

Dynamic Update with Probabilities

Johan van Benthem* Jelle Gerbrandy† Barteld Kooi‡

April 15, 2008

Abstract

Current dynamic-epistemic logics model different types of information change in multi-agent scenarios. We propose a way to generalize these logics to a probabilistic setting, obtaining a calculus for multi-agent update with different slots for probability, and a matching dynamic logic of information change that has a probabilistic character itself. We present a general completeness result that not only holds for the particular logical system set out in this paper, but for a larger class of dynamic probabilistic logics as well. Finally, we discuss how our basic update rule can be parameterized for different ‘update policies’.

1 Introduction

Conditional probabilities $P_i(\varphi \mid A)$ describe how agent i ’s probability distributions for propositions φ change as new information A comes in. The standard probabilistic calculus describing such changes revolves around Bayes’ Law in case the new information A is factual, concerning some actual situation under investigation. But there are also proposed mechanisms in the literature that deal with non-factual new information A , such as the Jeffrey Update Rule for probabilistic information of the form “ $P_i(A) = x$ ” and Dempster’s rule for combining evidence.

Current dynamic-epistemic logics manipulate formulas $[!A]K_i\varphi$ describing what an agent knows or believes after a proposition A has been publicly communicated or publicly observed. Here A may be either about the real world or about information that other agents have. More sophisticated modern update systems deal with many further scenarios, which involve partial observation and different information for different agents, as happens with whispers, or lies, or just seeing some situation from different angles. Thus, it seems of interest to combine the two perspectives to obtain one system for reasoning about interaction of knowledge and probability in a system where both may change.

This paper takes its point of departure in two earlier attempts at achieving such a combination. Kooi (2003) provides a complete dynamic logic of probabilistic update after public announcements, and van Benthem (2003) has one for probabilistic update after publicly observed events with known probabilities for their occurrence. These two aspects come together in non-trivial scenarios such as the ‘Quizmaster’ (or ‘Monty Hall’), where the participants in a quiz observe a door being opened by the quizmaster, and must recompute the probability that the car is behind the door they have chosen originally. In doing so, two probabilities play a crucial role. Their own prior probability for the car being behind any door matters, but so does their knowledge of the ‘process’, viz. the probability that the quizmaster would have opened this door, given his knowledge of the door behind which the prize car

*University of Amsterdam, Stanford University, johan@science.uva.nl

†Università di Torino, jelle@gerbrandy.com

‡University of Groningen, B.P.Kooi@rug.nl

is located. The main new contribution of this paper is a still more comprehensive view of probabilistic update from a dynamic-epistemic perspective, identifying not two, but *three* crucial probabilistic aspects of incoming information, including what we call *observation probabilities*. All three are then used to provide a generalized update rule, whose logic can be axiomatized completely.

The paper is organized as follows. The first two sections cover our point of departure: in Section 2 we present a static epistemic-probabilistic logic, and in Section 3 we give purely dynamic-epistemic logics. In Section 4 we give our full probabilistic update rule involving all three aspects: prior probability, occurrence probability, and observation probability. In Section 5 we turn to reasoning about these updates, and prove a general completeness result for dynamic-epistemic-probabilistic logic. Finally, in Section 6, we discuss how our update rule can be generalized still further so as to allow for different ‘policies’ or ‘agent types’. We provide a mechanism for weighing the different sources of probabilistic information available to us. While our presentation is self-contained, we refer to the extended on-line version (van Benthem, Gerbrandy, and Kooi 2006) for details and additional topics.

2 Static epistemic-probabilistic logic

Epistemic and probabilistic languages describing what agents know and believe plus the probabilities they assign were introduced by Halpern and Tuttle (1993) and further developed by Fagin and Halpern (1993). We take a simple instance of such a system as our starting point.

Definition 1 (Epistemic probability models) *Given is a set of agents Ag and a set of propositional variables At . An epistemic probability model is a structure $M = (S, \sim, P, V)$ such that*

- S is a non-empty set of states,
- \sim is a set of equivalence relations \sim_i on S for each agent $i \in Ag$,
- P is a set of probability functions P_i , which assign for each agent $i \in Ag$ and each state $s \in S$,
- V assigns a set of states to each propositional variable.

So in these models both the non-probabilistic information and the probabilistic information of the agent is represented (by \sim_i and P_i respectively). This is reflected in the semantics by two modal operators for these notions.

Definition 2 (Static epistemic-probabilistic language) *The static epistemic-probabilistic language is given by the following Backus-Naur form:*

$$\varphi ::= p \mid \neg\varphi \mid \varphi \mid K_i\varphi \mid P_i\varphi = k$$

where $p \in At$, $i \in Ag$, and $k \in \mathbb{Q}$.

This language allows for mixed formulas such as:

$$K_iP_j(\varphi) = k, \text{ or } P_i(K_j\varphi) = k.$$

In this way, we can talk about agents’ knowledge of each other’s probabilities, or about the probabilities they assign to the fact that someone knows some proposition. The semantics of this language are an extension of the semantics for epistemic logic, where the probability statements $P_i(\varphi) = k$ over propositions φ are evaluated by summing over those states in M where φ holds

Definition 3 (Semantics for epistemic-probabilistic logic)

$$\begin{array}{ll}
M, s \models p & \text{iff } s \in V(p) \\
M, s \models \neg\varphi & \text{iff } M, s \not\models \varphi \\
M, s \models \varphi \wedge \psi & \text{iff } M, s \models \varphi \text{ and } M, s \models \psi \\
M, s \models K_j\varphi & \text{iff for all } t \in S: \text{ if } s \sim_i t, \text{ then } M, t \models \varphi \\
M, s \models P_i(\varphi) = k & \text{iff } \sum_{t \text{ with } M, t \models \varphi} P_i(s)(t) = k
\end{array}$$

The definition of epistemic probability models leaves room for further constraints on the relation between the probability assignments and the knowledge of the agents as defined by \sim . For example, it may be reasonable to ask that probability assignments are *uniform* in the sense that if $P_i(s)(t)$ is positive, then $P_i(s) = P_i(t)$, or to assume that the probability assignments are related to the knowledge of the agents, e.g. by assuming that $P_i(s)$ assigns positive probabilities only to states that are in the \sim_i -equivalence class of s . Such assumptions define classes of models with different logics. For example, in many natural applications epistemically indistinguishable states get the same probability distribution. Thus, agents will know the probabilities they assign to propositions, and hence we have a valid principle

$$P_i(\varphi) = k \rightarrow K_i P_i(\varphi) = k$$

that is, ‘epistemic introspection’ holds for subjective probability. Fagin and Halpern (1993) and Halpern (2003) provide excellent overviews with completeness and complexity results for various systems.

What we will have to say about the dynamics of probabilistic update does not hinge on this type of decisions about the class of models, which is why we assume as little as possible. That said, our model could also have even been more general: we could have generalized $P_i(s)$ to be a σ -algebra over a subset of S , as in Fagin and Halpern (1993), or represent insecurity about probabilities by upper and lower bounds, as in Dempster-Shafer theory, also discussed in Halpern (2003). We believe our account extends to these cases as well, but as our main emphasis is on dynamic update phenomena, the preceding simple system suffices to make our point.

3 Dynamic-epistemic logics for non-probabilistic information update

Dynamic-epistemic logics describe information flow engendered by events. The simplest informative event, and a pilot case for much of the theory, is a public announcement $!A$ of some true proposition A to a group of agents. Updates for more complex communicative events can be described in terms of ‘update models’, which model more complex patterns of access that agents may have to the even currently taking place. While much of the theory has been developed with conversation and communication in mind, it is important, also for our later probabilistic applications, to stress that we are not doing some sort of formal linguistics. The formal systems we will be dealing with apply just as well to observation, experimentation, learning, or any sort of information-carrying event. The logics of both public and more private informational scenarios are discussed below.

3.1 Public announcements

The dynamic effect of a public announcement is to change some current (non-probabilistic) model $M = (S, \sim, V)$ to an updated model $M|A$, which is defined by restricting the states of M to just those where A is true.

$$\begin{aligned}
[!A]p &\leftrightarrow (A \rightarrow p) \\
[!A]\neg\varphi &\leftrightarrow (A \rightarrow \neg[!A]\varphi) \\
[!A](\varphi \wedge \psi) &\leftrightarrow ([!A]\varphi \wedge [!A]\psi) \\
[!A]K_i\varphi &\leftrightarrow (A \rightarrow K_i[!A]\varphi)
\end{aligned}$$

Table 1: Reduction axioms for public announcement logic.

A public announcement is usually very informative. Hence, the truth values of epistemic statements can change due to an announcement. E.g., I did not know that A before, but I do now, after I learned that B . These truth value changes can be quite subtle, witness the existence of self-refuting true statements, such as “You don’t know that p , but p is true”, which become false upon public announcement. Therefore we need a dynamic-epistemic language, whose logic helps us keep careful track of things over time.

First, one adds a ‘dynamic’ modal operator $[!A]$ to the epistemic language.

Definition 4 (Public announcement language) *The public announcement language is given by the following Backus-Naur form:*

$$\varphi ::= p \mid \neg\varphi \mid \varphi \wedge \varphi \mid K_i\varphi \mid [!A]\varphi$$

A formula $[!A]\varphi$ is read as ‘ φ holds after the announcement that φ ’. The resulting language is interpreted in standard models for epistemic logic $M = (S, \sim_i, V)$. (These models are epistemic probabilistic models from Definition 1, without the set of probability functions). The semantics for this language is the same as in Definition 3 in as far as the languages have the same operators. The public announcement operator has the following semantics,

Definition 5 (Semantics of public announcements) *Let an epistemic model $M = (S, \sim, V)$ be given, where $s \in S$.*

$$M, s \models [!A]\varphi \quad \text{iff} \quad M, s \models A \text{ implies } M|A, s \models \varphi$$

where $M|A$ is the model (S', \sim', V') such that, writing $\llbracket A \rrbracket$ for $\{t \in S \mid M, t \models A\}$:

- $S' = S \cap \llbracket A \rrbracket$,
- $\sim'_i = \sim_i \cap (S' \times S')$,
- $V'(p) = V(p) \cap S'$.

A complete dynamic-epistemic logic *PAL* for public announcement was first given by Plaza (1989), and was independently developed by Gerbrandy (1998). It exemplifies a typical set-up for dynamic-epistemic analysis. There is a complete set of axioms for the static base language over epistemic models – the logic *S5*, for example – and on top of that, a number of *reduction axioms* that analyze effects of informational events. These axioms describe the effects of an announcement by relating what is true *after* to what is true *before* an announcement.

The axioms given in Table 1 describe how public announcement operators interact with other logical operators. Note how these axioms move each logical operators of the static language outside the scope of the new operator $[!A]$. Thus, they perform a compositional analysis of the effects of receiving information. As a side-effect, working inside out in a stepwise manner, such a ‘recursion equation’ allows us to translate any sentence from the dynamic language into an equivalent sentence of the underlying static language – always provided the latter has enough expressive power to do the necessary ‘pre-encoding’.

This expressive harmony between the static and dynamic parts of the system is not always obvious, and we may have to redesign the static base language to achieve it. For instance, conditional probabilities are crucial for this purpose, and later on, we also need the ‘linear inequalities’ of Halpern and Tuttle (1993) and Fagin and Halpern (1993). Much more can be said about this methodology (cf. van Benthem, van Eijck, and Kooi 2006), but the main point for our paper is just this: Once the design is right, for any class of epistemic models with a complete set of axioms in the static language, a completeness result for the extended language comes ‘for free’. We just need to identify the right ‘recursion equations’, in the form of reduction axioms like the above as schemata, plus the rule of substitution of equivalents: from $\varphi \leftrightarrow \psi$, infer that $\chi \leftrightarrow \chi'$, where χ' is obtained from χ by replacing an occurrence of φ by ψ . Below, we will formulate such reduction axioms for a suitably designed dynamic probabilistic language, and obtain the same kind of dynamic completeness result.

3.2 Update models

Baltag, Moss, and Solecki (1998) first introduced more general update models.

Definition 6 (Update models) *Given is a set of agents Ag and a logical language \mathcal{L} . An update model is a structure $A = (E, \sim, \text{pre})$ such that*

- E is a non-empty finite set of events,
- \sim is a set of equivalence relations \sim_i on E for each agent $i \in Ag$,
- pre assigns a formula from \mathcal{L} to each event $e \in E$.

The ‘precondition function’ pre determines in which states the events can actually occur by assigning a formula (pre_e) to each event in E .

These models are quite similar to epistemic models, but instead of information about static situations, information about *events* is modeled. The indistinguishability relations \sim over events model insecurity about which event actually happens in the same way that the relations in the static models model ignorance about situations: $e \sim_i e'$ can be read as ‘given that event e occurs, it is consistent with agent i ’s information that event e' occurs.’ The result of an event represented by A occurring in a situation represented by M is modeled by means of a *product construction*.

Definition 7 (Update rule) *Let M be an epistemic model and let A be an update model. The product update model $M \times A = (S', \sim', V')$ is defined by setting:*

- $S' = \{(s, e) \mid s \in S, e \in E \text{ and } M, s \models \text{pre}_e\}$,
- $(s_1, e_1) \sim'_i (s_2, e_2)$ iff $s_1 \sim_i s_2$ and $e_1 \sim_i e_2$.
- $V'(p) = \{(s, e) \in S' \mid s \in V(p)\}$.

The indistinguishability relation in $M \times A$ is determined by the indistinguishability relations in M and A . An agent cannot distinguish a pair (s_1, e_1) from (s_2, e_2) in the new model if the agent could not distinguish s_1 from s_2 in the old model and could not distinguish event e_1 from e_2 . Note that truth values of propositional variables do not change due to an epistemic event: the propositional variables true in (s, e) are those true in s ¹.

Again, there is a dynamic-epistemic language and a matching complete dynamic logic to reason about product updates.

¹This mechanism can easily be generalized to include an account of factual change in the world: cf. (van Benthem, van Eijck, and Kooi 2006).

$$\begin{aligned}
[A, e]p &\leftrightarrow (\text{pre}_e \rightarrow p) \\
[A, e]\neg\varphi &\leftrightarrow (\text{pre}_e \rightarrow \neg[A, e]\varphi) \\
[A, e](\varphi \wedge \psi) &\leftrightarrow ([A, e]\varphi \wedge [A, e]\psi) \\
[A, e]K_i\varphi &\leftrightarrow (\text{pre}_e \rightarrow \bigwedge_{e' \sim_i e} K_i[A, e']\varphi)
\end{aligned}$$

Table 2: Reduction axioms for update models.

Definition 8 (Dynamic-epistemic language) *The syntax of the dynamic-epistemic language is given by the following Backus-Naur form:*

$$\varphi ::= p \mid \neg\varphi \mid \varphi \wedge \psi \mid K_i\varphi \mid [A, e]\varphi$$

In this language the update models are update models with respect to language defined above. There is a harmless simultaneous recursion here: update models are defined in terms of the language, and the dynamic language is defined in terms of update models, just as formulas and programs are defined simultaneously in PDL. This language is also interpreted on epistemic models.

Definition 9 (Semantics of update models) *Given an epistemic model $M = (S, \sim, V)$ with $s \in S$:*

$$M, s \models [A, e]\varphi \quad \text{iff} \quad M, s \models \text{pre}_e \text{ implies } M \times A, (s, e) \models \varphi$$

where $M \times A$ is the product update model.

A formula of the form $[A, e]\varphi$ states that, if event e can occur, then φ is true in the result. Next, there is the issue of valid reasoning. As before, we get a simple superstructure on top of whatever valid principles we had for the static base language – this may be multi-agent $S5$, but it does not have to be. The axiomatization is a straightforward generalization of the earlier one for the logic of public announcements. The reduction axioms are given in Table 2. A growing literature shows how this product update can model a wide variety of informational scenarios (Baltag 2002; Himmerson, Meyer, and de Vink 2005; van Benthem, van Eijck, and Kooi 2006; van Ditmarsch, van der Hoek, and Kooi 2007).

4 Modeling probabilistic information change

Extensions of dynamic epistemic logic with probabilistic information have been proposed, as mentioned earlier, in Kooi (2003), on probabilistic update after public announcement, and van Benthem (2003) on probabilistic update after publicly observed events with known probabilities for their occurrence. Both papers also provide dynamic update rules as well as matching complete logics, which are of independent interest. But in this paper, we forego such details, and move straight ahead to our new proposals mostly subsuming these. The first step is a more comprehensive view of three crucial probabilistic aspects of incoming information, which will feed into a generalized epistemic-probabilistic update rule.

4.1 Three sources of probability

Taken together, Kooi’s and van Benthem’s earlier approaches perform a two-fold ‘probabilization’ of *DEL*-style product update. The proposed update mechanism distinguishes (a) *prior probabilities of worlds* in the current epistemic-probabilistic model M , representing agents’ current informational attitudes, and (b) *occurrence probabilities for events* from the update model A , representing agents’ views on

what sort of process produces the new information. Both kinds of probability play a role in the well-known Monty Hall Puzzle, as indicated briefly in our Introduction.

These are certainly two natural places to look for probability. But there is also a third type of uncertainty, that may be called

- (c) *observation probability*, reflecting agents' uncertainty as to which event is currently being observed.

For a simple example, suppose that I see you reading a letter from our funding Agency, and I know it is either a rejection of your grant proposal or an acceptance. You know which event (reading 'yes', or reading 'no') is taking place, while I do not. If I know nothing more than this, and I have no idea about the frequency of rejection versus acceptance letters, pure epistemic product update might compute a new model, but it will not fully indicate *to what extent* I should consider a state in that model possible. But now suppose that there is additional information in my observation. Perhaps I saw a glimpse of your letter, or you looked smug, and I therefore assume that you were probably reading a letter of acceptance rather than a rejection. This would be a case of 'observation probability' in our sense. The notion of observation probability is well-known from scenario's used to motivate Jeffrey conditioning, where one is uncertain about the evidence one receives due to partial observation and such.

A simple but not unrealistic scenario where all three kinds of probability distinguished here come together is as follows.

Example 1 (The Hypochondriac) *Suppose that you are reading about some horrible disease on a website, and you start to wonder whether you have this disease. The chances of having this disease are very slight, say only 1 in 100.000. The website states that one of the symptoms of this disease is that a certain gland is swollen. If you have the disease, the chance that this gland is swollen is 97%, whereas if you do not have the disease, the chances are 0 that it is swollen. You immediately start to examine the gland. The problem is that it is hard to determine whether this gland is swollen or not. It is the first time you are actually examine the gland and – not being a physician – you do not know what its size ought to be. You are uncertain, but you think the chances are 50% that the gland is swollen. What are the chances you should assign to having the disease?*

We will now define a general mechanism for computing the answer.

4.2 Update models and probabilistic product update

For a start, our static epistemic-probabilistic models M are still the same as before, and so is our epistemic-probabilistic language. We will also continue using the earlier notation $[A, e]\varphi$ for the effects of executing an update model (A, e) in the current epistemic probabilistic model M . Our first task is to define appropriate probabilistic update models. For this purpose, we redefine the earlier update models, so as to make them look more like processes consisting of various events with uniformly specified occurrence and observation probabilities:

Definition 10 (Probabilistic update models) *Probabilistic update models are structures $A = (E, \sim, \Phi, \text{pre}, P)$ where:*

- E is a non-empty finite set of events,
- \sim is a set of equivalence relations \sim_i on E for each agent $i \in \text{Ag}$,
- Φ is a set of pairwise inconsistent sentences called preconditions,

- *pre* assigns to each precondition $\varphi \in \Phi$ a probability distribution over E (we write $\text{pre}(\varphi, e)$ for the probability that e occurs given φ),
- For each i , P_i assigns to each event e a probability distribution over E .

The formal language for the preconditions Φ is left open, but it will be formally defined in Section 5. Just as for ordinary update models there is a harmless simultaneous recursion here.

The definition should be understood as follows. Part of the models consists in the specification of ‘occurrence probabilities’ of a process which makes events occur with certain probabilities, depending on a set of conditions Φ . Such a process is captured by the function *pre*. Diseases and quizmasters are such processes, that follow rules of the form “if P holds, then do a with probability q ”. But one can also think of Markov Processes or other standard probabilistic devices. The ‘evidence spaces’ of Halpern and Pucella (2006) that connect hypotheses with a space of possible observations are very similar as well.

The second component of the models are the ‘observation probabilities’ represented by the functions P_i . The probability $P_i(e)(e')$ is the probability assigned by the observer i as to event e' taking place, given that e actually takes place. These probability functions add a probabilistic structure to the uncertainty relations \sim_i in much the same way as they do in the static models².

Our next task is defining a dynamic update rule using these models. Merging the input from all three sources of probability, computing an update is a straightforward generalization of earlier rules:

Definition 11 (Probabilistic Product Update Rule) *Let M be an epistemic-probabilistic model and let A be a probabilistic update model. If s is a state in M , write $\text{pre}(s, e)$ for the value of $\text{pre}(\varphi, e)$, where φ is the element of Φ that is satisfied in M, s . If no such φ exists, $\text{pre}(s, e) = 0$.*

Now, the product update model $M \times A = (S', \sim', P', V')$ is defined by setting:

- $S' = \{(s, e) \mid s \in S, e \in E \text{ and } \text{pre}(s, e) > 0\}$
- $(s, e) \sim'_i (s', e')$ iff $s \sim_i s'$ and $e \sim_i e'$
- $P'_i((s, e), (s', e')) :=$

$$\frac{P_i(s)(s') \cdot \text{pre}(s', e') \cdot P_i(e)(e')}{\sum_{s'' \in S, e'' \in E} P_i(s)(s'') \cdot \text{pre}(s'', e'') \cdot P_i(e)(e'')} \text{ if the denominator } > 0$$

and 0 otherwise

- $V'((s, e)) = V(s)$

So, the new space of states after the update consists of all pairs (s, e) such that event e occurs with a positive probability in s (as specified by *pre*). The indistinguishability relations are defined just as before.

The most interesting part is the definition of the new probability measures, and it reflects our earlier intuition of the Hypochondriac example. The new probabilities $P'_i(s, e)$ for (s', e') are the arithmetical *product* of the prior probability for s' , the

²One more argument in favor of distinguishing these various probabilities may be this. One might think of an occurrence probability more in terms of *objective* frequencies, and of observation probabilities more as *subjective* chances. Thus, our perspective allows for natural co-existence of both major views of probability within the same scenario.

probability that e' actually occurs in s' , and the probability that i assigns to observing e' . To obtain a proper probability measure in the resulting state, we normalize the computed product value. Taking a normalized product of probability measures is similar to Dempster's rule for combining beliefs (with belief functions as probability measures), but our rule does not grind all probabilities together, as usual – but rather separates out the process description and the observation probabilities, while it allows for indexes for many agents in a natural manner³.


Here is how our general update mechanism works out in practice.

Example: The Hypochondriac Again In our example of the hypochondriac feeling a certain gland, the initial hypothesis about the proposition p of having the disease is captured by a prior probability distribution


$$\frac{1}{100000} \quad \frac{99999}{100000}$$

$p \bullet \text{-----} \circ \neg p$

Then the hypochondriac feels whether the gland is swollen, assigning our probabilities regarding the disease (if he has the disease, he has a swollen gland with probability .97) and his power of observation (he thinks the gland is swollen with probability .5) as above, resulting in the following probabilistic update model:



The product of our initial state with this model is as follows:



This diagram is our new probabilistic information state after the whole episode. The probability the Hypochondriac should assign to having the disease is still 1 in 100.000. Since his observation was inconclusive he has not gained any information about whether he the disease or not. Had he found it more probable that the gland was swollen, the probability of having the disease would have been higher than 1 in 100.000, and had he found it more probable that it was not swollen, the probability would have been lower than 1 in 100.000.

More generally, the construction also illustrates another typical feature of product update. We do not just eliminate existing states or change prior probabilities,

³If the denominator in the definition of the new probability measure sums to 0, we just stipulate that the value of the whole division is 0. This means that the model $M \times A$ is not strictly speaking a probabilistic epistemic model: after an update, $P_i(s, e)(\cdot)$ may assign probability 0 to all states. From a strictly formal viewpoint, this is no problem (and the choice is certainly defensible from a probabilistic viewpoint as well, cf. e.g. Bacchus 1990), but for the reader who does not like this feature, there are straightforward ways of circumventing it.

but may also *construct new types of possibility*. Initially, we only considered options for one single aspect of reality: having the disease or not not. After the update, we consider more complex epistemic possibilities, including information about whether the gland is swollen or not. In this way, the number of options may increase even though we have obtained genuine information about the original situation.

4.3 Discussion and further developments

The successive effect of our epistemic probabilistic product update rule is this. We start from a simple probability space and, piece by piece, build up more complex probability spaces using descriptions of informational events given by update models. This control over successive spaces may be useful in practice, where it is the management of relevant spaces, rather than correct application of the calculus of probability which forms the main difficulty in reasoning with uncertainty. Over time, the new possibilities may be viewed as the set of all possible runs of some total informational process, linking up with more global descriptions in terms of epistemic probabilistic versions of temporal logics. The update rule provides a good modeling tool for the analysis of such complex scenarios.

As for connections with the existing literature, we just say this. We have emphasized a view of update models with occurrence probabilities as representing probabilistic *processes*. But there are alternative interpretations. Occurrence probabilities and the way we update with them are also very similar to what Shafer (1982) calls a *parametric model*, and Halpern and Pucella (2006) an *evidence space*. On the latter view, observations constitute evidence for certain hypotheses, and these observations are statistically related to the hypotheses in the way described by the model. Then, our preconditions take the place of hypotheses, our events correspond to observations, and our precondition function correspond to Halpern and Pucella (2006)'s likelihoods. Updating with evidence spaces goes back to Shafer (1976), and they are a special case of our update rule – for a single agent, and without observation probabilities.

We conclude this section with a few more technical points, showing how this setting relates to existing logical themes, while inviting a number of further ones.

Conservative extension. First, we have truly generalized the original non-probabilistic update models that we described before:

Fact 1 *For each non-probabilistic update model A there is a probabilistic update model B such that for each M , if M' is the non-probabilistic model obtained from removing the probability measures from M , then $M' \times A$ is the same model as $M \times B$ with its probability measures removed.*

For the simple, but somewhat laborious proof, we refer to the extended version (van Benthem, Gerbrandy, and Kooi 2006) of this paper.

Model theory and probabilistic bisimulation. Next, one further test of correct generalization is extension of existing system properties. Kooi (2003) proposed a notion of epistemic-probabilistic *bisimulation* which is adequate for our static language. It is easy to show that our product update rule respects such bisimulations between input models, and hence the model theory of our system is still like that of its predecessors.

Shifting loci of probabilistic information. As for new issues, one very natural question to ask is whether the three components of our system: prior world probabilities, and occurrence and observation probabilities on events are really independent. Intuitively, in modeling concrete scenarios one might choose where to

locate things. Indeed, (van Benthem, Gerbrandy, and Kooi 2006) provide a number of technical transformations on update models, showing how under suitable redefinitions of events, occurrence probabilities can absorb observation probabilities, and vice versa. However, in our view, such technical observations do not endanger the intuitive plausibility of our three-source scheme.

Modeling complex temporal protocols. Finally, our update mechanism is more powerful than might appear at first sight. It is sometimes claimed that ‘local updates’ in the dynamic-epistemic style cannot deal with the more complex global *protocols* found in epistemic-temporal logics, describing long-term features of informational processes. Such protocols also occur naturally in probabilistic settings: for instance, an agent whose assertions over time have a certain probabilistic reliability may be viewed as a protocol over time. Now, the relationship between dynamic-epistemic logic and epistemic-temporal logic is somewhat delicate (cf. van Benthem, Gerbrandy, and Pacuit 2007), but here is an example that much more can be done with update models than might be apparent. The key point here is the freedom we have in choosing what we take to be the relevant *events*.

Here is how one can update our information about the *kind of process* that we are observing, instead of just taking that to be a fixed piece of knowledge. Instead of showing this in formal detail, we give a simple example.

Suppose that we know that a coin is either fair, or has only heads. Then we can represent the observation of one throw of the coin with an initial model with two preconditions {fair, heads_only} and two events to be observed, heads and tails, related with the obvious probabilities: $\text{pre}(\text{fair}, \text{heads}) = \text{pre}(\text{fair}, \text{tails}) = \frac{1}{2}$, while $\text{pre}(\text{heads_only}, \text{heads}) = 1$ and $\text{pre}(\text{heads_only}, \text{tails}) = 0$. In this particular scenario, the *iteration* of successive observations is of interest as well – this can be modeled by simply repeating the update a number of times. The model predicts what would be expected: one observation of tails rules out the unfair coin as a possibility, while each observation of heads makes the probability of an unfair coin more likely.

In other scenarios, however, this simple trick will not do. Say, we meet a person telling us something in one of the ubiquitous ‘island puzzles’ beloved by logicians: who might be a truth-teller or a liar. We have to find out what is what (cf. Liu 2008, Chapter 5 for such scenarios in straight epistemic logic). In this setting, it seems much more natural to encode the options inside the update model, introducing more complex structured ‘pair events’

‘(Truth Teller, !A)’, ‘(Liar, !A)’

encoding both the assertion made, and the type of agent making it.

A general event construction with pairs ‘(process type, observed event)’ is found in (van Benthem, Gerbrandy, and Kooi 2006), which shows how, in this manner, the scope of our analysis extends into more general protocol logics⁴.

This concludes the first new contribution of this paper: the distinction between three sorts of input probability, and the definition of one simple epistemic-probabilistic update mechanism based upon it. We now pass to our second main result, the design of a complete dynamic logic for agents reasoning with, or about our Product Update Rule.

5 Dynamic logics of probabilistic update

In order to reason explicitly about probabilistic information change in a dynamic-epistemic format, we must extend existing epistemic probabilistic logics with ap-

⁴Even so, (van Benthem, Gerbrandy, and Pacuit 2007) do propose a full-fledged merge of dynamic-epistemic logic with explicit epistemic temporal protocols in the end.

appropriate dynamic reduction axioms. In this section, we show how this can be done for a logic based on our proposed product update semantics.

5.1 Adding probabilistic inequalities

As explained before, the crucial information about our Product Update Rule will be reflected in recursive ‘reduction axioms’, which state when propositions get certain probabilities after an epistemic event took place. Moreover, we already saw that such axioms express a certain harmony between the dynamic and static parts of an epistemic language. In order to obtain completeness in this style, we crucially need what might look like a mere technical feature of the system of Fagin and Halpern (Halpern and Tuttle (1993), Fagin and Halpern (1993)). They add linear inequalities to the language of epistemic-probabilistic logic:

$$\alpha_1 \times P_i(\varphi_1) + \dots + \alpha_n \times P_i(\varphi_n) \geq \beta$$

where $\alpha_1, \dots, \alpha_n, \beta$ are rational numbers.

Incorporating this feature, here is the total dynamic language that we will use.

Definition 12 (Dynamic-epistemic-probabilistic language) *The dynamic-epistemic-probabilistic language is given by the following Backus-Naur form:*

$$\varphi ::= p \mid \neg\varphi \mid \varphi \wedge \psi \mid K_i\varphi \mid [A, e]\varphi \mid \alpha_1 \cdot P_i(\varphi_1) + \dots + \alpha_n \cdot P_i(\varphi_n) \geq \beta$$

where A is a probabilistic update model, and e an event from the domain of A , while α_k, β stand for rational numbers.

Note again that there is a joint recursion hidden in this set-up: the formulas that define the preconditions in our probabilistic update models come from the same language that we are defining here, but through the dynamic modalities, such models themselves enter the language again. The semantics for $[A, e]\varphi$ is similar to the non-probabilistic case.

Definition 13 (Semantics of probabilistic updates models) *Let an epistemic probability model $M = (S, \sim, P, V)$ be given, with $s \in S$. The key truth clause is:*

$$M, s \models [A, e]\varphi \text{ iff if there is a } \varphi \in \Phi \text{ such that } M, s \models \varphi, \text{ then } M \times A, (s, e) \models \varphi$$

where $M \times A$ is the product update model.

5.2 A complete dynamic epistemic probabilistic logic

With all this in place, here is the main result of this Section:

Theorem 1 *The dynamic-epistemic probabilistic logic of update by probabilistic event models is completely axiomatizable, modulo some given axiomatization of the logic of the chosen class of static models.*

Proof We explain the numerical core idea first. To obtain a complete logic for product update, we must find the key axiom that relates formulas of the form $[A, e]\psi$ with ψ involving probabilities to static assertions with suitable probabilities in the original model (M, s) . The following calculation is the heart of our reduction.

Consider the probability value $P_i(\psi)$ of a formula ψ in a product model $(M, s) \times (A, e)$. In the equations below, we drop some subscripts, exchanging exactness for legibility. We will abbreviate $P_i(s)$ in the initial model with P^M , write $P^{M \times A}$ for the value of $P_i(s, e)$ in the product model, and write P^A for $P_i(e)$ in the action model. Furthermore, we write $\langle A, e \rangle \psi$ for $\neg[A, e]\neg\psi$.

If $\sum_{s'' \in S, e'' \in A} P^M(s'') \cdot \text{pre}(s'', e'') \cdot P^A(e'') > 0$, the following must hold:

$$\begin{aligned}
& P^{M \times A}(\psi) \\
&= \sum_{(s', e') \text{ in } M \times A : M \times A, (s', e') \models \psi} P^{M \times A}(s', e') \\
&= \sum_{s' \in S, e' \in E : M, s' \models \langle A, e' \rangle \psi} P^{M \times A}(s', e') \\
&= \frac{\sum_{s' \in S, e' \in E : M, s' \models \langle A, e' \rangle \psi} P^M(s') \cdot \text{pre}(s', e') \cdot P^A(e')}{\sum_{s'' \in S, e'' \in E} P^M(s'') \cdot \text{pre}(s'', e'') \cdot P^A(e'')} \\
&= \frac{\sum_{s' \in S, e' \in E, M, s' \models \langle A, e' \rangle \psi} P^M(s') \cdot \text{pre}(s', e') \cdot P^A(e')}{\sum_{s'' \in S, e'' \in E} P^M(s'') \cdot \text{pre}(s'', e'') \cdot P^A(e'')}
\end{aligned}$$

The numerator of this last equation can be written as

$$\sum_{\varphi \in \Phi, s' \in S, e' \in E, M, s' \models \varphi, M, s' \models \langle A, e' \rangle \psi} P^M(s') \cdot \text{pre}(\varphi, e') \cdot P^A(e')$$

which is equivalent to

$$\sum_{\varphi \in \Phi, e' \in E} P^M(\varphi \wedge \langle A, e' \rangle \psi) \cdot \text{pre}(\varphi, e') \cdot P^A(e')$$

We can analyze the denominator of the equation in a similar way, and rewrite it as

$$\sum_{\varphi \in \Phi, e'' \in E} P_M(\varphi) \cdot \text{pre}(\varphi, e'') \cdot P_A(e'')$$

In other words, we can rewrite the probability $P^{M \times A}(\psi)$ in the new model as a term of the following form:

$$P^{M \times A}(\psi) = \frac{\sum_{\varphi \in \Phi, e' \in E} P^M(\varphi \wedge \langle A, e' \rangle \psi) \cdot k_{\varphi, e'}}{\sum_{\varphi \in \Phi, e'' \in E} P^M(\varphi) \cdot k_{\varphi, e''}}$$

where, for each φ and f , $k_{\varphi, f}$ is a constant, namely the value $\text{pre}(\varphi, f) \cdot P^A(f)$.

This observation gives us a reduction axiom of sorts. Because both the set of preconditions Φ and the domain of A are finite, we can enumerate them as $\varphi_0, \dots, \varphi_n$ and e_0, \dots, e_m . We can rewrite a formula in which 'P' refers to the probabilities after the update of the form

$$\langle A, e \rangle P(\psi) = r$$

to an equation in which 'P' refers to probabilities in the prior model:

$$\frac{\sum_{1 \leq i \leq n, 1 \leq j \leq m} k_{\varphi_i, e_j} \cdot P(\varphi_i \wedge \langle A, e_j \rangle \psi)}{\sum_{1 \leq i \leq n, 1 \leq j \leq m} k_{\varphi_i, e_j} \cdot P(\varphi_i)} = r$$

which can be rewritten as a sum of terms:

$$\sum_{1 \leq i \leq n, 1 \leq j \leq m} k_{\varphi_i, e_j} \cdot P(\varphi_i \wedge \langle A, e_j \rangle \psi) + \sum_{1 \leq i \leq n, 1 \leq j \leq m} -r \cdot k_{\varphi_i, e_j} \cdot P(\varphi_i) = 0$$

To express these observations as one reduction axiom in our formal language, we need sums of terms to deal with single probability assignments after the update. A

language with only simple equalities cannot do this, and thus it is not in ‘expressive harmony’ in our terms. But our language with linear inequalities is up to the job.

Here are the principles that we need. Concentrating on the only part of our language that is new, we must achieve a reduction, not just for single probability assignments, but also for linear inequalities of these. In order to achieve this, we start with a formula of the form

$$[A, e](\alpha_1 \cdot P(\psi_1) + \dots + \alpha_k \cdot P(\psi_k) \geq \beta)$$

We can replace the separate terms $P(\psi_k)$ after the modal update operator by their equivalents as computed just before. We then obtain an expression of the form

$$\sum_{1 \leq h \leq k, 1 \leq i \leq n, 1 \leq j \leq m} \alpha_h \cdot k_{\varphi_i, e_j} \cdot P(\varphi_i \wedge [A, e_j]\psi_h) + \sum_{1 \leq i \leq n, 1 \leq j \leq m} -\beta \cdot k_{\varphi_i, e_j} \cdot P(\varphi_i) \geq 0$$

This is an expression in the language. To formulate the axiom, then, let us abbreviate this last inequality as χ . The above formulas are equivalent only under the condition that the denominator of the equation that is used to compute the posterior probabilities is greater than 0. The full axiom then becomes:

$$([A, e](\alpha_1 \cdot P(\psi_1) + \dots + \alpha_k \cdot P(\psi_k) \geq \beta) \leftrightarrow \left(\left(\sum_{1 \leq i \leq n, 1 \leq j \leq m} k_{\varphi_i, e_j} \cdot P(\varphi_i) \geq 0 \rightarrow \chi \right) \wedge \left(\sum_{1 \leq i \leq n, 1 \leq j \leq m} k_{\varphi_i, e_j} \cdot P(\varphi_i) = 0 \rightarrow 0 \leq \beta \right) \right))$$

The other reduction axioms are familiar from the non-probabilistic event updates. We only need to formulate the preconditions of an event in the object language. We can define $\text{pre}_{A,e}$ to be the sentence $\bigvee_{\varphi \in \Phi, \text{pre}(\varphi, e) \geq 0} \varphi$. We then have the following set of valid equivalences:

1. $[A, e]p \leftrightarrow (\text{pre}_{A,e} \rightarrow p)$ if p is an atomic formula
2. $[A, e]\varphi \wedge \psi \leftrightarrow [A, e]\varphi \wedge [A, e]\psi$
3. $[A, e]\neg\varphi \leftrightarrow (\text{pre}_{A,e} \rightarrow \neg[A, e]\varphi)$
4. $[A, e]K_i\varphi \leftrightarrow (\text{pre}_{A,e} \rightarrow \bigwedge_{e \sim_i f} K_i[A, f]\varphi)$

Our theorem now follows by the usual argument. Applying the reduction axioms, each formula in the extended dynamic epistemic probabilistic language is provably equivalent to one in the base language, and hence it suffices to prove its static equivalent in the complete language of Halpern and Fagin. \square

Note that our methodology via reduction axioms yields a relative, rather than an absolute axiomatization of the full dynamic language. One can take any base system of reasoning about probabilities for the chosen static models, and the reduction axioms will then also allow for reasoning about effects of dynamic actions on top of that. The conditions under which the theorem applies are relatively few. First of all, the base logic should be able to express the above type of linear inequalities. The factors in these inequalities should be able to capture the probability values in the update models, because these turn up as $k_{\varphi, s}$ in the axiom. Moreover, the base logic should be formulated carefully, because uniform substitution does not hold in the dynamic logic – given the special reduction axiom for atomic formulas. But that should be about it – any reasonable axiomatization for any subclass of our

probabilistic epistemic models leads automatically to a complete axiom system for the dynamic language for update over these models⁵.

This second main result of our paper shows that our framework can formulate rich probability logics, in which information change due to probabilistic events is described explicitly. Moreover, the preceding completeness argument allows us to analyze complex probability updates over a wide variety of static base logics with standard semantical and proof-theoretical tools.

6 Parametrizing the Update Rule

The third and final contribution of this paper is an analysis of possible policies and agent-diversity in epistemic-probabilistic update.

6.1 Inductive logic, policies and weights

Our analysis so far identified three component probabilities that drive information update. But this still leaves out one more major issue, having to do with legitimate diversity of update rules. In the earliest publications on *Inductive Logic* in the 1950s, Carnap (1952) pointed out that update requires another component, viz. a *policy* on the part of agents. We have a current probability distribution, encoded in the model M . We observe a new event, encoded in an update model A . The resulting model will now depend on how much weight agents assign to the two factors: ‘past experience’ versus ‘the latest news’. The result was Carnap’s famous ‘continuum of inductive methods.’ Diversity of update policies is also a key feature in modern Learning Theory (Kelly 1996), and belief revision theory (Gärdenfors and Rott 1995). See also (Liu 2008) on diversity of update policies inside non-probabilistic dynamic-epistemic logic, for agents with different memory capacities or different belief revision habits.

By contrast, our updates in Section 4 essentially assigned equal weight to all factors. For instance, if one thought that $\neg p$ was true with probability 0.9, and now observes an event a which indicates with probability 0.9 that p occurred, then the product rule will give both options (p, a) and $(\neg p, a)$ equal weights 0.81 in the product model $M \times A$. But surely, other policies are possible. A truly conservative person will stick to her prior probabilities, a radical modernist sheds the past whenever something new occurs, and merely follows the observation probabilities in the update model.

Carnap’s continuum of inductive methods modeled compromises between such extremes by assigning weights to the probabilities that go into the Update Rule. These weights seem an independent dimension when modeling updating agents, viz. how they use the evidence that is given by probabilistic update models, and we will make a proposal later on for a rule that allows for variation. But before doing so, let us first consider a radical alternative.

6.2 Jeffrey Update and ‘over-ruling’

Actually, there already exists a well-known alternative probabilistic update rule, which favors new evidence absolutely over the prior probabilities, the so-called *Jeffrey Update*. Consider this example, adapted from (Halpern 2003):

Example 2 (The Dark Room) *An object in a room has one of 5 possible colors, 3 of them light (red, yellow, green), 2 dark (brown, black). We have an initial*

⁵This relative style of axiomatization may even make special sense in quantitative probabilistic settings, since we can ‘factor out’ the possibly high complexity of the underlying numerical reasoning in standard mathematical structures.

probability distribution over these five cases, say, the equiprobability measure. Now we make an observation of the object, and we see that, with probability 3/4, the object must be dark. What are the new probabilities?

Jeffrey Update takes this scenario as an instruction of the following form. The new probability of the object being dark must become 3/4, and that of its being light 1/4. But within those zones, the *relative probabilities* of the five initial cases should remain the same. Thus, the radical intuition behind the Dark Room scenario tells us to do two things:

- Set the probability values of propositions in some partition according to some stipulated values coming from the new observation,
- Stick to the old probability ratios for states within the cells of that partition.

More precisely, the information contained in a Jeffrey Update is given by a pair (Φ, P) of a set of sentences partitioning the logical space and a probability distribution P over Φ . The Jeffrey Update of a probability measure P^{old} with this new information is defined as:

$$P^{\text{new}}(s) = P^{\text{old}}(s \mid \varphi) \cdot P(\varphi)$$

Thus, in this update scenario, the observation of the new signal completely *overrules* any prior information about the sentences in Φ .

It is interesting to compare Jeffrey Update with our Product Update scenario so far. To do so conveniently, we make things comparable by taking an update model with ‘*signal events*’ for the relevant propositions (only one such event can happen in each world), and then assigning them observation probabilities equal to the desired Jeffrey values. Formally, the information represented by (Φ, P) is then easily captured in an event model $A = (\Phi, \sim, \Phi, \text{pre}, P)$ as before with ‘signal events’ for partition members. Here we set $\text{pre}(\varphi, \psi) = 1$ iff $\varphi = \psi$. For instance, with the object in the Dark Room, we have two signals ‘Light’, ‘Dark’, with occurrence probabilities 1 and 0 only, and observation probabilities 1/4, 3/4.

Now, our earlier straight Product Update will not get the same effect here, and it is easy to see why. Its value for the probability that the object is dark will weigh *two* factors: the observation probability, but also the prior probability that the object was dark. This interpolates somewhere between 2/5 and 3/4. And indeed, there may be something to this. The way the Dark Room is described in (Halpern 2003), it is not so clear intuitively that one would *want* to discard the prior in Jeffrey’s manner.

Even so, Jeffrey Update is a widely accepted and interesting rule. It has natural counterparts in belief revision, where ‘lexicographic reordering’ of states according to plausibility on the basis of a new fact A makes all A -states better than all $\neg A$ -states, but inside these two zones, the old comparison order is retained.⁶

Before we do something about this, a methodological comment is in order concerning the scope of update stipulations of the ‘overruling’ kind. Jeffrey Update sets the probabilities of the elements of Φ to certain specified values. This will only work if Φ contains ‘factual’ sentences without probability operators or epistemic operators which are sensitive to model changes. Formulas containing information about current probabilities or epistemic possibilities do not in general remain constant over an update – as we have observed before. This observation highlights a

⁶Jeffrey Update with over-ruling is also the model for the general ‘Priority Update Rule’ of (Baltag and Smets 2007) So, the failure of our Product Update Rule in subsuming such a natural scenario seems a problem to be taken seriously.

matter of ‘temporal’ perspective. *DEL*-style systems describe update through ‘preconditions’: what we learn from observing an event is what was true in order for it to happen. The reduction axioms express this backward-looking feature, analyzing preconditions for assertions. By contrast, Jeffrey Update involves ‘forward-looking’ instructions of the type found in belief revision theory, or *STIT*-type action logics: ‘See to it that A ’, ‘Come to believe that A ’. Thus, the two perspectives toward new beliefs and probabilities are related, but have a somewhat different thrust.

6.3 General weighing: the *ABC* formula

Now suppose that we want to allow agents to give different weights to the three probability factors in our update scenario. This can be done in various ways, but a convenient one would work with three numbers α, β, γ from the interval $[0, 1]$. These numbers represent the respective strength of the three kinds of probabilities in the light of new evidence, with 0 meaning “does not count at all” and 1 representing the judgment that this evidence is at least as good as any other.

Before formulating our weighed update rule, we need to consider more precisely which prior probabilities actually change when we encounter the evidence represented in an update model. An update model is *about* something specific – it represents evidence about the probabilities of the set of preconditions Φ , and no more. Our update rule reflects this, as it essentially only changes probabilities of members of the Φ , and changes the probabilities of other propositions only in so far as it is necessary to accommodate this change. If we, to use our earlier example, choose to give the evidence of medical self-examination a high weight with respect to our prior beliefs, this is no reason to adapt our prior probabilities about unrelated information, say, about where we parked our car yesterday night. In this way, it is similar to Jeffrey Update we discussed above.

More precisely, with our Product Update Rule, we have the following property:

Fact 2 *If worlds s and t satisfy the same precondition, then for all e the ratio of the probability of the sets $\{(s, e) \mid e \in E\}$ and $\{(t, e) \mid e \in E\}$ is the same as the ratio of the probabilities of s and t before the update. More precisely, for propositional $\varphi \in \Phi$ and for each ψ that $P^{\text{new}}(\psi \mid \varphi) = P^{\text{old}}(\psi \mid \varphi)$.*

We want to preserve this property for our weighed update rule. If we assign a low weight to our prior probabilities, we should only do that with regard to the propositions in the relevant set Φ . This can be done by an equation with a numerator of the following form:

$$P_i(s)(s' \mid \varphi_{s'}) \cdot P_i(s)(\varphi_{s'})^\alpha \cdot \text{pre}(s', e')^\beta \cdot P_i(e)(e')^\gamma,$$

where we stipulate that $x^0 = 1$ for all x .

The complete statement of our parametrized update rule then becomes:

Definition 14 (Weighted Product Update Rule)

$$P'_i((s, e), (s', e')) := \frac{P_i(s)(s' \mid \varphi_{s'}) \cdot P_i(s)(s')^\alpha \cdot \text{pre}(s', e')^\beta \cdot P_i(e)(e')^\gamma}{\sum_{s'' \in S, e'' \in E} P_i(s)(s'' \mid \varphi_{s''}) \cdot P_i(s)(s'')^\alpha \cdot \text{pre}(s'', e'')^\beta \cdot P_i(e)(e'')^\gamma}$$

if the denominator > 0 and 0 otherwise

To understand the power of this mechanism, one can consider a number of special cases of interest. First of all, setting all three weighing factors to 1 gives our original product update.

Next, setting $\alpha, \beta, \gamma = (1, 0, 0)$ effectively ‘binarizes’ the new evidence: all events that can occur will occur with *equal* probability at each state. This does not mean

exactly that all new evidence is ignored – states that were eliminated by the unweighed update will still be eliminated. What it does mean that all *probabilistic* evidence is ‘flattened.’ As a special case, it follows that, if no state in the prior model is eliminated by the update, then it produces an epistemic product model $M \times A$ where the summed probability of worlds (s, e) in the product model is the same as the probability of s in M . This conservatively copies the prior onto the new model, and distributes the probability of s evenly over the new states (s, e) .

Also of interest is the case $\alpha, \beta, \gamma = (0, 0, 0)$. Here we ignore *all* evidence pertaining to Φ – not just the new evidence, but also the prior evidence pertaining to the elements of Φ (“Now that I have heard this, I don’t know what to think anymore”). If the update does not eliminate any states, then in the resulting product model, all propositions in Φ have become equally probable.

Finally, setting $\alpha, \beta, \gamma = (0, 0, 1)$ is the opposite, extremely radical, policy where the observation probabilities for e determine the probabilities for worlds (s, e) . This mimics (and indeed, it generalizes) the Jeffrey Update for preconditions that do not contain probability statements or epistemic operators.

Fact 3 *The weighted $\alpha, \beta, \gamma = (0, 0, 1)$ update rule is Jeffrey Update.*

Proof We compute as follows. If we omit the 0 factors, we have

$$P'_i((s, e), (s', e')) := \frac{P_i(s)(s' \mid \varphi_{s'}) \cdot P_i(e)(e')}{\sum_{s'' \in S, e'' \in E} P_i(s)(s'' \mid \varphi_{s''}) \cdot P_i(e)(e'')} \\ \text{if the denominator } > 0 \text{ and } 0 \text{ otherwise}$$

To see that this is like Jeffrey Update, consider the example with the dark room. There are five equally probable states (red, yellow, green, brown, black). The observation probabilities are 1/4 and 3/4 for observing a light color and a dark color respectively. We have that $P_i(s)(red \mid light) = 1/3$, and likewise for all light colors. And we have $P_i(s)(brown \mid dark) = 1/2$ and likewise for black. Going through the calculations reveals that after executing this update the probability of a light color is now 1/4 and for a dark color it is now 3/4. \square

As for the explicit *update logic* of our weighted update rule, as long as the weighed probabilities can be represented in the static language, it can be axiomatized along the lines of the previous section for the pure case. But a more interesting logical issue might be to have a language which can define various types of updating agent explicitly, and then analyze their *interaction*, such as learning about other agents’ types, and choosing optimal strategies for dealing with them.

7 Related work

In the logical literature, combinations of epistemic logics and probabilistic reasoning have been studied since the 1990s (cf. e.g., van der Hoek 1992). Fagin and Halpern (1993) and Halpern and Tuttle (1993) were our point of departure for the static case, and Kooi (2003) and van Benthem (2003) for the dynamic aspect. In addition, Halpern (2003) should be compared as a general study of probabilistic reasoning in an epistemic-temporal setting, and in particular also, the work by Grünwald and Halpern (2003) as a study of probabilistic update, including Jeffrey Update. We also mention the paper by Aucher (2005) which was developed independently in a dynamic-epistemic line. Some of Aucher’s conclusions seems similar to ours, whereas other features diverge (e.g., he also treats drastic forms of belief revision triggered by ‘surprise events’ of probability zero) – but we must leave detailed

comparisons to other times, places, and agents. Next, Baltag and Smets (2007) raise an interesting challenge to our parameterized *ABC* approach, by providing one uniform update rule for qualitative belief update, relocating the parameter information in the description of the signal. Moreover, their rule is like Jeffrey’s in allowing overruling of old plausibility comparison of worlds by new plausibilities among events observed. It would be of interest to see if our Product Update Rule can also be cast in this more uniform format, by changing the way we present our three probability factors as input to our rule.

Finally, other areas are relevant, too. We already mentioned inductive logic and learning theory as paradigms to be compared, with agents modifying their probability distributions over time. But maybe more pointedly, the foundations of *Bayesian statistics* seem close to what we have been discussing, and the concerns and insights of its practitioners (and also its critics, (cf. Fitelson 2001) seem very congenial to ours. Romeijn (2005) is a first attempt by a person from the latter tradition at a fruitful confrontation with dynamic-epistemic approaches.

8 Conclusions

We have presented an analysis of three major probabilistic aspects of observing an event in the framework of dynamic-epistemic logic. The resulting distinction of prior probabilities, occurrence probabilities, and observation probabilities seems to make general sense, and through our proposed new ‘product rule’, it allows for an explicit modular view of probabilistic update and the concomitant construction of successive new probability spaces. The resulting update logic merges ideas from multi-agent epistemic logic and probabilistic update in a harmonious fashion. In particular, we have shown how one can find complete logics for reasoning about and with these updates, provided the epistemic-probabilistic base language is made rich enough. Finally, we have shown how our approach can be parameterized to different update policies, representing different ways of responding to new evidence.

We believe that this is just a start. Throughout our paper new questions have come up, while we feel our system might also have uses in practice. In particular, our explicit calculus of model construction makes sense in analyzing well-known probabilistic scenarios, while qualitative versions of our product update rule might provide a richer view of the events that lead to belief revision.

9 Acknowledgments

This work goes back to various discussions between the authors over the past few years. We thank Tomasz Sadzik, and the audiences at the Edinburgh ESSLLI Workshop on Belief Revision (August 2005), the ILLC Probability and Update Workshop ‘Meet the Bayesians’ (September 2005), the ILLC-Stanford Update Day (September 2005), and the ILLC Graduate Update Seminar (October-November 2005), anonymous referees and audience at the LOFT conference (July 2006), as well as the FoTFS Conference on Logic and Probability at ILLC Amsterdam 2007, for their responses to ‘talk versions’ of this paper. Barteld Kooi would like to thank NIAS for the peace and quiet needed to work on this paper at the end of 2006.

References

- Aucher, G. (2005). How our beliefs contribute to interpret actions. 4th International Workshop of Central and Eastern Europe on Multi-Agent Systems

- (CEEMAS05). To appear in the Lecture Notes in Artificial Intelligence published by Springer-Verlag.
- Bacchus, F. (1990). *Representing and Reasoning with Probabilistic Knowledge, A Logical Approach to Probabilities*. The MIT Press.
- Baltag, A. (2002). A logic for suspicious players: Epistemic actions and belief-updates in games. *Bulletin of Economic Research* 54(1), 1–45.
- Baltag, A., L. S. Moss, and S. Solecki (1998). The logic of common knowledge, public announcements, and private suspicions. In I. Gilboa (Ed.), *Proceedings of the 7th conference on theoretical aspects of rationality and knowledge (TARK 98)*, pp. 43–56.
- Baltag, A. and S. Smets (2007). Probabilistic dynamic belief revision. In J. van Benthem, S. Ju, and F. Veltman (Eds.), *A Meeting of the Minds: Proceedings of the Workshop on Logic, Rationality and Interaction, Beijing*, Volume 8 of *texts in computer science*, London, pp. 21–39. College Publications.
- Carnap, R. (1952). *The Continuum of Inductive Methods*. Chicago: The University of Chicago Press.
- Fagin, R. and J. Y. Halpern (1993). Reasoning about knowledge and probability. *Journal of the ACM* 41(2), 340–367.
- Fitelson, B. (2001). *Studies in Bayesian Confirmation Theory*. Ph. D. thesis, Department of Philosophy, University of Wisconsin, Madison.
- Gärdenfors, P. and H. Rott (1995). Belief revision. In *Handbook of logic in artificial intelligence and logic programming (Vol. 4): epistemic and temporal reasoning*, pp. 35–132. Oxford, UK: Oxford University Press.
- Gerbrandy, J. D. (1998). *Bisimulations on Planet Kripke*. Ph. D. thesis, University of Amsterdam. ILLC Dissertation Series DS-1999-01.
- Grünwald, P. and J. Y. Halpern (2003). Updating probabilities. *Journal of AI Research* 19, 243–278.
- Halpern, J. Y. (2003). *Reasoning about Uncertainty*. The MIT Press.
- Halpern, J. Y. and R. Pucella (2006). A logic for reasoning about evidence. *Journal of Artificial Intelligence Research* 26, 1–34.
- Halpern, J. Y. and M. R. Tuttle (1993). Knowledge, probability, and adversaries. *Journal of the ACM* 40(4), 917–962.
- Hommersom, A., J.-J. Meyer, and E. de Vink (2005). Update semantics of security protocols. *Synthese* 142(2), 229–267.
- Kelly, K. (1996). *The Logic of Reliable Inquiry*. Oxford: Oxford University Press.
- Kooi, B. (2003). Probabilistic dynamic epistemic logic. *Journal of Logic, Language and Information* 12(4), 381–408.
- Liu, F. (2008). *Changing for the Better: Preference Dynamics and Agent Diversity*. Ph. D. thesis, University of Amsterdam. ILLC Dissertation Series DS-2008-02.
- Plaza, J. A. (1989). Logics of public communications. In M. L. Emrich, M. S. Pfeifer, M. Hadzikadic, and Z. W. Ras (Eds.), *Proceedings of the 4th International Symposium on Methodologies for Intelligent Systems*, pp. 201–216.
- Romeijn, J. W. (2005). Dutch books and epistemic events. Talk presented at the *Interfacing Probabilistic and Epistemic Update* workshop in Amsterdam.
- Shafer, G. (1976). *A Mathematical Theory of Evidence*. Princeton University Press.

- Shafer, G. (1982). Belief functions and parametric models. *Journal of the Royal Statistical Society* 44(3), 322–352.
- van Benthem, J. (2003). Conditional probability meets update logic. *Journal of Logic, Language and Information* 12(4), 409–421.
- van Benthem, J., J. Gerbrandy, and B. Kooi (2006). Dynamic update with probabilities. ILLC Prepublication PP-2006-21, <http://www.illc.uva.nl/Publications/ResearchReports/PP-2006-21.text.pdf>.
- van Benthem, J., J. Gerbrandy, and E. Pacuit (2007). Merging frameworks for interaction: DEL and ETL. In D. Samet (Ed.), *Proceedings of TARK 11*, pp. 72–82.
- van Benthem, J., J. van Eijck, and B. Kooi (2006). Logics of communication and change. *Information and Computation* 204(11), 1620–1662.
- van der Hoek, W. (1992). *Modalities for Reasoning about Knowledge and Quantities*. Ph. D. thesis, Vrije Universiteit Amsterdam.
- van Ditmarsch, H., W. van der Hoek, and B. Kooi (2007). *Dynamic Epistemic Logic*. Dordrecht: Springer.