Reasoning with Probabilities

Joshua Sack

August 5, 2013

Goal of the course

Probability is among the most popular tools for reasoning quantitatively about uncertainty.

Modal logic is a powerful tool for reasoning about computation and uncertainty.

Systems that incorporate both probability and modal operators arise in numerous fields:

- Philosophy
- Game theory
- Computer Science

The goal of this course is to explore and find common themes in different systems of modal probability logic.

Rough outline of course

- Day 1: Background in models of uncertainty: Aumann's agree to disagree, epistemic models, probabilities and measurability, Dutch book argument.
- Day 2: Probabilistic modal logic: probabilistic propositional logic, probabilistic modal logic, expressivity, proof systems, bisimulation, completeness and complexity.
- Day 3: Mixing qualitative and quantitative uncertainty: probabilistic epistemic logic, logics over probabilistic automata, probabilistic bisimulations, proof systems, complexity.
- Day 4: Probabilistic epistemic dynamics: probabilistic dynamic epistemic logic, reduction axioms
- Day 5: Mixed strategies, puzzles, and general remarks: modal logic for mixed strategies, probability puzzles.

What is not covered in this course

- Inductive logic
- First-order probability logic
- Probability-valued logic
- Fuzzy logic
- Defeasible reasoning
- Quantum probability logic

Agreeing to disagree

Theorem (informally stated)

Suppose n agents share a common prior and have different private information (updating using Bayes' Rule). If there is common knowledge in the group of the posterior probabilities, then the posteriors must be equal.

Robert Aumann. Agreeing to disagree. *Annals of Statistics* **4** (1976).

We will proceed with an illustration from

G. Bonanno and K. Nehring. *Agreeing to disagree: a survey*. (unpublished) 1997.

Basic outline of the example

1. Two scientists start off with a common prior over a set of 7 possibilities
2. They perform different experiments that to each scientist reveals the true state to be among a certain subset of possibilities
3. They revise their probabilities (Bayesian updates)
4. They reveal their beliefs about the probability of a certain event
5. They revise again
6. They repeatedly reveal their beliefs about the same event until they agree on the probability of the event.

Example

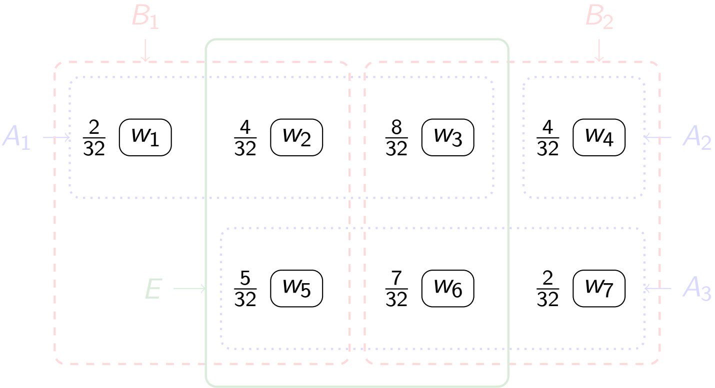

Both scientists share a common prior probability for all seven possibilities.

Example

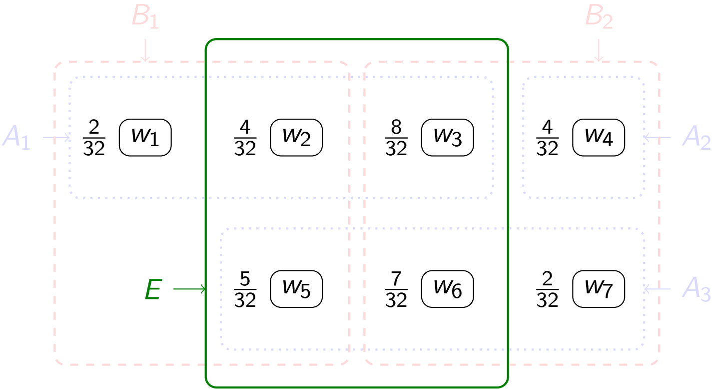

They are particularly interested in the set E and they both agree ($i \in \{1, 2\}$) that

$$\begin{aligned}\mu_i(E) &= \mu_i(w_2) + \mu_i(w_3) + \mu_i(w_5) + \mu_i(w_6) \\ &= \frac{4}{32} + \frac{8}{32} + \frac{7}{32} + \frac{5}{32} = \frac{24}{32}.\end{aligned}$$

Example

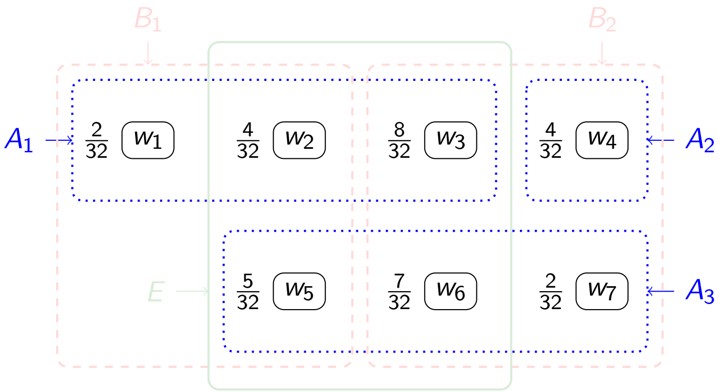

They agree that Scientist 1 performs an experiment given by the blue partition. The result will yield the partition that contains the actual world.

Example

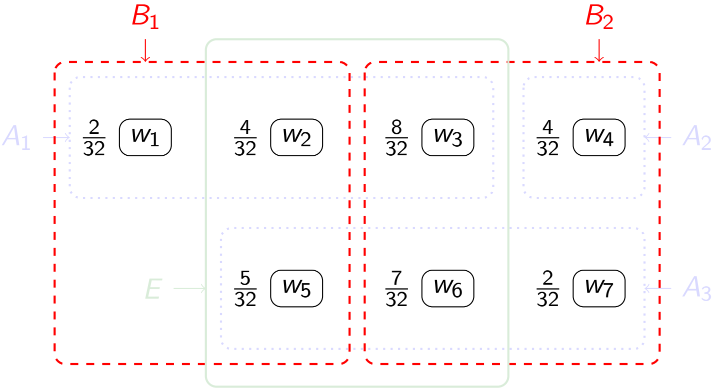

They agree that Scientist 1 performs an experiment given by the red partition.

Example

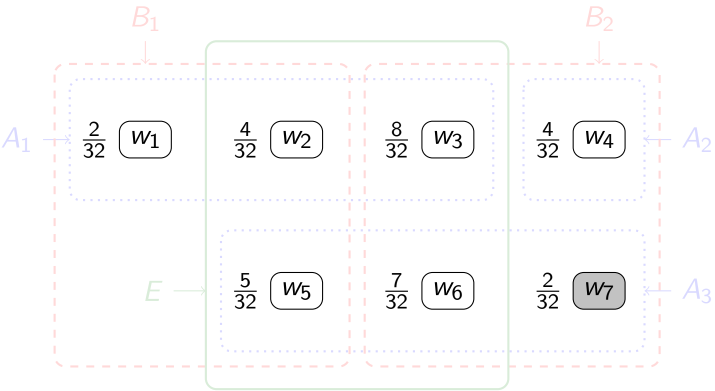

Suppose the actual state is w_7.

Example

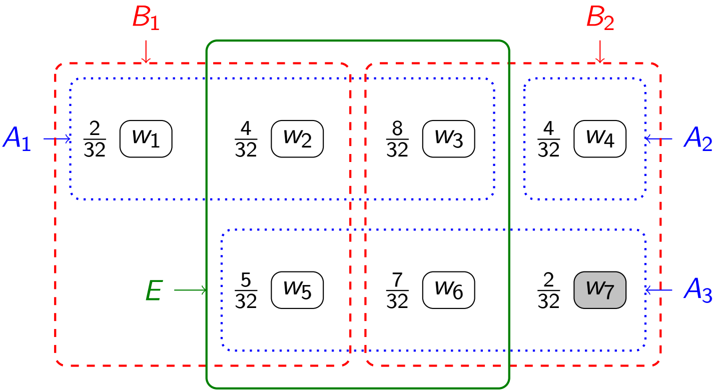

Their experiments yield A_3 and B_2.

$$\mu_1(E \mid A_3) = \frac{12}{14} \qquad \mu_2(E \mid B_2) = \frac{15}{21}$$

Example

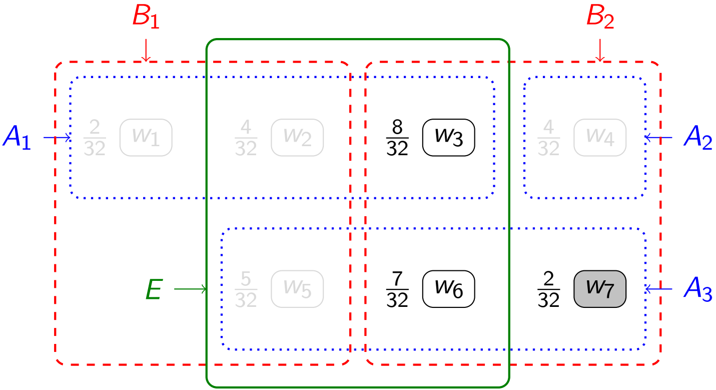

They email each other their new probabilities for E

$$\mu_1(E \mid A_3 \ \&\ \mu_2(E \mid X_2) = 15/21) = \frac{7}{9}$$

$$\mu_2(E \mid B_2 \ \&\ \mu_1(E \mid X_1) = 12/14) = \frac{15}{17}$$

where X_i is the result of i's experiment.

Example

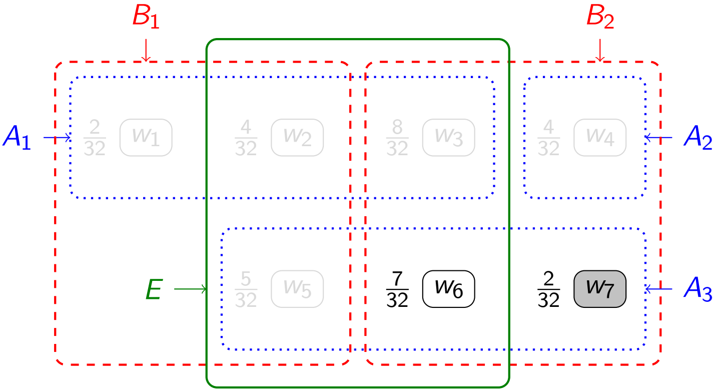

They email each other again their new probabilities for E

$$\mu_1(E \mid A_3 \,\&\, \mu_2(E \mid X_2) = \frac{15}{21} \,\&\, \mu_2(E \mid Y_2) = \frac{15}{17}) = \frac{7}{9}$$

$$\mu_2(E \mid B_2 \,\&\, \mu_1(E \mid X_1) = \frac{12}{14} \,\&\, \mu_1(E \mid Y_1) = \frac{7}{9}) = \frac{7}{9}$$

where Y_i is the information scientist i had just received.

Example

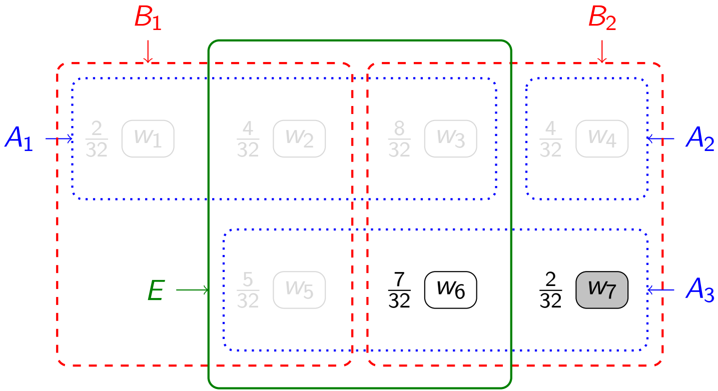

The scientists are now in agreement that the probability of E is $7/9$. Their revisions have converged to common knowledge of what each other believes the probability of E to be. And they agree.

Work related to Agreeing to Disagree

Implications

- No trade theorem:
  P. Milgrom and N. Stokey. Information, trade and common knowledge. *Journal of economic theory* (1982).
- From probabilities of events to aggregates:
  R. McKelvey and T. Page. Common knowledge, consensus, and aggregate information. *Econometrica* (1986).

Deeper understanding of Aumann's theorem

- How the posteriors become common knowledge:
  J. Geanakoplos and H. Polemarchakis. We can't disagree forever. *Journal of economic theory*. (1982).
- Pairwise communication:
  R. Parikh and P. Krasuchi. Communication, consensus, and Knowledge. *Journal of economic theory*. (1990)

(Qualitative) Epistemic language

Let Φ be a set of proposition letters, and Agt a set of agents.
Formulas:

$$\varphi ::= \top \mid p \mid \neg\varphi \mid \varphi \wedge \varphi \mid [i]\varphi$$

where $p \in \Phi$, $i \in Agt$.

- $[i]\varphi$ is read "agents i knows/believes φ
- $\langle i\rangle\varphi \equiv \neg[i]\neg\varphi$ is read "agent i considers φ possible.

Epistemic Models and Semantics

Let Φ be set of proposition letters and Agt a set of agents. An epistemic model is a tuple $M = (W, R, \| \cdot \|)$, where

- W is a set of possible worlds
- R is a collection of relations $R_i \subseteq W^2$ for each $i \in Agt$.
- $\| \cdot \| : \Phi \to \mathcal{P}(W)$.

Define $I_i : \mathcal{P}(W) \to \mathcal{P}(W)$ to be such that $I_i(X) = \{x \in W \mid R_i(x) \subseteq X\}$. The semantics is given by a function $[\![\cdot]\!]$ from formulas to subsets of W.

$$
\begin{array}{lcl}
[\![\top]\!] & = & W \\
[\![p]\!] & = & \|p\| \\
[\![\neg\varphi]\!] & = & W - [\![\varphi]\!] \\
[\![\varphi \wedge \psi]\!] & = & [\![\varphi]\!] \cap [\![\psi]\!] \\
[\![[i]\varphi]\!] & = & I_i([\![\varphi]\!])
\end{array}
$$

Commonly accepted axioms

- (normality) $[i](\varphi \to \psi) \to ([i]\varphi \to [i]\psi)$
- (knowledge; reflexivity) $[i]\varphi \to \varphi$
- (positive introspection; transitivity) $[i]\varphi \to [i][i]\varphi$
- (negative introspection; Euclidean) $\neg[i]\varphi \to [i]\neg[i]\varphi$

Probability space

Definition (Probability space)

A probability space is a tuple $(S, \mathcal{A}, \mu)$, where

- S is a set:
  S is called the "sample space", its elements "outcomes"
- $\mathcal{A} \subseteq \mathcal{P}(S)$ is a σ-algebra:
  a non-empty set of subsets of S closed under complements and countable unions.
- $\mu : \mathcal{A} \to [0, 1]$ is a probability measure:
  a function satisfying
  - $\mu(S) = 1$ (normalize to 1)
  - $\mu(\bigcup_{i=1}^{\infty} A_i) = \sum_{i=1}^{\infty} \mu(A_i)$ for every pairwise disjoint collection of sets $\{A_i\}_{i \in \mathbb{N}}$ in $\mathcal{A}$. (countable additivity)

The sets $A \in \mathcal{A}$ are called "events" or "measurable sets".

Measure space

Definition (Measure space)

A measure space is a tuple $(S, \mathcal{A}, \mu)$, where

- S is a set.
- $\mathcal{A}$ is a σ-algebra.
- $\mu : \mathcal{A} \to [0, \infty]$ is a measure: a function satisfying
  - $\mu(\emptyset) = 0$
  - $\mu(\bigcup_{i=1}^{\infty} A_i) = \sum_{i=1}^{\infty} \mu(A_i)$ for every pairwise disjoint collection of sets $\{A_i\}_{i \in \mathbb{N}}$ in $\mathcal{A}$. (countable additivity)

A probability measure is just a measure normalized to 1.

Definition (Measurable space)

Given any measure space $(S, \mathcal{A}, \mu)$, the pair $(S, \mathcal{A})$ is a measurable space.

Uniform probability distributions

Definition

The uniform probability distribution μ over an interval $[a, b]$ is given by

$$\mu(A) = \int_A \frac{1}{b-a} \, \mathrm{d}\lambda$$

where λ is the Lebesgue measure.

Important properties of Lebesgue measure:

- $\lambda(A) = \lambda(B)$, whenever $B = \{a + t : a \in A\}$ for some $t \in \mathbb{R}$. (translation invariance)
- $\lambda(\bigcup_{i=1}^{\infty} A_i) = \sum_{i=1}^{\infty} \lambda(A_i)$ for every pairwise disjoint collection of sets $\{A_i\}_{i \in \mathbb{N}}$ in $\mathcal{A}$. (countable additivity)

Vitali sets

One motivation for σ-algeras are Vitali sets.

- Define equivalence $\sim$ over $\mathbb{R}$, where $a \sim b$ iff $a - b \in \mathbb{Q}$.
- Denote an equivalence class with representative a by $[a]$; let $\mathcal{E}$ be set of equivalence classes.
- Let $f : \mathcal{E} \to [0,1]$ be any function for which $[f(E)] = E$.
- A Vitali set is

$$V = \{f(E) : E \in \mathcal{E}\}.$$

- Let $V_q = \{x + q : x \in V\}$ for each $q \in [-1,1] \cap \mathbb{Q}$.
- Then $[0,1] \subseteq \bigcup V_q \subseteq [-1,2]$.

Continuity of a set function

Definition (Continuity from above)

A set function $\mu : \mathcal{A} \to [0, \infty]$ is continuous from above if for every non-increasing sequence $A_1 \supseteq A_2 \supseteq \cdots \in \mathcal{A}$,

$$\mu\left(\bigcap_{n=1}^{\infty} A_n\right) = \lim_{n\to\infty} \mu(A_n).$$

We say that μ is continuous at $\emptyset$ if $\lim_{n\to\infty} \mu(A_n) = 0$, whenever $\bigcap A_n = \emptyset$.

Theorem

1. *Any measure is continuous from above*
2. *Any finitely additive set function $\mu : \mathcal{A} \to [0, \infty)$ that is continuous at $\emptyset$ is a measure (is countably additive).*

Proof of part 1

- Let $\{A_n\}$ be a decreasing sequence of sets.
- Let $A = \bigcap A_n$.
- Let $B_n = X - A_n$.
- Let $D_0 = B_1$ and $D_{n+1} = B_{n+1} - B_n$.

$$
\begin{aligned}
& 1 - \mu(A) = \mu(X) - \mu(A) \\
{\color{blue}=} \; & \mu(X - A) = \mu(\bigcup D_n) \\
{\color{red}=} \; & \sum \mu(D_n) = \lim \mu(B_n) \\
{\color{blue}=} \; & \lim(\mu(X) - \mu(A_n)) = 1 - \lim \mu(A_n)
\end{aligned}
$$

${\color{red}=}$ from countable additivity

${\color{blue}=}$ from finite additivity

Proof of part 2

- Let $\{A_n\}$ be a sequence of pairwise disjoint sets.
- Let $A = \bigcup_{i=1}^{\infty} A_i$
- Let $B_n = \bigcup_{i=1}^{n} A_i$
- Then $\bigcap(A - B_n) = \emptyset$
- Then $\lim \mu(A - B_n) = 0$ (by continuity at $\emptyset$)
- Then $\lim(\mu(A) - \mu(B_n)) = 0$ (by finite additivity)
- Then $\lim(\mu(A) - \sum_{i=1}^{n} \mu(A_i)) = 0$ (by finite additivity)
- Then $\lim \sum_{i=1}^{n} \mu(A_i) = \mu(A)$
- Thus $\sum_{i=1}^{\infty} \mu(A_i) = \mu(A) = \mu(\bigcup A_n)$

Product spaces

Definition

Given a family $(X_1, \mathcal{A}_1), \ldots, (X_n, \mathcal{A}_n)$ of measurable spaces, we define the product measurable space to be $(X, \mathcal{A})$, where

- $X = X_1 \times \cdots \times X_n$
- $\mathcal{A}$ is the σ-algebra generated by $\{A_1 \times \cdots \times A_n \mid A_i \in \mathcal{A}_i\}$.

If μ_i is a measure on $(X_i, \mathcal{A}_i)$, for each i, then we define the product measure μ to be

$$\mu(A) = \inf \left\{ \sum_{j=1}^{\infty} \prod_{i=1}^{n} \mu_i(A_i^j) \mid A_i^j \in \mathcal{A}_i \text{ and } A \subseteq \bigcup_{j=1}^{\infty} \prod_{i=1}^{n} A_i^j \right\}.$$

Outer measure

Definition (Outer measure)

Given a set S, an outer measure on S is a function $\mu : \mathcal{P}(S) \to [0, \infty]$, such that

- $\mu(\emptyset) = 0$
- $\mu(A_1) \leq \mu(A_2)$ whenever $A_1 \subseteq A_2$. (monotonicity)
- $\mu(\bigcup_{i=1}^{\infty} A_i) \leq \sum_{i=1}^{\infty} \mu(A_i)$ for every collection of sets $\{A_i\}_{i \in \mathbb{N}}$ in $\mathcal{P}(S)$. (countable subadditivity)

Measurable sets

Definition (μ-measurable sets)

Given an outer measure $\mu : \mathcal{P}(X) \to [0, \infty]$, a set $A \subseteq X$ is *μ-measurable* if for every set $B \subseteq X$

$$\mu(B) = \mu(B \cap A) + \mu(B - A).$$

Proposition

Given an outer measure μ, the set of μ-measurable sets forms a σ-algebra.

Definition (Measurable sets of a measure space)

Given a measure space $(S, \mathcal{A}, \mu)$, the set $\mathcal{A}$ is the set of *measurable sets* of the space.

Note the distinction between the measurable sets of a space and the μ-measurable sets of an outer measure μ.

From outer measure to measure

Given an outer measure $\mathcal{P}(S) \to [0, \infty]$,

- let $\mathcal{A}$ be the set of μ-measurable sets,
- let $\mu' : \mathcal{A} \to [0, \infty]$ such that $\mu'(A) = \mu(A)$ for all $A \in \mathcal{A}$.

Then $(S, \mathcal{A}, \mu')$ is a measure space.

Example: Lebesgue measure

First define the Lebesgue outer measure $\lambda^* : \mathcal{P}(\mathbb{R}^n) \to [0, \infty]$ by

$$\lambda^*(E) = \inf \left\{ \sum_{j=1}^{\infty} \prod_{i=1}^{n} (b_i^j - a_i^j) \mid E \subseteq \bigcup_{j=1}^{\infty} \prod_{i=1}^{n} [a_i^j, b_i^j] \right\}.$$

We define the Lebesgue measure λ to be the restriction of λ^* to the λ^*-measurable sets.

From measure to outer measure

Definition (Outer Measure Extension of a Measure)

If $\mathcal{A}$ is a σ-algebra over S and $\mu : \mathcal{A} \to [0, \infty]$ is a measure, then the *outer measure extension of* μ is defined to be $\mu^* : \mathcal{P}(S) \to [0, \infty]$ given by

$$\mu^*(A) = \inf\{\mu(B) : B \in \mathcal{A}, A \subseteq B\}$$

Proposition

Let $\mathcal{A}$ be a σ-algebra over S, let $\mu : \mathcal{A} \to [0, \infty]$ be a measure, and let $\mu^* : \mathcal{P}(S) \to [0, \infty]$ be the outer measure extension of μ. Then

- $\mu^*(A) = \mu(A)$ *for every* $A \in \mathcal{A}$ *(*μ^* *does indeed extend* μ*).*
- μ^* *is an outer measure.*
- *If* $\mathcal{A}'$ *consists of the* μ^**-measurable sets, then* $\mathcal{A} \subseteq \mathcal{A}'$*.*

Zero-one laws

Theorem (Kolmogorov's zero-one law)

Given a probability space $(X, \mathcal{A}, \mu)$ and a sequence $\mathcal{A}_k$ of independent σ-algebras ($A_i \in \mathcal{A}_i$ for each i implies $\mu(\bigcap_{i=1}^{\infty} A_i) = \prod_{i=1}^{\infty} \mu(A_i)$), each contained in $\mathcal{A}$, if for each k, B_k is the smallest σ algebra containing $\bigcup_{i=k}^{\infty} \mathcal{A}_i$, then for any $F \in \bigcap_{k=0}^{\infty} B_k$, $\mu(F)$ is either 0 or 1.

Theorem

The probability that a given first-order formula is true on finite models goes to either 1 or 0 with increasing domain size (and it is decidable which case).

R. Fagin. Probabilities on Finite Models. *Journal of Symbolic Logic* (1976).

Dutch Book and subjective probabilities

- Most of the probabilities in this course are subjective probabilities: probabilities agents assign to the likelihood of certain events.
- Subjective probabilities are often viewed in terms of an agent's willingness to bet.
- The Syncronic Dutch Book literature provides justification for the laws of probability using betting games.
- The Diachronic Dutch Book literature provides justification for Bayesian updating as a means for changing subjective probabilities using betting games.

Strategies for Synchronic Dutch Book

Let $(\Omega, \mathcal{A})$ be a finite measurable space. Consider three players α, β, and η

- α's strategy ("a system of beliefs") is a function $\mu : \mathcal{A} \to \mathbb{R}$
  View $\mu(A)$ as a price for a unit wager for event A
- β's strategy ("a system of bets") is a function $\nu : \mathcal{A} \to \mathbb{R}$
  View $\nu(A)$ as being the quantity β buys of unit wagers for event A
- η's strategy ("the actual outcome") is an $\omega \in \Omega$.

Payoffs

Fix a strategy profile (μ, ν, ω).

- α's payoff is

$$\sum_{\{A \in \mathcal{A} | \omega \notin A\}} \mu(A)\nu(A) + \sum_{\{A \in \mathcal{A} | \omega \in A\}} (\mu(A) - 1)\nu(A)$$

- β's payoff is

$$\sum_{\{A \in \mathcal{A} | \omega \notin A\}} -\mu(A)\nu(A) + \sum_{\{A \in \mathcal{A} | \omega \in A\}} (1 - \mu(A))\nu(A)$$

- η's payoff is 0 regardless of the strategies played

Synchronic Dutch Book

Definition (Dutch Book)

β's strategy is a Dutch Book with respect to α's strategy if regardless of η's strategy, β will receive a positive payoff.

Theorem (Synchronic Dutch Book Theorem)

If α's strategy μ is not a probability measure, then β has a strategy that is a Dutch book with respect to α's strategy.

Theorem (Converse Synchronic Dutch Book Theorem)

If α's strategy μ is a probability measure, then β has no strategy that is a Dutch book with respect to α's strategy.

Proof of Dutch Book Theorem

Possible violations of the laws of probability:

- $\mu(A) < 0$ for some A
- $\mu(S) > 1$
- $\mu(A \cap B) + \mu(A \cap \overline{B}) > \mu(A)$
- $\mu(A \cap B) + \mu(A \cap \overline{B}) < \mu(A)$

Proof continued

If $\mu(A) < 0$ for some A,

- $\nu(B) = 0$ for all $B \neq A$
- $\nu(A) = a$ for any positive number a (β buys a quantity of a unit wagers)
- This guarantees β at least $a|\mu(A)|$ (and at most $a|\mu(A)| + a$).

If $\mu(S) > 1$ for some A,

- $\nu(A) = 0$ for all $A \neq S$
- $\nu(S) = -a$ for any positive number a (β sells a quantity of a unit wagers)
- this guarantees β at least $a\mu(S)$ (and at most $a\mu(S) + a$).

Proof continued

If $\mu(A \cap B) + \mu(A \cap \overline{B}) > \mu(A)$,

- $\nu(A \cap B) = -1$
- $\nu(A \cap \overline{B}) = -1$
- $\nu(A) = 1$

Then

- $\omega \in A \cap B$ implies β's payoff is $\mu(A \cap B) - 1 + \mu(A \cap \overline{B}) - \mu(A) + 1 > 0$
- $\omega \in A \cap \overline{B}$ implies β's payoff is $\mu(A \cap B) + \mu(A \cap \overline{B}) - 1 - \mu(A) + 1 > 0$
- $\omega \in \overline{A}$ implies β's payoff is $\mu(A \cap B) + \mu(A \cap \overline{B}) - 1 - \mu(A) + 1 > 0$

Proof continued

If $\mu(A \cap B) + \mu(A \cap \overline{B}) < \mu(A)$,

- $\nu(A \cap B) = 1$
- $\nu(A \cap \overline{B}) = 1$
- $\nu(A) = -1$

The proof is the same as for $\mu(A \cap B) + \mu(A \cap \overline{B}) > \mu(A)$, but with every sign reversed.

Proof of converse Dutch Book Theorem

Let $\mathcal{B} = \{B_1, \ldots, B_n\}$ be the finest partition of Ω (finite) for which each $B_i \in \mathcal{A}$.

Fix a probability measure μ for α. Given ν, let ν' be given by

$$\nu'(A) = \begin{cases} \sum_{\{A' \mid A \subseteq A'\}} \nu(A') & A \in \mathcal{B} \\ 0 & \text{otherwise} \end{cases}$$

If $\omega \in B_\omega \in \mathcal{B}$, then β's payoff when ν is played:

$$\begin{aligned} & -\sum_{A \in \mathcal{A}} \mu(A)\nu(A) + \sum_{\{A \in \mathcal{A} \mid \omega \in A\}} \nu(A) \\ = & -\sum_{A \in \mathcal{A}} \sum_{\{B \in \mathcal{B} \mid B \subseteq A\}} \mu(B)\nu(A) + \sum_{\{A \in \mathcal{A} \mid B_\omega \subseteq A\}} \nu(A) \\ = & -\sum_{B \in \mathcal{B}} \sum_{\{A \in \mathcal{A} \mid B \subseteq A\}} \mu(B)\nu(A) + \sum_{\{A \in \mathcal{A} \mid B_\omega \subseteq A\}} \nu(A) \\ = & -\sum_{B \in \mathcal{B}} \mu(B)\nu'(B) + \nu'(B_\omega) \end{aligned}$$

which is β's payoff when ν' is played.

Proof continued

If $\nu'(B_i) \geq 0$ for all i, let $\nu'(B_M) = \max\{\nu'(B_i)\}$ and $\nu'(B_m) = \min\{\nu'(B_i)\}$. Then if $\omega \in B_M$, β's payoff is

$$\nu'(B_M) - \sum_{B \in \mathcal{B}} \mu(B)\nu'(B)$$

$$\geq \nu'(B_M) - \sum_{B \in \mathcal{B}} \mu(B)\nu'(B_M)$$

$$= \nu'(B_M) - \nu'(B_M) \sum_{B \in \mathcal{B}} \mu(B) = 0$$

and if $\omega \in B_m$, β's payoff

$$\nu'(B_m) - \sum_{B \in \mathcal{B}} \mu(B)\nu'(B)$$

$$\leq \nu'(B_m) - \sum_{B \in \mathcal{B}} \mu(B)\nu'(B_m) = 0$$

If $\nu'(B_i) \leq 0$ for all i, use the same reasoning as the case where $\nu'(B_i) \geq 0$.

Proof continued

If $\nu'(B_i) > 0$ for some i and $\nu'(B_i) < 0$ for some i, let

- $\nu'(B_M) = \max\{\nu'(B_i) > 0\}$
- $\nu'(B_N) = \max\{-\nu'(B_i) \mid \nu'(B_i) < 0\}$.

Using the same reasoning as for the cases with $\nu'(B_i) \geq 0$ for all B_i or $\nu'(B_i) \leq 0$ for all B_i,

- if $\omega \in B_M$, β's payoff is at least 0
- if $\omega \in B_N$, β's payoff is at most 0.

Strategies for Diachronic Dutch Book

Let

- $(\Omega, \mathcal{A})$ be a finite measurable space,
- $D_1, \ldots, D_n \in \mathcal{A}$ partition Ω,
- $\mathcal{A}_i = \{A \cap D_i : A \in \mathcal{A}\}$ for each i.

Consider three players α, β, and η

- α's strategy ("a system of beliefs") is a probability measure $\mu : \mathcal{A} \to \mathbb{R}$ for which $\mu(D_i) \neq 0$ for $1 \leq i \leq n$, together with probability measures $\{\mu_i : \mathcal{A}_1 \to \mathbb{R}\}_{i=1}^n$
- β's strategy ("a system of bets") is a function $\nu : \mathcal{A} \to \mathbb{R}$ together with functions $\{\nu_i : \mathcal{A}_1 \to \mathbb{R}\}_{i=1}^n$,
- η's strategy ("the actual outcome") is an $\omega \in \Omega$.

Payoffs

Fix a strategy profile $(\{\mu, \mu_1, \ldots, \mu_n\}, \{\nu, \nu_1, \ldots, \nu_n\}, \omega)$.
Define the function π by

$$\pi(\mathcal{X}, \mu, \nu) = \sum_{\{A \in \mathcal{X} | \omega \notin A\}} \mu(A)\nu(A) + \sum_{\{A \in \mathcal{X} | \omega \in A\}} (\mu(A) - 1)\nu(A)$$

- α's payoff is $\pi(\mathcal{A}, \mu, \nu) + \sum_{\{i | \omega \in D_i,\}} \pi(\mathcal{A}_i, \mu_i, \nu_i)$
- β's payoff is $-\pi(\mathcal{A}, \mu, \nu) - \sum_{\{i | \omega \in D_i\}} \pi(\mathcal{A}_i, \mu_i, \nu_i)$
- η's payoff is 0 regardless of the strategies played

Diachronic Dutch Book

Definition (Dutch Book)

β's strategy is a Dutch Book with respect to α's strategy if regardless of η's strategy, β will receive a positive payoff.

Theorem (Diachronic Dutch Book Theorem)

If there is an i and an $A \in \mathcal{A}$ such that $\mu_i(A \cap D_i) \neq \mu(A \cap D_i)/\mu(D_i)$, then β has a strategy that is a Dutch book with respect to α's strategy.

Theorem (Converse Diachronic Dutch Book Theorem)

If there for all i and $A \in \mathcal{A}$, it is the case that $\mu_i(A \cap D_i) = \mu(A \cap D_i)/\mu(D_i)$, then β has no strategy that is a Dutch book with respect to α's strategy.

Proof of Diachronic Dutch Book Theorem

Suppose $\mu_i(A \cap D_i) > \mu(A \cap D_i)/\mu(D_i)$. Let

- $\nu(D_i) = -\mu_i(A \cap D_i)$
- $\nu(A \cap D_i) = 1$
- $\nu_i(A \cap D_i) = -1$

Then

- $\omega \notin D_i$ implies β's payoff is $\mu(D_i)\mu_i(A \cap D_i) - \mu(A \cap D_i) > \mu(D_i)\mu(A \cap D_i)/\mu(D_i) - \mu(A \cap D_i) = 0$
- $\omega \in \overline{A} \cap D_i$ implies β's payoff is $\mu(D_i)\mu_i(A \cap D_i) - \mu_i(A \cap D_i) - \mu(A \cap D_i) + \mu_i(A \cap D_i) = \mu(D_i)\mu_i(A \cap D_i) - \mu(A \cap D_i) > 0$
- $\omega \in A \cap D_i$ implies β's payoff is the same as with $\omega \in \overline{A} \cap D_i$. Two new stakes must be payed (1 and -1).

Proof of Diachronic Dutch Book Theorem cont'd

Suppose $\mu_i(A \cap D_i) < \mu(A \cap D_i)/\mu(D_i)$. Let

- $\nu(D_i) = \mu_i(A \cap D_i)$
- $\nu(A \cap D_i) = -1$
- $\nu_i(A \cap D_i) = 1$

Then every term in the previous slide is negated, and the inequalities are not reversed.

Proof of Converse Dutch Book Theorem

Suppose $\mu, \mu_1, \ldots, \mu_n$ is α's strategy for which μ and each μ_i is a probability measure, and $\mu_i(A \cap D_i) = \mu(A \cap D_i)/D_i$ for each i and A. Given any strategy $\nu, \nu_1, \ldots, \nu_n$ for β, let $\nu_{i,D_i}(A) = 0$ for each $A \in \mathcal{A}$. For each $B \in \mathcal{A}_i$ with $B \neq D_i$ and $A \in \mathcal{A}$, let

$$\nu_{i,B}(A) = \begin{cases} \nu_i(B) & A = B \\ -\mu(B)\nu_i(B)/\mu(D_i) & A = D_i \\ 0 & \text{otherwise} \end{cases}$$

Let $\nu_i'(A) = 0$ for each i and A, and

$$\nu'(A) = \nu(A) + \sum_{i=1}^{n} \sum_{B \in \mathcal{A}_i} \nu_{i,B}(A)$$

Then the payoffs are the same if we replace β's strategy with $\nu', \nu_1', \ldots, \nu_n'$. As $\nu_i' = 0$, we can appeal to the converse synchronic Dutch Book theorem.