Conditional Probability and Update Logic

Second version, January 2003

Johan van Benthem, Amsterdam & Stanford

Abstract Dynamic update of information states is the dernier cri in logical semantics. And it is old hat in Bayesian probabilistic reasoning. This note brings the two perspectives together, and proposes a mechanism for updating probabilities while changing the informational state spaces.

1 Tree diagrams for probability

Many textbooks use a perspicuous tree format for simple probability spaces. Branches are histories of successive events. Going down the tree, actions generate new probability spaces, with the current space being more or less the current tree level. Arrows downward from a node are labeled with probabilities, summing to *1*. By way of illustration, take the perennial Monty Hall puzzle. First, Nature puts a car behind one of three doors (the quizmaster knows which, you do not), then you choose a door, and finally, the quizmaster opens a door not chosen by you which has no car behind it. This involves a tree-diagram like the following. Of course, which actions you put in *precisely* is a matter of picking the right level of detail.

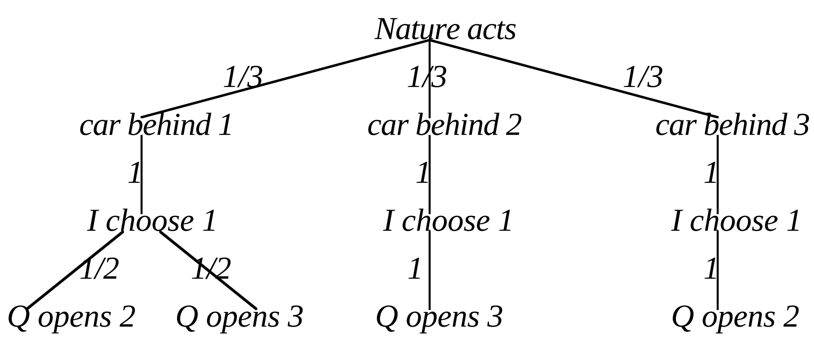

Let's say I chose door *1*, Monty opened door *3*. Should I switch or not? We must find the right conditional probability for the car being behind door *1*, given all that has passed. If we conditionalize on *'the car is not behind 3'*, we find a probability of *1/2*. But, if we do the job well, we will pick up the more informative true proposition *A = 'Monty opened door 3'* to compute *P('car behind 1'|A) = 1/3* – and conclude that we should switch. Is luck needed in picking the right *A*, or is there a systematic principle at work? In this note we will analyze this process, which is close to current dynamic update logics, with information flowing down the tree.

2 Update logic in a nutshell

Update logics are about changing information states as propositions are announced, or more general actions observed. In this note, information models for groups of agents G are standard epistemic structures $\boldsymbol{M}$ with a universe of possible worlds and equivalence relations $\sim_i$ between these modeling the uncertainty of agent $i \in G$. Each model has an *actual world* $\boldsymbol{s}$. These structures interpret the usual epistemic language with operators K_i for individual knowledge and C_G of common knowledge in the group. Formulas of this language describe the static properties of worlds s in a given information model $\boldsymbol{M}$ (Fagin, Halpern, Moses & Vardi 1995).

But epistemic *actions* change such models! E.g., truthful public announcement of a proposition ϕ removes all worlds from the current model where ϕ does not hold:

from

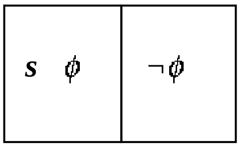

to

This is the dynamic effect of an answer "*Yes*" to a question " ϕ?". Successive state elimination is the simplest update procedure, yielding the familiar picture of shrinking sets representing ever stronger group knowledge about the actual world.

Despite this March of Progress, update steps may change the current truth value of assertions. Before you answered my question, I did not know if ϕ. Now I do, and so the ignorance statement has become false. The technical reason is that we have to re-evaluate formulas with epistemic operators in the new smaller models, which may affect their earlier truth values. The resulting pattern of changing truth and falsity can make even public update sequences surprising, witness puzzles like Muddy Children where repeated communication of ignorance leads to knowledge. In a more elaborate dynamic-epistemic logic, one records this in mixed assertions

$$[A!]\phi$$

saying that after public announcement of A, formula ϕ holds. This expresses things like $[A!]C_GA$: after public announcement of A, it has become common knowledge.

More complex epistemic actions do not just eliminate states, they may transform the worlds of the model. E.g., in the group of *{you, I, she}*, she and I do not know if ϕ is true, but you do. As a matter of fact, ϕ is true. We can draw a model like this:

$$\boldsymbol{\phi} \quad \overline{\underline{\quad I,\ she \quad}} \quad \neg\phi$$

Now I ask in public if ϕ, but you answer just to me: while she sees you answering. There are two relevant actions: "you says YES", "you say NO". Each of these has a public *precondition,* as it is common knowledge in our story that you speak truly. The first action requires that you know that ϕ, the second that you know that $\neg\phi$. You and I can distinguish the two, but she cannot. The result is a new epistemic model whose worlds are old worlds with an action attached whose precondition is satisfied – with the new uncertainties computed by the following rule of

> *Product Update* Let s, t be worlds in the current model, and a, b actions at s, t, resp., whose preconditions are satisfied there. Ordered pairs (s, a) encode the result of performing action a in state s. Uncertainty among new states can only come from existing uncertainty via indistinguishable actions:
>
> $$(s, a) \sim_i (t, b) \quad \text{iff} \quad \text{both } s \sim_i t \text{ and } a \sim_i b$$

In our example, this gives a new model where you and I know, while she does not, though she knows that we know (that we know is even common knowledge):

$$(\boldsymbol{\phi}, \textbf{\textit{say YES}}) \quad \overline{\quad she \quad} \quad (\neg\phi, \textit{say NO})$$

Product update can also blow up the size of an epistemic model. This would happen, e.g., when she is not sure whether you answered my question or not.

The general update process has two drivers: (a) an epistemic *information model* of all relevant possible worlds with agents' uncertainty relations, and (b) an *action model* of all relevant actions, again with agents' uncertainty relations between them.

Product update takes successive products of these two models. Again, truth values of propositions can change drastically in such transformations. Dynamic-epistemic logics record this explicitly. In a more refined version of the model, one might even impose global constraints on the possible runs of the update process, which record higher information like "sooner or later, she will tell me all she knows".

3 Stage setting: epistemic action update in Monty Hall

There is pure epistemics in setting the stage for the quiz puzzle. We start with this, adding probabilities only later. Let's compute the uncertainties of agents as we go:

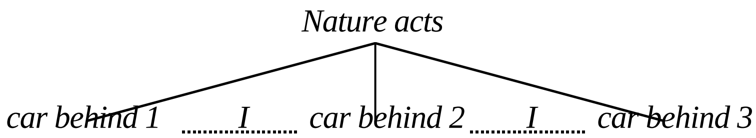

Nature has three actions, indistinguishable for me (*I*), but not for the quizmaster *Q*. The result is the three-world epistemic model at the second tree level. Now I choose a door. This is a public transparent action, but it would be tedious to represent all options. Let's just say that I publicly chose door *1*. The product update rule yields

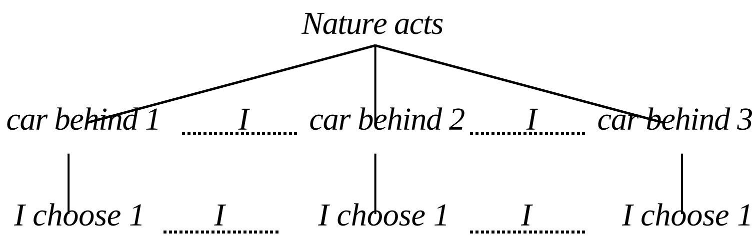

Next, we have three possible actions of *Q*'s publicly opening some door, with preconditions (a) I did not choose that door, and (b) *Q* knows that the car is not behind it. Product update takes only those pairs *(s, a)* where *s* satisfied the relevant precondition PRE_a, and computes uncertainties. The result is the usual Monty tree:

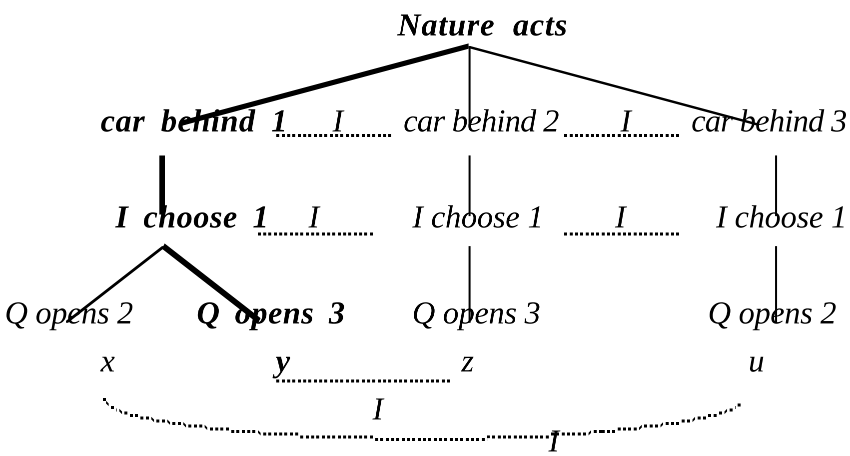

Let the car be behind door *1*, while *Q* opened door *3*: The actual world is *y*, reached via the bold-face branch. In the epistemic model at the bottom level, I know the world is either *y* or *z*. Through the tree, Quizmaster always knows exactly where he is. In other scenarios, both agents might have genuine uncertainties, resulting in a much more complex pattern of linked equivalence relations. So much for logic!

4 Conditional probability and update

Now for a quick review of Bayesian update. An agent's probability model is a set of worlds with a probability measure P on events, defined by propositions that can be true or false at worlds. And conditional probabilities $P(\phi \mid A)$ give the probability for ϕ given that A is the case, using P rescaled to the set of worlds satisfying A:

$$P(\phi \mid A) \; = \; P(\phi \wedge A) \,/\, P(A)$$

Bayes' Rule then helps compute such probabilities in forms like

$$P\,(\phi \mid A) \;\; = \;\; P\,(A \mid \phi)\; x\; P(A) \;\; / \, P(\phi)$$

and more elaborate versions of the same inversion idea.

Conditional probability looks like eliminative update. It zooms in on those worlds where the new information A holds, and then recomputes probabilities. This is like eliminating all $\neg A$-worlds, and re-evaluating epistemic formulas. And the binary format $P(\phi|A)$ storing all possible updates is like the above dynamic notation $[A!]\phi$.

But the epistemic perspective has two further features. First, it considers many agents together, with their mutual information. This would be like having my probability about your probabilities, etc. But even more importantly, product update does not just select subzones of the current information space, but it transforms the latter much more drastically as required by relevant information-carrying actions. Probabilistic theory speaks about *events* A on which we conditionalize, which seems a similar ambition. One wants to combine conditional probability with an account of how actions change the current probability model. Let's see how this works out by continuing with the earlier example of probability tree diagrams.

5 Computing updates on probabilities with public actions

The Monty Hall example is about public action, where we update probabilities in a transparent setting. The earlier epistemic update created information models at each tree level, and we expect that probabilities for an agent will give weights to her indistinguishable worlds in such models, giving fine-structure to her information. But, there is also a second sense of probability involved, working in the other driver of the story. Action diagrams may have indistinguishabilities between actions, and agents might also fine-structure their action alternatives numerically.

The latter, too, happens in textbook tree representations, which assign probability values to moves from tree nodes to their daughters. In the simplest case of public action, these move probabilities are the same from the viewpoint of every agent.

Digression: a subtlety of interpretation Action probabilities so far do not record uncertainties about *what action has taken place* once we observe that something has happened. They provide estimates for the likelihood that an action *will be taken* at the appropriate stage. But the two aspects can interfere in update, witness:

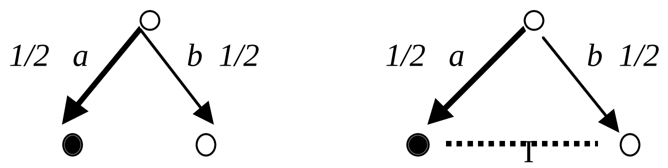

Suppose action *a* was in fact taken. If *a, b* are distinguishable for me, in the black dot on the left-hand side, I know exactly where I am through observation – though there is some 'ancient history' that with probability *1/2, b* might have been taken. But if *a, b* are indistinguishable for me, on the right, I do not know what happened, and the earlier probability induces a live option that I am in the white world below.

Here is another point which we can see in the Monty tree. First, one action type can have different probabilities at different nodes of the current tree level. E.g., *'opening door 2'* has probability *1/2* when the car is behind door *1*, but probability *1* with the car behind door *3*. Of course, by making descriptions of action tokens disjoint, we can make probabilities unique – but this seems less natural in practice.

Now for the *product update rule* along a branch. The usual textbook explanation makes probability of a branch a product of the probabilities of its actions. Recursively, this amounts to repeating the following step:

> Look at the current probabilities, and compute those for the next
> state by taking a suitable product with weighted available actions.

But the epistemic context matters. Probabilities need not sum to *1* at single nodes, or a whole horizontal tree level, but only on one 'information set': a maximal component of the uncertainty relation. Such components are the natural probability spaces from a given node at the current stage of the overall process. And in epistemic models, such components may be different for different agents. So we need probability functions relative to agents *i* and nodes *s*:

$P_{i,s}$ defined on the probability space $D_{i,s} = \{ t \mid t \sim_i s \}$

Then, the preceding considerations lead to the following update principle:

Product Rule For public actions a, and public local probabilities $P_{i,s}$:

(1)

$$P_{i,\ (s,a)}((t,a)) \quad = \quad \frac{P_{i,s}(t) \times P_t(a)}{\sum_{(u,a) \in D\, i,\, (s,a)} P_{i,s}(u) \times P_u(a)}$$

To keep complex notations in line, we will also write the denominator as follows:

$\Sigma\ \{P_{i,s}(u) \times P_u(a) \mid (u, a) \in D\ i, (s, a)\}$

where *{}* refers to a *multiset* counting occurrences of numbers.

A word of explanation may help. The rule computes probabilities per world. Its notation *(t, a)* presupposes that *t* satisfies the precondition for executing action *a*. Thus, the probability space from *(s, a)*'s perspective may have shrunk from the previous level *s*. The numerator is the obvious total product. The denominator renormalizes values to sum to *1* in the relevant space.

Finally, here is the general product rule for an arbitrary formula ϕ:

(2)

$$P_{i,\ (s,a)}(\phi) \quad = \quad \frac{\Sigma\{P_{i,s}(u) \times P_u(a) \mid (u,a) \in D_{i,(s,a)}\ \&\ (u,a) \models \phi\}}{\Sigma\{P_{i,s}(u) \times P_u(a) \mid (u,a) \in D_{i,(s,a)}\}}$$

This rule computes new probabilities after the action has taken place. But we can also describe it in terms of the *old situation* before the update! The index in the numerator ranges over all tuples *(u, a)* in $D_{i,(s,a)}$. But this amounts to looking at all *u* in $D_{i,s}$ satisfying the action precondition PRE_a. Thus, formula (2) is equivalent to

(3)

$$\frac{\Sigma\{P_{i,s}(u) \times P_u(a) \mid u \in D_{i,s}\ \&\ u \models PRE_a\ \&\ u \models [a]\phi\}}{\Sigma\{P_{i,s}(u) \times P_u(a) \mid u \in D_{i,s}\ \&\ u \models PRE_a\}}$$

This format may be viewed as a sort of *generalized conditional probability*

$$P_i^a\ ([a]\phi \mid PRE_a)$$

Essentially, we compute a standard conditional probability, but over a *new space* whose worlds are pairs *(u, a)* of old worlds and executable instances of the action *a*. This is precisely the combination of two mechanisms that we wanted.

6 A check on two examples

Monty Hall revisited Consider the earlier tree, now with probabilities indicated.

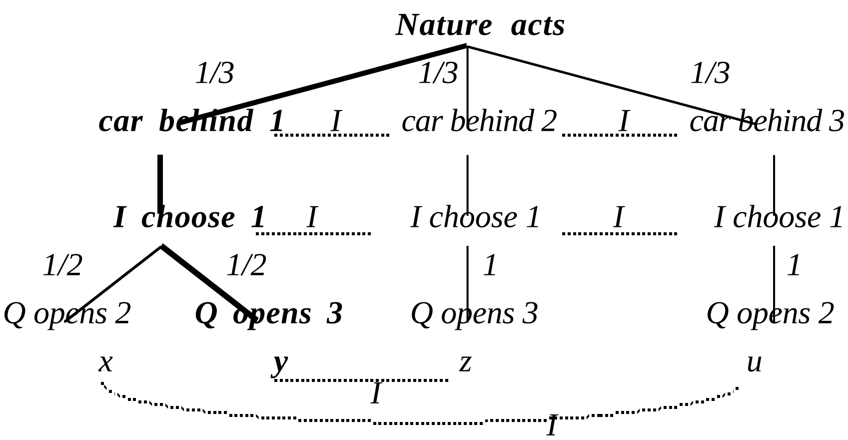

It is easy to check that the probabilities in my final set $\{x, y\}$ work out to

for y: $(1/3 \cdot 1/2) / (1/3 \cdot 1/2 + 1/3 \cdot 1) = 1/3$

for z: $(1/3 \cdot 1) / (1/3 \cdot 1/2 + 1/3 \cdot 1) = 2/3$

In this picture, we see our product rule at work, including its non-trivial features. E.g., we are now in world *y*, where we know *Q* has opened door *3*. Nevertheless, in computing the probability for being in *y* rather than *z*, we take the old probability *1/2* into account for opening *door 2* in the state preceding *y*. Why: now that we *know* this action was not taken? My intuitive response would be as follows:

"Counterfactual chances are still relevant. We observe an opening of door 3. What is the chance it lies on the left-hand branch, and not the middle one? Well, on the left-hand branch, there was a chance of 1/2 that the other door was opened, while on the middle branch, it was the only option. So, seeing door 3 opened provides more evidence for our being on the middle branch."

But to critics of the received view on Monty Hall, this line may sound circular... Thus, an update rule is not a neutral mathematical fact justifying dynamic-epistemic-probabilistic laws. It builds in such laws, as also shows in the axioms of Section 8.

Public announcement Product rules (1), (2) also specialize to the stipulation for public announcement *A!* of an assertion *A* in Kooij 2001. His setting is simpler than the Monty Hall tree, as the action probabilities for truly asserting propositions cannot vary per location: they either equal *1* or the action cannot be performed at all. Here is the stipulation, slightly adapted to our setting:

(4) $$P_{i,\ s, A!}(\phi) \quad = \quad \frac{\Sigma\{ P_{i,s}(u) \mid u \in D_{i,s} \ \&\ u \models A \ \&\ [A!]\phi \}}{\Sigma\{ P_{i,s}(u) \mid u \in D_{i,s} \ \&\ u \models A\}}$$

This is a special case of the general product rule (2). This shows also in that this rule for public announcement really computes a *standard conditional probability*

$$P^{old}{}_{i,s}([A!]\phi \mid A)$$

Kooij 2001 formulates the values in the update rule by referring to the 'old' situation straightaway – but (4) is more in the spirit of general epistemic updates.

7 General probabilistic product update

The product rule so far builds in special epistemic features. E.g., probabilities for worlds have uniform values across a whole information set, as seen from every vantage point. In terms of epistemic logic, this means the following:

> If agents know their probabilities of all propositions at some stage,
> product update will always lead to new probabilities which they know.

In the Monty Hall tree, probabilities are even common knowledge among *I* and *Q*. One might prove a *characterization* of product update via nice epistemic properties, on the lines of that given for the pure epistemic version in van Benthem 2001.

Here are some richer options, more in the spirit of the recent update literature (Baltag-Moss-Solecki 1999, van Ditmarsch 2001). Some require just hanging some subscripts in the above product rule at the right places, others involve new ideas.

First, action probabilities can be agent-dependent. To allow this, just replace terms $P_u(a)$ in formulas (1), (2) by $P_{i,u}(a)$. We might also let agents' probability functions vary within their uncertainty sets. Both might happen in Monty Hall: '*Q* does not know where the car is', '*I* think that *Q* prefers opening doors with lower numbers', etc. So, let's probabilize epistemic product update with general action diagrams $\boldsymbol{A}$. Here is a matching generalization (for further options, cf. van Benthem 2002):

(5) $$P_{i,\ (s,a)}(\phi) \quad =$$

$$\frac{\Sigma\{P_{i,s}(u) \times P_{i,u}(b) \mid u \sim_i s \ \&\ b \sim_i a \text{ in } \boldsymbol{A} \ \&\ u \models PRE_b \ \&\ (u, b) \models \phi\}}{\Sigma\{P_{i,s}(u) \times P_{i,u}(b) \mid (u, b) \in D_{i,(s,a)} \ \&\ b \sim_i a \text{ in } \boldsymbol{A}\}}$$

8 Probabilistic epistemic update logic

Update rules validate a logic for reasoning with knowledge and probability. This will be a combined language with modal operators of various sorts. With pure epistemic update, key axioms interchange update actions and our knowledge:

$$[A!]\, K_i\, \phi \qquad \leftrightarrow \qquad (A \rightarrow K_i\, (A \rightarrow [A!]\, \phi))$$

With general action diagrams $\boldsymbol{A}$ and product update, we get

$$[a]\, K_i\, \phi \qquad \leftrightarrow \qquad PRE_a \rightarrow \& \{K_i\, (PRE_b \rightarrow [b]\phi)) \mid b \sim_i a \text{ in } \boldsymbol{A}\}$$

This time, we get similar valid principles for public announcement with probability:

$$[A!]\, P_i\,(\phi) = k \quad \leftrightarrow \quad P_i\,(\,[A!]\phi \mid A\,) = k$$

General public actions a as in Section 5 validate a similar principle, but with a superscript ^{a} referring to the product conditional probability introduced there:

$$[a]\, P_i(\phi) = k \quad \leftrightarrow \quad P_i{}^a\,(\,[a]\phi \mid PRE_a\,) = k$$

With general epistemic action diagrams $\boldsymbol{A}$, we need a generalization to a suitable version of conditional probability involving the whole diagram, in the format:

$$[a]\, P_i(\phi) = k \quad \leftrightarrow \quad P_i{}^A(\bigvee_b [b]\phi \mid \bigvee_b PRE_b\} = k$$

$$\textit{with index } b \textit{ ranging over } \{\, b\sim_i a \text{ in } \boldsymbol{A}\,\}$$

Slightly neater formulations arise when we add notation to standard dynamic logic, such as explicit world-dependent function symbols $P_{i,\, s,\, [a]}$ for probabilities.

Given these observations, it becomes a routine exercise to generalize the complete axiom system for probabilistic public update in Kooij 2001 to the general setting of epistemic actions with product update of Baltag, Moss & Solecki 1999.

9 Comparison with Bayesian update

What becomes of the usual Bayesian calculus in this setting? The probabilistic notation $P_i\,(\phi \mid A)$ looks updatish as it is, and an update logic is nothing but a more systematic calculus for making this dynamics more explicit. Of course, there are also some differences. For a start, take the syntax. For us, ϕ was a static *proposition*, while *A* was more properly viewed as a dynamic *action "A"*. But that is just the point of a dynamic language, which handles diverse expressions like:

$$P_i\,(\phi) = k\;, \qquad P_i\,(\,["A"]\phi) = k\;, \qquad ["A"]\; P_i\,(\phi) = k$$

Our system relates these, by stating how to compute posterior probabilities in terms of prior probabilities before the action took place, using rules like

$$[a]\; P_i\,(\phi\,) = k \quad \leftrightarrow \quad P_i^{\,a}\,(\,[a]\phi \mid PRE_a\,) = k$$

Next, let us look at the way conditional probability works in practice. As an illustration, take again Monty Hall. Here is what most people would consider the canonical solution, by a simple appeal to the *standard* Bayes Rule:

$$P\,(B\,|A) \qquad = \qquad P\,(A\,|B) \bullet P(A)\;\,/\,P(B)$$

Set A = "The car is behind door 1", B = "The quizmaster opened door 3".

Then $\qquad P(A) = 1/3\;,\; P\,(A\,|B) = 1/2\;,\; P(B) = 1/2:$
et voilá: $\qquad P\,(B\,|A) = 1/3\;!$

But to analyze what these lines mean, one needs to specify the probability space, and justify the postulated *P*-values. Two of these are simple: *P(A)* is a given prior probability, and *P(B)* refers to the given probabilities of all possible actions. But what justifies the stated value *1/2* for *P(A|B)* ? The relevant probability space which most people seem to have in mind here (as in the earlier tree) is not a simple subspace of the initial one with car states. Its worlds are rather ordered pairs of

<*car state; action taken*>

But if so, our product update model is close to practice, since that is what it says!

More generally, the Product Rule of Section 5 is like computing a probabilistic update using a *prior* and a *likelihood function* over possible events (cf. Good 1965, Hirshleifer & Riley 1992). Thus, it converges with a powerful theoretical paradigm in probability theory. Against this background, the above analysis then adds a systematic view of dynamic-epistemic model transformations on probability spaces, with a well-understood logical mechanism for reasoning about it.

But what about Bayes' Law? Let us analyze some working principles of conditional probability in our update logic. First, here is the basic definition:

$$P(\phi \mid A) \bullet P(A) \quad = \quad P\,(\phi \,\&\, A)$$

This has the following counterpart in dynamic-probabilistic logic – assuming for simplicity that the proposition *A* is a precondition for *action* "*A*":

$$["A"]\, P_i(\phi) = k \ \ \&\ \ P_i(A) = l \quad \rightarrow \quad P_i(\, ["A"]\phi \,\&\, A) = k \bullet l$$

But Bayes' Rule itself, the main engine of probabilistic update, is more problematic, as it *inverts* the order of action:

$$P_i(\phi \mid A) \quad = \quad P_i(A \mid \phi) \bullet P_i(\phi) \ / \ P_i(A)$$

In a dynamic setting like update logic, public announcements of even pure epistemic statements can have different effects when they are made in different orders. Thus, order inversions stating when the announcement *A* is true after announcing its effect ϕ violate the spirit of updating! In fact, for general epistemic assertions, Bayes' Law fails. Here is an illustration, with a familiar formula from the literature:

Example Epistemic failure of Bayes' Law
Consider the following epistemic model with two agents:

$$\begin{tikzcd} \boldsymbol{p,\ q} \arrow[r, dash, "you"'] \arrow[d, dash, "me"] & \neg p,\ \neg q \\ \neg p,\ q \end{tikzcd}$$

The actual world has *p, q* both true. Now consider the assertions

A 'you do not know if *p* is the case',

which is true in the two uppermost worlds, but not at the bottom. Next, take

ϕ 'I know if *p* is the case'

which is only true in the world to the right. Thus, in this model $P(A) = 2/3$, while $P(\phi) = 1/3$. A public update with the assertion *A* takes this model to the new

$$\begin{tikzcd} \boldsymbol{p,\ q} \arrow[r, dash, "you"'] & \neg p,\ \neg q \end{tikzcd}$$

where ϕ holds everywhere:

$$P(\phi \mid A) = 1$$

An update with ϕ takes the initial model to the one-world model

$\neg p, \ \neg q$

where A is false everywhere:

$P(A \mid \phi) = 0$

Substituting, we see that Bayes' Law fails:

$P(\phi \mid A) = 1 \neq (0 \cdot 1/3) / 2/3$

The update-logical status of Bayes' Rule seems to be this. Order inversions are invalid in general, but they are admissible when the relevant assertions are simple enough – like non-epistemic announcements of atomic facts. Despite this slight snub, the Rule is widely useful, and it has lived happily for centuries without logical underpinning. Our analysis has put it in perspective, not called into doubt.

10 Conclusion: reasoning with probability

This note is about a dynamic take on probabilistic reasoning, going one step further in an existing line of research. What it proposes in a nutshell is that marrying epistemic product update with probabilistic conditionalization produces a more principled joint account of both model change and probability adjustment.

More generally, this analysis touches on two problems people have in probabilistic reasoning. The first has often been observed. Usually, people do not miscalculate, but they misidentify the relevant model. E.g., in Monty Hall, many people compute the conditional probability with respect to the fact that *the car is not behind door 3*, which yields probability 1/2 for its being behind door 1, and switching is useless. There is nothing wrong with this reasoning per se, as it is indeed the correct update for a public announcement or observation that the car is not behind door 3. The problem is rather the choice of the model. Keeping track of the right ambient models is made easier by update mechanisms. But there is always a non-automatic feature. This now becomes finding the relevant actions with their preconditions and probabilities. Some training with the pet examples of update logic might help here.

But good frameworks should not just moralize: they should also predict and *explain* reasoning failures. Here is an illustration, which also suggests a separation of concerns. If probabilistic update involves epistemic dynamics, one would expect that people's problems with it are a mixture of known purely dynamic difficulties and genuine probabilistic ones. Again in Monty Hall, intuitively, we tend to look at *postconditions* of observed actions: what holds once they have been performed.

Opening door 3 certainly reveals the car's not being there. But our update analysis (1), (2), (4) says it is rather the *preconditions* of the relevant actions which count: what has to hold beforehand for them to be executable at all. This may be a hard distinction for human agents. Difficulties in probabilistic reasoning might depend on our faint grasp of 'timing' and dynamics, rather than of probability per se.

Finally, a sweeping statement behind our proposal at a conceptual level. *Events* are a key term in probability theory, but their static modeling as sets of outcomes is *wrong*. They should be taken seriously as *dynamic* actions that change states!

Acknowledgment This note was written in June 2002, after a Stanford classroom discussion of Barteld Kooij's 2001 paper on probabilistic update for public announcement. I thank Robert van Rooy for useful critical comments.

References

A. Baltag, L. Moss & S. Solecki, 1998, 'The Logic of Public Announcements, Common Knowledge and Private Suspicions, *Proceedings TARK 1998*, 43–56, Morgan Kaufmann Publishers, Los Altos.

A. Baltag, L. Moss & S. Solecki, 2002, Updated version of the preceding, department of cognitive science, Indiana University, Bloomington, and department of computing, Oxford University.

J. van Benthem, 2001, 'Games in Dynamic Epistemic Logic', in G. Bonanno & W. van der Hoek, eds., *Bulletin of Economic Research* 53:4, 219–248.

J. van Benthem, 2002, 'Conditional Probability is Update Logic', lecture note, department of philosophy, Stanford University.

H. van Ditmarsch, 2000, *Knowledge Games*, dissertation DS-2000-06, Institute for Logic, Language and Computation, University of Amsterdam.

R. Fagin, J. Halpern, Y. Moses & M. Vardi, 1995, *Reasoning about Knowledge*, MIT Press, Cambridge (Mass.).

I. Good, 1965, *The Estimation of Probabilities*, MIT Press, Cambridge, MA.

J. Hirshleifer & J.G. Riley, 1992, *The Analytics of Uncertainty and Information*, Cambridge University Press, Cambridge.

J. Kemeny, J. Snell & G. Thompson, 1974, *Introduction to Finite Mathematics*, Prentice Hall, Englewood Cliffs (N.J.).

B. Kooij, 1999, The Monty Hall Dilemma, Master's thesis, department of philosophy, University of Groningen.

B. Kooij, 2001, 'Probabilistic Dynamic Epistemic Logic', Department of computer science, University of Groningen.

R. van Rooy, 2002A, 'Quality and Quantity of Information Exchange', ILLC Amsterdam. To appear in *Journal of Logic, Language and Information*, special issue on theories of information.

R. van Rooy, 2002, 'Utility, Informativity, and Protocols', in G. Bonanno et al., eds., *Proceedings Logic and the Foundations of the Theory of Games and Decisions, LOFT-5*, Torino.