Extending Probabilistic Dynamic Epistemic Logic

Joshua Sack

May 29, 2008

Probability Space

Definition

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

1. S is a set called the sample space.
2. $\mathcal{A} \subseteq \mathcal{P}(S)$ is a σ-algebra: a set of subsets of S 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
   - $\mu(S) = 1$ 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)

If $\mathcal{A} = \mathcal{P}(S)$, then the probability space is called "discrete", and the probability function μ can be viewed as mapping *each* element of S into $[0, 1]$.

How the σ-algebra helps us

The σ-algebra lets us restrict the domain of the probability measure.

- Restricting the domain of the probability measure lets us reflect uncertainty about the probability of individual elements of S.

Probabilistic Epistemic Logic offers a different way of handling uncertainty about probabilities; it lets us express uncertainty about probability spaces.

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
$\mathcal{M} = (X, \{\xrightarrow{i}\}_{i\in\mathbf{I}}, \|\cdot\|, \{\mathcal{P}_{i,x}\}_{i\in\mathbf{I},x\in X})$, 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 $\mathcal{P}_{i,x}$ is defined as the tuple $(S_{i,x}, \mathcal{A}_{i,x}, \mu_{i,x})$, where
  - $S_{i,x} \subseteq X$ 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 $S_{i,x}$

Adding probability formulas

Add to epistemic logic probability formulas of the form $P_i(\varphi) \geq q$ for a rational number q. Evaluate the truth of such a formula at a pointed model (M, x) (where M is a probabilistic epistemic model and x is a state in M) in the following way.

- $(M, x) \models P_i(\varphi) \geq q$ iff $\mu_{i,x}([\![\varphi]\!] \cap S_{i,x}) \geq q$, where $[\![\varphi]\!]$ is the set of states in M where φ is true.

Of course, this definition only works if $[\![\varphi]\!] \cap S_{i,x} \in \mathcal{A}_{i,x}$. We can ensure this by either

1. requiring each $\mathcal{A}_{i,x}$ be large enough so that $[\![\varphi]\!] \cap S_{i,x} \in \mathcal{A}_{i,x}$ is guaranteed for all φ,
2. extending the function $\mu_{i,x}$ to all subsets. *Inner* and *outer measures* are two such extensions, and they need not obey all conditions of a measure.

Inner and Outer Probabilities

Let $(S, \mathcal{A}, \mu)$ The inner and outer measures are defined on any set $Y \in \mathcal{P}(S)$ by

- (outer measure) $\mu^*(Y) = \inf\{\mu(B) : Y \subseteq B, B \in \mathcal{A}\}$
- (inner measure) $\mu_*(Y) = \sup\{\mu(B) : B \subseteq Y, B \in \mathcal{A}\}$

When the σ-algebra $\mathcal{A}$ is finite, these are equivalent to

- (outer measure) $\mu^*(Y) = \mu(\bigcap\{B : Y \subseteq B, B \in \mathcal{A}\})$
- (inner measure) $\mu_*(Y) = \mu(\bigcup\{B : B \subseteq Y, B \in \mathcal{A}\})$

The inner and outer measures are related by

$$\mu^*(Y) = 1 - \mu_*(\overline{Y})$$

Neither the inner nor the outer measure is in general a measure.

Observation about finite σ-algebras

- Any finite σ-algebra $\mathcal{A}$ can be characterized by an equivalence relation. The equivalence relation is such that the equivalence classes are the sets $\bigcap\{A \in \mathcal{A} : x \in A\}$ for each $x \in S$. Some of these sets produced are identical for different x's. Although S may be infinite, the there will only be finitely many equivalence classes.
- A σ-algebra can be generated by the equivalence classes of any equivalence relation.

Fagin, Halpern, and Tuttle example

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.

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). Until k performs the action s or d, agent i considers any of these four states possible.

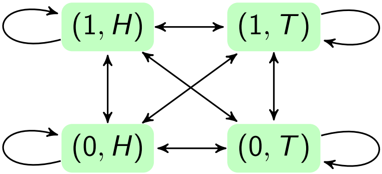

We indicate i's uncertainty between two states using a double arrow between the two states. In particular, an arrow from state x to state y indicates that i considers y possible if x is the actual state. (Before the bit is given, k's epistemic relation will be the same).

Here is a possibility for i's probability spaces. The sample space enclosed in a box, and the σ-algebra equivalence classes are enclosed in the dotted ovals.

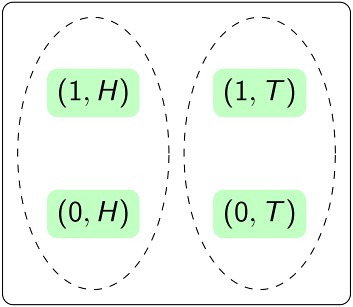

$\mathbf{M}_1$

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

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).

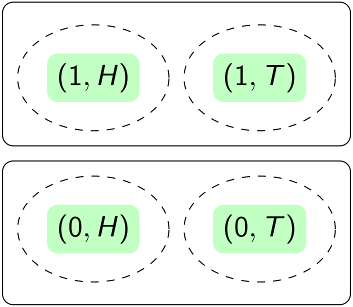

$\mathbf{M}_2$

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.

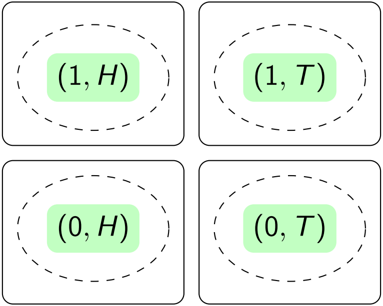

$\mathbf{M}_3$

Modeling a sequence of events

It is suggested that each of these models may reasonably represent i's probability spaces at a certain stage in the sequence of events (but to make to make better sense of the transition, we add a little more in parentheses that was not in the original statement of the example):

- M_1 with the time before the bit is given to k (suppose i does not yet know that k will perform action s or d).
- M_2 with the time after the bit is given to k, (after k tells i he will do either s or d depending on the coin toss,) but before the coin is flipped.
- M_3 with the time after the coin is tossed, (after k spontaneously offers i a bet about what action he will take,) but before k performs his action.

action model

Definition (Action Model)

An *action model* $(\Sigma, \{\xrightarrow{i}\}, \{\mathcal{P}_{i,\sigma}\}, \mathrm{pre})$ is a probabilistic epistemic model with the valuation function $\| \cdot \|$ replaced by a function *pre* which assigns to each $\sigma \in \Sigma$ a class of pointed probabilistic epistemic models. Each element $\sigma \in \Sigma$ is called an *action type*.

Update Product

We define a mechanical procedure called the *update product* for transforming one model into another model given an action (represented by an action model). We define the update product in two stages.

1. The first, called the *unrestricted product*, is to take a product that does not make use of the *pre* function.
2. The second, called the *relativization*, is to restrict the unrestricted product to a set of states characterized by the *pre* function.

finite product measure

Definition (Finite product space)

The product space of probability spaces $(S_1, \mathcal{A}_1, \mu_1)$ and $(S_2, \mathcal{A}_2, \mu_2)$ is $(S_3, \mathcal{A}_3, \mu_3)$, where

1. $S_3 = S_1 \times S_2$ is the Cartesian product.
2. $\mathcal{A}_3$ is the smallest σ-algebra containing

$$\{A \times B : A \in \mathcal{A}_1, B \in \mathcal{A}_2\}$$

3. The probability measure is defined as

$$\mu_3(A) = \sum_{k=1}^{n} \mu_1(B_k)\mu_2(C_k)$$

where $B_k \in \mathcal{A}_1$, $C_k \in \mathcal{A}_2$, and $A = \biguplus_{i=1}^{n} B_k \times C_k$

unrestricted product

Definition (unrestricted product)

The unrestricted product between a probabilistic epistemic model $\mathbf{M}$ and an action model $\mathbf{\Sigma}$ is $\mathbf{M} \otimes_U \mathbf{\Sigma}$ with the following components:

1. $X_{\otimes} = X \times \Sigma$
2. $(x, \sigma) \xrightarrow{i} (z, \tau)$ iff $x \xrightarrow{i} z$ and $\sigma \xrightarrow{i} \tau$
3. $\|p\|_{\otimes} = \|p\| \times \Sigma$
4. We define $\mathcal{P}_{i,(x,\sigma)}$ to be the finite product space between $\mathcal{P}_{i,x}$ and $\mathcal{P}_{i,\sigma}$

relativization

Definition (relativization)

The relativization of a probabilistic epistemic model **M** to $Y \subseteq X$ is given by $\mathbf{M} \otimes_R Y$ with the following components:

1. $X_Y = Y$
2. $x \xrightarrow{i}_Y z$ iff $x \xrightarrow{i} z$ and $x, z \in Y$
3. $\|p\|_Y = \|p\| \cap Y$
4. For $x \in Y$, if $\mu^*_{i,x}(Y) = 0$, then define $\mathcal{P}_{i,x}$ to be the trivial probability space on the singleton x. Otherwise
    1. $S_{Y\,i,x} = S_{i,x} \cap Y$
    2. $\mathcal{A}_{Y\,i,x}$ is the σ-algebra generated by $\{A \cap Y : A \in \mathcal{A}_{i,x}\}$
    3. The probability measure is defined by

$$\mu_{Y\,i,x}(A) = \frac{\mu^*_{i,x}(A \cap Y)}{\mu^*_{i,x}(Y)}$$

Establishing Additivity

Theorem

Let $(S, \mathcal{A}, \mu)$ be a probability space, such that $\mathcal{A}$ is a finite σ-algebra. If $A, B \in \mathcal{A}$ and $Y \subseteq S$, then

- $\mu^*(A \cap B \cap Y) = \mu^*(A \cap Y) + \mu^*(B \cap Y)$
- $\mu_*(A \cap B \cap Y) = \mu_*(A \cap Y) + \mu_*(B \cap Y)$

The proof of the outer measure part rests on the fact that $\widehat{Y} = \bigcap\{C : Y \subseteq C : C \in \mathcal{A}\} \in \mathcal{A}$, and that $\widehat{A \cap B \cap Y}$ is a disjoint union of $\widehat{A \cap Y}$ and $\widehat{B \cap Y}$.

The proof of the inner measure part follows similar reasoning.

The update product

Definition

Update Product Given a probabilistic epistemic model M and an action model Σ, let

$$Y = \{(x, \sigma) : (M, x) \in \mathrm{pre}(\sigma)\}$$

Then the update product between M and Σ, written $M \otimes \Sigma$ is

$$M \otimes \Sigma = (M \otimes_U \Sigma) \otimes_R Y$$

What action models should be used?

From M_1 to M_2, there are two events:

1. a semi-private announcement of the bit to k
2. a public announcement that k plans to do either action s or d.

From M_2 to M_3, there are two events:

1. a semi-private announcement to k of the result of the coin toss
2. a public announcement regarding k's bet offer

We first consider going from M_1 to M_2 using just one action model, and similarly from M_2 to M_3 with just one action model. We then consider going from M_1 to M_2 using a sequence of two action models, and similarly from M_2 to M_3.

Semi-private announcement

The relational structure of a semi-private announcement is given by

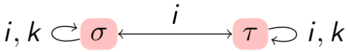

i and k's probability spaces:

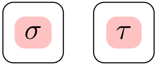

M_1 to M_2

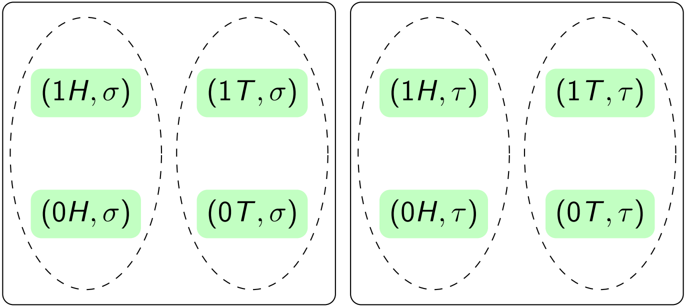

M_1 to M_2

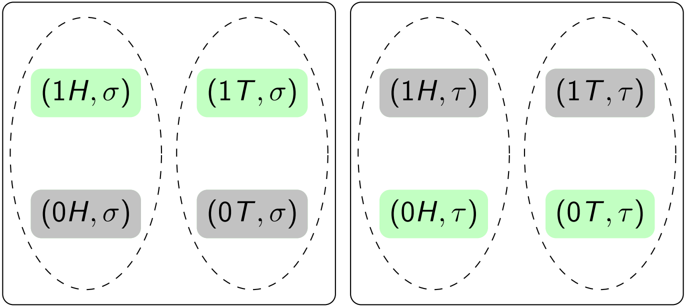

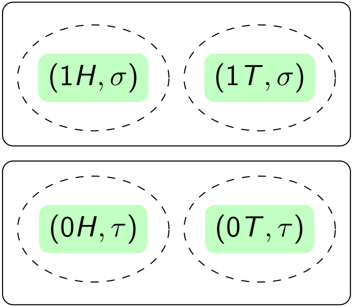
(1H, σ)
(1T, σ)
(0H, τ)
(0T, τ)

M_2 to M_3

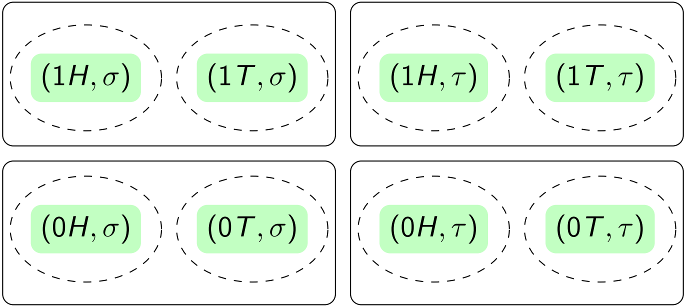

M_2 to M_3

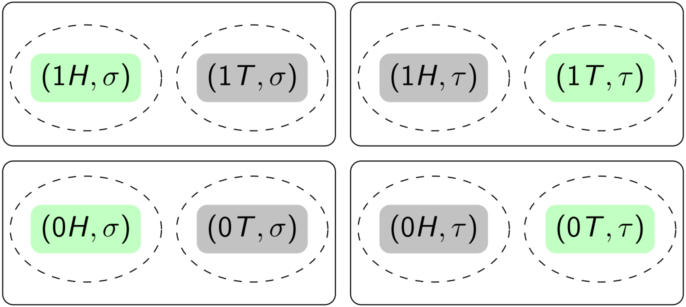

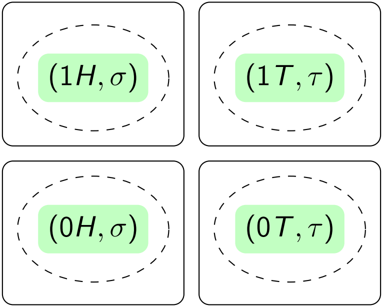
$(1H, \sigma)$
$(1T, \tau)$
$(0H, \sigma)$
$(0T, \tau)$

From M_1 to M_2 first stage: semi-private announcement

relational structure:

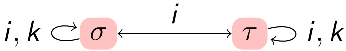

i's probability space:

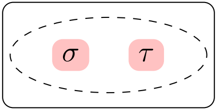

k's probability spaces:

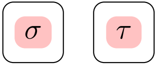

$pre(\sigma)$ includes states with 1, and $pre(\tau)$ includes states with 0.

From M_1 to M_2 second stage: public announcement

relational structure:

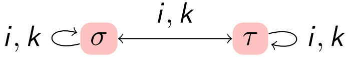

This is the public announcement "the precondition of σ or the precondition of τ" as long as no state satisfies both preconditions.
i and k's probability spaces:

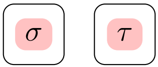

$pre(\sigma)$ includes states with 1, and $pre(\tau)$ includes states with 0.

From M_2 to M_3

The semi-private and public announcement action models are the same in all components except for the precondition function *pre*.

- Instead of 1, the precondition of σ is H
- Instead of 0, the precondition of τ is T.

M_1

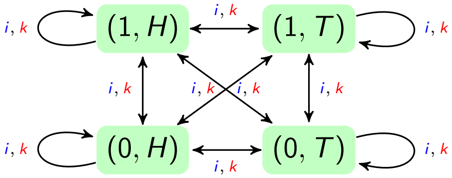

M_1

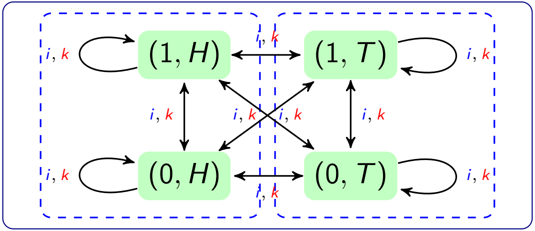

M_1

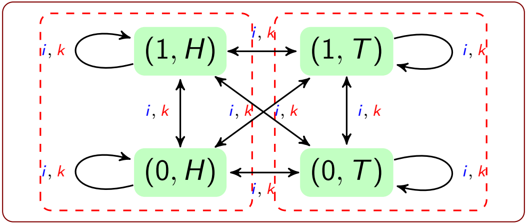

M_1

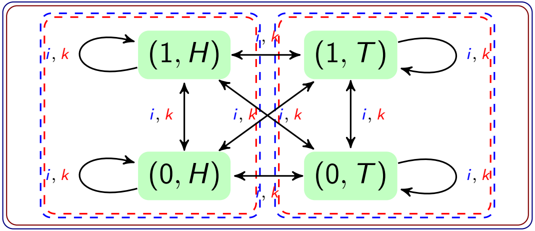

after 1st semi-private announcement

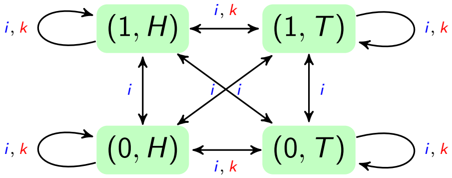

after 1st semi-private announcement

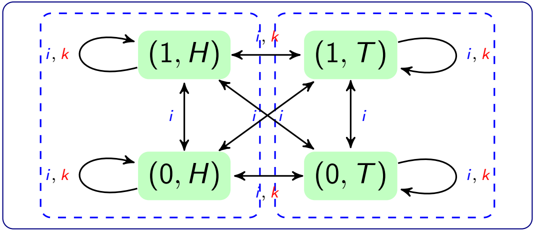

after 1st semi-private announcement

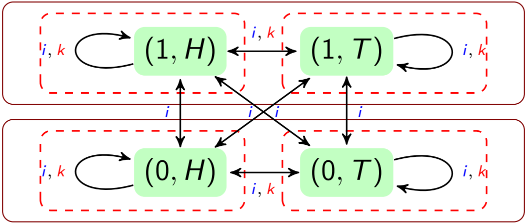

after 1st semi-private announcement

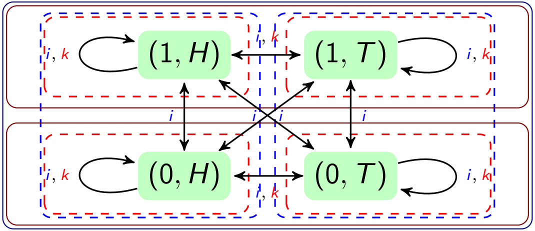

after first public announcement (M_2)

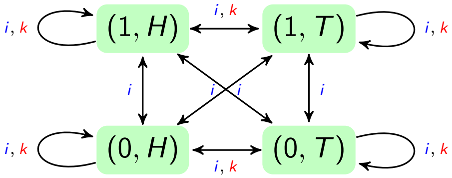

after first public announcement (M_2)

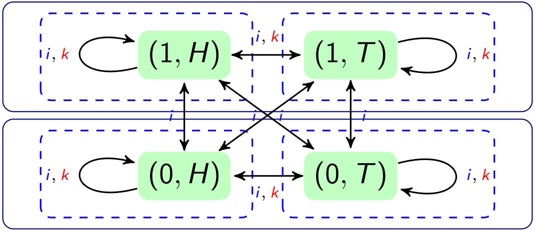

after first public announcement (M_2)

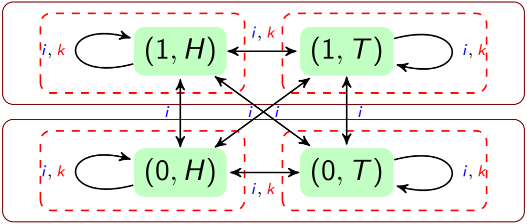

after first public announcement (M_2)

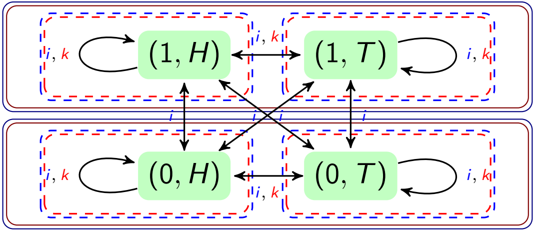

after 2nd semi-private announcement

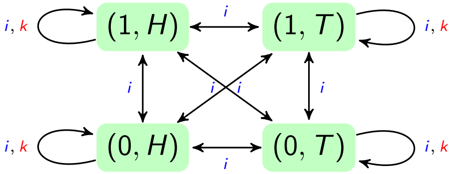

after 2nd semi-private announcement

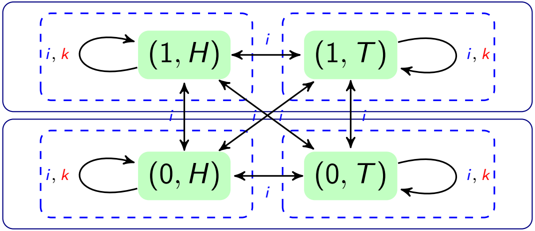

after 2nd semi-private announcement

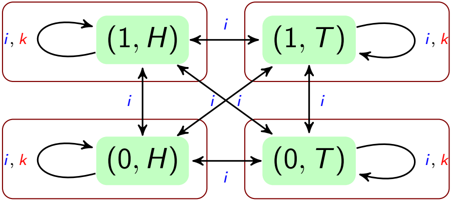

after 2nd semi-private announcement

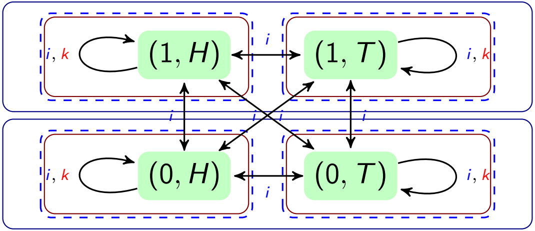

M_3

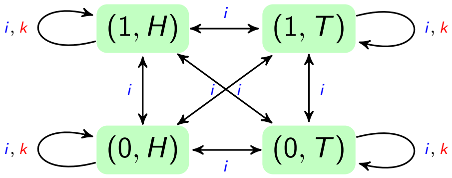

M_3

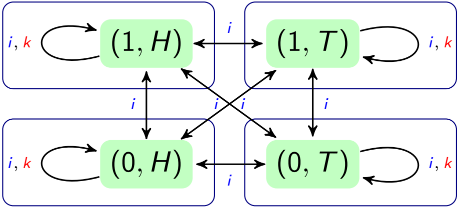

M_3

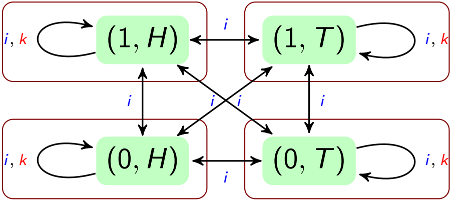

M_3

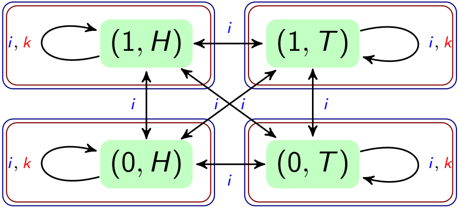

Recording the past

Here are two ways: list history or static history:

- Define a list history to be a list $(S_0, S_1, \ldots, S_n)$ of probabilistic epistemic model, for which for each j, $S_{j+1} = S_j \otimes A$ for some action model A. We may want to fix some underlying structure of A for technical reasons.
- Define a static history to be an augmented probabilistic epistemic model $(X, \{\stackrel{i}{\to}\}_{i\in\mathbf{I}}, \|\cdot\|, \{\mathcal{P}_{i,x}\}, Y)$, where Y is a binary relation over X, and a number of conditions are imposed in order to ensure that xYz can adopt the reading "x follows from z and some action." The goal is to ensure that the static history is structurally equivalent to a list history.

When proving completeness, we have found it easier to use the static histories.

Need for non-measurable sets in completeness proof?

The answer to this is still unknown. Here are some comments:

- Completeness for Probabilistic Epistemic Logic (which involves non-measurable sets) constructs a filtration that has discrete measures.
- A similar filtration can be constructed for a probabilistic dynamic epistemic logic with a past-time operator, but the filtration will guarantee all conditions needed to reflect update products.
- I suggest using truth-preserving model transformations to convert the filtration into a true static history.
- It is yet unknown whether we would benefit from being able to transform a model into one where probabilities are not discrete.