Dynamic Update with Probabilities

Johan van Benthem, Jelle Gerbrandy & Barteld Kooi

April 15, 2006

Abstract

We propose a way of adding probabilities to Dynamic Epistemic Logic (Baltag, Moss, and Solecki (1998)) to obtain a logic for reasoning about information change that is brought about by probabilistic information. We propose a general format for representing probabilistic information in a multi-agent setting, define the corresponding notion of epistemic update, and discuss the logic that results.

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 "$P_i(A) = x$".

Current dynamic-epistemic logics manipulate formulas $[!A]K_i\varphi$ describing what agents know or believe after a proposition A has been communicated. Here A may be either about the real world or about information that other agents have. And the most sophisticated modern update systems can even deal with a much greater variety of informative events, such as partial observation, whispers, or lies. Thus, it seems of interest to combine the two perspectives – for reasons of mutual benefit.

The paper is organized as follows. The first two sections cover our point of departure: in Section 2 we present static epistemic-probabilistic logic, and in Section 3 we present dynamic epistemic logics. Combinations of probabilistic and dynamic epistemic logics have been proposed before. The systems by Kooi and van Benthem are very shortly presented in Section 4. In Section 5 we present the way we would like to model probabilistic updates with varying degrees of generality. In Section 6 we turn to reasoning about probability.

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 just take some such system as our base in this paper, as our main emphasis is on the dynamic update phenomena. Epistemic probability models are structures

$$M = (S, \sim_i, P_i, V),$$

where S is a set of states, for each agent i, $\sim_i$ is an equivalence relation on S, and P_i is a probability measure on each equivalence class of $\sim_i$. Here, the above probabilities $P_i(\varphi)$ over propositions φ are generated by summing those over the states in M where φ holds. There are standard ways to generalize this definition to phenomena having to do with infinity.

The basic epistemic-probabilistic language that is interpreted on the preceding models allows 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 knowledge operators of this language are interpreted as usual in epistemic logic, while the probability statements are interpreted as follows:

$$M, s \models K_j\varphi \quad \text{iff} \quad \text{for all } t\text{: if } s \sim_i t, \text{ then } M, t \models \varphi$$

$$M, s \models P_i(\varphi) = k \quad \text{iff} \quad \sum_{t \sim_i s \text{ with } M,t \models \varphi} P_i(s)(t) = k$$

$$M, s \models P_i(\varphi \mid A) = k \quad \text{iff} \quad \frac{\sum_{t \sim_i s \text{ with } M,t \models A \& \varphi} P_i(s)(t)}{\sum_{t \sim_i s \text{ with } M,t \models A}} = k.$$

This system was studied by Halpern and others. Its logic is complete, and decidable modulo some numerical theory of the probability values.

Remark: Simplified Notation For many technical purposes it suffices to use notations $P(\varphi)$ without any index i, as the relevant agent will be clear from context when it matters. Going one step further, we can also drop the reference to equivalence classes of worlds in the relation $\sim_i$ for different agents, pretending we are just discussing one universe.

3 Dynamic-epistemic logics for non-probabilistic information update

Dynamic-epistemic logics describe information flow engendered by observed events. The simplest informative event, and a pilot case for much of the theory, is a public announcement $!A$ of some true proposition A in a group of agents. The dynamic effects of this is to change the current model $M = (S, \sim_i, V)$ to an updated model $M|A$ defined by restricting the worlds of M to just those where A holds. The effects of public announcements can be described in terms of epistemic statements true before them, or after them. In particular, truth values of epistemic statements can change because of an announcement. E.g., I did not know that A before, but now I do. 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. To keep track of all this, a dynamic epistemic language is needed, whose logic helps us keep careful track of things.

First, one adds a 'dynamic' modal operator $[!A]$ to the epistemic language. A formula $[!A]\varphi$ is read as 'φ holds after the announcement that A'. This language is interpreted in standard models for epistemic logic, with the following key clause, referring to the submodel $M \mid A$ of M consisting of all worlds where the formula A is true:

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

A complete dynamic-epistemic logic PAL for public announcement was first found in Plaza (1989), and independently in 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, and on top of that, a bunch of *reduction axioms* analyzing effects of informational actions compositionally. In particular, the following axioms describe how public announcement operators interact with other logical operators when added to multi-agent S5:

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

Updates for more complex communicative actions can be described in terms of 'action models', which stand for complex events that carry information for agents. Baltag, Moss, and Solecki (1998) define action models as structures

$$A = (E, \sim_i, \mathsf{Pre})$$

consisting of a set of events, an indistinguishability relation for each agent, and a 'precondition function' that determines in which worlds the events can actually occur. These models are quite similar to epistemic models, but instead of situations involving information, events involving information flow are modeled. The execution of an event represented by A in an epistemic model represented M is modeled by means of a *product construction* $M \times A$. The possible worlds in $M \times A$ are all pairs (w, e) of worlds in M and events in E such that the event can occur in the world.

$$\{(s, e) \in S \times E \mid (M, s) \models \mathsf{Pre}_e\}$$

The uncertainty relation in $M \times A$ is determined by the uncertainty relations in M and A. An agent cannot distinguish a pair (s, e) from (s', e') if the agent cannot distinguish s from s' in M and e from e' in A:

$$(s, e) \sim_i (s', e') \text{ iff } s \sim_i s' \text{ and } s \sim_i e'.$$

Finally, truth values of propositional variables do not change in executing an epistemic action: the propositional variables true in (s, e) are those true in s.

We refer to the growing literature for examples of how this mechanism models a wide variety of informational scenarios. Again, there is a simple dynamic-epistemic language reflecting this, with action models as modal operators. A formula of the form $[A, e]\varphi$ is interpreted as follows:

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

The axiomatization for this logic consists of a number of 'reduction axioms' for the new modal operators: 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 – often a multi-agent $S5$. The axiomatization is a straightforward generalization of the earlier one for PAL:

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

4 Some earlier combinations of epistemic and probabilistic update

4.1 Kooi's probabilistic PAL

Kooi (2003) presents a complete probabilistic dynamic-epistemic logic of truthful public announcements $!A$. Kooi's static base language involves both absolute probability assignments, and the numerical inequalities mentioned earlier. The 'model-crossing' truth condition for the dynamic modalities is as stated before.

Kooi's system can model many update scenarios. He also develops some epistemic-probabilistic model theory, using a notion of probabilistic bisimulation. The system allows for subtle comparison between different assertions involving conditional probabilities, update modalities, and knowledge.

The earlier methodology still works in this extended setting: completeness and decidability are obtained via reduction axioms. The key axiom for absolute probability statements is the following:

$$[!A]P_i(\varphi) = k \leftrightarrow (A \to P_i([!A]\varphi \mid A) = k) \tag{1}$$

Kooi's reduction involves a conditional probability. Hence, he also has a reduction axiom for conditional probability statements resulting from announcements:

$$[!A]P_i(\varphi \mid \psi) = k \leftrightarrow (A \to P_i([!A]\varphi \mid A\&([!A]\psi)) = k \tag{2}$$

4.2 Van Benthem's probabilistic DEL

Kooi's system describes the effects of announcements $!A$ of propositions that are true now, as happens in straightforward conversations where we accept what the source is telling us. Thus, one updates prior probabilities with non-probabilistic new information. But what if the source is unreliable, and hence the new information itself is probabilistic? In this case, we must distinguish two sorts of probability that factor into the update:

(a) prior probabilities of worlds,

recording our informative experience so far, but also what may be called 'process descriptions' for the event being observed, which require

(b) occurrence probabilities for new events in worlds.

Van Benthem's update rule runs as follows. Product models $M \times A$ are defined just as for DEL in Section 2. The only new feature to be defined are the new probabilities in the product model, i.e. the values for $P_{i,(s,a)}((t,b))$ for 'accessible' worlds (t,b) in $M \times A$, where we have that $t \sim_i s$ in M and $b \sim_i a$ in A, while $M, t \models \mathsf{Pre}_b$:

Definition 1 (Probability Product Update Rule)

$$P_{i,(s,a)}((t,b)) = \frac{P_{i,s}(t) \times P_{i,t}(b)}{\sum_{(u,c)\in D_{i,(s,a)}} P_{i,s}(u) \times P_{i,u}(c)} \tag{3}$$

This rule looks forbidding, but the idea is very simple, and it comes out clearly for one agent with uniform probabilities. In that case, subscripts of P operators for agent and state are superfluous, and the update rule becomes:

$$P((t,b)) = \frac{P(t) \times P_t(b)}{\sum_{(u,c)\in D} P(u) \times P_u(c)} \tag{3}$$

The new probabilities are computed in the numerator, taking a product of the prior probability for the initial world s and the occurrence probability for the event a. The denominator just normalizes the probabilities assigned by taking the sum of all numerator values in all relevant cases.

As for probabilistic reasoning in this setting, e.g., that involved in solving Monty Hall-style puzzles, van Benthem does not provide a complete logic. But he points out that the crucial reduction axiom can be designed by redescribing the Probability Product Update Rule in terms of the model M before the update.

A novel and potentially useful feature is the systematic construction of successive probability spaces. The standard picture in Bayesian conditioning is that of a current space which gets smaller and smaller as new information comes in. Product update, however, can modify the current model in more complex ways, 'branching' currently single worlds into different continuations. This seems helpful in practice, as the main difficulty which people have with probabilistic reasoning does not seem to be the mathematical calculus per se, but the selection of the right probability model to put this calculus to work.

5 Modelling probabilistic information change

The main point of this paper is that the preceding systems still leave out further essential aspects of probability update and its logic. In this section, we identify a third crucial probabilistic aspect, and incorporate this aspect in to model to provide a generalized update rule.

5.1 Three sources of probability

The model of Section 4.2 performs only a partial 'probabilization' of DEL-style product update. But there is also a third type of uncertainty that we have not modeled yet: concerning the event that is being observed. I see you read the letter from the Agency, and I know it is either a rejection of your grant proposal or an acceptance. You know which event ("reading YES", "reading NO") is taking place, while I do not. Epistemic product update in its non-probabilistic version computes the resulting differential uncertainties for you and me, which may be quite intricate. But the latter, too, can be analyzed in a more fine-grained quantitative manner in the present setting. To do that, we distinguish not two, but *three* sorts of probability:

(a) prior probabilities of worlds,

(b) occurrence probabilities for new events in worlds.

as well as

(c) observation probabilities for the current event.

Perhaps I saw a glimpse of your letter, or you looked smug: and thus, I might assign probability 0.7 to your reading an acceptance rather than an rejection. A simple scenario where all three kinds of probability come together is this:

Example: The Mumbling Liar You do not know whether p – Mary, your three year old daughter, took a cookie – is true or not. You assign probability 0.5 to either case. You ask Mary whether she took the cookie: you know that if she did, she will speak the truth with probability .3. (If she did not steal the cookie, she will, of course, not lie about that.)

Mary mumbles something – you did not quite get whether her answer was 'yes' or 'no.' You think that you observed a "Say-yes"-event with probability 0.8, but there is a 0.2 chance that it was "Say-no".

What should be the new probability when all this is taken into account? There are four possible worlds in the product update model:

$$((\neg)p, \text{“Say } (\neg)p\text{”}),$$

and it seems reasonable to compute their probabilities by multiplying all three relevant probabilities: (a) the prior for the world with p or $\neg p$, (b) the occurrence probability for Mary's assertion, and (c) the observation probability that the event in the pair is in fact what we observed.

5.2 Event models and general product update

For a start, our static epistemic-probabilistic models M are still the same as before, and so is our epistemic-probabilistic language. Indeed, we can also use the earlier DEL notation $[A, e]\varphi$ for the effects of executing an epistemic event (A, e). But we will now redefine what we mean by such models, representing uniformly defined epistemic-probabilistic 'events':

Definition 2 (Probabilistic event models) *Probabilistic event models are structures $A = (E, \Phi, \mathsf{Pre}, P)$ where Φ is a set of mutually inconsistent and jointly exhaustive sentences called* preconditions*,* Pre *assigns to each $\varphi \in \Phi$ a probability distribution over E (we write $\mathsf{Pre}(\varphi, a)$ for the probability that a occurs given that φ is true). Finally, P is a probability function over E.*

This should be understood as follows. Part of the models consists in the specification of an *observed process* which makes events occur with certain probabilities, depending on a set of conditions Φ which cover all eventualities in a unique manner. Liars or Quizmasters are such processes, with instructions 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 language for these conditions Φ is left open here, but in an eventual clean set-up, it would be the formal language of our dynamic-epistemic-probabilistic system, used with the right set of simultaneous recursions.

The second component of the models is the probability function P, which stands for probabilities assigned by the observer as to which event actually took place. The latter feature encodes a property of the *observation process.*

With these models, computing an update involving all three sources of probability is a straightforward generalization of earlier rules.

Definition 3 (Full Product Update Rule) *For any static epistemic-probabilistic model M and probabilistic event model A, the product update model $M \times A$ is defined exactly as for DEL on the epistemic part, but now with a probability measure P over the set $S \times A$ of pairs of states from M and events from A. Let $\Phi(t)$ be the unique element of Φ that is true in t. We set:*

$$P((t, b)) = \frac{P(s) \times \mathsf{Pre}(\Phi(t), b) \times P(b)}{\sum_{u \in S, c \in E} P(u) \times \mathsf{Pre}(\Phi(u), c) \times P(c)} \tag{4}$$

Agent-indexed versions of this, indicating also the relevant models where probabilities are taken, have a numerator in the following format:

$$P^{M \times A}_{i,(s,a)}((t, b)) = P^{M}_{i,s}(t) \times \mathsf{Pre}^{A}(\Phi(t), b) \times P_{i,a}(A, b),$$

while the denominator would only sum over pairs (world, event) that are epistemically accessible in $M \times A$ from the vantage point (s, a).

Here is how this update mechanism works out in practice.

Example: The Mumbling Liar Again In our example of Mary taking a cookie, our initial hypothesis about the proposition p of taking the cookie is captured by a prior probability distribution

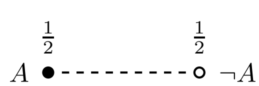

We then observe Mary mumbling her answer, assigning our probabilities regarding her character (if she has a reason, she lies with probability .7) and our powers of observation (we think she said p with probability .8) as above, resulting in the following probabilistic event model:

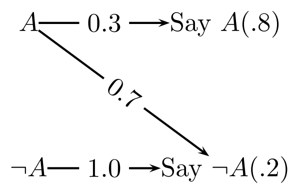

The product of our initial state with this model is as follows:

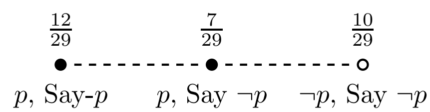

This diagram is our new probabilistic information state after the whole episode. Note how it also illustrates another typical feature of product update. We do not just eliminate existing worlds or change prior probabilities, but may also construct new types of possibility. Initially, we only considered options for one single aspect of reality (Mary took the cookie, or not), after the update, we consider more complex epistemic possibilities, including information about whether she lied about it or not. In this way, the number of options may increase even if we have in fact obtained information. The new possibilities may be viewed as possible runs of the total process represented by our probabilistic event model.

6 Dynamic logics of probabilistic update

In order to reason about probabilistic information change in a dynamic-epistemic format, we must extend existing axiom systems like Kooi's with appropriate reduction axioms. In this section, we show how this can be done for our general mechanism of update with all three probability factors.

The crucial information about our Product Update Rule will be reflected in our reduction axioms, which state when propositions get certain probabilities after an epistemic event took place. Moreover, we know already that reduction axioms express a certain harmony between the dynamic and static parts of an epistemic language. E.g., absolute probabilities after public announcement called for a language with conditional probabilities. Some experimentation with putative axioms after products with event models will reveal the need for something even stronger, viz. additive probability statements of the form

$$\alpha_1 \times P(\varphi_1) + \cdots + \alpha_k \times P(\varphi_k) = x$$

or even the earlier-mentioned *linear inequalities* of probability statements, replacing the equality by a 'less than' sign.

More precisely, our *dynamic-epistemic-probabilistic language* can be defined in the following inductive format:

$$p \mid \neg\varphi \mid \varphi \wedge \psi \mid \varphi \vee \psi \mid K_i\varphi \mid [(A,e)]\varphi \mid$$
$$\alpha_1 \times P(\varphi_1) + \cdots + \alpha_k \times P(\varphi_k) \leq \beta \times P(\psi)$$

where (A, e) is a probabilistic event model, and α_i, β are rational numbers. There is a joint recursion hidden in this set-up, somewhat like that for purely epistemic product update with event models. The formulas that define the partition in our probabilistic event models come from the same language that we are defining here, but through the clause for the dynamic modalities, such models themselves enter the language again.

The key reduction axiom for reasoning about product update must relate formulas of the form $[(A, e)]\psi$ with ψ involving probabilities to assertions in terms of probabilities in the original model (M, s).

For a start, we analyze the probability value $P(\psi)$ of a formula ψ in a product model $(M, s) \times (A, e)$. The worlds in the product model where ψ holds can be divided into a finite partition consisting of sets of worlds t such that $(t, f) \models \psi$ in the product model, where f runs over all events in A. Fix such an event f. We must have (t, f) epistemically indistinguishable from (s, e). Moreover, in such a pair, we clearly have that $[A, f]\psi$ holds in M at t. Finally, the given partition Φ in A assigns a unique proposition $\varphi(t)$ to the world t for which A provides an occurrence probability of f in t. By our Product Update Rule, the probability of (t, f) is then the product of three probability terms:

$$\text{prior-prob}(t) \times \text{occurrence-prob}(\varphi(t), f) \times \text{observation-prob}(f)$$

or, more precisely:

$$P_{(M,s)\times(A,e)}(t, f) := P_M(t) \times \mathsf{Pre}_A(\varphi(t), f) \times P_A(f)$$

Summing all this in an obvious way, we have that:

$$P_{(M,s)\times(A,e)}(\psi) = \Sigma_{f\in A,\varphi\in\Phi}(P_M(\varphi \wedge [A, f]\psi) \times \mathsf{Pre}_A(\varphi, f) \times P(f))$$

Using this observation, we automatically get a reduction axiom in standard DEL style for dynamic-probabilistic assertions of the form:

$$[A, e](P(\psi) = k)$$

in which 'P' refers to the probabilities after the update, as a sum of terms referring to probabilities in the initial model M of the form:

$$\alpha_0 \times P(\varphi_0 \wedge [(A, f_0)]\psi) + \ldots + \alpha_n \times P(\varphi_n \wedge [(A, f_n)]\psi) = k$$

Now the language is not yet in expressive harmony here, since we need sum terms in the static language to deal with single probability assignments after the update. However, it is easy to see that this can be solved once we provide the static language with either linear equalities or linear inequalities of probability statements.

With these observations, we have the following, quite general, result.

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.*

This can be proven as follows. When analyzing a statement

$$[(A, e)](\alpha_1 \times P(\varphi_1) + \cdots + \alpha_k \times P(\varphi_k) \leq \beta \times P(\psi))$$

we can replace the separate terms $P(\varphi_i), P(\psi)$ after the modal update operator by their equivalents as computed just before. The result is a linear inequality where the terms of the main sum may themselves contain sums. Rearranging terms, this can be brought into the form of one big sum with probabilities for formulas in M with suitable coefficients. In particular, no essential multiplication is introduced for the probability terms.

This general result shows that with the DEL framework one can formulate rich probability logics, in which transformations of models are systematically described in the object language itself. The completeness result allows us to analyze of complex probability updates with standard semantical and proof theoretical tools.

Our methodology via reduction axioms yields a relative, rather than an absolute axiomatization of the full dynamic language. In particular, we cannot say without further information whether the total system will be decidable. 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.

7 Related work

Combinations of epistemic logics and probabilistic reasoning have been studied since the 1990s (cf. van der Hoek (1992)). Fagin and Halpern (1993) and Halpern and Tuttle (1993) were our point of departure for the static case. We have also reviewed two earlier DEL-style attempts by Kooi, and van Benthem. In addition, Halpern should be mentioned as a general study of probabilistic reasoning in an epistemic-temporal setting, and in particular, 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. Some of his conclusions seems similar to ours, whereas other features diverge (e.g., he also treats drastic forms of belief revision) – but we must leave detailed comparisons to other times, places, and agents. Finally, we mention the tradition of foundations of Bayesian reasoning and its critics (Jeffrey, Glymour, Fitelson) whose concerns and results seem very congenial to ours. Romeijn (2005) provides 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 it allows for a richer modular view of probabilistic update and the concomitant construction of successive new probability spaces. The resulting update system has a model theory much like that of existing ones for purely epistemic or doxastic settings. We have also shown how one can find complete logics with reduction axioms, provided the epistemic-probabilistic base language is made rich enough.

In the extended version of this paper (which has been 'prepublished' as van Benthem, Gerbrandy, and Kooi (2006)) we show how our approach can be parametrized for different types of agent assigning different weights to these factors, thus allowing for the sort of diversity widely encountered in studies of probabilistic agents. Adding weights to the framework allows us to cover a wider range of probability updates. In particular, updates in which the incoming probabilities completely overrule the initial probabilities – so-called 'Jeffrey Update' – can be modeled as a case in which observation and occurrence probability factors get a high weight, while the initial probabilities get none.

9 Acknowledgments

This work goes back to various discussions between the authors over the past few years. In particular, we thank 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) for their responses to 'talk versions' of this paper.

References

Aucher, G. (2005). How our beliefs contribute to interpret actions. Accepted as a full paper at the 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.

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.

Fagin, R. and J. Y. Halpern (1993). Reasoning about knowledge and probability. *Journal of the ACM 41*(2), 340–367.

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. and M. R. Tuttle (1993). Knowledge, probability, and adversaries. *Journal of the ACM 40*(4), 917–962.

Kooi, B. (2003). Probabilistic dynamic epistemic logic. *Journal of Logic, Language and Information 12*(4), 381–408.

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.

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 der Hoek, W. (1992). *Modalities for Reasoning about Knowledge and Quantities*. Ph. D. thesis, Vrije Universiteit Amsterdam.