Bayesianism and Reliability

Models and Simulations in Philosophy
April 14th, 2014

Review

Last Class:

- Difference between individual and social epistemology
- Why simulations are particularly useful for social epistemology and philosophy of science
- Intro. to Agent-Based Models (ABMs)

Today's Class

Today: Two theories of what is rational for an individual to believe

- Bayesianism
- Logical Reliability

Question: Why discuss individual epistemology given that the focus of this course is social epistemology and philosophy of science?

Answer: To

- Model how group members do or ought to change their beliefs.
- Develop criteria of group rationality.

Outline

Bayesianism

Bayesianism is the conjunction of two theses:

- Beliefs = Probabilities
- Beliefs updated by conditionalization

The Sample Space

In probability, there is a set Ω called the sample space that represents all the possible outcomes of an experiment.

For instance:

- Experiment 1: Roll a die.
- $\Omega = \{1, 2, 3, 4, 5, 6\}$.
- Experiment 2: Flip a coin twice.
- $\Omega = \{\langle H, H\rangle, \langle T, T\rangle, \langle H, T\rangle, \langle T, H\rangle\}$

Events

Subsets of the sample space are called events.

Events

Experiment 1: Roll a die

Example Event: The event the die lands on an odd number is $A = \{1, 3, 5\}$.

Events

Experiment 2: Flip a coin twice

Example Event: The event the coin lands heads exactly once is:

$$A = \{\langle H, T\rangle, \langle T, H\rangle\}$$

Impossible Event

In both experiments, $\emptyset$ represents the "impossible" event:

- Experiment 1: The die does not land on any face.
- Experiment 2: The coin lands on neither heads nor tails.

Probability Axioms

A probability measure is a function that assigns every event A some number $P(A)$ between 0 and 1 (inclusive) such that

- $P(\Omega) = 1$, where Ω is the entire sample space, and
- **Finite Additivity:** If $A \cap B = \emptyset$, then $P(A \cup B) = P(A) + P(B)$.

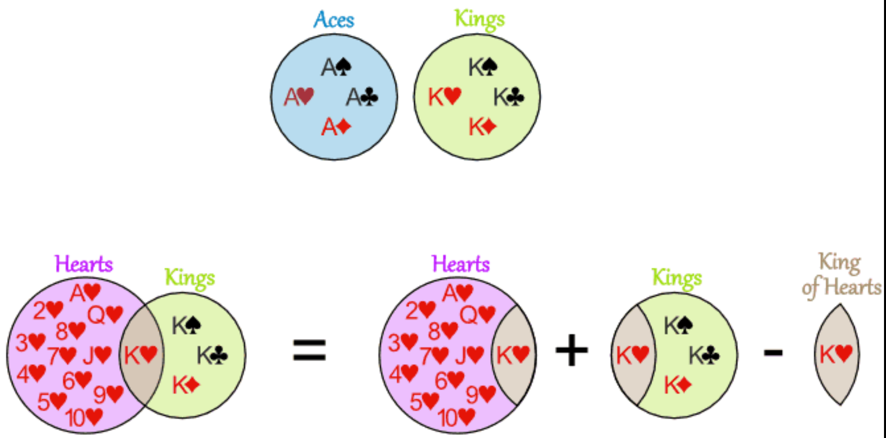

Probability Axioms

Scientists and statisticians often assume

Countable Additivity: If $E_1, E_2, \ldots$ is a countable sequence of events and $E_i \cap E_j = \emptyset$ whenever $i \neq j$, then

$$P(\bigcup_n E_n) = \sum_n P(E_n).$$

Belief and Probability

The first thesis of Bayesianism is that degrees of belief are representable by probabilities.

Properties of Probability

Probabilities are always comparable.

- Numbers can be compared: either $x > y$, or $y > x$, or $x = y$.
- Are degrees of belief always comparable? E.g., What's more likely: (i) aliens invade Earth before 2200 or (ii) monkeys increase in intelligence and become our overlords within the next 3000 years?

Properties of Probability

Probabilities are always quantitatively comparable.

For example, .642 is 32.1 times greater than .02.

Quantitative Comparisons

Tremendously Awful

Really awful (but better?)

Properties of Probability

Probabilities are always quantitatively comparable.

- Comparisons of numbers can be **quantified**. E.g. .642 is 32.1 times greater than .02.
- Compare: Clearly, Nickelback is worse than Creed. Are there 32.1 times as bad?
- Can assessments of degrees of belief be made so precise?

Properties of Probability

Probabilities are bounded (between 0 and 1)

- Compare: The length of objects can be assigned a number, but there is no maximum length.
- Why is it that there is some maximum degree of certainty? And some minimum degree of confidence?

Properties of Probability

Probabilities are additive.

- Compare: Ice cream is delicious. So is sausage. Is the deliciousness of sausage with a side of ice cream the sum of the degrees to which others are delicious?
- Moreover, why didn't you multiply the two numbers? What's so special about addition?

Properties of Probability

To answer these questions, let's consider how we might measure beliefs . . .

Measuring Probability

How can degrees of belief be measured?

Idea: Betting behavior (with small sums)

Odds

- On April 16th, the "football" team FC Bayern will play Kaiserlautern.
- Several bookies are offering around 3 : 2 odds on Kaiserlautern.

Odds

That is, you pay \$2 for a bet such that

- You win \$3 if Kaiserlautern wins.
- You get nothing otherwise.

Odds

A bookie who offers you such a bet clearly thinks that FC Bayern will likely beat Kaiserlautern.

Can we quantify how much so?

Fair Odds

Say the bookie considers her odds $a : b$ on an event E to be fair if she is willing to **both** sell and buy the bet $a : b$ on E.

Odds and Degrees of Belief

- Define the bookie's degree of belief $P(E)$ in an event E to be $\frac{b}{b+a}$ where $a : b$ are her fair betting odds for E.
- Notice that her degrees of belief are numbers between 0 and 1.

Odds and Degrees of Belief

Question: Would it be prudent for the bookie's degrees of beliefs to satisfy the remaining probability axioms?

Dutch Book

- Let E be the event that FC Bayern wins its next match
- Let E^c the event they don't.
- Suppose the bookie posts $2:1$ odds on E and $2:1$ odds on E^c.

Dutch Book

Notice the bookies degrees of belief are not probabilities as

$$P(E) = P(E^c) = \frac{1}{2+1} = \frac{1}{3}$$

and hence,

$$P(\Omega) = P(E \cup E^c) = P(E) + P(E^c) = \frac{2}{3}.$$

Dutch Book

- You are very wise and decide to buy both bets from the bookie.
- So you pay \$1 for a bet that pays \$2 if FC Bayern wins, and
- You pay \$1 for a bet that pays \$2 if FC Bayern does not win.
- Regardless of what happens, you win \$1, i.e., you are **guaranteed** to take money from the bookie.

Dutch Book Theorem

Theorem (Dutch Book Theorem)

There is no series of bets against the bookie that ensures that she loses money for sure if and only if her degrees of belief obey the probability axioms.

Dutch Book Theorem and Bayesianism

The Dutch book theorem is one (of several) arguments for the Bayesian's first thesis:

Beliefs = Probabilities

Bayesianism

Bayesianism is the conjunction of two theses:

- Beliefs = Probabilities
- Update by conditionalization

Updating Degrees of Belief

- Suppose your degree of belief at time t that my sibling is male is $P_t(M)$.
- You learn at time $t+1$ that my sibling's name is "Evan."
- What should your degrees of belief be then?

Conditionalization

Conditionalization is the thesis that upon learning E, you ought to revise your degrees of belief in M as follows:

$$P_{t+1}(M) = P_t(M|E) := \frac{P_t(M\&E)}{P_t(E)}$$

Conditionalization seems most plausible when P_t are **frequencies**.

Diachronic Dutch Books

There are Dutch Book theorems for conditionalization as well, but we'll skip them.

Normative Arguments

Together, these theorems seem to provide a normative argument for the theses:

- Your degrees of belief ought to act like probabilities, and
- Your ought to update your degrees of belief by conditionalization.

This is what Strevens [2006] calls the "*a priori*" justification for Bayesianism.

Descriptive Models

Sometimes Bayesianism is also descriptive of how humans in fact reason [Gopnik and Wellman, 2012].

Bayesianism in this Course

Question: Why will we care about Bayesianism in this course?

Answer:

- Many of the models we consider will assume group members have probabilistic beliefs.
- Some assume they update those beliefs by conditionalization.

Outline

Bayesianism and Truth

- Notice the arguments for Bayesianism had the following structure:
- If you want to avoid sure loss, then
  - Beliefs = Probabilities
  - Update = Conditionalization
- In particular, neither argument showed a relationship between your beliefs and what is true.

Truth and Rationality

> *I argue that epistemic states and changes of such states as well as the rationality criteria governing epistemic dynamics can be, and should be, formulated independently of the factual connections between the epistemic inputs and the outer world. One consequence of this position is that the concept of truth is irrelevant ...*

Gaerdenfors [1988], page 9.

Bayesianism and Truth

There are theorems that connect Bayesian updating with developing true beliefs, but they may not be as widely applicable as many imagine [Belot, 2013].

Logical Reliability Defined

Kelly [1996] argues that the method we use to update our beliefs ought be logically reliable, i.e., it is "guaranteed to converge in every possible circumstance consistent with . . . background assumptions."

Logical Reliability Defined

> *The logical reliabilist conceives of inductive problems the way a computer scientist conceives of computational problems. A solution to a computational problem (an algorithm) is supposed to be guaranteed by its mathematical structure to output a correct answer on every possible input. The logical perspective stands in sharp contrast with the received view among inductive methodologists, who are often more interested in whether a belief is justified by evidence . . . Logical reliabilism simply demands of inductive methods what is routinely required of algorithms.*

Kelly [1996], page 4.

Logical Reliability in this Course

Question: Why will we care about logical reliability in this course?

Answer: Many of the models we consider will evaluate the rationality of an organizational structure by asking whether all group members' beliefs tend towards the truth.

Bounded Rationality

Kelly [1996] also discusses another important issue in this course: bounded rationality.

For Kelly, one should not prescribe methods for updating beliefs that cannot not, in principle, be carried out by a computer.

Bounded Rationality

> *Traditional methodological analysis has centered on quixotic "ideal agents" who have divine cognitive powers ... This idealization is harmless only so long as what is good for ideal agents is good for their more benighted brethren. It will be seen, however, that rules the seem harmless and inevitable for ideal agents are in fact disastrous for computationally bounded agents. A prime example is the simple requirement that a hypothesis be dropped as soon as it is inconsistent with the data, a rule endorse by almost all inductive methodologies [e.g., Bayesianism]. It turns out the there are problems for which computable methods can be consistent in this sense, and some computable method is a reliable solution, but no computable, consistent method is a reliable solution.*

Kelly [1996], page 6.

Logical Reliability in this Course

Question: Why will we care about bounded rationality in this course?

Answer:

- Some of the models we consider will evaluate ideal agents. We should carefully consider whether the models apply to non-ideal ones.
- Other models consider boundedly rational agents. We should consider whether agents' limitations are realistic and normatively defensible.

Structure of the Course

Upcoming Weeks:

- Three Units: Disagreement, Diversity, and Testimony
- Each unit has two parts:
  1. An introduction to a "traditional" problem in epistemology of philosophy of science
  2. Analysis of computer models aimed at answering said question.

Structure of the Course

Course Website:

- Printable slides
- Sample code from class

NetLogo

Today we will discuss:

- Data types in NetLogo
- Declaring and modifying variables
- Arithmetic, Boolean, and List Operations
- Global vs. Local Variables

References I

Belot, G. (2013). Bayesian orgulity. *Philosophy of Science*, 80(4):483–503.

Gaerdenfors, P. (1988). *Knowledge in flux: Modeling the dynamics of epistemic states.* MIT press.

Gopnik, A. and Wellman, H. M. (2012). Reconstructing constructivism: Causal models, bayesian learning mechanisms, and the theory theory. *Psychological bulletin*, 138(6):1085—1108.

Kelly, K. T. (1996). *The logic of reliable inquiry*. Oxford University Press, New York.

Strevens, M. (2006). Notes on bayesian confirmation theory.