Probabilistic Dynamic Epistemic Logic

Barteld P. Kooi
University of Groningen, Department of Computing Science, P.O. Box 800, 9700 AV Groningen, The Netherlands barteld@cs.rug.nl

Abstract. In this paper I combine the dynamic epistemic logic of Gerbrandy (1999) with the probabilistic logic of Fagin and Halpern (1994). The result is a new probabilistic dynamic epistemic logic, a logic for reasoning about probability, information, and information change that takes higher order information into account. Probabilistic epistemic models are defined, and a way to build them for applications is given. Semantics and a proof system is presented and a number of examples are discussed, including the Monty Hall Dilemma.

Keywords: dynamic logic, epistemic logic, probability, updates, higher order information

1. Introduction

Epistemic logic is a modal logic used to reason about information, including higher order information. Dynamic epistemic logics are extensions of epistemic logic which can be used to reason about information and information change. In probability theory Bayesian updating can be seen as a model for information change, but higher order information is overlooked. This is a problem when one wants to formalize inferences about changing probabilistic higher order information.

In this paper I combine probabilistic logic with dynamic epistemic logic yielding a new logic, PDEL, that deals with changing probabilities and takes higher order information into account. In section 3 probabilistic epistemic models are introduced. The language of PDEL can be interpreted on these models. In section 4 I give a method for making models for specific situations and I provide a sound and complete proof system for PDEL. Bisimulation for probabilistic epistemic models is introduced in section 5. In section 6 some examples of application are discussed. Finally, in section 7 some conclusions are drawn and some directions for further research are indicated. But first I want to make clear why I develop PDEL in the first place.

2. Motivation

In this section I give some of my views on (dynamic) epistemic logic and probability theory, which will make clear why a combination of these is worthwhile and what the scope of this paper is.

2.1. EPISTEMIC LOGIC AND DYNAMIC EPISTEMIC LOGIC

Epistemic logic was initially developed by Hintikka (1962). His main goal was a conceptual analysis of knowledge and belief. In this paper I take the term "epistemic" broader, applying to belief and other ways an agent might have information as well. This concurs with much of the literature in this area. Epistemic logic typically deals with what an agent considers to be possible given his current information. This information also contains information about information other agents have, because epistemic logic is suited to deal with situations involving more than one agent. In this way epistemic logic also deals with *higher order information*, i.e. information about information.

Consider the following example from Van Ditmarsch 2000. There are three players, 1, 2, and 3 and three cards: red, white, and blue. The cards are distributed among the players such that red, white, and blue are held by 1, 2, and 3 respectively. Let us assume that the players can only see their own cards and that they all have the information that the cards are distributed among them such that each has one card. With the language of epistemic logic we can formalize complicated statements such as 'player 1 knows that player 2 does not know which card player 3 has.' This kind of higher order information is dealt with by epistemic logic. But there are even more complicated forms of higher order information such as *common knowledge*. In the example above for instance, it is common knowledge that there are exactly three cards and it is common knowledge which these are. Epistemic logic also provides a good conceptual analysis of common knowledge. In modern introductions to epistemic logic such as (Fagin et al., 1995) and (Meyer and van der Hoek, 1995) we find that epistemic logic has a very wide range of application.

Although epistemic logic provides a good analysis of higher order information, information change is not included into its scope. Dynamic epistemic logics are extensions of epistemic logic that deal with information change. On the one hand the development of these systems was inspired by the semantics of natural language, where the meaning of a sentence is viewed as the way it changes the information of those who hear the sentence. On the other hand their development was inspired by the study of games where information exchange occurs and higher order information plays an important role. Several systems have been proposed over the years, notably those by Gerbrandy and Groeneveld (1997) (which was inspired by Veltman (1996)), Baltag et al. (1998), van Ditmarsch (2000), and Ten Cate 2002.

Let me remark beforehand that the kind of information change that is studied in dynamic epistemic logic is not the dramatic kind that is studied in philosophy of science where entire theories are replaced, or the kind of information change that is studied in belief revision where things that were not considered possible by an agent do occur and world views are subsequently changed substantially. The focus is on everyday calm and quiet information change, such as it is studied in probability theory and game theory, where small pieces of information are processed in a piecemeal fashion. The main objective is to give an account of how an agent's information changes due to actions, including the change of higher order information. Therefore in many systems change of the world itself is not studied at all.

In dynamic epistemic logic, incorporating new information is called an update. This should not be confused with the notion of updating as it is used in the belief revision paradigm. The simplest example of an update is where an agent learns that a certain sentence φ holds. An update with a sentence in this case means that alternatives that the agent considers possible where φ does not hold are removed, yielding a new set of alternatives. There are much more complicated forms of updates, than just updates with sentences. In a multi-agent setting, for example, different agents may have different access to the new information and the information the agents have about each other also plays a role. Consider the situation mentioned above. Suppose player 1 shows his card to player 2. Meanwhile player 3 can see that this is going on, but she cannot see which card 1 shows. The players' information then changes as a result: player 2 knows which card player 1 has, player 3 knows that player 2 knows which card player 1 has, but not which one that is, player 1 knows that player 3 knows that. See van Ditmarsch (2000) for an extensive discussion of this example.

2.2. PROBABILITY THEORY AND PROBABILISTIC LOGIC

Probability theory is a well-studied area. It is, however, not a logic. To me it seems important to try to make a logic out of it or to incorporate it in a logic, because it is often presented as a theory that models reasoning. A logical approach has much to offer if one wants to study reasoning: a formal language with which inferences can be represented; a clear distinction between syntax and semantics; various notions of validity; decision procedures, and so on. In my view logic is the best way to study reasoning.

In the philosophy of probability, a distinction is made between objective theories of probability – probability is taken to be a part of physical reality – and epistemic theories of probability – they see probability as degree of belief. (See Gillies (2000) for an overview.) This distinction seems related to a distinction made by Bacchus (1990) between *statistical* and *propositional* probabilities. A statistical probability statement is about the proportion of individuals that have a certain property. A propositional probability statement expresses the probability that a certain individual has a certain property. From the viewpoint of possible world semantics, a statistical probability is an attribute of a possible world. The proportion of individuals with a certain property simply *is* or *is not* a specific rational number. In that sense a statistical probability is a part of physical reality. Propositional probability, on the other hand, seems to be a *modal* notion which involves more possible worlds. If an individual has different properties in different possible worlds, then the probability it has some property is defined as the ratio of the set of worlds where that individual has that property. In that sense a propositional probability is not a part of physical reality. If the accessibility of these possible worlds is interpreted epistemically, then propositional probabilities express the degree of belief in that proposition. When Kripke wrote:

> 'Possible worlds' are little more than the miniworlds of school probability blown large. (Kripke, 1980, p. 18)

clearly he had propositional (or epistemic) probabilities in mind. As this paper is mainly concerned with epistemic concepts such as information and information change, the focus is on propositional probability.

Note that in this paper I am interested in logics that deal with probabilities explicitly, i.e. one can express statements involving probability in the language of the logic. This is quite different from logics that use probabilistic semantics to investigate non-monotonic inference relations (see Kyburg (1994)). This paper is concerned with reasoning about uncertainty rather than uncertain reasoning.

For my purposes, the logic developed by Fagin and Halpern (1994) is particularly promising. For it is not only a logic of propositional probability, but an epistemic logic as well: there are probability operators and epistemic operators in the language. In this paper I will call this logic PEL for *probabilistic epistemic logic*. The benefit is that in PEL one can distinguish an event that is highly unlikely in the sense that its probability is zero from an event that is epistemically impossible. For example, when a coin is flipped repeatedly until heads comes up, the infinite sequence of tails has probability zero, but it is possible. This is quite different from the coin landing heads and tails simultaneously, which is impossible and also has probability zero.

Another feature of PEL is that it is also able to model a form of uncertainty that is non-probabilistic. Suppose that an agent knows a coin lands heads in one third of the cases or is fair, but she does not know which of these is the case. This ignorance can be modeled as such. There need not be a probability distribution for the coin being fair or not. See Fagin and Halpern (1994) for more discussion about why this is a desirable feature.

Just as in the case of epistemic logic, in probability theory the matter of incorporating newly acquired information has been investigated. In probability theory this is done by taking *posterior* probabilities instead of *prior* probabilities, i.e. the conditional probabilities given the new information, which is also called Bayesian updating. Posterior probabilities can be calculated using Kolmogorov's definition (Kolmogorov, 1956):

$$\mathbf{P}(X|Y) = \frac{\mathbf{P}(X \cap Y)}{\mathbf{P}(Y)} \qquad \text{if } \mathbf{P}(Y) > 0$$

The idea is that this rule gives one the probability of X after one gets the information that Y is the case. So posterior probability can be used to model information change.

2.3. A COMBINATION

Although the distinction between improbable and impossible events and ignorance about probabilities make PEL an appealing system for reasoning about probability, the main motivation for using PEL as the basis for PDEL is that probabilistic higher order information can be studied in PEL, and by making a dynamic version of PEL we can study probabilistic higher order information change. It is interesting to note that both in dynamic epistemic logic and in probability theory, the incorporation of new information is studied. But they seem to come up with different answers to how this is properly done. The difference is that in dynamic epistemic logic more kinds of information change are distinguished that explicitly take higher order information into account. The intriguing question that pops up is what these two fields could learn from each other with respect to information change.

Fortunately a formal connection between the two areas has been established (see Bacchus (1990) and Halpern (1991) for details), where we can see that probabilistic logic can be seen as an extension of (non-dynamic) epistemic logic. The language of epistemic logic can be seen as a fragment of the language of probabilistic logic. This is done by relating belief and certainty. Standard probability theory can be seen as an extension of KD45. Now let us focus on certainty and conditional certainty: $\mathbf{P}(\varphi) = 1$, which is abbreviated by $cert(\varphi)$ and $\mathbf{P}(\varphi \mid \psi) = 1$, which is abbreviated by $cert(\varphi \mid \psi)$. We get the following sentence for conditional certainty:

$$\mathbf{P}(\psi) > 0 \rightarrow (cert(\varphi \mid \psi) \leftrightarrow cert(\psi \rightarrow \varphi))$$

The consequent of this implication is very much like the Knowledge-Update axiom (also called the generalized Ramsey axiom) of dynamic epistemic logic (see figure 2 on page 16).

$$[\psi]\Box\varphi \leftrightarrow \Box(\psi \rightarrow [\psi]\varphi)$$

The only difference, apart from notation (see section 3 for the definitions), is that instead of $\psi \rightarrow \varphi$ in the case of probabilistic logic, we have $\psi \rightarrow [\psi]\varphi$ in dynamic epistemic logic. This crucial difference is due to a difference in perspective on information change: in dynamic epistemic logic learning that ψ can change the truth value of φ. In probabilistic logic this is assumed not to be the case. This difference in perspective only becomes apparent when one is interested in higher order information. Assuming that learning something does not change facts (i.e. truth values of propositional variables), the truth value of φ can only change if an agent learns that ψ, if φ somehow involves a statement about the information the agent has. In this paper I develop a probabilistic dynamic epistemic logic that does take into account that the truth value of sentences can change due to information change. To keep this paper simple I limit updates to *public announcements*, i.e. all agents simultaneously get the same information and it is common knowledge that they receive it. This simple dynamic epistemic logic is introduced by Plaza (1989). I use the version introduced by Gerbrandy and Groeneveld (1997) and I will call it DEL in this paper. I think that probability theory could greatly benefit from the theory of information change provided by dynamic epistemic logic.

3. Probabilistic Epistemic Models

The standard models in the semantics of epistemic logic are multi-agent Kripke models: models with accessibility relations for all agents being considered. The accessibility relations can be interpreted epistemically; a world is accessible to an agent iff the state of affairs in that world is consistent with the information the agent has. The question is how to add probability to these models. Fagin and Halpern define probabilistic epistemic models. For those readers familiar with probability theory, in their probabilistic epistemic models a probability space is assigned to each agent in each world. In this paper I limit this to models, where the σ-algebra of measurable sets is always the powerset of the sample space. Therefore the definition is a bit simpler. Most of the results in this paper however equally apply to the more general notion of probabilistic epistemic models.

DEFINITION 1. (Probabilistic epistemic models). *Let a countable set of propositional variables $\mathcal{P}$ and a finite set of agents $\mathcal{A}$ be given. A probabilistic epistemic model M is a quadruple (W, R, V, P) such that:*

1. $W \neq \emptyset$*; a set of possible worlds;*

2. $R : \mathcal{A} \to 2^{W \times W}$*; assigns an accessibility relation to each agent;*

3. $V : \mathcal{P} \to 2^W$ *; assigns a set of worlds to each propositional variable;*

4. $P : (\mathcal{A} \times W) \to (W \rightharpoonup [0, 1])$*; such that*

$$\forall a \in \mathcal{A} \forall w \in W \sum_{v \in dom(P(a,w))} P(a, w)(v) = 1$$

assigns a probability function to each agent at each world such that its domain is a non-empty subset of the set of possible worlds. ($\rightharpoonup$ means that it is a partial function; some worlds may not be in the domain of the function.)

Below we often use the notion of a pointed model (M, w). This is a model with a designated world, called its point, which is taken to be the actual world. We also want to generalize the probability function to sets of worlds. If E is a subset of $dom(P(a, w))$, then $P(a, w)(E) = \sum_{v \in E} P(a, w)(v)$.

The advantages of having separate probability functions and accessibility relations are, as one noted earlier, that one can distinguish events with probability zero from events that are impossible. Moreover one can have *ignorance about probabilities*, not just *ignorance in terms of probabilities*. All this becomes clear in the examples. There are restrictions that can be imposed on the probability function. For example one might want the domain of the probability function assigned to a world to be a subset of the set of accessible worlds. Restrictions, such as probabilistic versions of positive and negative introspection, are discussed extensively by Fagin and Halpern (1994). In this paper there are no restrictions on probability functions.

We can interpret the language of PEL, which is introduced in Fagin and Halpern (1994), on these models. Here, we extend this language with update operators from DEL, thus obtaining a new language for reasoning about probability and information change.

DEFINITION 2. (Language of PDEL). *Let a countable set of propositional variables $\mathcal{P}$ and a finite set of agents $\mathcal{A}$ be given. The language of* PDEL $\mathcal{L}^{\mathbf{P}[\cdot]}_{\mathcal{PA}}$ *is given by the following rule in extended Backus-Naur form :*

$$\varphi ::= p \mid \neg\varphi \mid \varphi_1 \wedge \varphi_2 \mid \Box_a\varphi \mid [\varphi_1]\varphi_2 \mid q_1\mathbf{P}_a(\varphi_1) + \cdots + q_n\mathbf{P}_a(\varphi_n) \geq q$$

where $p \in \mathcal{P}$, $a \in \mathcal{A}$ and $q_1, \ldots, q_k$ and q are rationals. Besides the usual abbreviations, we have the following.

$$\begin{array}{lcl}
\sum_{i=1}^{n} q_i\mathbf{P}_a(\varphi_i) \geq q & : & q_1\mathbf{P}_a(\varphi_1) + \cdots + q_n\mathbf{P}_a(\varphi_n) \geq q \\
q_1\mathbf{P}_a(\varphi) \geq q_2\mathbf{P}_a(\psi) & : & q_1\mathbf{P}_a(\varphi) - q_2\mathbf{P}_a(\psi)) \geq 0 \\
\sum_{i=1}^{n} q_i\mathbf{P}_a(\varphi_i) \leq q & : & \sum_{i=1}^{n} - q_i\mathbf{P}_a(\varphi_i) \geq -q \\
\sum_{i=1}^{n} q_i\mathbf{P}_a(\varphi_i) < q & : & \neg(\sum_{i=1}^{n} q_i\mathbf{P}_a(\varphi_i) \geq q) \\
\sum_{i=1}^{n} q_i\mathbf{P}_a(\varphi_i) > q & : & \neg(\sum_{i=1}^{n} q_i\mathbf{P}_a(\varphi_i) \leq q) \\
\sum_{i=1}^{n} q_i\mathbf{P}_a(\varphi_i) = q & : & (\sum_{i=1}^{n} q_i\mathbf{P}_a(\varphi_i) \leq q) \wedge (\sum_{i=1}^{n} q_i\mathbf{P}_a(\varphi_i) \geq q)
\end{array}$$

The language of PEL $\mathcal{L}^{\mathbf{P}}_{\mathcal{PA}}$ *consists of those sentences of* $\mathcal{L}^{\mathbf{P}[\cdot]}_{\mathcal{PA}}$ *in which no update operators occur.*

A sentence of the form $\Box_a\varphi$ is to be read as 'a believes that φ'. A sentence of the form $\mathbf{P}_a(\varphi) \geq q$ can be read as 'the probability a assigns to φ is greater than or equal to b'. Note that in this language higher order probability statements can be expressed, such as $\mathbf{P}_a(\mathbf{P}_b(\varphi) \geq q_1) \geq q_2$. This expresses that the probability a assigns to the sentence that the probability b assigns to φ is greater or equal to q_1, is greater or equal to q_2. This is higher order in the sense that it expresses what information an agent has about the information of an(other) agent, completely analogous to the case in epistemic logic where sentences such as $\Box_a\Box_b p$ express that a has information about b's information.

A sentence of the form $[\varphi]\psi$ can be read as 'ψ is the case, after everyone simultaneously and commonly learns that φ is the case'. In order to interpret this language we have to give two definitions simultaneously, i.e. a truth-definition and a definition of updated models. These definitions are interdependent, but not circular.

DEFINITION 3. (Semantics for $\mathcal{L}^{\mathbf{P}[\cdot]}_{\mathcal{PA}}$). *Let a probabilistic epistemic model* $M = (W, R, V, P)$ *and a world* $w \in W$ *be given.*

$$\begin{array}{ll}
(M,w) \models p & \textit{iff } w \in V(p) \\
(M,w) \models \neg\varphi & \textit{iff } (M,w) \not\models \varphi \\
(M,w) \models (\varphi \wedge \psi) & \textit{iff } (M,w) \models \varphi \textit{ and } (M,w) \models \psi \\
(M,w) \models \Box_a \varphi & \textit{iff } (M,v) \models \varphi \textit{ for all } v \textit{ such that } wR(a)v \\
(M,w) \models [\varphi]\psi & \textit{iff } (M_\varphi, w_\varphi) \models \psi \textit{ (see definition 4)} \\
(M,w) \models \sum_{i=1}^n q_i \mathbf{P}_a(\varphi_i) \geq q & \textit{iff } \sum_{i=1}^n q_i P(a,w)(\varphi_i) \geq q
\end{array}$$

where $P(a,w)(\varphi_i) = P(a,w)(\{v \in dom(P(a,w)) \mid (M,v) \models \varphi_i\})$.

DEFINITION 4. (Semantics for updates). *Let a probabilistic epistemic model* $M = (W, R, V, P)$ *and a world* $w \in W$ *be given. The* updated *model* $M_\varphi = (W_\varphi, R_\varphi, V_\varphi, P_\varphi)$ *is defined as follows.*

$$\begin{array}{lll}
W_\varphi & = & W \\
R_\varphi(a) & = & \{(u,v) \mid (u,v) \in R(a) \textit{ and } (M,v) \models \varphi\} \\
V_\varphi & = & V \\
dom(P_\varphi(a,u)) & = & \begin{cases} dom(P(a,u)) & \textit{if } P(a,u)(\varphi) = 0 \\ \{v \in dom(P(a,u)) \mid (M,v) \models \varphi\} & \textit{otherwise} \end{cases} \\
P_\varphi(a,u)(v) & = & \begin{cases} P(a,u)(v) & \textit{if } P(a,u)(\varphi) = 0 \\ \dfrac{P(a,u)(v)}{P(a,u)(\varphi)} & \textit{otherwise} \end{cases}
\end{array}$$

For a pointed model (M,w) *the updated model is* (M_φ, w) *(i.e.* $w_\varphi = w$*).*

Announcing φ yields an updated model which is a copy of the original model. It is not an identical copy of the original model, for the accessibility relations and the probability functions differ. Worlds where φ does not hold are no longer accessible to any of the agents. The probability functions are treated similarly to accessibility relations. Worlds where φ does not hold are no longer in the domain of the function. Note that the announcement of φ does not presume that φ is actually true. Consequently an update can always be executed[1], i.e. it holds in general that $\langle\varphi\rangle\top$.

However an update only changes the probability functions of those agents who assign non-zero probability to φ. There are some approaches in probability theory for updating with sentences that have probability zero. The most common one is to leave it undefined. If we would take this approach in case of probabilistic logic, there would be truth value

[1] There are other dynamic epistemic logics which limit public announcements to *truthful* public announcements where only true announcements can be made.

gaps, which would make it very difficult to give a complete proof system. Another approach is to assign probability zero to everything after an update with a sentence that has probability zero. This approach is found in Bacchus (1990), but also in probability theory, for example in Prohorov and Rozanov (1969). This would seem to go against the laws of modal logic; after learning a sentence with probability zero, even the truth would be assigned probability zero. By analogy to *ex falso sequitur quodlibet* it would be more appropriate to assign probability one to everything in that case. This on the other hand would go against the laws of probability theory. So both choices would make it difficult to provide a complete proof system. There are more advanced approaches to updating with sentences with probability zero (see Halpern (2001) for an overview of the different approaches and references). All these approaches handle updating with sentences that have a non-empty set of worlds where that sentence holds. However updating with a sentence that does not hold in any world such as the absurdity remains a problem, and would still result in truth value gaps.

Dynamic epistemic logic cannot deal well with updates with inconsistent information as well. Typically, the accessibility relation become empty after an inconsistent update. A method of *revision* such as it is studied in belief revision is not available here. In PDEL too, we must also deal with updates with information that has probability zero in a way that is not intuitively appealing. The approach given in definition 4 is simply to ignore the information. This is to ensure that one does not divide by zero. There is no compelling philosophical reason for this choice, except maybe that the agent would just not believe the information received, and would therefore leave things as they were. This makes the proof system relatively simple.

LEMMA 1. *If* (M, w) *is a probabilistic epistemic model, then* (M_φ, w_φ) *is a probabilistic epistemic model too.*

Proof. The only difficulty lies in whether P_φ assigns a probability function to each agent in each world. Take a world u and an agent a. If $P(a,u)(\varphi) = 0$, then $P_\varphi(a,u) = P(a,u)$ and therefore it is a probability function. If $P(a,u)(\varphi) \neq 0$, then the domain of $P_\varphi(a,u)$ is exactly the set of worlds in the original domain where φ holds, therefore:

$\sum_{v \in dom(P_\varphi(a,u))} P_\varphi(a,u)(v)$

$=\{$definition of the domain$\}$

$\sum_{v \in dom(P(a,u)) \text{ and } (M,v)\models\varphi} P_\varphi(a,u)(v)$

$=\{$definition of the probability$\}$

$\sum_{v \in dom(P(a,u)) \text{ and } (M,v) \models \varphi} \frac{P(a,u)(v)}{P(a,u)(\varphi)}$

$=\{\text{algebra}\}$

$$\frac{\sum_{v \in dom(P(a,u)) \text{ and } (M,v) \models \varphi} P(a,u)(v)}{P(a,u)(\varphi)}$$

$=\{\text{definition } P(a,u)(\varphi)\}$

$$\frac{P(a,u)(\varphi)}{P(a,u)(\varphi)}$$

$=\{P(a,u)(\varphi) \neq 0\}$

1

Moreover $P(a,u)(v) \in [0,1]$ for all $v \in dom(P_\varphi(a,u))$ and $P(a,u)(v) \leq P(a,u)(\varphi)$. Therefore $P_\varphi(a,u)(v) \in [0,1]$.

As one can see this notion of updating is quite similar to Bayesian updating. In fact for many sentences it holds that

$$\mathbf{P}_a(\varphi|\psi) = q \text{ iff } [\psi]\mathbf{P}_a(\varphi) = q$$

Here notation is abused by adding conditional probabilities to the language, where conditional probability is defined as $\mathbf{P}_a(\varphi|\psi) = \frac{\mathbf{P}_a(\varphi \wedge \psi)}{\mathbf{P}_a(\psi)}$. The equivalence holds if the truth value of φ is not changed by learning that ψ. However the equivalence above does not hold in general. An example of this failure is in the case of an *unsuccessful update*. A successful update with φ will result in a state where the agents believe that φ. But for example when you get the information that 'you do not know that it is raining and it is raining', afterwards you will not believe that you do not know that it is raining[2]. In Gerbrandy (1999) this topic is discussed more extensively, including the *muddy children puzzle*, where an interesting example of unsuccessful updating occurs. The probabilistic version of this is quite similar. Suppose I flip a fair coin, such that you cannot see the outcome, but I can. Then I tell you that the probability you assign to heads is not zero and that the outcome is tails. After that update you do assign probability zero to the outcome being heads.

$$[\mathbf{P}_a(\text{heads}) > 0 \wedge \text{tails}]\mathbf{P}_a(\mathbf{P}_a(\text{heads}) > 0 \wedge \text{tails}) = 0$$

[2] Some people argue that the occurrence of unsuccessful updates is due to the fact that these sentences are not properly labeled with time indices. In that case propositions can never change in truth value. In the context of dynamic logic however it seems more useful to have a notion of propositions that can change truth values.

However

$$\mathbf{P}_a(\mathbf{P}_a(\text{heads}) > 0 \wedge \text{tails}|\mathbf{P}_a(\text{heads}) > 0 \wedge \text{tails}) = 1$$

This is due to the difference of perspective on information change in probability theory and dynamic epistemic logic as was explained in the introduction. Although after a public announcement it is common knowledge that φ is true at the time of the announcement, it need not be common knowledge that φ after the announcement, because φ may involve statements about information.

One of remaining open question about DEL is to give a syntactic characterization of those sentences which may lead to an unsuccessful update. This is also unsolved for PDEL. Some progress has been made for a syntactic characterization of those sentences which always lead to successful updates by van Benthem (2002b).

Although acquiring new information can now be modeled in a way that takes higher order information into account, the language is not very sophisticated yet with regard to how an update came about. When one models games, public announcements are made by the players using some sort of strategy. Perhaps they do not reveal all they know, or perhaps their actions depend on very complex protocols. For the kind of updates we are considering we have to assume that "the announcement that φ" came about by some process of which the result was either an announcement of φ or an announcement of $\neg\varphi$. For example when it is an answer to a question regarding φ. Consider the following example by (Albers, 2003, chapter 1). A fair die is thrown and one agent a can see the outcome, whereas b cannot. Now b can inform a about the outcome by saying either that the outcome is odd, or that the outcome is even, or that it is a multiple of three. Now if b truthfully states that the outcome is even, what is the probability that the outcome is 6? The answer depends on b's strategy or protocol. This kind of update cannot be dealt with in PDEL. For more sophisticated extensions of PDEL see section 7.

4. Reasoning about probability

In probability theory, inferences are often justified by making a model of the situation that is being investigated. Then the relevant propositions are analyzed in that model. In logic, inferences are usually shown to be valid by translating them into a formal language and showing that the conclusion can be deduced from the premises in a formal proof system. In this section I provide a way to make models of particular situations and a formal proof system for probabilistic dynamic epistemic logic.

In section 6 an example of application of each of these approaches is given.

4.1. Building a Model

Although I introduced probabilistic epistemic models in section 3, it is still not immediately clear how one can model a specific situation. In Halpern and Tuttle (1993) an approach for this is given. In this section I give a similar approach, which differs from the approach in Halpern and Tuttle (1993) in the sense that I introduce purely probabilistic models. From that perspective one could say that in Halpern and Tuttle (1993) only purely probabilistic models are considered that are $\mathsf{S5}$ and connected. The interesting feature that these models have is that the agents have a *common prior*, which means that if they were to forget everything they have learned, then they would agree on all the probabilities. It is still an open question whether this class of models can be characterized by a sentence in the language of PDEL. The importance of having a common prior is that it is often assumed in game theory (Aumann, 1976).

As before, suppose an agent knows that a coin lands heads in one third of the cases or is fair, but she does not know which of these is the case. It is not easy to make a model of this at once. In this section I show how to construct a probabilistic epistemic model from two models: one for the *non-probabilistic* information (i.e. propositional and epistemic information) and another for the *probabilistic* information. It is often easier to think about these domains of information separately. The idea is to multiply an epistemic model with what I call a purely probabilistic model.

DEFINITION 5. (Purely probabilistic models). *Let a nonempty set E and a finite set of agents $\mathcal{A}$ be given. A purely probabilistic epistemic model* M *is a triple* $(\mathsf{W}, \mathsf{R}, \mathsf{P})$ *such that:*

$-\mathsf{W} \neq \emptyset$

$-\mathsf{R} : \mathcal{A} \to 2^{\mathsf{W} \times \mathsf{W}}$

$-\mathsf{P} : \mathsf{W} \to \{P \mid P$ *is a probability functions with domain* E

Thus a probability function is assigned to each world and the domain of all of these is E. I call these models *purely* probabilistic, because there are no propositional variables in them, but probability functions have a similar role. Nevertheless the accessibility relations will be interpreted epistemically.

Given an *epistemic* model M and a *purely probabilistic* model M we can make a *probabilistic epistemic* model $\mathfrak{M}$. Both models must be defined with respect to the same agents, and the set of possible worlds W of the epistemic model must be the domain of all probability spaces in the range of P (i.e. $E = W$). Worlds in the purely probabilistic model provide prior probability distributions over the set of worlds of the epistemic model. The probability an agent assigns to a set of worlds is its prior probability conditionalized on the agent's knowledge.

DEFINITION 6. (Multiplication). *Given an epistemic model M and a purely probabilistic model M, such that $M = (W, R, V)$ and $\mathsf{M} = (\mathsf{W}, \mathsf{R}, \mathsf{P})$*

$$M \otimes \mathsf{M} = \mathfrak{M} = (\mathfrak{W}, \mathfrak{R}, \mathfrak{V}, \mathfrak{P})$$

iff

$$\begin{array}{lcl} \mathfrak{W} & = & W \times \mathsf{W} \\ \mathfrak{R}(a) & = & \{((w, \mathsf{w}), (v, \mathsf{v})) \mid wR(a)v \wedge \mathsf{w}\mathsf{R}(a)\mathsf{v}\} \\ \mathfrak{V}(p) & = & V(p) \times \mathsf{W} \\ dom(\mathfrak{P}(a, (w, \mathsf{w}))) & = & \{v \mid wR(a)v\} \times \{\mathsf{w}\} \\ \mathfrak{P}(a, (w, \mathsf{w}))(v, \mathsf{w}) & = & \dfrac{\mathsf{P}(\mathsf{w})(v)}{\sum_{(u,\mathsf{w}) \in dom(\mathfrak{P}(a,(w,\mathsf{w})))} \mathsf{P}(\mathsf{w})(u)} \end{array}$$

The domain of the probability function that is assigned to an agent at a pair (w, w) contains those pairs (v, w) such that v is accessible to the agent in the epistemic model from w. So the domain is a probabilistic copy of the set of worlds accessible to the agent in the epistemic model. The probability assigned to a world by an agent is its conditional probability given that it is in the domain of the agent's probability function (disregarding the second element of the pair).

Now we can deal with the initial example: a coin is tossed and an agent a does not know the outcome. So she cannot distinguish worlds where the outcome is heads from worlds where it is tails. She knows the coin is fair or that it lands heads one third of the times, but she does not know which is the case. Hence she cannot distinguish worlds where the probability of heads is $\frac{1}{2}$ from worlds where it is $\frac{1}{3}$. I can make an *epistemic* model for a's information about the outcome and a *purely probabilistic* model for a's information about the coin. These two models and the result of multiplying these models are shown in figure 1. Now we can see that a sentence $\mathbf{P}_a(\varphi) \geq q$ should *not* be read as 'the probability a assigns to φ is greater than or equal to q,' because there need not be a *unique* probability q assigns to φ. In the example a cannot distinguish two probability distributions. $\mathbf{P}_a(\varphi) \geq q$ should be read as 'the probability a *should* assign to φ is greater than or equal to q, given the "actual" probability distribution over the worlds and given

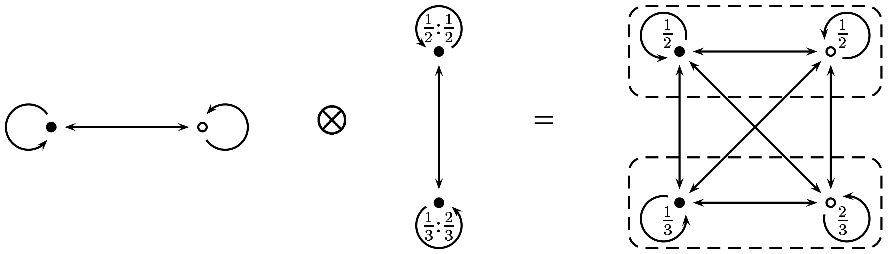

Figure 1. An example of multiplication. The epistemic model is on the left, the purely probabilistic model is in the middle, and the probabilistic epistemic model is on the right. In the probabilistic epistemic models the solid nodes indicate that the outcome is heads, and an open node indicates the outcome is tails. The dashed boxes indicate the domains of the probability functions.

a's other information.' Hence we should be interested in sentences of the form $\Box_a(\mathbf{P}_a(\varphi) \geq q)$. Such a sentence holds iff a knows the probability she should assign to φ is greater than or equal to q.

There is one requirement the underlying models should meet for multiplication to work: the sets of worlds accessible to the agents should have non-zero probabilities. This ensures that the probability functions are well-defined, because it ensures that the set of worlds accessible to an agent is not empty and that no division by zero occurs.

4.2. Proof system, soundness and completeness

The proof system PDEL provided in this section is based on the proof system DEL in Gerbrandy (1999) for dynamic epistemic logic, and the proof system AX_{MEAS} in Fagin and Halpern (1994) for probabilistic epistemic logic. These systems two systems joined with the axioms Probability-Update 1 and Probability-Update 2 constitute the proof system for probabilistic dynamic epistemic logic.

DEFINITION 7. (Proof System). *The proof system of probabilistic dynamic epistemic logic,* PDEL*, is provided in Figure 2.*

Let us call the axioms and rules for propositional logic, epistemic logic and update logic without the Probability-Update axioms DEL and the axioms and rules for propositional logic, epistemic logic, linear inequalities and probability logic PEL. In Gerbrandy and Groeneveld (1997) the soundness and completeness of DEL is proved, although it is proved with respect to non-well-founded objects, the correspondence between these and epistemic models implies it is also sound and complete for epistemic models. In Fagin and Halpern (1994) the soundness and completeness of PEL is proved. Although it is proved for a more general class of

Propositional Logic	
PC	$\vdash \varphi$ where φ is an instance of a propositional tautology
Epistemic Logic	
$\Box_a$-distribution	$\vdash \Box_a(\varphi \to \psi) \to (\Box_a\varphi \to \Box_a\psi)$
$\Box_a$-necessitation	From $\vdash \varphi$, infer $\vdash \Box_a\varphi$
Update Logic	
$[\varphi]$-distribution	$\vdash [\varphi](\psi \to \chi) \to ([\varphi]\psi \to [\varphi]\chi)$
Functionality	$\vdash \neg[\varphi]\psi \leftrightarrow [\varphi]\neg\psi$
Atomic Permanence	$\vdash p \leftrightarrow [\varphi]p$
Knowledge-Update	$\vdash [\varphi]\Box_a\psi \leftrightarrow \Box_a(\varphi \to [\varphi]\psi)$
Probability-Update 1	$\vdash \mathbf{P}_a(\varphi) > 0 \to (([\varphi]\sum_{i=1}^n q_i\mathbf{P}_a(\varphi_i) \geq q) \leftrightarrow (\sum_{i=1}^n q_i\mathbf{P}_a(\varphi \wedge [\varphi]\varphi_i) \geq q\mathbf{P}_a(\varphi)))$
Probability-Update 2	$\vdash \mathbf{P}_a(\varphi) = 0 \to (([\varphi]\sum_{i=1}^n q_i\mathbf{P}_a(\varphi_i) \geq q) \leftrightarrow (\sum_{i=1}^n q_i\mathbf{P}_a([\varphi]\varphi_i) \geq q))$
$[\varphi]$-necessitation	From $\vdash \psi$, infer $\vdash [\varphi]\psi$
Linear Inequalities	
0 terms	$\vdash \sum_{i=1}^n q_i\mathbf{P}_a(\varphi_i) \geq q \leftrightarrow (\sum_{i=1}^n q_i\mathbf{P}_a(\varphi_i)) + 0\mathbf{P}_a(\varphi_{k+1}) \geq q$
Permutation	$\vdash \sum_{i=1}^n q_i\mathbf{P}_a(\varphi_i) \geq q \to \sum_{i=1}^n q_{j_i}\mathbf{P}_a(\varphi_{j_i}) \geq q$ where $j_1, \ldots, j_k$ is a permutation of $1 \ldots k$
Addition	$\vdash \sum_{i=1}^n q_i\mathbf{P}_a(\varphi_i) \geq q \wedge \sum_{i=1}^n q_i'\mathbf{P}_a(\varphi_i) \geq q' \to \sum_{i=1}^n (q_i + q_i')\mathbf{P}_a(\varphi_i) \geq (q + q')$
Multiplication	$\vdash (\sum_{i=1}^n q_i\mathbf{P}_a(\varphi_i) \geq q) \leftrightarrow (\sum_{i=1}^n dq_i\mathbf{P}_a(\varphi_i) \geq db)$ where $d > 0$
Dichotomy	$\vdash (t \geq q) \vee (t \leq q)$
Monotonicity	$\vdash (t \geq q) \to (t > q')$ where $q > q'$
Probability Logic	
Nonnegativity	$\vdash \mathbf{P}_a(\varphi) \geq 0$
Probability of truth	$\vdash \mathbf{P}_a(\top) = 1$
Additivity	$\vdash \mathbf{P}_a(\varphi \wedge \psi) + \mathbf{P}_a(\varphi \wedge \neg\psi) = \mathbf{P}_a(\varphi)$
Equivalence	From $\vdash \varphi \leftrightarrow \psi$, infer $\vdash \mathbf{P}_a(\varphi) = \mathbf{P}_a(\psi)$

Figure 2. The proof system PDEL for probabilistic dynamic epistemic logic

models than the models of definition 1, the models of definition 1 form a subclass, therefore PEL is still sound. Furthermore, in the completeness proof, the countermodels are probabilistic epistemic models in the sense of definition 1 (their system has the finite model property). Therefore this logic is also complete for these models.

The axiom Probability-Update 1 clarifies the relationship between conditional probability and the notion of updating probabilities in this paper. The relationship is best captured by the following equivalence, which was pointed out to me by Johan van Benthem:

$$[\psi](\mathbf{P}_a(\varphi) = q) \qquad \text{iff} \qquad \mathbf{P}_a([\psi]\varphi \mid \psi) = q$$

Note that it is not a normal modal logic, because we do not have universal substitution. This is due to the existence of unsuccessful updates. For example, $\vdash [p]\Box_a p$ is a theorem, but $\vdash [\neg\Box_a p \wedge p]\Box_a(\neg\Box_a p \wedge p)$ is not, although it is a substitution instance. There are more principles in dynamic epistemic logics which are valid but not derivable schematically (see (van Benthem, 2002b)).

THEOREM 1. (Soundness). *If* $\vdash \varphi$ *then* $\models \varphi$.

Proof. The soundness of the axioms of DEL and PEL I do not prove. The proof of their soundness can be found in Gerbrandy (1999) or Gerbrandy and Groeneveld (1997), and Fagin and Halpern (1994).

For Probability-Update 1, first of all note that

$$\begin{array}{l} \{v \mid (M,v) \models [\varphi]\psi \wedge \varphi \text{ and } v \in dom(P(a,w))\} \\ = \\ \{v \mid (M_\varphi, v_\varphi) \models \psi \text{ and } v_\varphi \in dom(P_\varphi(a,w_\varphi))\} \end{array} \qquad (1)$$

Suppose $(M,w) \models \mathbf{P}_a(\varphi) > 0$. Now the following equivalences hold:
$(M,w) \models [\varphi]\sum_{i=1}^{n} q_i \mathbf{P}_a(\varphi_i) \geq q)$
$\equiv$ {truth definition}
$(M_\varphi, w_\varphi) \models \sum_{i=1}^{n} q_i \mathbf{P}_a(\varphi_i) \geq q)$
$\equiv$ {truth definition}
$\sum_{i=1}^{k} q_i P_\varphi(a,w_\varphi)(\varphi_i) \geq q$
$\equiv$ {By (1), the definition of updates, and $(M,w) \models \mathbf{P}_a(\varphi) > 0$}
$$\sum_{i=1}^{k} q_i \frac{P(a,w)([\varphi]\varphi_i \wedge \varphi)}{P(a,w)(\varphi)} \geq q$$
$\equiv$ {algebra}
$\sum_{i=1}^{k} q_i P(a,w)([\varphi]\varphi_i \wedge \varphi) \geq qP(a,w)(\varphi)$
$\equiv$ {truth definition}
$(M,w) \models \sum_{i=1}^{n} q_i \mathbf{P}_a([\varphi]\varphi_i \wedge \varphi) \geq q\mathbf{P}_a(\varphi))$

The soundness of probability-update 2 is immediate from the definition of update, because if φ has probability zero, nothing happens to the domain of the probability function after updating with φ.

To prove completeness I provide a translation of the sentences of probabilistic dynamic epistemic logic to the sentences of probabilistic epistemic logic. Given that PEL is complete for probabilistic epistemic logic, it then suffices to show that a sentence is provably equivalent in PDEL to its translation. This is the same proof method as is used in (Gerbrandy, 1999). One uses the axioms of the proof system to translate the language to another language for which a complete proof system is available.

DEFINITION 8. (Translation from $\mathcal{L}_{\mathcal{PA}}^{\mathbf{P}[\cdot]}$ to $\mathcal{L}_{\mathcal{PA}}^{\mathbf{P}}$). *The translation* $t : \mathcal{L}_{\mathcal{PA}}^{\mathbf{P}[\cdot]} \to \mathcal{L}_{\mathcal{PA}}^{\mathbf{P}}$ *is defined as follows:*

1. $t(p) = p$
2. $t(\neg\varphi) = \neg t(\varphi)$
3. $t(\varphi \wedge \psi) = t(\varphi) \wedge t(\psi)$
4. $t(\Box_a \varphi) = \Box_a t(\varphi)$
5. $t(\sum_{i=1}^{n} q_i \mathbf{P}_a(\varphi_i) \geq b) = (\sum_{i=1}^{n} q_i \mathbf{P}_a(t(\varphi_1)) \geq q)$
6. $t([\varphi]p) = p$
7. $t([\varphi]\neg\psi) = \neg t([\varphi]\psi)$
8. $t([\varphi](\psi \wedge \chi)) = t([\varphi]\psi) \wedge t([\varphi]\chi)$
9. $t([\varphi]\Box_a \psi) = \Box_a (t(\varphi) \to t([\varphi]\psi))$
10. $t([\varphi](\sum_{i=1}^{n} q_i \mathbf{P}_a(\varphi_i) \geq q)) =$

$$(\mathbf{P}_a(t(\varphi)) > 0 \wedge (\textstyle\sum_{i=1}^{n} q_i \mathbf{P}_a(t(\varphi) \wedge t([\varphi]\varphi_i)) \geq q\mathbf{P}_a(t(\varphi))))$$
$$\vee$$
$$(\mathbf{P}_a(t(\varphi)) = 0 \wedge \textstyle\sum_{i=1}^{n} q_i \mathbf{P}_a(t([\varphi]\varphi_i)) \geq q)$$

Note that although the update operator has an infinitary character (it has effects for the entire model), when evaluating a sentence the effect only needs to be given for the finite intention depth (the number of stacked modal operators).

LEMMA 2. *For every sentence φ of* PDEL*, the translation of that sentence $t(\varphi)$ is a sentence in probabilistic epistemic logic to which it is provably equivalent in* PDEL.

THEOREM 2. (Completeness). *If $\models \varphi$, then $\vdash \varphi$.*

Proof. The axiom system PEL is complete with respect to the semantics of $\mathcal{L}_{\mathcal{PA}}^{\mathbf{P}}$. Therefore the proof of any sentence $\varphi \in \mathcal{L}_{\mathcal{PA}}^{\mathbf{P}}$ can be obtained by only using axioms and rules of the axiom system PEL. From lemma 2 we get that every sentence $\varphi \in \mathcal{L}_{\mathcal{P}a}^{\mathbf{P}[\cdot]}$ is provably equivalent in PDEL to a sentence in $t(\varphi) \in \mathcal{L}_{\mathcal{PA}}^{\mathbf{P}}$. Therefore the axiom system is complete for probabilistic dynamic epistemic logic.

COROLLARY 1. *The language of probabilistic dynamic epistemic logic is just as expressive as the language of probabilistic epistemic logic.*

COROLLARY 2. *The validity problem for probabilistic dynamic epistemic logic is decidable.*

Proof. This follows directly from the decidability result in Fagin and Halpern (1994).

As complexity is concerned, the validity problem for probabilistic epistemic logic is complete for polynomial space. However the translation from definition 8 is exponential in space in the depth of probabilistic operators after an update, i.e. sentences of the form $[\varphi](\mathbf{P}_a(\mathbf{P}_a(\ldots)))$. We can of course conclude that polynomial space is a lower bound on complexity and exponential space is an upper bound on complexity. The complexity problem is also still open for DEL, because the size of the translation is also exponential in space in the depth of sequences of knowledge operators after update operators.

5. Bisimulation for probabilistic dynamic epistemic logic

Bisimulation is a useful notion in modal logic. It generally holds that if two structures are bisimilar, then they are behaviorally indistinguishable. In the case of probabilistic epistemic models, behaviorally indistinguishable means satisfying the same sentences. A well-known result in modal logic is that if two pointed models are bisimilar, then they satisfy the same sentences (see for example Blackburn et al. (2001) for a textbook explanation of this notion.) In this section I show that such a result holds for probabilistic dynamic epistemic logic as well.

DEFINITION 9. (Bisimulation). *I use the following abbreviations.*

$$\begin{aligned} \mathsf{forth}(E, E') &:= \forall x \in E \exists y \in E'(xBy) \\ \mathsf{back}(E, E') &:= \forall y \in E' \exists x \in E(xBy) \end{aligned}$$

Let two probabilistic epistemic models M and M' be given. A relation $B \subseteq W \times W'$ is a bisimulation iff for all $w \in W$ and $w' \in W'$, if wBw', then for all $n \in \mathcal{A}$ the following hold:

atoms $w \in V(p)$ *iff* $w' \in V'(p))$ *for every* $p \in \mathcal{P}$

forth $\mathsf{forth}(\{v \mid wR(a)v\}, \{v' \mid w'R'(a)v'\})$

back $\mathsf{back}(\{v \mid wR(a)v\}, \{v' \mid w'R'(a)v'\})$

pforth *For every $E \subseteq dom(P(a,w))$ there is an $E' \subseteq dom(P'(a,w'))$ such that*

$$P(a,w)(E) \leq P'(a,w')(E') \text{ and } \mathsf{back}(E,E')$$

pback *For every $E' \subseteq dom(P'(a,w'))$ there is an $E \subseteq dom(P(a,w))$ such that*

$$P'(a,w')(E') \leq P(a,w)(E) \text{ and } \mathsf{forth}(E,E')$$

I write $(M,w)\leftrightarroweq(M',w')$, if there is a bisimulation between M and M' linking w and w'.

Atoms, **forth** and **back** are the usual conditions for bisimulation. I added **pforth** and **pback** to accommodate probabilistic sentences. For those readers familiar with probability theory, this definition can easily be extended to the more general notion of probabilistic epistemic models given in Fagin and Halpern (1994) with probability spaces, where one takes the inner measure instead of the probability function in **pforth** and **pback**. The theorem below also holds for these models.

THEOREM 3. *For all models (M,w) and (M',w') and for all sentences φ, if $(M,w)\leftrightarroweq(M',w')$, then $(M,w) \models \varphi$ iff $(M',w') \models \varphi$*

Proof. By induction on φ. Suppose $(M,w)\leftrightarroweq(M',w')$. The base case and cases for conjunction, negation and individual epistemic operators $\Box_a$ are straightforward. By lemma 2, we get the case for updates for free.

Suppose uBu' and $(M,u) \models \sum_{i=1}^{n} q_i \mathbf{P}_a(\varphi_i) \geq q$. Let

$$E_i = \{v \in dom(P(a,u)) \mid (M,v) \models \varphi_i\}$$

and

$$E_i' = \{v' \in dom(P'(a,u')) \mid (M',v') \models \varphi_i\}$$

If we show that $P(a,u)(E_i) \leq P'(a,u')(E_i')$ we are done. From uBu' and **pforth** it follows that there is an $S' \subseteq dom(P'(a,w'))$ such that

$$P(a,w)(E_i) \leq P'(a,w')(S') \text{ and } \mathsf{back}(E_i,S')$$

The induction hypothesis together with $\mathsf{back}(E_i,S')$ imply that $(M',v') \models \varphi_i$ for every $v' \in S'$. Therefore $S' \subseteq E_i'$ and therefore $P'(a,u')(S') \leq P'(a,u')(E_i')$. Now we conclude that

$$P(a,u)(E_i) \leq P'(a,w')(S') \leq P'(a,u')(E_i')$$

The case for right to left is analogous. Which gives as an additional result that $P(a,u)(E_i) = P'(a,u')(E_i')$.

Therefore for all models (M, w) and (M', w'), if $(M, w) \leftrightarroweq (M', w')$, then for all sentences φ: $(M, w) \models \varphi$ iff $(M', w') \models \varphi$

The converse of theorem 3 also holds when the models are finite or when one uses an infinitary language which allows conjunctions over arbitrary sets of sentences.

The notion of bisimulation presented in this paper can also be applied to probability spaces, which can be seen as special cases of probabilistic epistemic Kripke models, and therefore it is also interesting for probability theory, to see whether two models of the same experiment are equivalent. It would be worthwhile to investigate the mathematics of this further.

There are richer languages for reasoning about probability which are able to distinguish bisimilar models. Dependent upon the language which is used in reasoning about probability, one might wonder whether there is information being modeled which is not needed. This leads to the question whether one can define minimal models. In modal logic one can define a minimal model with respect to an arbitrary Kripke model by identifying all bisimilar worlds. The result for modal logic seems to be folklore. This can also be done for probabilistic epistemic models.

DEFINITION 10. (Minimal models). *Let a probabilistic epistemic model $M = (W, R, V, P)$ be given. The minimal model associated with M is the model $M' = (W', R', V', P')$, where:*

- $W' = \{E \subseteq W \mid \textit{for all } w, v \in E : (M, w) \leftrightarroweq (M, v)\}$

- $R'(a) = \{(E, E') \subseteq (W' \times W') \mid \textit{there is a } w \in E \textit{ and a } v \in E' \textit{ such that } wR(a)v\}$

- $V'(p) = \{E \mid \textit{there is a } w \in E \textit{ such that } w \in V(p)\}$

- $dom(P'(a, E)) = \{E' \subseteq W' \mid \textit{there is a } w \in E \textit{ and a } v \in E' \textit{ such that } v \in dom(P(a, w))\}$

- $P'(a, E)(E') = \sup\{q \in \mathbb{R} \mid \textit{there is a } w \in E \textit{ such that } q = P(a, w)(E' \cap dom(P(a, w)))\}$

where in the last clause $E' \in dom(P'(a, E))$.

LEMMA 3. *Every model M is bisimilar to the minimal model M' associated with it.*

Proof. Let $M = (W, R, V, P)$ and $M' = (W', R', V', P')$ as in definition 10. Now we will show that the $\in$ relation on $W \times W'$ is a bisimulation. The case for **atoms**, **forth**, and **back** are straightforward.

For the case for **pforth** assume that $w \in E$, where $E \in W'$. Suppose $S \subseteq dom(P(a,w))$. Now we have to show that there is a subset $\mathbb{S}$ of the domain of $P'(a,E)$ such that the probability assigned to it is greater or equal to the probability assigned to S and $\mathsf{back}(S,\mathbb{S})$. Let $\mathbb{S} =$

$$\{E' \in W' \mid \text{there is a } v \in S \text{ such that } v \in E'\}$$

From the definition of $dom(P'(a,E))$ it follows that $\mathbb{S} \subseteq dom(P'(a,E))$. It is also easily seen that $\mathsf{back}(S,\mathbb{S})$. From the definition of $P'(a,E)$ it follows that $P(a,w)(E' \cap dom(P(a,w))) \leq P'(a,E)(E')$, for every $E' \in \mathbb{S}$. Therefore $P(a,w)(S) \leq P'(a,E)(\mathbb{S})$.

For the case of **pback** assume that $w \in E$, where $E \in W'$. Suppose $\mathbb{S} \subseteq dom(P'(a,E))$. Now we have to show that there is a subset S_w of the domain of $P(a,w)$ such that the probability assigned to it is greater or equal to the probability assigned to $\mathbb{S}$ and $\mathsf{forth}(S,\mathbb{S})$. Let $S_w =$

$$\{v \mid \text{there is an } E' \in \mathbb{S}' \text{ such that } v \in E' \text{ and } v \in dom(P(a,w))\}$$

It is obviously the case that $S_w \subseteq dom(P(a,w))$ and $\mathsf{forth}(S_w,\mathbb{S})$. Suppose that $P'(a,E)(\mathbb{S}) > P(a,w)(S_w)$. Therefore there is a $u \in E$ such that for the set $S_u =$

$$\{v \mid \text{there is an } E' \in \mathbb{S}' \text{ such that } v \in E' \text{ and } v \in dom(P(a,v))\}$$

$P(a,u)(S_u) > P(a,w)(S_w)$. But because both w and v are in E they must be bisimilar. From **pback** it follows that for S_u there is a set $S \subseteq dom(P(a,w))$ such that $P(a,u)(S_u) \leq P(a,w)(S)$ and $\mathsf{forth}(S,S_u)$. Therefore $S \subseteq S_w$. But this leads to a contradiction, because it now follows that $P(a,w)(S) \leq P(a,w)(S_w) < P(a,u)(S_u) \leq P(a,w)(S)$.

Therefore $M \leftrightarroweq M'$.

In the literature on probabilistic transition systems, notions of probabilistic bisimulation have also been put forward: notably those by Larsen and Skou (1991), who introduce a notion of bisimulation for discrete systems, and de Vink and Rutten (1999), which is a generalization of Larsen en Skou's approach to general probabilistic transition systems. There are some small differences between these notions of probabilistic bisimulation and the notion presented in this paper, and the question whether the notions coincide, or one is more general than the other requires further investigation. However as far as I know, the result that bisimilarity of two probabilistic epistemic models implies that they have the same probabilistic dynamic epistemic theory is new, as well as the result about minimal models are new.

6. Examples

In this section two examples are given that apply the theory presented in this paper. Section 6.1 links up with the approach of building models of section 4.1. It shows that when higher order information is involved it can be useful to formalize the inferences that are involved in the example with the language introduced in this paper. Section 6.2 shows how the Monty Hall Dilemma can be formalized. Here a syntactical approach to the problem is given instead of the usual semantic approach. It shows how PDEL can be used to do the usual reasoning about conditional probabilities in a rather elegant way.

6.1. Coins

Let us look at the following game, which is based on an example by Van Rooy 2003, which is in turn based on an example by (Hirshleifer and Riley, 1992, p. 220). Suppose two players, a and b, are playing the following game. A coin is tossed and the players have to guess the outcome. Player b guesses first, and after hearing player b's guess player a guesses the outcome. If they guess the same outcome both players receive a payoff of 12 euros regardless of the outcome, otherwise, when the guesses differs, the player who guessed the outcome correctly receives 30 euros.

Suppose that it is not known to player a whether the coin is fair, or whether the coin lands heads with probability one third, but player b does know, and this is common knowledge. Consequently player a does not really know which game she is playing: it is a game of incomplete information (see Binmore, 1992, chapter 11). One can construct a probabilistic epistemic model for this situation in the way described in section 4.1 (see Figure 3). A game theoretical analysis of this situation tells us that when the coin is fair and player a is risk neutral, the best strategy for a is to guess the opposite of player b's guess. The expected payoff is 15 euros. But when the coin is not fair and player a is risk neutral, the best strategy for player a is to guess the same outcome as player b. This strategy guarantees an outcome of 12 euros. If player a would guess differently, the expected outcome is only 10 euros, because player b's best strategy is to guess the outcome that is most likely (in this case tails). In that the probability that player a wins 30 euros is $\frac{1}{3}$.

Now suppose a public announcement is made as to which probabilities player b assigns to the outcome tails. Afterwards player a will know what strategy to follow. Note that the announcement does not say anything about what the actual outcome is. Player a only learns

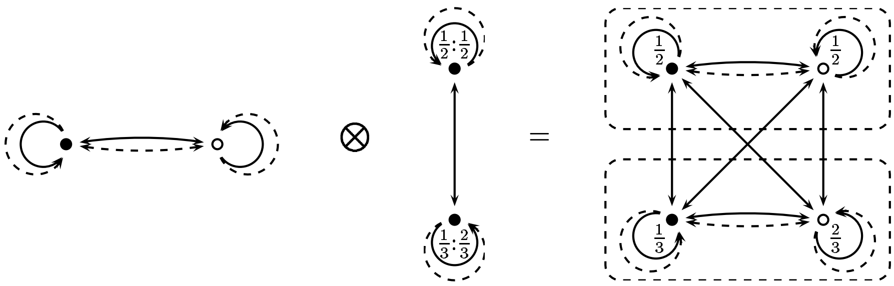

Figure 3. The construction of a probabilistic epistemic model for the situation where player a does not know whether the coin is fair or not, but player b does. In the (probabilistic) epistemic models the solid nodes indicate that the outcome is heads, and an open node indicates the outcome is tails. The solid lines represent the accessibility relation of player a, the dashed lines represent the accessibility relation of player b. The dashed boxes indicate the domains of the probability functions.

about the probabilities player b assigns to the outcomes. This higher order information determines what the best strategy is.

6.2. The Monty Hall Dilemma

In this section I discuss a puzzle that often leads to furious discussions. It received worldwide attention after Marilyn vos Savant discussed it in her column 'Ask Marilyn' in *Parade Magazine* (Vos Savant, 1990), where she answers questions sent in by the readers.

> Suppose you're on a game show, and you're given the choice of three doors. Behind one door is a car, behind the others, goats. You pick a door, say number 1, and the host, who knows what's behind the doors, opens another door, say number 3, which has a goat. He says to you, "Do you want to pick door number 2?" Is it to your advantage to switch your choice of doors?
>
> Craig. F. Whitaker
> Columbia, MD

This is the Monty Hall Dilemma. It got its name from the American game show host Monty Hall (see Selvin (1975)). Vos Savant, who, reportedly, is listed in the Guinness Book of World Records for the highest IQ, argued that it is to your advantage to switch. If you switch, you get a goat in one third of the cases and win the car in two third of the cases. This could be argued as follows. Suppose you initially pick the door with the car, then you should not switch. This happens in one third of the cases. Suppose on the other hand you initially pick a door that contains a goat, which happens in two third of the cases. Monty Hall cannot open the door with the car and he cannot open the door you picked. He has to open the other door with a goat. So, if you pick a door with a goat, Monty Hall only has one option. After he opens that door, the remaining unopened door you did not pick must contain the car. Therefore, if you initially pick a door with a goat, switching will guarantee that you win the car. You pick such a door in two third of the cases. Hence by switching you lose in one third of the cases and you win in two third of the cases.

The Monty Hall Dilemma is a puzzle for which intuitions fail many people. The best way to show that the counterintuitive results are correct is to use some formal method. PDEL provides such a method. (A comparison to other methods is given in section 7.) In fact it is quite easy to represent the inference in the language of probabilistic dynamic epistemic logic.

In the following analysis I show what happens to the contestant's information, by formalizing the information changes that occur. I prove that Vos Savant's inference is valid by giving a formal proof.

I take the set of agents to be $\mathcal{A} = \{c, m\}$ (the contestant and Monty Hall). The set of propositional variables $\mathcal{P}$ is the union of the three sets $A = \{A_1, A_2, A_3\}$ (where A_i means that the car is behind door number i), $C = \{C_1, C_2, C_3\}$ (where C_i means that the contestant initially chooses door number i), and $O = \{O_1, O_2, O_3\}$, (where O_i means that door number i is opened by Monty Hall).

Now I characterize the rules of the game. One of the rules is that there is only one car behind the doors, the contestant may only choose one door, and Monty Hall may only open one door.

$$\begin{array}{lcl} \mathsf{onecar} & = & \bigoplus A \\ \mathsf{onechoice} & = & \bigoplus C \\ \mathsf{oneopen} & = & \bigoplus O \end{array}$$

Where $\bigoplus$ means exclusive or. I assume that the contestant should assign a probability of $\frac{1}{3}$ to the car being behind a particular door. This is an assumption that has to be made to get Vos Savant's answer. Moreover I assume the contestant does not learn anything about the location of the car by picking a door. Therefore the contestant should still assign a probability of $\frac{1}{3}$ after picking a door: the contestant's choice is independent of where the car is.

$$\begin{array}{lcl} \mathsf{equal} & = & \bigwedge_{i \in \{1,2,3\}} \mathbf{P}_c(A_i) = \frac{1}{3} \\ \mathsf{independentAC} & = & \bigwedge_{j \in \{1,2,3\}} [C_j]\mathsf{equal} \end{array}$$

This is a nice way of expressing independence. This assumption remains implicit in most other analyses I found of the Monty Hall Dilemma. The crucial part of the analysis of the Monty Hall Dilemma is to see under what conditions Monty Hall opens a door. He opens exactly one door such that the contestant did not pick it and the car is not behind it.

$$\mathsf{conditions} = \bigwedge_{i,j=\{1,2,3\}} [C_i](O_j \leftrightarrow (\neg A_j \wedge \neg C_j \wedge \bigwedge_{k\in\{1,2,3\},k\neq j} \neg O_k))$$

Let us use $\mathsf{initial}$ as an abbreviation for the conjunction of onecar, $\mathsf{onechoice}$, $\mathsf{oneopen}$, equal, $\mathsf{independentAC}$, and $\mathsf{conditions}$.

The question is whether the contestant should switch or not:

$$\mathsf{switch} = [C_1][O_3]\mathbf{P}_c(A_1) \leq \mathbf{P}_c(A_2)$$

If this sentence is true, then the chances that the contestant wins the car do not decrease by switching. It turns out that $\mathsf{initial}$ is not enough to deduce this result. What is needed is that the contestant is informed about the game: $\mathbf{P}_c(\mathsf{initial}) = 1$. We also need two other very natural assumptions, namely that the probability that the contestant chooses door number one is greater than zero according to the contestant: $\mathbf{P}_c(C_1) > 0$, and that after the contestant chooses door number one the probability that Monty Hall opens door number three is greater that zero according to the contestant: $[C_1]\mathbf{P}_c(O_3) > 0$. This suffices to deduce switch.

The $\mathsf{independentAC}$ assumption implies that $[C_1]\mathbf{P}_c(A_1) = \frac{1}{3}$, and therefore:

$$[C_1]\mathbf{P}_c(O_3 \wedge A_1) \leq \tfrac{1}{3}$$

By $\mathsf{conditions}$ and $\mathsf{onechoice}$ we get $[C_1]\mathbf{P}_c(A_2 \to O_3) = 1$. Some probabilistic reasoning gives us that $[C_1]\mathbf{P}_c(O_3 \wedge A_2) = \mathbf{P}_c(A_2)$. This, together with $[C_1]\mathbf{P}_c(A_2) = \frac{1}{3}$ (from $\mathsf{independentAC}$), allows us to infer that $[C_1]\mathbf{P}_c(O_3 \wedge A_2) = \frac{1}{3}$, which yields

$$[C_1]\mathbf{P}_c(O_3 \wedge A_1) \leq \mathbf{P}_c(O_3 \wedge A_2)$$

By atomic permanence we get

$$[C_1]\mathbf{P}_c(O_3 \wedge [O_3]A_1) \leq \mathbf{P}_c(O_3 \wedge [O_3]A_2)$$

Then by probability update 1, 0-terms and $[C_1]\mathbf{P}_c(O_3) > 0$ we have:

$$[C_1][O_3]\mathbf{P}_c(A_1) \leq \mathbf{P}_c(A_2)$$

Thus far we have made no assumptions about the strategy used by the contestant or Monty Hall. We do not need this to deduce switch, but we do need to assume something about the strategy of Monty Hall if we want to deduce that the probability that the contestant wins the car by switching equals one third. Then we need to assume that if Monty Hall can choose between opening two doors (if the door the contestant picked is the same as where the car is), then the probability he opens one door is the same as the probability he opens the other door. This boils down to:

$$\mathsf{equalopen} = \bigwedge_{\{i,j,k\}=\{1,2,3\}} [C_i]\mathbf{P}_c(O_j) = \mathbf{P}_c(O_k)$$

With this we can deduce:

$$[C_1][O_3]\mathbf{P}_c(A_1) = \tfrac{1}{3} \wedge \mathbf{P}_c(A_2) = \tfrac{2}{3}$$

7. Conclusion and further research

In this paper I presented a probabilistic dynamic epistemic logic, which can be used to reason about probability, information, and information change. The difference between information change as it is modeled in this logic and as it is modeled in probability theory is that higher order information is taken into account. Besides semantics I have provided a method to build models and a sound and complete proof system. Moreover the notion of bisimulation has also been defined for this logic. It can be applied to game situations such as card games with public announcements and the Monty Hall Dilemma.

The principal advantage PDEL has with respect to probability theory is that it can be used to formalize inferences into a formal language, such that standard logical tools can be used to see whether it is a good inference. Therefore it is very suitable to model reasoning. In probability theory probabilities are assigned to sets of worlds. These sets appear in the 'language' of probability theory, which means one is always working with a specific model. This makes it difficult to assess whether inferences hold in all models, which is exactly what logic provides. Moreover, by having a language PDEL can explicitly deal with higher order information.

Concerning information change, PDEL provides a novel approach to probabilistic updating. Updating in probability theory and PDEL are very similar. In probability theory new information is usually represented as a set of possible worlds. One learns that the actual world is an element of that set. In PDEL new information is represented as a sentence. By definition 3 every sentence is associated with a set of possible worlds. One learns that the actual world is an element of that set, just as in probability theory. But by having this linguistic component that can express higher order information, we can take into account that the truth value of sentences can change due to an update. Updating with the same sentence twice may yield different results than updating once. Updating with one sentence and then another may be different from updating the other way around. These phenomena are not taken into account in probability theory, where receiving the same information twice is always the same as receiving it once, and the order in which information is received does not matter.

Now let us turn our attention to further research. Publicly learning a sentence is not the only way one can acquire new information. There are changes in information that cannot be modeled with PDEL. In the future we want to model game actions such as: one player showing another player a card, while a third player can see this is going on, but cannot see which card is being shown. To be able to handle these kinds of actions, we need to bring more of dynamic epistemic logic into PDEL, by making the dynamic operators more program-like (in the style of PDL, see Harel et al. (2000)). An extension with test, non-deterministic choice, sequential composition, and subgroup updates does not involve many difficulties. Subgroup updates are updates where some agents get new information, whereas the other agents do not get that information. The same proof technique, i.e. by a translation, for completeness applies.

There are more phenomena we would like to capture such as common knowledge, because common knowledge plays an important role in many game situations. This poses some problems on the proof system. In Baltag et al. (1998) a complete proof system for dynamic epistemic logic with common knowledge is provided, which gives good hopes that it can also be added to PDEL.

Another direction for further research is to develop a logic along the lines of Baltag (2000), where epistemic actions are viewed as epistemic action models that can be multiplied with epistemic models, yielding the result of executing the action. All these models could be made into probabilistic models. A step in this direction is made in van Benthem (2002a).

In probability theory there are other, more complex ways of incorporating new information, such as Jeffrey's rule of conditioning (Jeffrey, 1983), Dempster's rule of combination (Dempster, 1967), and cross entropy (Kullback and Leibler, 1951). All these ideas were born out of different kinds of dissatisfaction with conditional probability as a model for incorporating new information. It would also be interesting to investigate to what kinds of dynamic epistemic updates these kinds of information change correspond.

This paper is a first step in combining epistemic logic with probability theory and there are many more steps to make.

8. Acknowledgments

I would like to thank Gerard Renardel, Rineke Verbrugge, Johan van Benthem, Erik Krabbe, Theo Kuipers, Jeanne Peijnenburg, and Jan Willem Romeyn for their comments and suggestions on this paper. I would also like to thank four anonymous reviewers for their comments on earlier versions of this paper.

References

Albers, C. J.: 2003, 'Distributional Inference: The Limits of Reason'. Ph.D. thesis, Institute for Mathematics and Computing Science, Groningen.

Aumann, R. J.: 1976, 'Agreeing to Disagree'. *Annals of Statistics* **4**(6), 1236 – 1239.

Bacchus, F.: 1990, *Representing and Reasoning with Probabilistic Knowledge, A Logical Approach to Probabilities.* Cambridge, Massachusetts: The MIT Press.

Baltag, A.: 2000, 'A Logic of Epistemic Actions (Extended Version)'. To appear as CWI Technical report.

Baltag, A., L. Moss, and S. Solecki: 1998, 'The Logic of Public Announcements and Common Knowledge'. In: *TARK 98.*

Binmore, K.: 1992, *Fun and Games.* D.C. Heath and Company.

Blackburn, P., M. de Rijke, and Y. Venema: 2001, *Modal Logic*, Vol. 53 of *Cambridge Tracts in Theoretical Computer Science.* Cambridge: Cambridge University Press.

de Vink, E. and J. Rutten: 1999, 'Bisimulation for probabilistic transition systems: a coalgebraic approach'. *Theoretical Computer Science* **221**, 271–293.

Dempster, A. P.: 1967, 'Upper and lower probabilities induced by a multivalued mapping'. *Annals of Mathematical Statistics* **38**, 325–329.

Fagin, R. and J. Halpern: 1994, 'Reasoning About Knowledge and Probability'. *Journal of the Association for Computing Machinery* **41**(2), 340–367.

Fagin, R., J. Halpern, Y. Moses, and M. Vardi: 1995, *Reasoning About Knowledge.* Cambridge, Massachusetts: MIT Press.

Gerbrandy, J.: 1999, *Bisimulations on Planet Kripke.* Amsterdam: ILLC Dissertation Series.

Gerbrandy, J. and W. Groeneveld: 1997, 'Reasoning about Information Change'. *Journal of Logic, Language, and Information* **6**, 147 –196.

Gillies, D.: 2000, *Philosophical Theories of Probability*, Philosophical issues in science. London and New York: Routledge.

Halpern, J. and M. Tuttle: 1993, 'Knowledge, Probability, and Adversaries'. *Journal of the Association for Computing Machinery* **40**(4), 917–962.

Halpern, J. Y.: 1991, 'The relationship between knowledge, belief, and certainty'. *Annals of Mathematics and Artificial Intelligence* **4**, 301–322.

Halpern, J. Y.: 2001, 'Lexicographic probability, conditional probability, and nonstandard probability'. In: *Proceedings of the Eighth Conference on Theoretical Aspects of Rationality and Knowledge.* To appear.

Harel, D., D. Kozen, and J. Tiuryn: 2000, *Dynamic Logic*, Foundation of Computing. Cambrigde, Massachussetts: MIT Press.

Hintikka, J.: 1962, *Knowledge and Belief, An Introduction to the Logic of the Two Notions.* Ithaca & London: Cornell University Press.

Hirshleifer, J. and J. G. Riley: 1992, *The analytics of uncertainty and information*, Cambridge surveys of economic literature. Cambridge: Cambridge University Press.

Jeffrey, R. C.: 1983, *The logic of decision.* Chicago London: University of Chicago Press.

Kolmogorov, A. N.: 1956, *Foundations of the Theory of Probability.* New York: Chelsea Publishing Company, second edition. Translation edited by Nathan Morrison with an added bibliography by A.T. Bharucha-Reid.

Kripke, S.: 1980, *Naming and Necessity.* Oxford: Blackwell Publishers Ltd, revised and enlarged edition.

Kullback, S. and R. A. Leibler: 1951, 'On information and sufficiency'. *Annals of Mathematical Statistics* **22**, 79–86.

Kyburg, H. E.: 1994, 'Uncertainty logics'. In: D. M. Gabbay, C. J. Hogger, and J. A. Robinson (eds.): *Handbook of Logic in Artificial Intelligence and Logic Programming*, Vol. 3, Nonmonotonic reasoning and uncertain reasoning. Oxford: Clarendon Press.

Larsen, K. and A. Skou: 1991, 'Bisimulation through probabilistic testing'. *Information and Computation* **94**(1), 1–28.

Meyer, J.-J. and W. van der Hoek: 1995, *Epistemic Logic for AI and Computer Science.* Cambridge: Cambridge University Press.

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.

Prohorov, Y. V. and Y. A. Rozanov: 1969, *Probability Theory, basic concepts, limit theorems, random processes*, Vol. 157 of *Die Grundlehren der mathematischen Wissenschaften in Einzeldarstellungen mit besonderer Berücksichtigung der Anwendungsgebiete.* Berlin, Heidelberg, New York: Springer-Verlag. Translation of *Teorija Verojatnostej* by K. Krickeberg and H. Urmitzer.

Selvin, S.: 1975, 'A Problem in Probability'. *The American Statistician* **29**(1), 67.

ten Cate, B. D.: 2002, 'Internalizing Epistemic Actions'. In: M. Martinez (ed.): *Proceedings of the NASSLLI 2002 student session.* Stanford University.

van Benthem, J. F. A. K.: 2002a, 'Conditional probability is update logic'. unpublished.

van Benthem, J. F. A. K.: 2002b, "One is a lonely number': on the logic of communication'. To be published in the proceedings of the *Logic Colloquium 2002.*

van Ditmarsch, H. P.: 2000, *Knowledge Games.* Amsterdam: ILLC Dissertation Series DS-2000-06. Available online at `http://www.illc.uva.nl/Publications/`.

van Rooy, R.: 2003, 'Quality and quantity of information exchange'. *Journal of Logic Language and Information.* To appear.

Veltman, F.: 1996, 'Defaults in Update Semantics'. *Journal of Philosophical Logic* **25**, 221–261.

Vos Savant, M.: 1990, 'Ask Marilyn'. *Parade Magazine* p. 15.