Iterated Belief Change in MAS

Jan-Willem Roorda
Computer Science Chalmers
University of Technology
Sweden
jwr@cs.chalmers.se

Wiebe van der Hoek
Computer Science
Utrecht University
University of Liverpool
The Netherlands & U.K.
wiebe@cs.uu.nl

John-Jules Meyer
Computer Science
Utrecht University
The Netherlands
jj@cs.uu.nl

ABSTRACT

We give a model for iterated belief change in multi-agent systems. The formal tool we use for this is a combination of modal and dynamic logic. Two core notions in our model are the *expansion of the knowledge and beliefs of an agent*, and *the processing of new information.* An expansion has been defined as the change in the knowledge and beliefs of an agent when it decides to believe an incoming formula while holding on to its current propositional beliefs. To prevent our agents from forming inconsistent beliefs they do not expand with every piece of information they receive. Instead, our agents remember their original beliefs and every piece of information they receive. After every receipt of information they decide which (consistent) subset of the received information should be incorporated into their original beliefs. This procedure is called the processing of new information. We show that our model of belief update behaves in an intuitive way and that it is not sensitive to criticism on comparable models.

Categories and Subject Descriptors

I.2.4 [**Artificial Intelligence**]: Knowledge Representation Formalisms and Models—*modal logic*;

General Terms

[**Theory, Design**]:

Keywords

iterated belief revision, knowledge, formalisms

1. INTRODUCTION

We describe a model for the dynamic doxastic and epistemic behaviour of rational agents. Our main goal is to give a model for a multi-agent system which describes how the knowledge and beliefs of agents should change when they iteratively receive new information. The formal tool that we use to describe the behaviour of agents is a combination of modal and dynamic logic (cf. [7]). To see how these logics can be used to describe multi-agent systems we refer to [8].

Our model of knowledge and belief is based on the notion of possible worlds. We say that the agent a believes the formula ϕ (notation $\mathsf{B}_a\phi$) if ϕ is true in all the worlds it considers doxastically possible. An agent a is said to know the formula ϕ (notation $\mathsf{K}_a\phi$) if ϕ is true in all the worlds it considers epistemically possible. The difference between the notions of knowing and believing is that we require that known formulas are true, while an agent can have false beliefs. The reason that we equip our agents with beliefs, is that they are needed for (common sense) reasoning. If our agents would only be equipped with knowledge, it would be impossible for them to reason with information from sources that are not hundred percent reliable. And because in real world applications it is not often the case that agents are confronted with absolutely reliable information, it would pose severe restrictions on their reasoning capabilities.

One of the reasons that we equip our agents with knowledge is that it can act as a certain lower bound for beliefs, which can be used to filter incoming information. For instance, when we know that in the world the agent is going to live in certain formulas are, and will remain, true (like the fact in RoboCup that there is at most one ball in the field), we can incorporate them in the knowledge of our agent. When the agent receives incoming information that contradicts these formulas (balls are sensed at two different places), it will not use them to update his beliefs, because it knows that the incoming information cannot be true.

1.1 Multi-agent belief expansions

A core notion in our model of belief update is *the expansion of the knowledge and beliefs of an agent by a propositional formula* ϕ. We define an *expansion* as the change in the knowledge and the beliefs of an agent when it decides to believe an incoming formula while holding on to its current propositional beliefs. To illustrate the idea of expansion we give an example; for simplicity reasons the example is only concerned with beliefs and not with knowledge. Suppose we have two agents, say a and b. Both agents have no beliefs about whether the proposition p is true. Furthermore, each agent believes that the other agent has no beliefs about whether p is true etc. This situation can be modelled by the Kripke-world (M, w) described in the picture below. In Figure 1 the doxastic accessibility relation of agent a is represented by the continuous arrows, the doxastic accessibility relation of agent b by the dotted arrows.

AAMAS'02, July 15-19, 2002, Bologna, Italy.

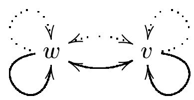

Figure 1: $w \models p, v \models \neg p$

Suppose agent a receives (without agent b being aware of it) the message that p is true. Furthermore, suppose that the agent considers the source to be trustworthy and therefore decides to expand his beliefs with the information that p is true. After the update we expect that agent a believes that p is true. Furthermore, we expect that it believes that it believes that p is true, etc. But we expect that the beliefs of agent b, which was not aware of the receipt of information by a, have not changed. The Kripke world (M', w'), described in Figure 2 fulfils these expectations. We say that (M', w') is the Kripke-world resulting from the Kripke-world (M, w) by expanding the beliefs of agent a with the formula p.

In this paper we describe a function $[\cdot,\cdot]$, such that for a Kripke world (M, s), a formula of propositional logic ϕ, and a label of an agent a, $(M, s)[\phi, a]$ yields the Kripke world equal to (M, s) expect for the beliefs of agent a that have been expanded with the formula ϕ. In recent literature several approaches to belief expansions in modal logic have been presented. Examples are the work of Gerbrandy and Groeneveld [6] who have added belief-update operators to a system based on non-well-founded sets, the work of Lomuscio and Ryan [10] using Kripke-trees, van Ditmarsch' semantical approach to ([3]) on *Knowledge Games* and the work of Baltag et al [1], who use standard Kripke semantics (but also a 'semantical language').

In essence, our belief expansion function is equal to Baltag's 'conscious update operator'. However, a difference is that our belief expansion function works on models of knowledge and belief, while Baltag's update operators are only concerned with knowledge. Another contribution of our approach here is a proof that our expansion function is the *unique* function that satisfies a set of desired characteristics of belief expansion.

1.2 Iterated Belief Change

One of the characteristics of our belief expansion function is that after an expansion of the beliefs of agent a with ϕ, agent a believes ϕ: $(M, s)[\phi, a] \models \mathsf{B}_a\phi$. Another fact is that an expansion does not decrease the propositional beliefs of an agent. So for every two propositional formula ϕ and ψ the following holds: $(M, s) \models \mathsf{B}_a\psi \Rightarrow (M, s)[a, \phi] \models \mathsf{B}_a\psi$.

From these two characteristics follows that repeated expansion can lead to inconsistent beliefs, for instance an expansion of the beliefs of agent a with ϕ followed by an expansion with $\neg\phi$ yields inconsistent beliefs for agent a. Thus: $((M, s)[a, \phi])[a, \neg\phi] \models \mathsf{B}_a\bot$.

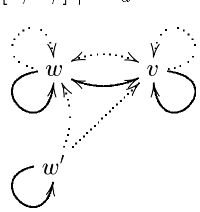

Figure 2: $w, w' \models p, v \models \neg p$

In real world situations agents receive information from a wide variety of sources. The information that these sources provide need not be true. So it can be the case that an agent receives contradicting information. If an agent should expand his beliefs by every piece of information it receives, it could end up with inconsistent beliefs. Of course this is unwanted. The question how to deal with information that is inconsistent with current beliefs is known in literature as the problem of (iterated) *belief revision*.

The best known formal approach towards belief revision is the so called AGM framework ([5]). In AGM, beliefs are represented by sets of propositional sentences. The framework gives a number of postulates that describe how the beliefs of an agent should change upon receiving new information (that may be inconsistent with its current beliefs). One of these postulates states that a revision of a belief set K with a formula ϕ should lead to believing the formula ϕ (even if $\neg\phi$ is part of the original belief set K). Thus: $\phi \in K \star \phi$.

So, the AGM framework gives priority to incoming information: the last piece of information should always be believed, even if it contradicts the current beliefs of the agent. The principle of giving priority to incoming information has been criticised by several people. Halpern and Friedman [4], for instance, argue that the principle leads to counter intuitive results if the incoming information is received from sources that are not completely reliable. In that case the beliefs of an agent which receives contradicting information, are completely dependent on the order in which the information is received. Halpern and Friedman argue that this is wrong, they state that instead of only relying on the order of the incoming information, incoming information should compete for acceptance against what is already believed. In order to describe a model in which this it the case, they state that it is needed to 'describe how to combine a given strength of belief in the observation with the strengths of beliefs in the original epistemic state'.

Revisions in AGM are so-called *contractions* followed by expansions: first beliefs are removed to enable new beliefs to be accommodated. Contractions thus play a predominant role in AGM belief revision. In our approach to belief revision we have chosen to abandon the principle of giving priority to incoming information. Moreover, in our setting contractions do not play a direct role. Instead of giving priority to the last piece of information an agent receives we let the agent decide which pieces of received information it will use to expand its beliefs. This decision can for instance be based on the trust the agent has in the sources of the received information. An essential aspect of this decision procedure is that if an agent chooses to use a certain piece of information, it is able to review this choice later on.

For instance, suppose an agent starts its life with believing p but having no beliefs about q. Now, if an agent receives a piece of information say $\neg q$, it can decide to expand his beliefs with $\neg q$, resulting in believing p and $\neg q$. But, if later on it receives another piece of information say $p \rightarrow q$ from a more trustworthy source, it can decide that the expansion with $\neg q$ was not wanted after all, and to expand its original beliefs (the beliefs he had before expanding with $\neg q$) with $p \rightarrow q$, resulting in believing p and q.

So, in our approach agents do not always perform a belief expansion when they receive a piece of new information. In general a newly received piece of information competes with previously received pieces of information to be expanded in the beliefs of the agent. Furthermore, an expansion with a piece of information can be undone later on when an agent receives information from a more trustworthy source contradicting the expanded information.

The contribution of our research can be split up into several parts. The first part is a model of belief expansion in multi agent epistemic logic. Although our belief expansion function does not differ much in spirit from Baltag's 'conscious update operator', we give a solid defence of the correctness of our approach, which is missing in Baltag's article. Furthermore, Baltag's approach is only concerned with knowledge, whereas our approach describes the change in both knowledge and beliefs. The second, more novel part, of our contribution is our model of iterated belief change. Our model replaces the paradigm of 'giving priority to incoming' information by the paradigm that gives priority to information received from more trustworthy sources. Such an approach has not been formally described before. We will also not employ direct contraction operators like in the AGM setting. Instead of contracting beliefs we use different expansions of the information obtained so far (which we record in histories). We might thus characterize our approach succinctly by the slogan "backtracking expansions rather than performing contractions". Finally, we show how an instantiation of our framework in which agents give priority to more recent information meets the intuition behind the AGM postulates for belief revision and the postulates for iterated belief change of Darwich and Pearl ([2]).

The remainder of this paper is organised as follows. First, in Section 2 we describe our formal model for representing knowledge an beliefs. In Section 3 we give our belief update function which describes how the knowledge and beliefs of an agent should change upon expanding. We give a number of desired characteristics of such a function and prove that our function is the unique function satisfying these criteria. In Section 4 we give our model for iterated belief change in multi modal epistemic logic. In Section 5 we compare our framework with the AGM approach to belief change. We end with conclusions and suggestions for further research.

2. KNOWLEDGE AND BELIEF

Given a set of agents $\mathcal{A}$ and a set of propositional symbols $\mathcal{P}$, our language $\mathcal{L}$ is the smallest set containing $\mathcal{P}$ such that:
(1) $\phi, \psi \in \mathcal{L} \Rightarrow \neg\phi, \phi \wedge \psi \in \mathcal{L}$
(2) $i \in \mathcal{A}, \phi \in \mathcal{L} \Rightarrow \mathsf{K}_i\phi, \mathsf{B}_i\phi \in \mathcal{L}$
Furthermore, $\mathcal{L}_0$ is the propositional subset of $\mathcal{L}$.

The class $\mathcal{M}$ of Kripke models contains tuples $M = (S, \pi, \mathrm{R}, \mathrm{B})$, where S is a set of possible states, $\pi : S \to \mathcal{P} \to \{0, 1\}$ is a function that assigns truth values to the propositional symbols in the possible states, $\mathrm{R} : \mathcal{A} \to S \times S$ is a function that assigns to each agent its epistemic accessibility relation, and, finally, $\mathrm{B} : \mathcal{A} \to S \times S$ is a function that assigns to each agent its doxastic accessibility relation.

To achieve the right properties of knowledge and belief we restrict the class of models $\mathcal{M}$ to $\mathcal{M}^*$. The class $\mathcal{M}^*$ of knowledge and belief models is the greatest subset of the class of Kripke models $\mathcal{M}$ such that for every $M = (S, \pi, \mathrm{R}, \mathrm{B}) \in \mathcal{M}^*$ holds that (1) $\mathrm{R} : \mathcal{A} \to S \times S$ is an equivalence relation, (2) $\mathrm{B} : \mathcal{A} \to S \times S$ is a serial, transitive and Euclidean relation, (3) for every $i \in \mathcal{A}$ holds $\mathrm{B}_i \subseteq \mathrm{R}_i$ and (4) $\forall i \in \mathcal{A}$, $\forall u, s, t \in S$: $\mathrm{R}_i(s, t) \Rightarrow (\mathrm{B}_i(s, u) \Leftrightarrow \mathrm{B}_i(t, u))$.

We use the same system for our models of knowledge and belief as [9]. The first two requirements above state that we use the standard S5-model for knowledge and the standard KD45-model for belief. This ensures that known facts are true, (an agent can however believe formulas that are false, but it does not believe contradictory information), and that our agents have positive and negative introspection for both knowledge and belief. The third requirement ensures that everything that is known also is believed. Finally, the fourth requirement ensures that agents know about their own beliefs.

Definition 1. *(Truth in Kripke models)*
Let $M = (S, \pi, \mathrm{R}, \mathrm{B})$ be a Kripke model and $s \in S$. Then:

$$\begin{array}{lcl}
(M, s) \models p \in \mathcal{P} & \Leftrightarrow & \pi(s)(p) = 1 \\
(M, s) \models \phi \wedge \psi & \Leftrightarrow & (M, s) \models \phi \text{ and } (M, s) \models \psi \\
(M, s) \models \neg\phi & \Leftrightarrow & (M, s) \not\models \neg\phi \\
(M, s) \models \mathsf{K}_i\phi & \Leftrightarrow & \forall t \in S[\mathrm{R}_i(s, t) \Rightarrow (M, t) \models \phi)] \\
(M, s) \models \mathsf{B}_i\phi & \Leftrightarrow & \forall t \in S[\mathrm{B}_i(s, t) \Rightarrow (M, t) \models \phi)]
\end{array}$$

Two Kripke worlds (M, s) and (M', s'), are called elementary equivalent, if, and only if, for every $\phi \in \mathcal{L}$ one has $(M, s) \models \phi \Leftrightarrow (M', s') \models \phi$.

3. THE BELIEF EXPANSION FUNCTION

Let (M, s) be a Kripke world, ϕ a formula of propositional logic, and a a label of an agent. By $(M, s)[\phi, a]$ we denote the Kripke world in which the beliefs of agent a are updated with the formula ϕ. We can think of $[\cdot, \cdot]$ as a function of type: $\mathcal{M} \times \mathcal{S} \times \mathcal{L}_0 \times \mathcal{A} \to \mathcal{M} \times \mathcal{S}$.

We now describe what characteristics a belief expansion function $[\cdot, \cdot]$ should have. Suppose that we update the knowledge and beliefs of agent a in the Kripke world (M, s) with the propositional formula ϕ. First of all, the function should change *knowledge* and *beliefs*, not reality, so the truth of formulas consisting of propositional variables should not be changed. Furthermore, the function should only change the knowledge and beliefs of agent a, and not those of other agents. So the truth of formulas $\mathsf{B}_b\alpha$ and $\mathsf{K}_b\alpha$ should not be changed for $a \neq b$.

The set of worlds which agent a considers doxastically possible should be changed. Agent a should consider worlds doxastically possible only if ϕ is true in them. Furthermore, the beliefs of agent a should be updated in the doxastically possible worlds too. So, it should be the case that in the updated Kripke world the formula $\mathsf{B}_a\alpha$ is true if, and only if, for all states t holds, [if (t is a doxastically possible state in the original Kripke world (M, s) and (M, t) verifies ϕ) then ((M, t) expanded with ϕ verifies α)].

Updating the beliefs of agent a should not change the set of worlds it considers epistemically possible. But of course the agent should know that its beliefs are updated in the worlds it considers epistemically possible. So the updated model should satisfy the formula $\mathsf{K}_a\alpha$ if, and only if, for all states t that agent a considers epistemically possible, (M, t) expanded with ϕ models α. Finally, the function should behave naturally towards $\wedge$ and $\neg$. We will call a function which satisfies these criteria *proper*. All this is summarised below.

Definition 2. *A belief expansion function $[\cdot, \cdot] : \mathcal{M} \times \mathcal{S} \times \mathcal{L}_0 \times \mathcal{A} \to \mathcal{M} \times \mathcal{S}$ is called* proper *if*

$$
\begin{array}{lll}
(M,s)[\phi,a] \models p & \Leftrightarrow & (M,s) \models p \\
(M,s)[\phi,a] \models \alpha \wedge \beta & \Leftrightarrow & (M,s)[\phi,a] \models \alpha \& \\
 & & (M,s)[\phi,a] \models \beta \\
(M,s)[\phi,a] \models \neg\alpha & \Leftrightarrow & (M,s)[\phi,a] \not\models \alpha \\
(M,s)[\phi,a] \models \mathsf{B}_b\alpha & \Leftrightarrow & (M,s) \models \mathsf{B}_b\alpha \ \ (a \neq b) \\
(M,s)[\phi,a] \models \mathsf{B}_a\alpha & \Leftrightarrow & \forall t((\mathrm{B}_a(s,t) \ \& \ (M,t) \models \phi) \Rightarrow \\
 & & (M,t)[\phi,a] \models \alpha) \\
(M,s)[\phi,a] \models \mathsf{K}_b\alpha & \Leftrightarrow & (M,s) \models \mathsf{K}_b\alpha \ \ (a \neq b) \\
(M,s)[\phi,a] \models \mathsf{K}_a\alpha & \Leftrightarrow & \forall t(\mathrm{R}_a(s,t) \Rightarrow (M,t)[\phi,a] \models \alpha)
\end{array}
$$

For all proper belief expansion functions $[\cdot,\cdot]$, agents a, b with $a \neq b$, propositional formulas ϕ, ψ and formulas χ holds:

$$
\begin{array}{lll}
(M,s)[\phi,a] \models \mathsf{B}_a\phi & & \\
(M,s)[\phi,a] \models \mathsf{B}_a\mathsf{B}_a\phi \wedge \mathsf{K}_a\mathsf{B}_a\phi & & \\
(M,s)[\phi,a] \models \mathsf{K}_a\psi & \Leftrightarrow & (M,s) \models \mathsf{K}_a\psi \\
(M,s)[\phi,a] \models \mathsf{B}_b\chi & \Leftrightarrow & (M,s) \models \mathsf{B}_b\chi
\end{array}
$$

The first clause states that after the expansion of the beliefs of agent a with ϕ it believes that ϕ. The second clause expresses that after the expansion the agent knows and believes that it believes ϕ. The third clause states that the knowledge of propositional formulas is not changed by expansions. Finally, the fourth close states that the beliefs of agent b are not changed by an expansion of the beliefs of agent a. A nice property of our definition of proper belief expansion functions is that it defines a unique function.

Theorem 1. *There is, up to elementary equivalence, at most one proper belief expansion function.*

We define *the* proper belief expansion function below. This definition is based on the ideas in [1]. We extended their definition for the concept of knowledge. The intuition behind this definition is, when expanding with ϕ, we make two copies of the original Kripke world. The one containing the states $\mathrm{old}(S)$ is an exact copy of the original Kripke world. Because, for all agents not equal to a, all from s' epistemically and doxastically accessible worlds, lie in this 'old' copy, they see no difference between the original and the expanded Kripke world. Agent a, on the other hand, is able to see the 'new' copy, where worlds not satisfying ϕ are no longer doxastically reachable.

Definition 3. *(Proper belief expansion)*
$[\cdot,\cdot]_p$ is defined as follows: Given a world $((S,\pi,\mathrm{R},\mathrm{B}),s)$, a label of an agent a, and a propositional formula ϕ:
$(M,s)[\phi,a]_p = ((S',\pi',\mathrm{R}',\mathrm{B}'),s')$, where

- $S' = \mathrm{old}(S) \bigcup \mathrm{new}(S)$,
- $s' = \mathrm{new}(s)$,
- $\pi'(old(u))(p) = \pi'(new(u))(p) = \pi(u,p)$
- *and the B'_i and R'_i are minimal such that:*

$$
\begin{array}{lll}
\mathrm{B}'_i(\mathrm{old}(u),\mathrm{old}(v)) & \Leftrightarrow \mathrm{B}_i(u,v) & i \in \mathcal{A} \\
\mathrm{B}'_a(\mathrm{new}(u),\mathrm{new}(v)) & \Leftrightarrow \mathrm{B}_a(u,v) \ \& \ (M,v) \models \phi & \\
\mathrm{B}'_b(\mathrm{new}(u),\mathrm{old}(v)) & \Leftrightarrow \mathrm{B}_b(u,v) & (a \neq b)
\end{array}
$$

$$
\begin{array}{lll}
\mathrm{R}'_i(\mathrm{old}(u),\mathrm{old}(v)) & \Leftrightarrow \mathrm{R}_i(u,v) & i \in \mathcal{A} \\
\mathrm{R}'_a(\mathrm{new}(u),\mathrm{new}(v)) & \Leftrightarrow \mathrm{R}_a(u,v) & \\
\mathrm{R}'_b(\mathrm{new}(u),\mathrm{old}(v)) & \Leftrightarrow \mathrm{R}_b(u,v) & (a \neq b)
\end{array}
$$

Example 1. *We show for a specific Kripke world the result of an expansion via $[\cdot,\cdot]_p$. To make things easier we treat an example with only doxastic relations. Let (M,w) be the Kripke world described in Figure 1. The expansion of the beliefs of agent a with p in state w yields the Kripke world (M',w') of Figure 3. Note that there is no $\cup_i\mathrm{B}_i$-path from w' to v'. So we can delete the state v' without changing the meaning of the Kripke world (M',w'), yielding exactly the situation denoted in Figure 2.*

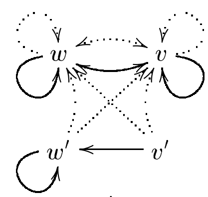

Figure 3: $w, w' \models p; v, v' \models \neg p$

Theorem 2. *$[\cdot,\cdot]_p$ is proper and well defined on the set of Kripke worlds modulo elementary equivalence*

4. PROCESSING NEW INFORMATION

As explained in the previous sections iteratively expanding with every piece of information received by an agent can lead to inconsistent beliefs. To prevent forming inconsistent beliefs our agents remember their original beliefs (the beliefs they had before receiving any information) and every piece of information they receive. After every receipt of information they decide which (consistent) subset of the received information should be expanded into their original beliefs. The decision which subset should be used is based on the trust the agents have in the sources of the received information. We call this procedure the process new information (pni) action. To model this action we shift from a static modal logic to a dynamic modal logic. We describe its syntax and semantics in this section.

Definition 4. *(Language)*
Given finite sets $\mathcal{A} = \{1,\ldots,n\}$ of agents and $\mathcal{O} = \{I, II, III,\ldots,m\}$ of sources, and a set $\mathcal{P}$ of propositional symbols we define the language $\mathcal{L}_d$ and the class of actions Ac by mutual induction as follows.
$\mathcal{L}_d$ is the smallest set containing $\mathcal{P}$ such that

- $\phi, \psi \in \mathcal{L}_d \Rightarrow \neg\phi, \phi \wedge \psi \in \mathcal{L}_d$
- $i \in \mathcal{A}, \alpha \in \mathsf{Ac}, \phi \in \mathcal{L}_d \Rightarrow \mathsf{K}_i\phi, \mathsf{B}_i\phi, [i,\alpha]\phi \in \mathcal{L}_d$

and Ac is the smallest set such that
$\phi \in \mathcal{L}_0, t \in \mathcal{O} \Rightarrow \mathsf{pni}(\phi,t) \in \mathsf{Ac}$
$\alpha_1, \alpha_2 \in \mathsf{Ac} \Rightarrow (\alpha_1;\alpha_2) \in \mathsf{Ac}$

Definition 5. *(Extended Kripke models)*
The class $\mathcal{E}$ of Extended Kripke models contains tuples $E = (M_c, (M_o, s_o), H, Sf)$, where

- *M_c and M_o are Kripke models as defined before,*
- *M_c is called the current Kripke-model, it models the current knowledge and beliefs of the agents.*

- $M_o = (S_o, \pi_o, \mathrm{R}_o, \mathrm{B}_o)$ *is called the original Kripke model, it models the knowledge and beliefs of the agents when they start their lives, we demand that* $\mathrm{R}_{o(i)} = \mathrm{B}_{o(i)}$ *for all* i*;* $s_o \in S_o$*,*

- H *is the received information function, of type:* $\mathcal{A} \to [\mathcal{L}_0 \times \mathcal{O}]$*. The received information function gives for every agent the list of the formulas it received in the past, together with their sources,*

- Sf *is the information selection function of type* $\mathcal{A} \to \mathcal{M} \times \mathcal{S} \to [\mathcal{L}_0 \times \mathcal{O}] \to [\mathcal{L}_o \times \mathcal{O}]$*, it selects for an agent which information it uses to form its knowledge and beliefs. Of course we demand* $Sf(i)(M, s)(U) \subseteq U$*.*

By demanding that $\mathrm{R}_{o(i)} = \mathrm{B}_{o(i)}$ for all i we ensure that the original beliefs of agents do not compromise proper beliefs. By proper believing a formula we mean that the agent believes, but does not know, the formula. The reason that we use this assumption, is that we consider precisely those formulas that an agents knows as unquestionable. Every other propositional formula that an agent believes should have some kind of justification, in our model this justification consists of the fact that the formula is believed because of an expansion with a certain piece of incoming information. This way, we ensure that our agents are always able to make the judgement whether a proper believed formula should be given up in favour of an incoming formula.

Given an extended Kripke model $E = (M_c,\ (M_o, s_o),\ H,\ Sf)$ where $M_c = (S, \pi, \mathrm{R}, \mathrm{B})$, and $s \in S$, we define:

$$\begin{array}{lcl}
(E,s) \models p \in \mathcal{P} & \Leftrightarrow & \pi(s,p) = 1 \\
(E,s) \models \phi \wedge \psi & \Leftrightarrow & (E,s) \models \phi \text{ and } (E,s) \models \psi \\
(E,s) \models \neg\phi & \Leftrightarrow & \text{not } (E,s) \models \phi \\
(E,s) \models \mathsf{K}_i\phi & \Leftrightarrow & \forall t \in S[(s,t) \in \mathrm{R}_i \Rightarrow (E,t) \models \phi)] \\
(E,s) \models \mathsf{B}_i\phi & \Leftrightarrow & \forall t \in S[(s,t) \in \mathrm{B}_i \Rightarrow (E,t) \models \phi)] \\
(E,s) \models [i,\alpha]\phi & \Leftrightarrow & r(i,\alpha)(E,s) \models \phi
\end{array}$$

where r is defined as follows:

$$\mathrm{r}(i, \alpha_1; \alpha_2)(E, s) \quad = \quad \mathrm{r}(i, \alpha_2)(\mathrm{r}(i, \alpha_1)(E, s))$$

The definition of the semantics above is incomplete: there is no definition for the function r applied to the process new information action. We will shortly complete the definition of the semantics by defining the effect of this action. First we give a few technical definitions.

Given a Kripke world, (M, s), a received information function H and a selection function Sf, define:

$$\mathrm{Filter}(M,s)(Sf)(H) = \\ [(\phi, i) | \exists t \in \mathcal{O}(\phi, t) \in Sf(i)(M, s)(H(i)), i \in \mathcal{A}]$$

Filter applied to a Kripke world, a selection function and a history function yields a list of (formula,agent) tuples. The meaning of this list is the following: if (ϕ, i) is an element of this list then agent i has decided to expand its beliefs by ϕ.

The use of Filter in combination with a selection function Sf is responsible for the 'contractive' effect of our model of belief update. For example, if agent a receives the formula p from source I, the selection function may tell the agent to use p to update its beliefs. In that case the received information function H and the selection function Sf are such that $H(a) = [(p, I)]$, and $\mathrm{Filter}(M_o, s_o)(Sf)(H) = [(p, a)]$. Thus, after the receival of p the beliefs of agent a are modelled by $(M_o, s_o)[p, a]$. If later on the agent receives $\neg p$ from the more reliable source II the selection function may demand that the agent backtracks to its original beliefs and uses only $\neg p$ to form its beliefs. In that case the following holds: $H(a) = [(\neg p, II), (p, I)]$, and $\mathrm{Filter}(M_o, s_o)(Sf)(H) = [(\neg p, a)]$. Therefore, after the receival of p and $\neg p$ the beliefs of agent a are modelled by $(M_o, s_o)[\neg p, a]$.

Thus, the output of the selection function will be used as input for the belief expansion function.

Given a Kripke world (M, s) and a list of (formula,agent)-tuples U, define[1]:

$$\begin{array}{lcl}
(M,s)[\,[\,]\,] & = & (M,s), \\
(M,s)[U] & = & ((M,s)[\phi,i])[U'], \text{ if } U = (\phi,i) : U'
\end{array}$$

Definition 6. *(Processing new information) Given an extended Kripke model* $E = (M_c, (M_o, s_o),\ H,\ Sf)$*, a state* $s \in S$*, a propositional formula* $\phi \in \mathcal{L}_0$*, an agent* $i \in \mathcal{A}$ *and a source* $t \in \mathcal{O}$*. We define:* $r(i, \mathsf{pni}(\phi, t))(E, s) = ((M_c', (M_o, s_o), H', Sf), s')$*, where*

- $H'(i) = (\phi, t) : H(i)$ *and* $H'(j) = H(j)$ *for* $j \neq i$*.*
- $(M_c', s') = (M_o, s_o)[\,\mathrm{Filter}(M_o, s_o)(Sf)(H')\,]$

The intuition behind this definition is the following. After agent i has performed the $\mathsf{pni}(\phi, t)$ action, the information (ϕ, t) will be added to its list of received information $H(i)$; this yields $H'(i)$. By means of the selection function Sf, it will decide which elements of the list of received information $H'(i)$ it will use to form his knowledge and beliefs; this yields $Sf(i)(M_o, s_o)(H'(i))$. The model which reflects the new beliefs of agent i is the original model (which reflects its original knowledge and beliefs) expanded with the information in $Sf(i)(M_o, s_o)(H'(i))$. Because the new model also must reflect the knowledge and beliefs of the other agents (which have not changed) we also expand with the elements from $Sf(j)(M_o, s_o)(H(j))$ for $j \neq i$.

For this intuition to make sense, it must be the case that the knowledge and beliefs of the agents are the result of expansions of their original knowledge and beliefs with the selected information from their received information list. This leads to the following definition.

Definition 7. *(Proper extended Kripke worlds)*
An extended Kripke world $((M_c, (M_o, s_o), H, Sf), s)$ *is called proper iff* $(M_c, s) = (M_o, s_o)[\,\mathrm{Filter}(M_o, s_o)(Sf)(H)\,]$*.*

Below we give some constraints on selection functions. First we give two concrete examples of selection functions: the naive and the trivial selection function. The trivial selection function forbids agents to use any incoming information. Given a Kripke world (M_o, s_o), a label of agent i, a list of formula-source tuples U, we define the trivial selection function, $Sf_{trivial}$, by: $Sf_{trivial}(i)(M_o, s_o)(U) = [\,]$.

The naive selection lets agents expand their beliefs with every piece of information they receive, even if the pieces of information contradict each other or contradict the original beliefs. Given a Kripke world (M_o, s_o), a label of agent i, a list of formula-source tuples U, we define the naive selection function, denoted Sf_{naive}, by: $Sf_{naive}(i)(M_o, s_o)(U) = U$.

In models with a naive selection function agents can form inconsistent beliefs. For example, in a model E with the

[1] $[\,]$ denotes the empty list, : constructs a new list from an element and a list.

naive selection function we have:
$(E,s) \models [i, \mathsf{pni}(\phi, t_1); \mathsf{pni}(\neg\phi, t_2)]\mathsf{B}_i\bot$.
Of course this is unwanted. A natural constraint on selection functions is that agents cannot form inconsistent beliefs by expansion. This leads to the definition of proper selection functions.

DEFINITION 8. *(Proper selection functions)* *A selection function Sf is called proper if for every proper extended Kripke world $(E,s) = (M_c, (M_o, s_o), H, Sf)$ and agent i the following holds:*
$(M_o, s_o) \not\models \mathsf{B}_i\bot$ & $Sf(i)(M_o, s_o)(U) = U'$
$\Rightarrow (M_o, s_o)[U', i] \not\models \mathsf{B}_i\bot$
where $(M,s)[U,i]$ is defined by $(M,s)[\,[\,]\,,i] = (M,s)$ and $(M,s)[(\phi,t) : U, i] = ((M,s)[\phi, i])[U,i]$.

In an extended Kripke model E containing a proper selection function we have:
$(E,s) \models \neg\mathsf{B}_i\bot \rightarrow [i, \mathsf{pni}(\phi, t_1); \mathsf{pni}(\neg\phi, t_2)]\neg\mathsf{B}_i\bot$
Demanding that a selection function is proper is needed, but not enough to let our agents work in an intuitive way. To illustrate this, the trivial selection function is proper, but by using this selection function our agents will never update their beliefs. This leads to the definition of eager selection functions. We call selection functions which let the agents use every piece of information as long as that does not lead to a contradiction, eager functions.

A proper selection function Sf is called eager, if for every Kripke world (M_o, s_o) and agent i the following holds:
$(M_o, s_o)[(H)(i), i] \not\models \mathsf{B}_i\bot \;\; \Rightarrow Sf(i)(M_o, s_o)(H(i)) = H(i)$
For a proper extended Kripke model E, that contains an eager selection function, $\phi \in \mathcal{L}_0$, $t \in \mathcal{O}$ the following holds:

- $(E,s) \models \neg\mathsf{B}_i\neg\phi \rightarrow [i, \mathsf{pni}(\phi, t)]\mathsf{B}_i\phi$
- $(E,s) \models \mathsf{B}_i\phi \wedge \neg\mathsf{B}_i\bot \rightarrow [i, \mathsf{pni}(\neg\phi, t)]\neg\mathsf{B}_i\bot$

The first clause states that if agent a does not believe $\neg\phi$, before receiving ϕ, it will believe ϕ afterwards and the second clause states that if agent a does believe ϕ, receiving $\neg\phi$ will not lead to inconsistent beliefs.

The definition of eager selection functions does not say anything about which information should be selected when an agent receives contradicting information. So demanding that a selection function is eager still does not ensure that the selection function works in an intuitive way. A natural constraint on selection functions is that they let agents update their beliefs with a 'maximal consistent' subset of the received information. This is expressed in the definition of maximal selection functions.

A proper selection function Sf is called maximal, if for every Kripke world (M_o, s_o) and agent i the following holds[2]:
$Sf(i)(M_o, s_o)(H(i)) = U \Rightarrow$
for every $h \in H(i) \backslash U$ holds $(M_o, s_o)[U +\!\!+ [h], i] \models \mathsf{B}_i\bot$

We say that U is a maximal consistent (w.r.t. (M_0, s_0)) sublist of $H(i)$. In general there are many such sublists of the received information list. A way of further restricting the allowed sublists is to use the notion of 'trust'. Given an ordering $<_{trust}$ on the set of sources $\mathcal{O}$ such that $t_2 <_{trust} t_1$ means that the agent considers source t_1 more trustworthy than source t_2, we can define the class of trust-based selection functions. The intuition behind the class of trust-based selection functions is that a selection function in this class prefers (maximal consistent) sublists that contain more information from more trustworthy sources.

Given a list U of (formula,source)-tuples, we define $U(t)$ as the sublist of U containing only the $(formula, source)$-tuples with the source equal to t.

A selection function Sf is called a trust-based maximal selection function (w.r.t. to a total ordering $<_{trust}$ on $\mathcal{O}$), if for every Kripke world (M_o, s_o) and agent i the following holds:
If $s_1 >_{trust} s_2 >_{trust} \ldots >_{trust} s_m$ then $Sf(i)(M_o, s_o)(H(i)) = \bigcup_{1 \le t \le m} U_t$, where each U_t is a maximal consistent sublist of $H(i)(t)$ w.r.t. $(M_o, s_o)[\bigcup_{1 \le k \le (t-1)} U_k, i]$

For a model E with a trust-based selection function based on an ordering with $I >_{trust} II$, we have:
$$(E,s) \models \neg\mathsf{B}_i\neg\phi \wedge \neg\mathsf{B}_i\phi \rightarrow [i, \mathsf{pni}(\phi, I); \mathsf{pni}(\neg\phi, II)]B_i\phi$$
In some applications agents do not consider the source of the information to be of importance, instead priority is given to more recent information. This intuition is reflected in the definition of the time based selection function.

DEFINITION 9. *The time based selection function Sf_{time} is defined by:*
$Sf_{time}(i)(M_o, s_o)([\,]) = [\,]$, *and*
$$Sf_{time}(i)(M_o, s_o)((\phi,t) : U) = \begin{cases} (\phi,t) : Sf_{time}(i)((M_o, s_o)[\phi, i])(U) & \textit{if } (M_o, s_o) \not\models \mathsf{B}_i\neg\phi \\ Sf_{time}(i)((M_o, s_o))(U) & \textit{otherwise} \end{cases}$$

For every proper extended Kripke model E where agent i has no original beliefs about p and q and containing the time-based selection function and $t \in \mathcal{O}$ the following holds:

- $(E,s) \models [i, \mathsf{pni}(\neg p, t); \mathsf{pni}(p, t)]\mathsf{B}_i p$
- $(E,s) \models [i, \mathsf{pni}(q, t); \mathsf{pni}(\neg p \wedge \neg q, t)]\mathsf{B}_i(\neg p \wedge \neg q)$
- $(E,s) \models [i, \mathsf{pni}(p, t); \mathsf{pni}(\neg p \wedge \neg q, t); \mathsf{pni}(q, t)]\mathsf{B}_i(p \wedge q)$

The first two clauses demonstrate that the time-based selection function gives priority to more recent information. In the third clause p is used by the agent because it does not contradict its original beliefs, $\neg p \wedge \neg q$ is not selected because it contradicts p, but q is selected because is does not contradict the expansion of the original beliefs with p. Therefore, p and q are expanded into the original beliefs of the agent, resulting in believing $p \wedge q$.

5. COMPARISON WITH AGM

The best known formal approach towards belief revision is the AGM framework proposed by Gärdenfors et al [5]. The framework gives a number of postulates that describe how the beliefs of an agent should change upon receiving new information (that may be inconsistent with its current beliefs). A characteristic of the AGM framework is that it gives priority to incoming information: the last piece of information should always be believed, even if it contradicts the current beliefs of the agent. In this section we compare our approach using the time-based selection function, which gives priority to most recent information, with the AGM framework.

In the AGM framework beliefs are represented by a (deductively closed) set of propositional formulas. Contrary to our framework no modalities or dynamics are allowed to occur in the beliefs of the agents; there are no beliefs about

[2] $+\!\!+$ denotes list concatination.

the agents' own belief or about the beliefs of others, and furthermore there are no beliefs about possible results of actions. The AGM framework describes three types of belief change:

Expansion: A new sentence ϕ is added to a belief set B together with the logical consequences of the addition. The belief set resulting from the expansion of the belief set B with the propositional formula ϕ is denoted $B+\phi$. Revision: A new formula ϕ, that may be inconsistent with the belief set B, is added. The result of a revision, denoted $B\star\phi$ should be consistent unless the incoming formula is a contradiction. Contraction: A formula ϕ is retracted from the belief set B, the new set $B-\phi$ should be deductively closed and should not contain ϕ.

Below we define a revision function in our framework as the process new information action using the time based selection function. We check to which extend the AGM postulates hold for the propositional beliefs of our agents using this revision function. Because the beliefs of our agents are much more complex than the beliefs in the AGM-framework, we first define a simple theory that describes the change in the propositional beliefs of our agents when processing new information.

The dynamics of the propositional beliefs of our agents can be modelled by the concept of belief triples.

DEFINITION 10. *(Belief triples)*
A belief triple consists of a set of original propositional beliefs, a list of received information and a selection function.

- *A belief triple B is a triple of a theory $B_0 \subseteq \mathcal{L}_0$, a list of (sentence,source) tuples $H : [\mathcal{L}_o \times \mathcal{O}]$, and a function $S : \mathbb{P}(\mathcal{L}_0) \rightarrow [\mathcal{L}_0 \times \mathcal{O}] \rightarrow [\mathcal{L}_0 \times \mathcal{O}]$. In accordance with the intuitions of the previous sections we call B_o the original beliefs of an agent, H its history of received information function, and S its selection function.*
- *We say that a belief tuple $B = (B_o, H, S)$ models the sentence α, notation $B \models \alpha$ iff $B_o \cup S(B_o)(H) \models \alpha$*
- *We define $|B| = \{\alpha \in \mathcal{L}_0 | B \models \alpha\}$.*

The time-based selection function S_{time} for belief triples is the equivalent of the time based selection Sf_{time} for extended Kripke worlds. In the definition of the time-based selection function we omit the sources of the information: we write ϕ instead of (ϕ, t).

DEFINITION 11. *(Time based selection function)*
The time based selection function $S_{time} : \mathbb{P}(\mathcal{L}_0) \rightarrow [\mathcal{L}_0] \rightarrow [\mathcal{L}_0]$ is defined by:
$S_{time}(B_o)([\,]) = [\,]$, *and* $S_{time}(B_o)(\phi : H)$
$$= \begin{cases} \phi : S_{time}(\mathrm{Cn}(B_o \cup \{\phi\}))(H) & \textit{if } B_0 \not\models \neg\phi \\ S_{time}(B_o)(H) & \textit{otherwise} \end{cases}$$

The revision- and expansion functions for belief triples are defined as follows:

- The revision of the belief triple $B = (B_o, H, S)$ with the (formula,source)-tuple (ϕ, t) is defined by: $B \star (\phi, t) = (B_o, (\phi, t) : H, S)$
- The expansion of the belief triple $B = (B_o, H, S)$ with the sentence ϕ is defined by: $B + \phi = (\mathrm{Cn}(B_o \cup \{\phi\}), H, S)$

The following proposition states that the dynamics of the propositional beliefs of our agents can be modelled by the concept of belief triples.

THEOREM 3.
For a proper extended Kripke world (E, s) with $E = (M_c, (M_o, s_o), H, Sf)$, an agent i, and a belief triple $B = (B_o, H(i), S)$ such that:
(1) $B_o = \{\alpha \in \mathcal{L}_0 | (M_o, s_o) \models \mathsf{B}_i\alpha\}$,
(2) *for all Kripke models (M', s'), theories B' and received information functions H' holds*
$B' = \{\alpha \in \mathcal{L}_0 | (M', s') \models \mathsf{B}_i\alpha\} \Rightarrow$
$S(B')(H'(i)) = Sf(i)(M', s')(H')$,

it holds that for every $\alpha, \phi \in \mathcal{L}_0$ and $t \in \mathcal{O}$:

- $B \models \alpha \Leftrightarrow (E, s) \models \mathsf{B}_i\alpha$
- $B \star (\phi, t) \models \alpha \Leftrightarrow (E, s) \models [i, \mathsf{pni}(\phi, t)]\mathsf{B}_i\alpha$

Below, we state an adaption of the AGM postulates for belief triples. Because we only check the validity of the postulates in belief triples with a time based selection we can omit the source of incoming information: we write $B \star \phi$ instead of $B \star (\phi, t)$.

Given a belief triple B, and propositional formulas ϕ and ψ, the AGM postulates for belief revision in terms of belief triples are

1. $|B \star \phi| = \mathrm{Cn}(|\mathsf{B} \star \phi|)$
2. $\phi \in |B \star \phi|$
3. $|B \star \phi| \subseteq |B + \phi|$
4. $\neg\phi \in |B| \Rightarrow |B \star \phi| \subseteq |B + \phi|$
5. $|B \star \phi| = \mathcal{L}_o \Leftrightarrow \vdash \neg\phi$
6. $\vdash \phi \leftrightarrow \psi \Rightarrow |B \star \phi| = |B \star \psi|$
7. $|B \star \phi \wedge \psi| \subseteq |B \star \phi + \psi|$
8. $\neg\psi \notin |B \star \phi| \Rightarrow |B \star \phi + \psi| \subseteq |B \star \phi \wedge \psi|$

THEOREM 4. *AGM postulates 1,3,4,6 and 7 hold for belief triples based on the time-based selection function, postulates 2,5 and 8 do not.*

One of the differences between the AGM and our approach is that we include the notion of original, (unquestionable, unchangeable) beliefs. Our agents ignore incoming information that contradicts the original beliefs. This way of handling information makes sense: in real world settings incoming information is very seldom hundred percent reliable, so when an incoming piece of information contradicts something that is know to be true, the wise thing is to just ignore the information.

Even if we set the original beliefs to be empty ($B_o = \emptyset$) this difference between the AGM- and our approach remains valid: after receiving the formula $\perp$ AGM agents believe everything while our agents simply ignore the formula. If we set $B_o = \emptyset$ and disregard the case of incoming formulas equivalent to $\perp$ all AGM postulates are valid in our approach. Darwich and Pearl propose four additional postulates, that in combination with the AGM postulates, should take care of intuitively acceptable iterated revision.

Given a belief triple B and propositional formulas α and μ, the postulates of Darwich and Pearl for belief revision are, in our terminology:

1. $\alpha \models \mu \Rightarrow |B \star \mu \star \alpha| = |B \star \alpha|$
   When two pieces of evidence arrive, the second being more specific than the first, the first is redundant.

2. $\alpha \models \neg\mu \Rightarrow |B \star \mu \star \alpha| = |B \star \alpha|$
   When two contradictory pieces of evidence arrive, the last one prevails.

3. $B \star \alpha \models \mu \Rightarrow |B \star \mu \star \alpha| \models \mu$
   An evidence μ should be retained after accommodating a more recent evidence α that implies μ given the current beliefs.

4. $B \star \alpha \not\models \neg\mu \Rightarrow |B \star \mu \star \alpha| \not\models \neg\mu$
   No evidence can contribute to its own demise.

THEOREM 5. *Under the assumption that α is consistent with the original beliefs in the first two postulates, the postulates of Darwich & Pearl are valid for belief triples using the time based selection function. Therefore, if we set $B_o = \emptyset$ and disregard the case of incoming formulas equivalent to $\perp$ all postulates of Darwich and Pearl are valid in our approach.*

6. CONCLUSION

In this paper we have given a model of belief update which describes how the knowledge and beliefs of agents change when they receive new information. The formal tool that we have used for this description is a combination of modal and dynamic logic. Two core notions in our ontology of belief update are the *expansion of the knowledge and beliefs of an agent with a propositional formula*, and *the processing of new information*. An expansion has been defined as the change in the knowledge and beliefs of an agent when it decides to believe an incoming formula while holding on to its current propositional beliefs.

To prevent our agents from forming inconsistent beliefs they do not expand their knowledge and beliefs with every piece of information they receive. Instead, they remember their original beliefs and every piece of information they receive. After every receipt of information they decide which (consistent) subset of the received information should be expanded into their original beliefs. The decision which subset should be used is determined by a the information selection function.

An example of a selection function is the time based selection function. If agents using the time based selection function receive contradicting pieces of information they give priority to the more recent information. We proved that if we disregard revisions with $\perp$, our framework with the time based selection function meets the AGM postulates for belief revision and the postulates of Darwich and Pearl for iterated belief revision. Revisions with $\perp$ are treated differently between the two approaches: the AGM-approach dictates that a revision with $\perp$ leads to inconsistent beliefs of agents, in our approach agents simply ignore incoming information equivalent to $\perp$.

The AGM framework gives priority to incoming information: the last piece of information should always be believed, even if it contradicts the current beliefs of the agent. In our approach we have dropped this principle, and instead incoming information is competing for acceptance against what is already believed. If an agent (proper) believes something, it is the result of a belief expansion with information from a certain source. If an agent receives incoming information that contradicts its current beliefs, it can decide by means of a selection function which subset of the received information should be used to form its beliefs. An intuitively attractive class of selection function is the class of trust-based selection functions. Agents using these selection functions give priority to information from more trustworthy sources.

One may conceive our approach using selection functions over lists of newly received information, as a layer on top of the dynamic-epistemic approaches of Baltag, van Ditmarsch, Gerbrandy, Lomuscio and others (cf. [1, 3, 6, 10]). Where their approaches give (slightly different) accounts of how to update with one newly learned piece of information, our framework for updating with new information and ways for backtracking those updates offers a way to deal with *streams* of new information, which one encounters when agents are involved in continuous sensing, or when they are encountered in a dialogue.

7. REFERENCES

[1] A. Baltag, L. Moss, and S. Solecki. The logic of public announcements, common knowledge and private suspicions. In *Proc. of TARK*, pages 43–56. Morgan Kaufmann Publishers, 1998.

[2] A. Darwiche and J. Pearl. On the logic of iterated belief revision. *Artificial Intelligence*, 89(1,2):1–29, 1997.

[3] H.P. van Ditmarsch. Descriptions of game actions, 2002. Accepted for the *Journal of Logic, Language and Information*.

[4] Nir Friedman and Joseph Y. Halpern. Belief revision: A critique. In Luigia Carlucci Aiello, Jon Doyle, and Stuart Shapiro, editors, *KR'96*, pages 421–431. Morgan Kaufmann, 1996.

[5] P. Gärdenfors. *Knowledge in Flux: Modeling the Dynamics of Epistemic States*. MIT Press, 1988.

[6] J. Gerbrandy and W Groeneveld. Reasoning about information change. *Journal of Logic, Language, and Information*, 6(2):147–169, 1997.

[7] R. Goldblatt. *Logics of Time and Computation*. University of Chicago Press, 1992.

[8] W. van der Hoek, B. van Linder, and J.-J.Ch. Meyer. An integrated modal approach to rational agents. In M. Wooldridge and A. Rao, editors, *Foundations of Rational Agency*, volume 14 of *Applied Logic Series*, pages 133–168. Kluwer, Dordrecht, 1998.

[9] B. van Linder, W. van der Hoek, and J.-J. Ch. Meyer. Actions that make you change your mind. In A. Laux and H. Wansing, editors, *Knowledge and Belief in Philosophy and AI*, pages 103–146. 1995.

[10] A. Lomuscio and M. Ryan. An algorithmic approach to knowledge evolution. In *Artificial Intelligence for Engineering Design, Analysis and Manufacturing (AIEDAM)*, 1999.