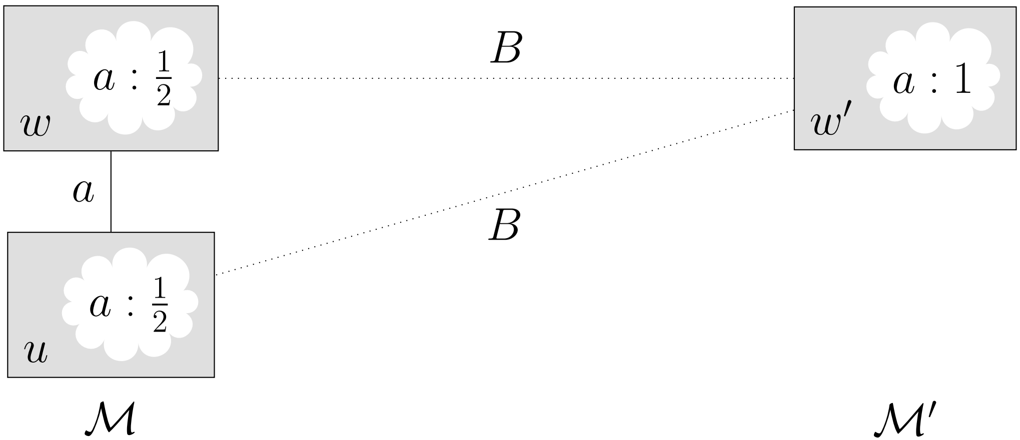

Example Bisimulation

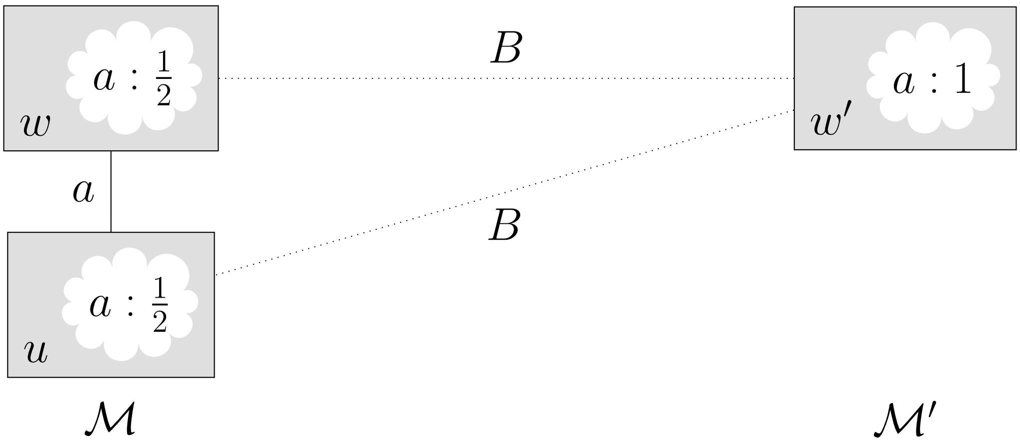

This notion of bisimulation is well-behaved (compare Giovanni's talk yesterday):

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

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.

$$
\begin{tikzcd}
hf\ 2 \arrow[r, dash] \arrow[d, dash, bend right] \arrow[d, dash, dashed, bend left] & h\overline{f}\ 3 \arrow[d, dash, dashed, bend right] \arrow[d, dash, bend left] \\
\overline{h}f\ 2 \arrow[r, dash] & \overline{hf}\ 1
\end{tikzcd}
$$

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.

$$\begin{tikzcd}
hf\ 2 \arrow[r, dash] \arrow[d, dash, bend right] \arrow[d, dash, dashed, bend left] & h\overline{f}\ 3 \arrow[d, dash, dashed, bend right] \arrow[d, dash, bend left] \\
\overline{h}f\ 2 \arrow[r, dash] & \overline{h}\overline{f}\ 1
\end{tikzcd}$$

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.

Normalized Version

Give each agent its own weight, and normalize the weight functions using the epistemic accessibilities.

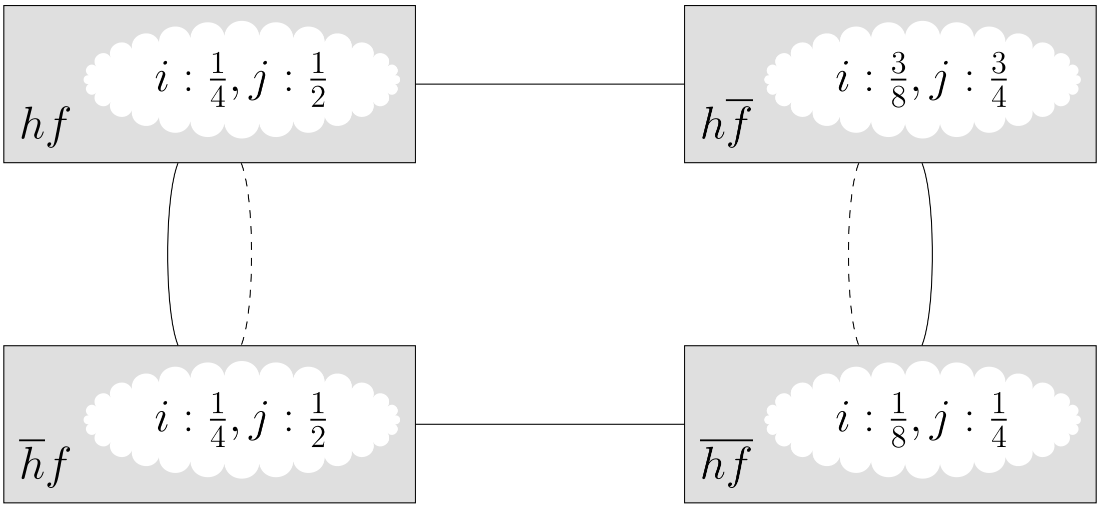

CALCULUS FOR EPISTEMIC COMPARISON LOGIC

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 $\vdash \varphi \leq_i \psi$

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.

Completeness of SC Calculus

Theorem 2 *Every consistent formula φ determines a canonical epistemic weight model $\mathcal{M}_\varphi$.*

Let Φ be the set of all subformulas of φ, closed under single negations. (Define subformulas of $\Psi \leq_i \Xi \ldots$)

Canonical model $\mathcal{M}_\varphi = (W, V, R, L)$.

W is the set of all maximal consistent subsets of Φ.

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

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

Relations: $\mathbf{w} R_i \mathbf{u}$ iff $\mathrm{sat}(\mathbf{w})$ and $\mathrm{sat}(\mathbf{u})$ contain the same i-comparison formulas. Clearly, all R_i are equivalence relations.

Definition of Weight Function

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.

It is possible to transform these i-comparison formulas in a system of linear inequalities that is consistent.

If $\mathbf{u} \in W$, write $\varphi_{\mathbf{u}}$ for the conjunction of all formulas in $\mathbf{u}$. We can prove:

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

Definition of Weight Function (2)

Now let $\Psi \leq_i \Xi$ be any i-comparison formula of **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.

Thus we get a consistent system of linear inequalities made up of i-inequalities in **w**. Let $(x_u^*)_{\mathbf{u}\in R_i(\mathbf{w})}$ be a solution, and define $\mathbb{L}_i(\mathbf{u}) = x_u^*$.

From Epistemic Weight 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 For any epistemic weight model $\mathcal{M}$ it holds that $\mathcal{M}^\bullet$ is a neighbourhood model.

Translating Knowledge and Belief

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

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

Theorem 3 *For all KB 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 4 *Let $\vdash$ denote derivability in the neighbourhood calculus for KB. Let $\vdash'$ denote derivability in the epistemic comparison calculus. Then $\vdash \varphi$ implies $\vdash' \varphi^\bullet$.*

Updates: Public Announcement

Public Announcement: restriction to φ worlds. $[!\varphi]\psi$.

EXAMPLE: REASONING ABOUT DISEASE

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.

A Weight Model for the Disease Problem

Let's use p for the outcome of the test and q for having the disease:

```
test = Pr p
disease = Pr q
```

```
dmodel :: EpistWM Int
dmodel = WMo
  [0..3]
  [a,b]
  [(0,[p,q]),(1,[q]),(2,[p]), (3,[])]
  [(a,[[0,1,2,3]]),(b,[[0,1,2,3]])]
  [(a,[(0,0.9),(1,0.1),(2,0.2*99),(3,0.8*99)]),
   (b,[(0,0.9),(1,0.1),(2,0.2),(3,0.8)])]
  [0..3]
```

Probability of having the disease

According to a:

```
*DEMO> sprob dmodel 0 a disease
1 % 100
```

According to b:

```
*DEMO> sprob dmodel 0 b disease
1 % 2
*DEMO> isTrue dmodel 0 (bb b disease)
False
*DEMO> isTrue dmodel 0 (bd b disease)
True
```

Update with the test result

```
dmodel' = upd_wpa dmodel test
```

```
DEMO> dmodel'
WMo [0,2] [a,b] [(0,[p,q]),(2,[p])]
[(a,[[0,2]]),(b,[[0,2]])]
[(a,[(0,9 % 10),(2,99 % 5)]),
 (b,[(0,9 % 10),(2,1 % 5)])] [0,2]
*DEMO> sprob dmodel' 0 a disease
1 % 23
*DEMO> sprob dmodel' 0 b disease
9 % 11
*DEMO> isTrue dmodel' 0 (bb b disease)
True
```

Compare with Applying Bayes' Rule

Compare 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 $P(D|T) = \frac{1}{23}$.

Generic Update

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 probabilistic update model is like a probabilistic epistemic model, but with the valuation function replaced by a function that assigns preconditions and actions (substitutions) to events.

Update models for S5 are like epistemic S5 models, but with their valuations replaced by precondition/action functions. A PA function assigns a Precondition formula and an Action (i.e., a substitution) to each state.

Update Model

An update model for probabilistic epistemic logic is a tuple

$$(W, P, R, L)$$

where

- W is a non-empty set of events
- P is a function $W \times L \to \mathcal{L} \times \mathbf{S}$ that assigns a pair (φ, S) consisting of a $\mathcal{L}$ formula φ (the precondition) and a substitution S (the action) to each world w.
- R is a function that maps each agent a to an equivalence R_a on W.
- L is a function from agents to $\mathbb{L}$-functions.

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.

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 [Gar81]?

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 $p + q$.

Initial Situation: Blissful Ignorance

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]
```

This gives:

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

Private Communication

An update model for telling a the value of p, while b does not learn this fact.

```
um1 :: FUM Prp
um1 = \ ags -> UM
    ags
    [0,1]
    [(0,(p,[])),(1,(Ng p,[]))]
    ((a,[[0],[1]]) :
        [ (x,[[0,1]]) | x <- ags \\ [a] ])
    [0,1]
    [(x,[(0,1/2),(1,1/2)]) | x <- ags ]
```

The result of updating with this

```
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]
[(a,[(0,1 % 2),(1,1 % 2)]),
 (b,[(0,1 % 2),(1,1 % 2)])]
```

Putting a second white marble in the urn

This can be implemented as 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]
[(a,[(0,1 % 2),(1,1 % 2)]),
 (b,[(0,1 % 2),(1,1 % 2)])]
```

Removing a white marble

An update model for removing either p or q from the urn. Nobody 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.

```
um2 :: FUM Prp
um2 = \ ags -> UM
    ags
    [0,1]
    [(0,(p,[(P 0,Ng Top)])),
     (1,(q,[(Q 0,Ng Top)]))]
    [ (x,[[0,1]]) | x <- ags ]
    [0,1]
    [(x, [(0,1/2),(1,1/2)]) | x <- ags ]
```

The 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]
[(a,[(0,1 % 3),(1,1 % 3),(2,1 % 3)]),
 (b,[(0,1 % 3),(1,1 % 3),(2,1 % 3)])]
```

Probability that the other marble is also white?

In our setting: what is the probability of $p \vee q$?

Well, it is different for a and b. Here is the whole story:

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

Risk Versus Uncertainty

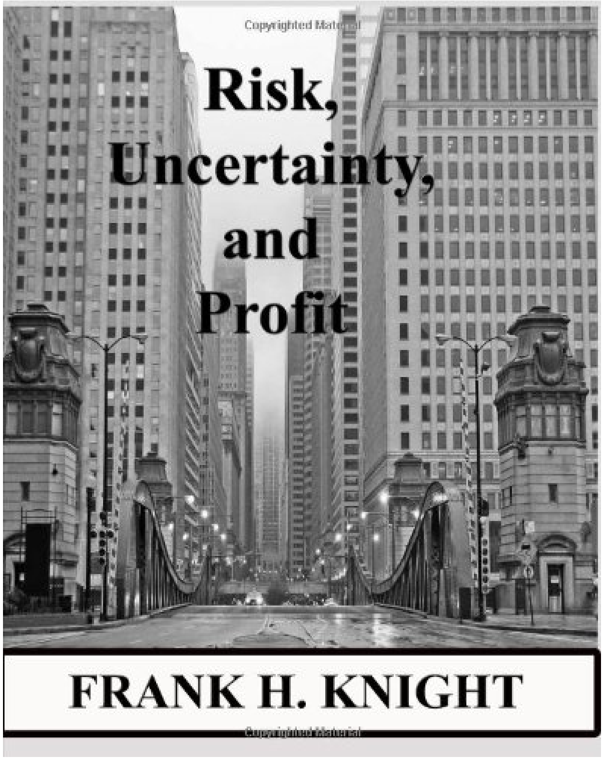

Knight's Distinction [Kni21]

Risk Choices involving known probabilities.

Uncertainty Choices involving unknown probabilities.

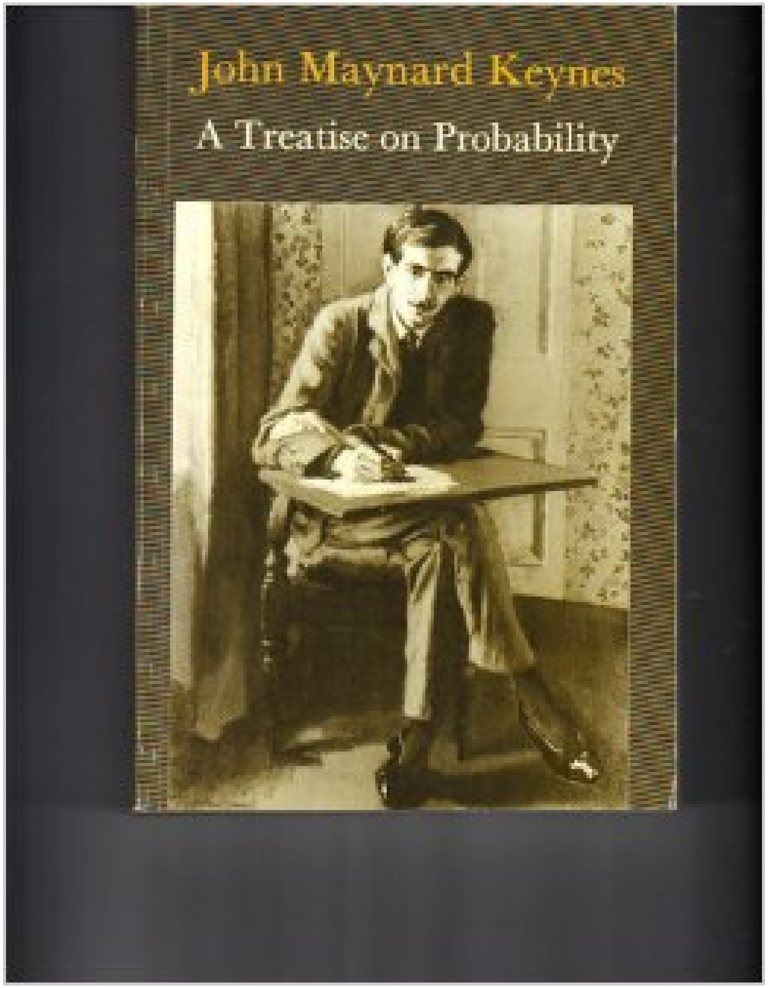
John Maynard Keynes
A Treatise on Probability

Keynes About the Distinction Between Risk and Uncertainty

Take a cue from how people actually deal with uncertainty. E.g. from the insurance trade:

> In fact underwriters themselves distinguish between risks which are properly insurable, either because their probability can be estimated within comparatively numerical limits or because it is possible to make a “book” which covers all possibilities, and other risks which cannot be dealt with in this way and which cannot form the basis of a regular business of insurance, — although an occasional gamble may be indulged in. [Key21, p. 21]

Or look at the practice of lawyers:

> A distinction, interesting for our present purpose, between probabilities, which can be estimated within somewhat narrow limits, and those which cannot, has arisen in a series of judicial decisions respecting damages. [Key21, p. 21]

Follows a case where a breeder of racehorses tries to recover damages for breach of a contract ...

Models with Imprecise Weights

Models with Imprecise Weights

- There are propositions that agents are not willing to bet on, because they do not trust their own estimations of the probability.

Models with Imprecise Weights

- There are propositions that agents are not willing to bet on, because they do not trust their own estimations of the probability.
- Traditional way to tackle this: distinguish between lower and upper probabilities. If these are wide apart, the agent is not willing to bet at all [Sha76, Hal03].

Models with Imprecise Weights

- There are propositions that agents are not willing to bet on, because they do not trust their own estimations of the probability.
- Traditional way to tackle this: distinguish between lower and upper probabilities. If these are wide apart, the agent is not willing to bet at all [Sha76, Hal03].
- 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.

Models with Imprecise Weights

- There are propositions that agents are not willing to bet on, because they do not trust their own estimations of the probability.
- Traditional way to tackle this: distinguish between lower and upper probabilities. If these are wide apart, the agent is not willing to bet at all [Sha76, Hal03].
- 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.
- Replace weights by weights with margin. 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 $y = 0$.

Interpretation of Comparison Formulas in Imprecise Weight Models

- $\mathcal{M}, w \models \Phi \leq_i \Psi$: The sum of the higher i-weights of Φ is less than or equal to the sum of the lower i-weights of Ψ.
- $\mathcal{M}, w \models \Phi <_i \Psi$: The sum of the higher i-weights of Φ is strictly less than the sum of the lower i-weights of Ψ.
- Note that $\Phi \leq_i \Psi$ and $\Phi <_i \Psi$ are no longer interdefinable.

Imprecise Probabilities and Neighbourhood Models

Imprecise Probabilities and Neighbourhood Models

- The lower probability of φ in state w is calculated as the ratio between A equal to the sum of the lower weight values for the states $v \in [w]_i$ with $M, v \models \varphi$, and B equal to A plus the sum of the lower and margin weight values for the points $v \in [w]_i$ with $M, v \models \neg\varphi$.

Imprecise Probabilities and Neighbourhood Models

- The lower probability of φ in state w is calculated as the ratio between A equal to the sum of the lower weight values for the states $v \in [w]_i$ with $M, v \models \varphi$, and B equal to A plus the sum of the lower and margin weight values for the points $v \in [w]_i$ with $M, v \models \neg\varphi$.

- This gives that the lower value for φ always equals 1 minus the higher value of $\neg\varphi$. Hence a bet of φ is safe iff the lower value of φ is greater than one half (or, equivalently, the higher value of $\neg\varphi$ is smaller than one half).

Imprecise Probabilities and Neighbourhood Models

- The lower probability of φ in state w is calculated as the ratio between A equal to the sum of the lower weight values for the states $v \in [w]_i$ with $M, v \models \varphi$, and B equal to A plus the sum of the lower and margin weight values for the points $v \in [w]_i$ with $M, v \models \neg\varphi$.

- This gives that the lower value for φ always equals 1 minus the higher value of $\neg\varphi$. Hence a bet of φ is safe iff the lower value of φ is greater than one half (or, equivalently, the higher value of $\neg\varphi$ is smaller than one half).

- This perspective is connected to what is sometimes called imprecise probability theory (See Walley [Wal91]).

Further Work: Imprecise Epistemic Comparison Logic

Further Work: Imprecise Epistemic Comparison Logic

- The distinction between lower and higher probabilities makes the probabilistic belief logic weaker.

Further Work: Imprecise Epistemic Comparison Logic

- The distinction between lower and higher probabilities makes the probabilistic belief logic weaker.
- Does this mean that now we can get a precise match between imprecise weight models and epistemic doxastic neighbourhood models?

Further Work: Imprecise Epistemic Comparison Logic

- The distinction between lower and higher probabilities makes the probabilistic belief logic weaker.
- Does this mean that now we can get a precise match between imprecise weight models and epistemic doxastic neighbourhood models?
- Give an axiomatisation of imprecise epistemic probability logic.

Further Work: Imprecise Epistemic Comparison Logic

- The distinction between lower and higher probabilities makes the probabilistic belief logic weaker.
- Does this mean that now we can get a precise match between imprecise weight models and epistemic doxastic neighbourhood models?
- Give an axiomatisation of imprecise epistemic probability logic.
- What is the appropriate notion of generic imprecise update?

The

UNPERSUADABLES

ADVENTURES *with the* ENEMIES *of* SCIENCE

WILL STORR

AUTHOR OF *WILL STORR VS. THE SUPERNATURAL*

OVERLOOK

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.

[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.

[Eij13] Jan van Eijck. Prodemo — implementation of model checking for epl. available from `homepages.cwi.nl:~/jve/software/prodemo`, 2013.

[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.

[FHM90] Ronald Fagin, Joseph Y Halpern, and Nimrod Megiddo.

A logic for reasoning about probabilities. Information and computation, 87(1):78–128, 1990.

[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.

[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.

[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.

[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.