Logic for probability, belief, and change.

Joshua Sack

Sep 15, 2010

Monty Hall Puzzle (Probability Dynamics)

A game show host presents to you three closed doors and says that behind one of them is a prize, and behind the other two there is nothing. He then does the following:

1. He lets you pick a door.
2. He then opens a door to an empty room that you did not select.
3. He then asks if you would like to switch.

Question: Should you switch?

Monty Hall Puzzle (Probability Dynamics)

A game show host presents to you three closed doors and says that behind one of them is a prize, and behind the other two there is nothing. He then does the following:

1. He lets you pick a door.
2. He then opens a door to an empty room that you did not select.
3. He then asks if you would like to switch.

Question: Should you switch?

Monty Hall Puzzle (Probability Dynamics)

A game show host presents to you three closed doors and says that behind one of them is a prize, and behind the other two there is nothing. He then does the following:

1. He lets you pick a door.
2. He then opens a door to an empty room that you did not select.
3. He then asks if you would like to switch.

Question: Should you switch?

Monty Hall Puzzle (Probability Dynamics)

A game show host presents to you three closed doors and says that behind one of them is a prize, and behind the other two there is nothing. He then does the following:

1. He lets you pick a door.
2. He then opens a door to an empty room that you did not select.
3. He then asks if you would like to switch.

Question: Should you switch?

Monty Hall Puzzle (Probability Dynamics)

A game show host presents to you three closed doors and says that behind one of them is a prize, and behind the other two there is nothing. He then does the following:

1. He lets you pick a door.
2. He then opens a door to an empty room that you did not select.
3. He then asks if you would like to switch.

Question: Should you switch?

What is Probability Exactly?

Definition

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

1. Ω is a set called the sample space.
2. $\mathcal{A} \subseteq \mathcal{P}(\Omega)$ is a σ-algebra: a set of subsets of Ω containing $\emptyset$, which is closed under complements and countable unions and intersections.
3. $\mu : \mathcal{A} \to [0, 1]$ is a probability measure, that is
   - $\mu(\Omega) = 1$ and $\mu(\emptyset) = 0$
   - If $\{A_1, A_2, \ldots\}$ is a countable set of pairwise disjoint elements of $\mathcal{A}$, then $\mu(\bigcup_{j=1}^{\infty} A_j) = \sum_{j=1}^{\infty} \mu(A_j)$. (Countable additivity)

$(\Omega, \mathcal{A})$ is a called a *measurable space*. Sets in the σ-algebra $\mathcal{A}$ are called *measurable*.

Probability Model

Definition (Probability Model)

A probability model is a tuple $(\mathbf{P}, \|\cdot\|)$, where

- $\mathbf{P} = (\Omega, \mathcal{A}, \mu)$ is a probability space.
- $\|\cdot\|$ is a function from Φ to the set $\mathcal{A}$.

Probability logic is generated by the rule:

$$\varphi ::= p \mid \neg\varphi_1 \mid \varphi_1 \wedge \varphi_2 \mid [\geq a]\varphi \mid [> a]\varphi \mid [= a]\varphi$$

where $a \in \mathbb{Q}$.

Semantics (selected components)

$$\begin{aligned} [\![p]\!] &= \|p\| \\ [\![[\geq a]\varphi]\!] &= \begin{cases} \Omega & \mu([\![\varphi]\!]) \geq a \\ \emptyset & \text{otherwise} \end{cases} \end{aligned}$$

Suppose you select door A...

Let

$$\begin{aligned}\Phi &= \{(x, y) \mid x \in \{A, B, C\}, y \in \{B, C\}, x \neq y\} \\ &= \{(A, B), (A, C), (B, C), (C, B)\}\end{aligned}$$

where

- x represents the door that opens to a prize, and
- y represents the door the host will open.

Define

- $isA \equiv \bigvee\{(x, y) \in \Phi \mid x = A\}$, and similarly for isB and isC.
- $openB \equiv \bigvee\{(x, y) \in \Phi \mid y = B\}$, and similarly for $openC$.

Then

- $[\![[= \frac{1}{3}] isA]\!] = [\![[= \frac{1}{3}] isB]\!] = [\![[= \frac{1}{3}] isC]\!] = \Omega$ that is $\mu([\![(A, B) \vee (A, C)]\!]) = \mu([\![(B, C)]\!]) = \mu([\![(C, B)]\!]) = \frac{1}{3}$.
- $[\![[= \frac{1}{2}] openB]\!] = [\![[= \frac{1}{2}] openC]\!] = \Omega$.

probability distribution

We end up with the following probabilities:

	$x = A$	$x = B$	$x = C$
$y = B$	$\mu([\![(A, B)]\!]) = 1/6$		$\mu([\![C, B]\!]) = 2/6$
$y = C$	$\mu([\![(A, C)]\!]) = 1/6$	$\mu([\![B, C]\!]) = 2/6$	

Probability Dynamics

Start with a probability model $M = ((\Omega, \mathcal{A}, \mu), \| \cdot \|)$ where $\mu(A) \neq 0$ for every $A \in \mathcal{A}$.

- Input: A subset $Y \subseteq \Omega$
- Update: $M[Y] = ((Y, \mathcal{A}^{M[Y]}, \mu^{M[Y]}, \| \cdot \|^{M[Y]})$, where
  - $\mathcal{A}^{M[Y]} = \{A \cap Y \mid A \in \mathcal{A}\}$
  - $\mu^{M[Y]} = \mu(\cdot \mid Y)$ maps $B \in \mathcal{A} \cap Y$ to $\mu(B)/\mu(Y)$.
  - $\|p\|^{M[Y]} = \|p\| \cap Y$.

Add formulas of the form $[!\varphi]\psi$ to the language:

$$[\![[!\varphi]\psi]\!] = [\![\neg\varphi]\!] \cup [\![\psi]\!]_{M[[\![\varphi]\!]]}$$

Using dynamics

Recall

	$x = A$	$x = B$	$x = C$
$y = B$	$\mu([\![(A,B)]\!]) = 1/6$		$\mu([\![C,B]\!]) = 2/6$
$y = C$	$\mu([\![(A,C)]\!]) = 1/6$	$\mu([\![B,C]\!]) = 2/6$	

Note that

$$\mu^{M[\![openB]\!]}([\![isC]\!]) = \frac{\mu([\![isC]\!])}{\mu([\![openB]\!])} = \frac{\mu([\![(C,B)]\!])}{\mu([\![(A,B)]\!] \cup [\![(C,B)]\!])} = \frac{\frac{2}{6}}{\frac{3}{6}} = \frac{2}{3}$$

Hence

$$[\![[!openB][=2/3]isC]\!] = [\![\neg openB]\!] \cup [\![[=2/3]isC]\!]_{M[\![openB]\!]}$$
$$= [\![\neg openB]\!] \cup [\![openB]\!] = \Omega.$$

Probability vs Epistemics

Probabilities may offer more detail than qualitative epistemics, but numerical details might not be appropriate. Should we:

- assign a probability to a computer outputting a bit 1?
- assign a probability to a coin flip resulting in heads?

We may wish to involve *both* qualitative and quantitive uncertainty.

One of many options

One may define

$$[\![\text{Knows}\,\varphi]\!] = [\![[=1]\varphi]\!]$$
$$[\![\text{Poss}\,\varphi]\!] = [\![[>0]\varphi]\!]$$

The next example shows the utility of more flexibility in the relationship between epistemics and probability.

Probability vs Epistemics

Probabilities may offer more detail than qualitative epistemics, but numerical details might not be appropriate. Should we:

- assign a probability to a computer outputing a bit 1?
- assign a probability to a coin flip resulting in heads?

We may wish to involve *both* qualitative and quantitive uncertainty.

One of many options

One may define

$$[\![\text{Knows}\,\varphi]\!] = [\![[=1]\varphi]\!]$$
$$[\![\text{Poss}\,\varphi]\!] = [\![[>0]\varphi]\!]$$

The next example shows the utility of more flexibility in the relationship between epistemics and probability.

Probability vs Epistemics

Probabilities may offer more detail than qualitative epistemics, but numerical details might not be appropriate. Should we:

- assign a probability to a computer outputting a bit 1?
- assign a probability to a coin flip resulting in heads?

We may wish to involve *both* qualitative and quantitive uncertainty.

One of many options

One may define

$$\llbracket \mathsf{Knows}\,\varphi \rrbracket = \llbracket [=1]\varphi \rrbracket$$
$$\llbracket \mathsf{Poss}\,\varphi \rrbracket = \llbracket [>0]\varphi \rrbracket$$

The next example shows the utility of more flexibility in the relationship between epistemics and probability.

Probability vs Epistemics

Probabilities may offer more detail than qualitative epistemics, but numerical details might not be appropriate. Should we:

- assign a probability to a computer outputing a bit 1?
- assign a probability to a coin flip resulting in heads?

We may wish to involve *both* qualitative and quantitive uncertainty.

One of many options

One may define

$$
\begin{aligned}
[\![\mathsf{Knows}\,\varphi]\!] &= [\![[=1]\varphi]\!] \\
[\![\mathsf{Poss}\,\varphi]\!] &= [\![[>0]\varphi]\!]
\end{aligned}
$$

The next example shows the utility of more flexibility in the relationship between epistemics and probability.

Coin and bit example (Probability and Epistemics)

Suppose there are two agents i and k.

1. k is first given a bit 0 or 1. k learns he has this bit, i is aware that k received a bit, but i does not know what bit k received.
2. k flips a fair coin and looks at the result. i sees k look at the result, but does not what the result is.
3. k performs action s if the coin agrees with the bit (given that heads agrees with 1 and tails agrees with 0), and performs action d otherwise.

This example is from

- R. Fagin & J. Halpern (1994) Reasoning about Knowledge and Probability. *Journal of the ACM* 41:2, pp. 340–367.

Coin and bit example (Probability and Epistemics)

Suppose there are two agents i and k.

1. k is first given a bit 0 or 1. k learns he has this bit, i is aware that k received a bit, but i does not know what bit k received.
2. k flips a fair coin and looks at the result. i sees k look at the result, but does not what the result is.
3. k performs action s if the coin agrees with the bit (given that heads agrees with 1 and tails agrees with 0), and performs action d otherwise.

This example is from

- R. Fagin & J. Halpern (1994) Reasoning about Knowledge and Probability. *Journal of the ACM* 41:2, pp. 340–367.

Coin and bit example (Probability and Epistemics)

Suppose there are two agents i and k.

1. k is first given a bit 0 or 1. k learns he has this bit, i is aware that k received a bit, but i does not know what bit k received.
2. k flips a fair coin and looks at the result. i sees k look at the result, but does not what the result is.
3. k performs action s if the coin agrees with the bit (given that heads agrees with 1 and tails agrees with 0), and performs action d otherwise.

This example is from

- R. Fagin & J. Halpern (1994) Reasoning about Knowledge and Probability. *Journal of the ACM* 41:2, pp. 340–367.

Coin and bit example (Probability and Epistemics)

Suppose there are two agents i and k.

1. k is first given a bit 0 or 1. k learns he has this bit, i is aware that k received a bit, but i does not know what bit k received.
2. k flips a fair coin and looks at the result. i sees k look at the result, but does not what the result is.
3. k performs action s if the coin agrees with the bit (given that heads agrees with 1 and tails agrees with 0), and performs action d otherwise.

This example is from

- R. Fagin & J. Halpern (1994) Reasoning about Knowledge and Probability. *Journal of the ACM* 41:2, pp. 340–367.

Probabilistic Epistemic Model

Definition

Let Φ be a set of proposition letters, and $\mathbf{I}$ be a set of agents.
A probabilistic epistemic model is a tuple
$\mathcal{M} = (X, \{\xrightarrow{i}\}_{i \in \mathbf{I}}, \|\cdot\|, \{\mathbf{P}_{i,x}\}_{i \in \mathbf{I}, x \in X})$, where

- X is a finite set
- $\xrightarrow{i}$ (a subset of X^2) is an epistemic relation for each agent $i \in \mathbf{I}$, that is $x \xrightarrow{i} y$ if i considers y possible from x
- $\|\cdot\|$ is a function assigning to each proposition letter p the set of states where it is true.
- for each agent i and state x, the probability space $\mathbf{P}_{i,x}$ is defined as the tuple $(\Omega_{i,x}, \mathcal{A}_{i,x}, \mu_{i,x})$, where
  - $\Omega_{i,x} \subseteq X$ is the sample space (finite because X is finite)
  - $\mathcal{A}_{i,x}$ is a σ-algebra
  - $\mu_{i,x} : \mathcal{A}_{i,x} \to [0,1]$ is a probability measure over $\Omega_{i,x}$

Discussion

There are four possible sequences of events: $(1, H), (1, T), (0, H), (0, T)$ (note that the action s or d is determined from the first two events). Until k performs the action s or d, agent i considers any of these four states possible.

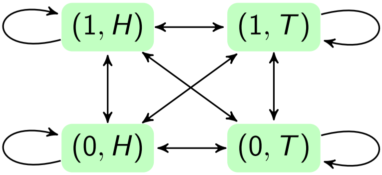

We indicate i's uncertainty between two states using a bidirectional arrow between the two states. In particular, an arrow from state x to state y indicates that i considers y possible if x is the actual state.

Note

It is reasonable that i can calculate the probability of d and s

Discussion

There are four possible sequences of events: $(1, H), (1, T), (0, H), (0, T)$ (note that the action s or d is determined from the first two events). Until k performs the action s or d, agent i considers any of these four states possible.

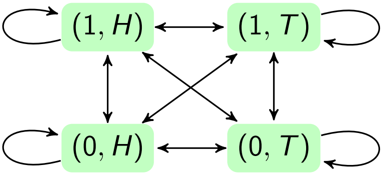

We indicate i's uncertainty between two states using a bidirectional arrow between the two states. In particular, an arrow from state x to state y indicates that i considers y possible if x is the actual state.

Note

It is reasonable that i can calculate the probability of d and s

Here is a possibility for i's probability spaces.

- sample space is enclosed in a box,
- σ-algebra equivalence classes are enclosed in the dotted ovals.

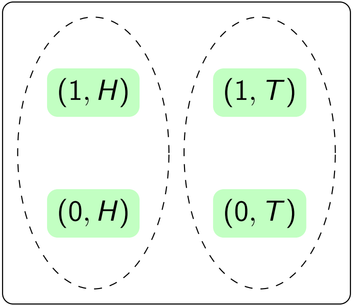

$\mathbf{M}_1$

- sample space is the same as set of states i considers possible.
- Individual states cannot be measurable (otherwise 0 or 1 must be assigned a probability).

Consequence: s and d are not measurable.

Here is a possibility for i's probability spaces.

- sample space is enclosed in a box,
- σ-algebra equivalence classes are enclosed in the dotted ovals.

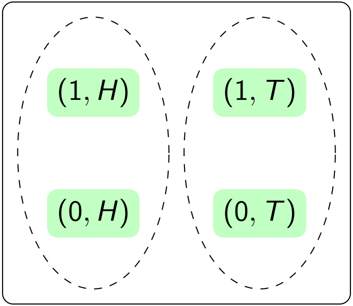

$\mathbf{M}_1$

- sample space is the same as set of states i considers possible.
- Individual states cannot be measurable (otherwise 0 or 1 must be assigned a probability).

Consequence: s and d are not measurable.

Another possibility has a sample space containing only the states with the correct bit (but recall that i considers all states possible and both sample spaces possible).

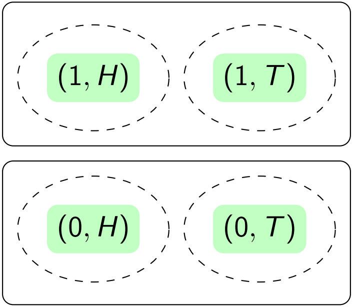

$\mathbf{M}_2$

Without assigning probability to the bit, i can now assign a probability to the actions s and d.

Another possibility has a sample space containing only the states with the correct bit (but recall that i considers all states possible and both sample spaces possible).

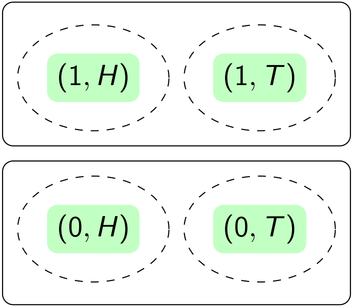

$\mathbf{M}_2$

Without assigning probability to the bit, i can now assign a probability to the actions s and d.

Here i is uncertain among 4 probability spaces.

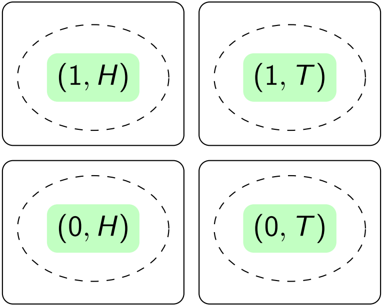

$\mathbf{M}_3$

Main point of example

The action s has probability $1/2$ (as does d). The only model to let us reason about the probability of s as having probability $1/2$ is M_2.

Observations about M_2:

- Epistemics is not defined in terms of probability
- In fact epistemics reflects uncertainty about what probability distribution there actually is:

  the set of (epistemic) possibilities is larger than the (probabilistic) sample spaces.

Main point of example

The action s has probability $1/2$ (as does d). The only model to let us reason about the probability of s as having probability $1/2$ is M_2.

Observations about M_2:

- Epistemics is not defined in terms of probability
- In fact epistemics reflects uncertainty about what probability distribution there actually is:

  the set of (epistemic) possibilities is larger than the (probabilistic) sample spaces.

Related Papers

1. D. Gerbrandy. (1998) *Bisimulations on Planet Kripke*. Dissertation, ILLC.
   (Discusses Muddy Children Puzzle and Surprise Exam Puzzle)
2. B. Kooi: Probabilistic Dynamic Epistemic Logic. (2003) *Journal of Logic Language and Information*, 12(4): 381–408.
   (Adds probability to Public Announcement Logic)
3. J. Sack. (2008) Temporal Languages for Epistemic Programs. *Journal of Logic Language and Information*. 17(2): 183–216.
   (Adds temporal operators to Dynamic Epistemic Logic)
4. J. Sack. (2009) Extending Probabilistic Dynamic Epistemic Logic. *Synthese*,1 69:2, pp. 241–257.
   (Adds previous-time operator and σ-algebras to Probabilistic Dynamic Epistemic Logic)

Books and Textbooks Relating to Topic

1. H. van Ditmarsch, W. van Der Hoek, B. Kooi (2008) *Dynamic Epistemic Logic*. Springer Synthese Library 337.
2. D. Harel, D. Kozen, J. Tiuryn (2000) *Dynamic Logic*. Foundations of Computing.
3. J.-J. Ch. Meyer and W. van Der Hoek (1995) *Epistemic Logic for AI and Computer Science*. Cambridge Tracts in Theoretical Computer Science 41.

Thank you!