

Belief updates in Multiple Agent Systems

Jan-Willem Roorda

UU-CS-1999-25

Belief Updates in Multiple Agent Systems

Jan-Willem Roorda*

Abstract

We give a model for a multi-agent system which describes how the knowledge and beliefs of agents should change when they receive new information. The formal tool we use for this description is a combination of modal and dynamic logic. Two core notions in our model of belief update are the *expansion of the knowledge and beliefs of an agent*, and the *processing of new information by an agent*. 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 of that, 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 incorporated 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. 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 the critique on comparable models.

1 Introduction

In this paper 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 receive new information.

There are two reasons why a formal description of such behaviour of rational agents is useful. First of all, real world implementations of agent systems will most probably be very complex. Since tests never can guarantee the correct working of such systems, it is desirable to prove their correctness in a formal way. In order to do so, a formal description of their behaviour is essential. Secondly, agents are equipped with features, like knowledge and belief, of which

*e-mail: jwroorda@math.uu.nl

no consensus about their exact meaning exists. In order to unambiguously define their characteristics, a formal specification is needed.

The formal tool that we use to describe the behaviour of agents is a combination of modal and dynamic logic. We assume that the reader is familiar with both these logics and furthermore, that he knows how they can be used to describe a multi-agent system. If not, the book of [Goldblatt, 1992] is suggested as an introduction to modal and dynamic logic. A good introduction to how these logics can be used to describe multi-agent systems is found in [Hoek *et al.*, 1999].

A well known model of belief update in a model context is given by [Hoek *et al.*, 1999] and [Linder *et al.*, 1995]. This model is part of the KARO-framework for the specification of the behaviour of rational agents. Critique on this model, given by others, has served as inspiration for this paper. We give an alternative model that is not sensitive to this critique.

As stated before, our main goal is to describe how in a multi-agent system, the knowledge and beliefs of agents should change when they receive new information. In order to attain this goal we first need to explain how we model knowledge and beliefs of agents. We do so in chapter two.

After we have described how we model knowledge and belief, we give a precise description of what we mean by updating the beliefs of an agent. Referring to the definition of [Halpern and Friedman, 1998], we call this the *ontology* of our model.

Two core notions in our ontology are *the expansion of the knowledge and beliefs of an agent by a propositional formula ϕ* and *the processing of new information by an agent*. 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 one to its current propositional beliefs. In chapter four we describe formally how the knowledge and beliefs of an agent should change upon expanding.

It is important to stress that we believe that an agent not always should perform a belief expansion when receiving a piece of new information. In general a newly received piece of information should compete with previously received pieces of information to be expanded in the beliefs of the agent. We will describe the processing of new information in chapter five.

In chapter six we restate critique that has been given by others on the KARO-model, we show that our model is not sensitive to it.

We end this article with our conclusions and suggestions for further research.

2 Models of knowledge and belief

To model the knowledge and beliefs of the agents in our multi-agent system we use Kripke models. In this section we give the syntax and semantics of our formal system. Furthermore we give some useful definitions and lemmas.

2.1 Syntax

Definition 2.1 (Language)

Given a finite set $\mathcal{A} = \{1, \dots, n\}$ of agents and a set \mathcal{P} of propositional symbols we define the language \mathcal{L} as follows.

\mathcal{L} is the smallest set containing \mathcal{P} such that

- $\phi, \psi \in \mathcal{L} \Rightarrow \neg\phi, \phi \wedge \psi \in \mathcal{L}$
- $i \in \mathcal{A}, \phi \in \mathcal{L} \Rightarrow \mathbf{K}_i\phi, \mathbf{B}_i\phi \in \mathcal{L}$

\mathcal{L}_0 is the propositional subset of \mathcal{L} .

2.2 Semantics

Definition 2.2 (Kripke models)

The class \mathcal{M} of Kripke models contains tuples $M = (S, \pi, \mathbf{R}, \mathbf{B})$.

where

- S is a set of possible states,
- $\pi : S \rightarrow \mathcal{P} \rightarrow \{0, 1\}$ is a function that assigns truth values to the propositional symbols in the possible worlds,
- $\mathbf{R} : \mathcal{A} \rightarrow S \times S$ is a function that assigns to each agent its epistemic accessibility relation,
- $\mathbf{B} : \mathcal{A} \rightarrow S \times S$ 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}^* .

Definition 2.3 (Knowledge and belief models)

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 $M = (S, \pi, \mathbf{R}, \mathbf{B}) \in \mathcal{M}^*$ holds

- $\mathbf{R} : \mathcal{A} \rightarrow S \times S$ is an equivalence relation,
- $\mathbf{B} : \mathcal{A} \rightarrow S \times S$ is a serial, transitive and euclidean relation,
- for every $i \in \mathcal{A}$ holds $\mathbf{B}_i \subseteq \mathbf{R}_i$
- for every $i \in \mathcal{A}$ and $u, s, t \in S$ holds $\mathbf{R}_i(s, t) \Rightarrow (\mathbf{B}_i(s, u) \Leftrightarrow \mathbf{B}_i(t, u))$.

We use the same system for our models of knowledge and belief as [Linder *et al.*, 1995], whose system in turn is based on [Kraus and Lehmann, 1988].

We make a distinction between the formulas that an agent *knows* and the formulas that an agent *believes*. We require that known facts are true, (an agent can however believe formulas that are false), and that our agents have positive and negative introspection for both knowledge and belief. Furthermore, we require that the agents know their own beliefs. Finally, we require that every known formula is believed.

We achieve these goals by, first of all, requiring that the epistemic accessibility relations are equivalence relations; we use the standard S5-model for knowledge. Furthermore, we demand that the doxastic accessibility relations are serial, transitive and euclidean relations, so that we end up the standard KD-45 model for belief. Finally, to achieve the wanted relation between knowledge and beliefs, we require that the doxastic relations are included in the epistemic relations, and that for every $i \in \mathcal{A}$ and $u, s, t \in S$ holds $R_i(s, t) \Rightarrow (B_i(s, u) \Leftrightarrow B_i(t, u))$

Definition 2.4 (Truth in Kripke models)

Let $M = (S, \pi, R, B)$ be a Kripke model and $s \in S$. We define:

$$\begin{aligned} (M, s) \models p \in \mathcal{P} &\Leftrightarrow \pi(s)(p) = 1 \\ (M, s) \models \phi \wedge \psi &\Leftrightarrow (M, s) \models \phi \text{ and } (M, s) \models \psi \\ (M, s) \models \neg\phi &\Leftrightarrow (M, s) \not\models \phi \\ (M, s) \models K_i\phi &\Leftrightarrow \forall t \in S [R_i(s, t) \Rightarrow (M, t) \models \phi] \\ (M, s) \models B_i\phi &\Leftrightarrow \forall t \in S [B_i(s, t) \Rightarrow (M, t) \models \phi] \end{aligned}$$

2.3 Elementary equivalence and bisimulation

Definition 2.5 (Depth of a formula)

The depth of a formula $\phi \in \mathcal{L}$, notation $\text{depth}(\phi)$, is defined as follows:

$$\begin{aligned} \text{depth}(p) &= 0 \\ \text{depth}(\neg\phi) &= \text{depth}(\phi) \\ \text{depth}(\phi \wedge \psi) &= \max(\text{depth}(\phi), \text{depth}(\psi)) \\ \text{depth}(K_i\phi) &= 1 + \text{depth}(\phi) \\ \text{depth}(B_i\phi) &= 1 + \text{depth}(\psi) \end{aligned}$$

Definition 2.6 (Elementary equivalence)

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

Two Kripke models, say (M, s) and (M', s') , are called n -elementary equivalent, notation $(M, s) \equiv_n (M', s')$ if, and only if, for every $\phi \in \mathcal{L}$ with $\text{depth}(\phi) \leq n$ $(M, s) \models \phi \Leftrightarrow (M', s') \models \phi$

Definition 2.7 (Bounded bisimulation)

Let (M, s) and (M', s') be two Kripke worlds, say $M = (S, \pi, R, B)$ and $M' = (S', \pi', R', B')$. Let n be a natural number. We call (M, s) and (M', s') 0-bisimilar, notation $(M, s) \simeq_0 (M', s')$, if, and only if, $\pi(s) = \pi'(s')$. We call (M, s) and (M', s') $n + 1$ -bisimilar, notation $(M, s) \simeq_{n+1} (M', s')$, if, and only if

- $\pi(s) = \pi'(s')$
- $\forall i \forall t [R_i(s, t) \Rightarrow \exists t' (R'_i(s', t') \wedge (M, t) \simeq_n (M', t'))]$ (forward choice for R)
- $\forall i \forall t' [R'_i(s', t') \Rightarrow \exists t (R_i(s, t) \wedge (M, t) \simeq_n (M', t'))]$ (backward choice for R)

- $\forall i \forall t [B_i(s, t) \Rightarrow \exists t' (B'_i(s', t') \wedge (M, t) \simeq_n (M', t'))]$ (forward choice for B)
- $\forall i \forall t' [B'_i(s', t') \Rightarrow \exists t (B_i(s, t) \wedge (M, t) \simeq_n (M', t'))]$ (backward choice for B)

Proposition 2.1 $(M, s) \equiv_n (M', s') \Leftrightarrow (M, s) \simeq_n (M', s')$

Proof: Like in [Gerbrandy, 1998].

3 Ontology

In their article "Belief Revision: A Critique" [Halpern and Friedman, 1998] describe the *ontology* of a belief change process as a careful description what it means for something to be believed [or known] by an agent, and what the status is of incoming information. In this section we will give a description of the ontology of our model of belief update.

3.1 Believing and knowing

Our model of knowledge and belief is based on the notion of possible worlds. We say that an agent believes, respectively knows, the formula ϕ if ϕ is valid in all the worlds it considers doxastically, respectively epistemically, possible.

The difference between the the notions of knowing and believing is that we require that known formulas are true, whilst 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, we can incorporate them in the knowledge of our agent. When the agent receives incoming information that contradicts these formulas, it will not use them to update his beliefs, because it knows that the incoming information cannot be true.

3.2 Incoming information

Our agents live in a system in which they receive information from different sources. They will use the information they receive from these sources to try to attain a more precise picture of the world they live in.

Two core notions in our ontology are *the expansion of the knowledge and beliefs of an agent by a propositional formula ϕ* and *the processing of new information*

by an agent. We define expansion 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.

In chapter four we will describe in an exact way how the beliefs and the knowledge set of an agent should change upon expanding. But for the remainder to make sense, it is needed to give some characteristics of expanding.

So, let us concentrate for a moment on what should happen with the knowledge and beliefs of an agent, when it decides to believe an incoming propositional formula ϕ while holding on to its current propositional beliefs. Because the agent decides to believe the formula ϕ it should be the case that that agent no longer considers worlds doxastically possible in which ϕ is not true. However, because incoming information need not to be true, and our notions of knowledge incorporate that an agent should only have true knowledge, these worlds should remain epistemically possible. A restriction on the doxastically possible worlds is not enough, because we would like that in every world that the agent considers (doxastically or epistemically) possible, it believes that ϕ is true. This ensures that we retain positive introspection. We will not go further into these questions now, for the remainder of this section it is enough to know that after an expansion with ϕ the agent believes that ϕ , but not necessarily knows that ϕ .

An essential assumption in our ontology is that when an agent receives new information, it has the choice whether it will use this information to expand his beliefs. The reason for this is the following. Our agents live in a world in which they 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 knowledge and beliefs by every piece of information it receives, it could end up with inconsistent beliefs. Of course this is unwanted.

To prevent an agent from forming inconsistent beliefs, we let the agent decide which pieces of received information it will use to expand its beliefs. This decision will 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 agent chooses to use a certain piece of information, he is able to review this choice later on. For instance, suppose an agent starts its life with knowledge K and beliefs B . Now, if an agent receives a piece of information, say ϕ , it can choose to expand his knowledge and beliefs with ϕ , which yields, say $K \star \phi$ and $B \star \phi$. But, if it later on receives an other piece of information, say ψ , it can decide that the update of ϕ was not wanted after all, and use only ψ to update his *original* knowledge K and *original* beliefs B , which yields $K \star \psi$ and $B \star \psi$.

Another core assumption in our ontology is that an agent starts its life without 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 propositional 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 the formula is believed because of an expansion with a certain piece of incoming information. If repeated expansion leads to inconsistent beliefs, the agent should decide by means of a selection function which of the performed expansion actions should be recovered.

4 The belief expansion function

4.1 Notation

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: $\mathcal{M} \times \mathcal{S} \times \mathcal{L}_0 \times \mathcal{A} \rightarrow \mathcal{M} \times \mathcal{S}$.

4.2 Desired characteristics

In this section we describe what characteristics a belief expansion function $[\cdot, \cdot]$ should have. We describe these characteristics in terms of the formulas an expanded Kripke world makes true.

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 $\mathbf{B}_b\alpha$ and $\mathbf{K}_b\alpha$ 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 $\mathbf{B}_a\alpha$ is true if, and only if, for all states t holds, [if $[t$ is a doxastically possible world in the original Kripke world (M, s) and (M, t) models ϕ] then $[(M, t)$ expanded with ϕ models α].

Updating the beliefs of agent a should not change the 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 $\mathbf{K}_a\alpha$ 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 summarized below.


Definition 4.1 (Proper belief expansion functions)

A belief expansion function $[\cdot, \cdot] : \mathcal{M} \times \mathcal{S} \times \mathcal{L}_0 \times \mathcal{A} \rightarrow \mathcal{M} \times \mathcal{S}$ is called proper if

$$\begin{aligned}
(M, s)[\phi, a] \models p &\Leftrightarrow (M, s) \models p \\
(M, s)[\phi, a] \models \alpha \wedge \beta &\Leftrightarrow (M, s)[\phi, a] \models \alpha \text{ and } (M, s)[\phi, a] \models \beta \\
(M, s)[\phi, a] \models \neg\alpha &\Leftrightarrow (M, s)[\phi, a] \not\models \alpha \\
(M, s)[\phi, a] \models \mathbf{B}_b\alpha &\Leftrightarrow (M, s) \models \mathbf{B}_b\alpha \\
(M, s)[\phi, a] \models \mathbf{B}_a\alpha &\Leftrightarrow \forall t((\mathbf{B}_a(s, t) \wedge (M, t) \models \phi) \Rightarrow (M, t)[\phi, a] \models \alpha) \\
(M, s)[\phi, a] \models \mathbf{K}_b\alpha &\Leftrightarrow (M, s) \models \mathbf{K}_b\alpha \\
(M, s)[\phi, a] \models \mathbf{K}_a\alpha &\Leftrightarrow \forall t((\mathbf{R}_a(s, t) \Rightarrow (M, t)[\phi, a] \models \alpha)
\end{aligned}$$

4.3 An example

Let (M, w) be the Kripke world described by the following picture. To keep things simple, the world is inhabited by only one agent. In the picture the doxastic accessibility relation of this agent is reflected by the continuous arrows, and its epistemic accessibility relation by the dotted arrows.



$$w \models p, v \models \neg p$$

Suppose we update the beliefs of the agent in world w with p . According to our definition the new Kripke world $(M, w)[p, a]$ models the following formulas:

- $\mathbf{B}p$
- $\mathbf{B}\mathbf{B}p$
- $\neg\mathbf{K}p$
- $\mathbf{K}\mathbf{B}p$

So we have that after the update agent a believes that p , it believes that it believes that p . It does not know that p , but it knows that it believes that p .

4.4 Properties of proper belief expansion functions

In this section we give some properties of proper belief expansion functions, they serve to strengthen the opinion that proper belief expansion functions behave in an intuitive way.

Proposition 4.1 (Properties of proper belief expansion functions)

For all Kripke models $(M, s) = ((S, \pi, \mathbf{R}, \mathbf{B}), s)$, propositional formulas ϕ, ψ , agents $a, b \in \mathcal{A}$ and proper belief expansion functions $[\cdot, \cdot]$ the following holds:

$$\begin{aligned}
(M, s)[\phi, a] &\models \mathbf{B}_a\phi \\
(M, s)[\phi, a] &\models \mathbf{B}_a\mathbf{B}_a\phi \\
(M, s)[\phi, a] &\models \mathbf{K}_a\mathbf{B}_a\phi \\
(M, s)[\phi, a] &\models \mathbf{K}_a\psi \quad \Leftrightarrow \quad (M, s) \models \mathbf{K}_a\psi
\end{aligned}$$

The first clause states that after the expansion of the beliefs of agent a with ϕ it believes that ϕ . The second and third clause express that after the expansion the agent knows and believes that it believes ϕ . The fourth clause states that the knowledge of propositional formulas is not changed by expansions.

Proof:

We only prove the first two clauses. By definition of proper expansion functions the first clause states $\forall t((\mathbf{B}_a(s, t) \wedge (M, t) \models \phi) \Rightarrow (M, t)[\phi, a] \models \phi)$. Because ϕ is propositional the righthand side is equivalent to $(M, t) \models \phi$. By using the definition of proper belief expansion functions again, we see that the second clause states that $\forall t((\mathbf{B}_a(s, t) \wedge (M, t) \models \phi) \Rightarrow (M, t)[\phi, a] \models \mathbf{B}_a\phi)$. But the righthand side is just the first clause.

4.5 Explicit definition of the proper belief expansion function

A nice property of our definition of proper belief expansion functions is that it defines a unique function. This is proved below.

Proposition 4.2 *There is, up to elementary equivalence, at most one proper belief expansion function*

Proof:

Let $[\cdot, \cdot]$ and $[\cdot, \cdot]'$ be two proper belief expansion functions. We have to prove that for all Kripke worlds (M, s) , for all propositional formulas ϕ , for all agents a that $(M, s)[\phi, a] \equiv (M, s)[\phi, a]'$.

Suppose a propositional formula ψ is given. We will prove by induction on the complexity of ψ that for all Kripke worlds (M, s) $(M, s)[\phi, a] \models \psi \Leftrightarrow (M, s)[\phi, a]' \models \psi$. The basic case is trivial. Suppose that the proposition holds for formulas α and β . The induction steps for $\psi = \alpha \wedge \beta$, $\psi = \neg\alpha$, $\psi = \mathbf{B}_b\alpha$ and $\psi = \mathbf{K}_b\alpha$ are trivial. Suppose $\psi = \mathbf{B}_a\alpha$, then $\forall t((\mathbf{B}_a(s, t) \wedge (M, t) \models \phi) \Rightarrow (M, t)[\phi, a] \models \alpha)$. The i.h. states that $(M, t)[\phi, a] \models \alpha \Leftrightarrow (M, t)[\phi, a]' \models \alpha$. So $(M, s)[\phi, a]' \models \mathbf{B}_a\alpha$. Suppose $\psi = \mathbf{K}_a\alpha$, then $\forall t(\mathbf{R}_a(s, t) \Rightarrow (M, t)[\phi, a] \models \alpha)$. The i.h. states that $(M, t)[\phi, a] \models \alpha \Leftrightarrow (M, t)[\phi, a]' \models \alpha$. So $(M, s)[\phi, a]' \models \mathbf{K}_a\alpha$.

We define *the* proper belief expansion function below.

Definition 4.2 (The proper belief expansion function)

Define $[\cdot, \cdot]_p$ as follows:

$$(M, s)[\phi, a]_p = (M', s')$$

where

$$\begin{aligned} M &= (S, \pi, \mathbf{R}, \mathbf{B}) \\ M' &= (S', \pi', \mathbf{R}', \mathbf{B}') \\ \text{new} : S &\rightarrow S \times \{0\} = \lambda s.(s, 0) \\ \text{old} : S &\rightarrow S \times \{1\} = \lambda s.(s, 1) \\ S' &= \text{old}(S) \cup \text{new}(S) \\ s' &= \text{new}(s) \\ \pi' &= \{(\text{old}(u), p, \pi(u, p)) \mid u \in S, p \in \mathcal{P}\} \cup \\ &\quad \{(\text{new}(u), p, \pi(u, p)) \mid u \in S, p \in \mathcal{P}\} \end{aligned}$$

and the \mathbf{B}_i and \mathbf{R}_i are such that:

$$\begin{aligned} \mathbf{B}'_i(\text{old}(u), \text{old}(v)) &\Leftrightarrow \mathbf{B}_i(u, v) && i \in \mathcal{A} \\ \mathbf{B}'_a(\text{new}(u), \text{new}(v)) &\Leftrightarrow \mathbf{B}_a(u, v) \wedge (M, v) \models \phi \\ \mathbf{B}'_b(\text{new}(u), \text{old}(v)) &\Leftrightarrow \mathbf{B}_b(u, v) && (a \neq b) \\ \mathbf{B}'_a(\text{new}(u), \text{old}(v)) &\Leftrightarrow \perp \\ \mathbf{B}'_i(\text{old}(u), \text{new}(v)) &\Leftrightarrow \perp && i \in \mathcal{A} \\ \mathbf{B}'_b(\text{new}(u), \text{new}(v)) &\Leftrightarrow \perp \\ \\ \mathbf{R}'_i(\text{old}(u), \text{old}(v)) &\Leftrightarrow \mathbf{R}_i(u, v) && i \in \mathcal{A} \\ \mathbf{R}'_a(\text{new}(u), \text{new}(v)) &\Leftrightarrow \mathbf{R}_a(u, v) \\ \mathbf{R}'_b(\text{new}(u), \text{old}(v)) &\Leftrightarrow \mathbf{R}_b(u, v) && (a \neq b) \\ \mathbf{R}'_a(\text{new}(u), \text{old}(v)) &\Leftrightarrow \perp \\ \mathbf{R}'_i(\text{old}(u), \text{new}(v)) &\Leftrightarrow \perp && i \in \mathcal{A} \\ \mathbf{R}'_b(\text{new}(u), \text{new}(v)) &\Leftrightarrow \perp \end{aligned}$$


The intuition behind this definition is, when expanding with ϕ , we make two copies of the original Kripke world. The one containing the states $\text{old}(S)$ 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.

This definition is based on the ideas in [Baltag *et al.*]. We extended their definition for the concept of knowledge.

4.6 An example


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 by the picture below. In the picture the doxastic

accessibility relation of agent a is reflected by the continuous arrows, the doxastic accessibility relation of agent b by the dotted arrows.



$$w \models p, v \models \neg p$$

The expansion of the beliefs of agent a with p in state w yields the following Kripke world (M', w') .



$$w, w' \models p; v, v' \models \neg p$$

4.7 Properties of $[\cdot, \cdot]_p$

Proposition 4.3 $[\cdot, \cdot]_p$ is proper

Proof:

Recalling definition 4.1 we should prove that for given propositional variable $p \in \mathcal{P}$, agents a and b such that $b \neq a$, formulas $\phi \in \mathcal{L}_0$ and formulas $\alpha, \beta \in \mathcal{L}$:

$$\begin{aligned} (M, s)[\phi, a]_p \models p &\Leftrightarrow (M, s) \models p \\ (M, s)[\phi, a]_p \models \alpha \wedge \beta &\Leftrightarrow (M, s)[\phi, a]_p \models \alpha \text{ and } (M, s)[\phi, a]_p \models \beta \\ (M, s)[\phi, a]_p \models \neg \alpha &\Leftrightarrow (M, s)[\phi, a]_p \not\models \alpha \\ (M, s)[\phi, a]_p \models \mathbf{B}_b \alpha &\Leftrightarrow (M, s) \models \mathbf{B}_b \alpha \\ (M, s)[\phi, a]_p \models \mathbf{B}_a \alpha &\Leftrightarrow \forall t ((\mathbf{B}_a(s, t) \wedge (M, t) \models \phi) \Rightarrow (M, t)[\phi, a]_p \models \alpha) \\ (M, s)[\phi, a]_p \models \mathbf{K}_b \alpha &\Leftrightarrow (M, s) \models \mathbf{K}_b \alpha \\ (M, s)[\phi, a]_p \models \mathbf{K}_a \alpha &\Leftrightarrow \forall t ((\mathbf{R}_a(s, t) \Rightarrow (M, t)[\phi, a]_p \models \alpha) \end{aligned}$$

The first item follows directly from the definitions of s' and π' . The second and third item follow directly from the semantical definitions of \wedge and \neg in modal logic.

The fourth claim is justified as follows. " \Rightarrow " Suppose $(M, s)[\phi, a]_p \models \mathbf{B}_b \alpha$. Let $(M, s)[\phi, a]_p = (M', s')$, as in the definition, then $(M', s') \models \mathbf{B}_b \alpha$. Hence $\forall t \in S' (\mathbf{B}'_b(s', t) \Rightarrow (M', t) \models \alpha)$. And thus by the definition of s' $\forall t \in S' (\mathbf{B}'_b(\text{new}(s), t) \Rightarrow (M', t) \models \alpha)$. By observing that $\mathbf{B}'_b(\text{new}(s), x)$ only is realized by $x \in \text{old}(S)$ it follows that $\forall t \in S (\mathbf{B}'_b(\text{new}(s), \text{old}(t)) \Rightarrow (M', \text{old}(t)) \models \alpha)$. Again by the definition of \mathbf{B}'_b $\forall t \in S (\mathbf{B}_b(s, t) \Rightarrow (M', \text{old}(t)) \models \alpha)$. Observe that

$(M', \text{old}(t)) \equiv (M, t)$. So $\forall t \in S(\mathbf{B}_b(s, t) \Rightarrow (M, t) \models \alpha)$. So $(M, s) \models \mathbf{B}_b\alpha$. All the proof steps are reversible, so this is also a proof for " \Leftarrow ".

The fifth claim is justified as follows. " \Rightarrow " Suppose $(M, s)[\phi, a]_p \models \mathbf{B}_a\alpha$. Then, by definition, $(M', s') \models \mathbf{B}_a\alpha$. Hence $\forall t \in S'(\mathbf{B}'_a(s', t) \Rightarrow (M', t) \models \alpha)$. And thus by the definition of s' $\forall t \in S'(\mathbf{B}'_a(\text{new}(s), t) \Rightarrow (M', t) \models \alpha)$. By observing that $\mathbf{B}'_a(\text{new}(s), x)$ only is realized by $x \in \text{new}(S)$ it follows that $\forall t \in S(\mathbf{B}'_a(\text{new}(s), \text{new}(t)) \Rightarrow (M', \text{new}(t)) \models \alpha)$. Using the definition of \mathbf{B}'_a we find that $\forall t \in S((\mathbf{B}_a(s, t) \wedge (M, t) \models \phi) \Rightarrow (M', \text{new}(t)) \models \alpha)$. And thus $\forall t \in S((\mathbf{B}_a(s, t) \wedge (M, t) \models \phi) \Rightarrow (M, t)[\phi, a]_p \models \alpha)$. This yields $(M, s) \models \mathbf{B}_a\alpha$. All proof steps are reversible, so this is also a proof for " \Leftarrow ".

The sixth claim can be proven in the same as the fourth claim has been proved.

Finally, we prove the last claim. " \Rightarrow " Suppose $(M, s)[\phi, a]_p \models \mathbf{K}_a\alpha$. Then, by definition, $(M', s') \models \mathbf{K}_a\alpha$. Hence $\forall t \in S'(\mathbf{R}'_a(s', t) \Rightarrow (M', t) \models \alpha)$. And thus by the definition of s' $\forall t \in S'(\mathbf{R}'_a(\text{new}(s), t) \Rightarrow (M', t) \models \alpha)$. By observing that $\mathbf{R}'_a(\text{new}(s), x)$ only is realized by $x \in \text{new}(S)$ it follows that $\forall t \in S(\mathbf{R}'_a(\text{new}(s), \text{new}(t)) \Rightarrow (M', \text{new}(t)) \models \alpha)$. Using the definition of \mathbf{R}'_a , we find that $\forall t \in S(\mathbf{R}_a(s, t) \Rightarrow (M', \text{new}(t)) \models \alpha)$. And thus $\forall t \in S(\mathbf{R}_a(s, t) \Rightarrow (M, t)[\phi, a]_p \models \alpha)$. This yields $(M, s) \models \mathbf{K}_a\alpha$. Again, all proof steps are reversible, so this is also a proof for " \Leftarrow ".

Proposition 4.4 $[\cdot, \cdot]_p$ is welldefined on the set of Kripke worlds modulo elementary equivalence

Proof:

We will prove by induction on n that for every $n \in \mathbb{N}$, $(M, s), (M', s') \in \mathcal{M} \times \mathcal{S}$, $\phi \in \mathcal{L}_0$ and $a \in \mathcal{A}$ holds $[(M, s) \equiv_n (M', s') \Rightarrow (M, s)[\phi, a]_p \equiv_n (M', s')[\phi, a]_p]$.

The case $n = 0$ is trivial.

The induction hypothesis is that for every $(M, s), (M', s'), \phi, a$ holds $[(M, s) \equiv_n (M', s') \Rightarrow (M, s)[\phi, a]_p \equiv_n (M', s')[\phi, a]_p]$. (*)

We have to prove that for every $(M, s), (M', s'), \phi, a$ holds $[(M, s) \equiv_{n+1} (M', s') \Rightarrow (M, s)[\phi, a]_p \equiv_{n+1} (M', s')[\phi, a]_p]$.

Suppose $(M, s), (M', s'), \phi, a$ are given, s.t. $(M, s) \equiv_{n+1} (M', s')$, together with a formula ψ with $\text{depth}(\psi) \leq n+1$. We have to prove that $(M, s)[\phi, a]_p \models \psi \Leftrightarrow (M', s')[\phi, a]_p \models \psi$. We prove this by induction on the structure of ψ .

The basic case is trivial.

The induction hypotheses is that $(M, s)[\phi, a]_p \models \alpha \Leftrightarrow (M', s')[\phi, a]_p \models \alpha$ and $(M, s)[\phi, a]_p \models \beta \Leftrightarrow (M', s')[\phi, a]_p \models \beta$.

We have to prove that for $\psi = \alpha \wedge \beta, \neg\alpha, \mathbf{B}_b\alpha, \mathbf{K}_b\alpha, \mathbf{B}_a\alpha, \mathbf{K}_a\alpha$ holds $(M, s)[\phi, a]_p \models \psi \Leftrightarrow (M', s')[\phi, a]_p \models \psi$

The cases for $\psi = \alpha \wedge \beta, \neg\alpha, \mathbf{B}_b\alpha, \mathbf{K}_b\alpha$ are trivial.

We will prove the last two cases. Because $\text{depth}(\psi) \leq n+1$, we have $\text{depth}(\alpha) \leq n$.

Suppose $(M, s)[\phi, a]_p \models \mathbf{B}_a\alpha$. Then $\forall t \in S(\mathbf{B}_a(s, t) \wedge (M, t) \models \phi \Rightarrow (M, t)[\phi, a]_p \models \alpha)$ (**). We have to prove that $(M', s')[\phi, a]_p \models \mathbf{B}_a\alpha$, i.e. $\forall t' \in S'(\mathbf{B}'_a(s', t') \wedge (M', t') \models \phi \Rightarrow (M', t')[\phi, a]_p \models \alpha)$.

Suppose that for some $t' \in S'$ $\mathbf{B}'_a(s', t') \wedge (M', t') \models \phi$ holds. We know $(M, s) \equiv_{n+1} (M', s')$, so by proposition 1.1 $(M, s) \simeq_{n+1} (M', s')$. So, by backward choice for \mathbf{B}_a , there is a $t \in S$ such that $\mathbf{B}_a(s, t)$ and $(M, t) \simeq_n (M', t')$. So, for this t we have $(M, t) \equiv_n (M', t')$ (***) . And for this t , we have $(M, t) \models \phi$. So by (**) we have $(M, t)[\phi, a]_p \models \alpha$. By (*) and (***) we have $(M, t)[\phi, a]_p \equiv_n (M', t')[\phi, a]_p$. So $(M', t')[\phi, a]_p \models \alpha$. So indeed $(M', s')[\phi, a]_p \models \mathbf{B}_a\alpha$.

Suppose $(M, s)[\phi, a]_p \models \mathbf{K}_a\alpha$. Then $\forall t \in S(\mathbf{R}_a(s, t) \Rightarrow (M, t)[\phi, a]_p \models \alpha)$ (**2). We have to prove that $(M', s')[\phi, a]_p \models \mathbf{K}_a\alpha$, i.e. $\forall t' \in S'(\mathbf{R}'_a(s', t') \Rightarrow (M', t')[\phi, a]_p \models \alpha)$.

Suppose that for some $t' \in S'$ $\mathbf{R}'_a(s', t')$ holds. We know $(M, s) \equiv_{n+1} (M', s')$, so by proposition 1.1 $(M, s) \simeq_{n+1} (M', s')$. So, by backward choice for \mathbf{R}_a , there is a $t \in S$ such that $\mathbf{R}_a(s, t)$ and $(M, t) \simeq_n (M', t')$. So, for this t we have $(M, t) \equiv_n (M', t')$ (**2). By (**2) we have $(M, t)[\phi, a]_p \models \alpha$. By (*) and (**2) we have $(M, t)[\phi, a]_p \equiv_n (M', t')[\phi, a]_p$. So $(M', t')[\phi, a]_p \models \alpha$. So indeed $(M', s')[\phi, a]_p \models \mathbf{K}_a\alpha$.

The above result is nice, but actually we want more. We want the result of an expansion of a knowledge and belief world to result in a knowledge and belief world. (Or at least to result in a world that is elementary equivalent to a knowledge and belief world.) I.e. we want that for every $(M, s) \in \mathcal{M}^* \times \mathcal{S}$, $\phi \in \mathcal{L}_0$, $a \in \mathcal{A}$: there exist $(M', s') \in \mathcal{M}^* \times \mathcal{S}$ such that $(M, s)[\phi, a] \equiv (M', s')$. Unfortunately, this is not the case, for example, suppose $(M, s) \in \mathcal{M}^* \times \mathcal{S}$, such that $(M, s) \models \mathbf{K}_b\neg\mathbf{B}_a\phi$, then $(M, s)[\phi, a] \models \mathbf{K}_b\neg\mathbf{B}_a\phi$ and $(M, s)[\phi, a] \models \mathbf{B}_a\phi$, i.e. agent b has false knowledge! A solution to this problem would be to restrict the knowledge and belief models to models in which agents do not have knowledge about knowledge and beliefs of other agents that is subject to change by expansions. Unfortunately, we have not succeeded in characterizing this class of models in terms of restrictions on the accessibility relations.

5 The process new information action

As described in the previous sections, our agents do not expand with every piece of information they receive. Instead of that, they 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 (**p-n-i**) action. An agent can use the forget information (**f-i**) action to undo the effects of a process new information action.

To model this actions we shift from a static modal logic to a dynamic modal logic. We describe its syntax and semantics in this section.

5.1 Dynamic Modal Logic

Definition 5.1 (Language)

Given a finite sets $\mathcal{A} = \{1, \dots, n\}$ of agents and $\mathcal{SC} = \{I, II, III, \dots\}$ of sources, and sets \mathcal{P} of propositional symbols and At of atomic actions, we define the language \mathcal{L} and the class of actions Ac by mutual induction as follows.

\mathcal{L} is the smallest set containing \mathcal{P} such that

- $\phi, \psi \in \mathcal{L} \Rightarrow \neg\phi, \phi \wedge \psi \in \mathcal{L}$
- $i \in \mathcal{A}, \alpha \in \text{Ac}, \phi \in \mathcal{L} \Rightarrow \mathbf{K}_i\phi, \mathbf{B}_i\phi, \langle \text{do}_i(\alpha) \rangle\phi, \mathbf{A}_i(\alpha) \in \mathcal{L}$

Ac is the smallest set containing At such that

- $\phi \in \mathcal{L} \Rightarrow \text{confirm } \phi \in \text{Ac}$
- $\alpha_1, \alpha_2 \in \mathcal{L} \Rightarrow (\alpha_1; \alpha_2) \in \text{Ac}$
- $\phi \in \mathcal{L}, \alpha_1, \alpha_2 \in \text{Ac} \Rightarrow (\text{if } \phi \text{ then } \alpha_1 \text{ else } \alpha_2 \text{ fi}) \in \text{Ac}$
- $\phi \in \mathcal{L}, \alpha \in \text{Ac} \Rightarrow (\text{while } \phi \text{ do } \alpha \text{ od}) \in \text{Ac}$
- $\phi \in \mathcal{L}_0, t \in \mathcal{SC} \Rightarrow \mathbf{p_n}\mathbf{j}(\phi, t) \in \text{Ac}$
- $\phi \in \mathcal{L}_0, t \in \mathcal{SC} \Rightarrow \mathbf{f}\mathbf{j}(\phi, t) \in \text{Ac}$

Definition 5.2 (Kripke models)

The class \mathcal{M} of Kripke models contains tuples $M = (S, \pi, \mathbf{R}, \mathbf{B}, (M_o, s_o), \mathbf{H}, \mathbf{r}, \mathbf{c}, \text{Sel})$.

where

- $S, \pi, \mathbf{R}, \mathbf{B}$ are as in definition 2.3,
- $\mathbf{r} : \mathcal{A} \times \text{At} \rightarrow \mathcal{M} \times \mathcal{S} \rightarrow \mathcal{M} \times \mathcal{S}$ is such that $\mathbf{r}(i, a)(M, s)$ yields the Kripke world transition in (M, s) caused by the event $\text{do}_i(a)$,
- $\mathbf{c} : \mathcal{A} \times \text{At} \rightarrow \mathcal{M} \times \mathcal{S} \rightarrow \{0, 1\}$ is the capability function such that $\mathbf{c}(i, a)(M, s)$ indicates whether the agent i is capable of performing the action a in s ,
- $M_o = ((S_o, \pi_o, \mathbf{R}_o, \mathbf{B}_o), s_o)$ is called the original Kripke world, it models the knowledge and beliefs of the agents when they start their lives. The original Kripke world is a Kripke world as in definition 2.3, to ensure that the agents start their lives without proper beliefs, we demand that $\mathbf{R}_{o(i)} = \mathbf{B}_{o(i)}$ for all i ,
- \mathbf{H} is the received information function, of type: $\mathcal{A} \rightarrow [\mathcal{L}_0 \times \mathcal{SC}]$. The received information function gives for every agent a list of the formulas it received in the past, together with their sources,
- Sel is the information selection function of type $\mathcal{M} \times \mathcal{S} \rightarrow \mathcal{A} \rightarrow [\mathcal{L}_0 \times \mathcal{SC}] \rightarrow [\mathcal{L}_o]$, it selects for an agent which information it will use to form its knowledge and beliefs.

We assume that the original knowledge and beliefs 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 propositional formulas that an agent knows as unquestionable. Every other propositional formula that an agent believes should have some kind of justification, in our model this justification will consist of the fact the formula is believed because of an expansion with a certain piece of incoming information. By this we ensure that our agents are always able to make the judgement whether a proper believed formula should be given up in favor of an incoming formula.

Definition 5.3 (Truth in Kripke models)

Let $M = (S, \pi, R, B, (M_o, s_o), H, r, c, \text{Sel})$ be a Kripke model and $s \in S$. We define:

$$\begin{aligned}
(M, s) \models p \in \mathcal{P} &\Leftrightarrow \pi(p) = 1 \\
(M, s) \models \phi \wedge \psi &\Leftrightarrow (M, s) \models \phi \text{ and } (M, s) \models \psi \\
(M, s) \models \neg\phi &\Leftrightarrow \text{not}(M, s) \models \phi \\
(M, s) \models \mathbf{K}_i\phi &\Leftrightarrow \forall t \in S[(s, t) \in R_i \Rightarrow (M, t) \models \phi] \\
(M, s) \models \mathbf{B}_i\phi &\Leftrightarrow \forall t \in S[(s, t) \in B_i \Rightarrow (M, t) \models \phi] \\
(M, s) \models [\text{do}_i(\alpha)]\phi &\Leftrightarrow \forall (M', s')((M', s') \in r(i, \alpha)(M, s) \Rightarrow (M', s') \models \phi) \\
(M, s) \models \mathbf{A}_i\alpha &\Leftrightarrow c(i, \alpha) = 1
\end{aligned}$$

where r and c are extended as follows:

$$\begin{aligned}
r(i, \text{confirm } \phi)(M, s) &= \{(M, s)\} \text{ if } (M, s) \models \phi \text{ and } \emptyset \text{ otherwise} \\
r(i, \alpha_1; \alpha_2)(M, s) &= (r(i, \alpha_2) \circ r(i, \alpha_1))(M, s) \\
r(i, \text{if } \phi \text{ then } \alpha_1 \text{ else } \alpha_2 \text{ fi})(M, s) &= r(i, \alpha_1)(M, s) \text{ if } (M, s) \models \phi \\
&\quad \text{and } r(i, \alpha_2)(M, s) \text{ otherwise} \\
r(i, \text{while } \phi \text{ do } \alpha_1 \text{ od})(M, s) &= \{(M', s') \mid \exists k \in \mathbb{N}, \exists M_0, \dots, M_k \in \mathcal{M}, \\
&\quad \exists s_0, \dots, s_k \in \mathcal{S}((M_0, s_0) = (M, s), \\
&\quad (M_k, s_k) = (M', s') \\
&\quad \&\forall j < k (M_{j+1}, s_{j+1}) \in \\
&\quad r(i, \text{confirm } \phi; \alpha_1)(M_j, s_j)) \\
&\quad \&(M', s') \models \neg\phi\} \\
c(i, \text{confirm } \phi)(M, s) &= 1, \text{ if } (M, s) \models \phi \text{ and } 0 \text{ otherwise} \\
c(i, \alpha_1; \alpha_2)(M, s) &= c(i, \alpha_1) \text{ and } c(i, \alpha_2)(r(i, \alpha_1)(M, s)) \\
c(i, \text{if } \phi \text{ then } \alpha_1 \text{ else } \alpha_2 \text{ fi})(M, s) &= c(i, \text{confirm } \phi; \alpha_1) \text{ or } c(i, \text{confirm } \neg\phi; \alpha_2) \\
c(i, \text{while } \phi \text{ do } \alpha_1 \text{ od})(M, s) &= 1 \text{ if } \exists k \in \mathbb{N} (c(i, (\text{confirm } \phi; \alpha_1)^k; \\
&\quad \text{confirm } \neg\phi)(M, s) = 1) \\
&\quad \text{and } 0 \text{ otherwise}
\end{aligned}$$

In the next section we complete the definition of the semantics by defining the process new information action.

5.2 The process new information action

Before we can define the semantics of the process new information action, we need to make a few technical definitions.

Definition 5.4 (Belief expansion on the new Kripke worlds)

We define the working of the belief expansion function on our new Kripke worlds as follows. Let $(M, s) = ((S, \pi, R, B, (M_o, s_o), H, r, c, \text{Sel}), s)$ be a Kripke world.

$$(M, s)[\phi, a] = ((S', \pi', R', B', (M_o, s_o), H, r, c, \text{Sel}), s')$$

where S', π', R', B' and s' are such that $((S, \pi, R, B), s)[\phi, a] = ((S', \pi', R', B'), s')$.

Definition 5.5 (Retrieving the for expansion selected information)

Given (M, s) , H and Sel , define:

$$\text{Filter}(M, s)(\text{Sel})(H) = [(\phi, i) | \phi \in \text{Sel}(M, s)(i)(H(i)), i \in \mathcal{A}]$$

Filter applied to a Kripke world, a selection function and a received information function yields a list of (agent, formula) 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 knowledge and beliefs by ϕ . This is purely a technical definition. In general the output of this function will be used as input for the belief expansion function, see the following definition.

Definition 5.6 (The belief expansion function on lists)

Let $U : [\mathcal{L}_0 \times \mathcal{A}]$, define:

$$\begin{aligned} (M, s)[\] &= (M, s), \\ (M, s)[U] &= ((M, s)[\phi, t])[U'], \text{ if } U = (\phi, t) : U' \end{aligned}$$

Definition 5.7 (The process new information action)

Let $M = (S, \pi, R, B, (M_o, s_o), H, r, c, \text{Sel})$ be a Kripke model. Let $s \in S$, $\phi \in \mathcal{L}_0$, $i \in A$ and $t \in \mathcal{SC}$. We define:

$$\begin{aligned} r(i, \text{p_n_i}(\phi, t))(M, s) &= \\ \{((S_o, \pi_o, R_o, B_o, (M_o, s_o), H', r, c, \text{Sel}), s_o)[\text{Filter}(M, s)(\text{Sel})(H')]\}. \end{aligned}$$

where H' is such that $H'(i) = (\phi, t) : H(i)$ and $H'(j) = H(j)$ for $j \neq i$.

furthermore:

$$c(i, \text{p_n_i}(\phi, t))(M, s) = 1 \text{ for all } (M, s).$$

The intuition behind this definition is the following. After agent i has performed the $\text{p_n_i}(\phi, t)$ 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 Sel , it will decide which elements of the list of received information $H'(i)$ it will use to form his knowledge and beliefs; this yields $\text{Sel}(M, s)(i)(H'(i))$. 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 $\text{Sel}(M, s)(i)(H'(i))$. 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 $\text{Sel}(j)(M, s)(H(j))$ for $j \neq i$.

For this intuition to make sense, it must be 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 5.8 (Proper Kripke worlds)

A Kripke world $((S, \pi, R, B, (M_o, s_o), H, r, c, \text{Sel}), s)$ is called proper iff.

$$((S, \pi, R, B, (M_o, s_o), H, r, c, \text{Sel}), s) = ((S_o, \pi_o, R_o, B_o, (M_o, s_o), H, r, c, \text{Sel}), s_o)[\text{Filter}(M, s)(\text{Sel})(H)]$$

In our definition of the process new information action we have made a rather arbitrary choice about the order in which we expand the Kripke world with the (agent, formula) tuples. The next proposition shows that this choice does not influence the meaning of the resulting Kripke world.

Proposition 5.1 $(M, s)[\phi, i][\psi, j] \equiv (M, s)[\psi, j][\phi, i]$

Proof:

The case for $i \neq j$ is trivial. The case for $i = j$ can be proved, by showing, by induction on the structure of α , that $(M, s)[\phi, i][\psi, i] \models \alpha \Leftrightarrow (M, s)[\phi \wedge \psi, i] \models \alpha$.

Definition 5.9 (The forget information action)

Let $M = (S, \pi, R, B, (M_o, s_o), H, r, c, \text{Sel})$ be a Kripke model. Let $s \in S, \phi \in \mathcal{L}_0, i \in A$ and $t \in \mathcal{SC}$ s.t. $(\phi, t) \in H(i)$. We define:

$$r(i, f_{\perp}(\phi, t))(M, s) = \{((S_o, \pi_o, R_o, B_o, (M_o, s_o), H', r, c, \text{Sel}), s_o)[\text{Filter}(M, s)(\text{Sel})(H')]\}.$$

where H' is such that $H'(i) = H(i) - (\phi, t)$ and $H'(j) = H(j)$ for $j \neq i$.

furthermore:

$$c(i, f_{\perp}(\phi, t))(M, s) = 1 \text{ if, and only if, } (\phi, t) \in H(i).$$

The intuition behind this action is similar to the intuition behind the p_{\perp} -action. Only in this case a piece of information is deleted from the list of received information $H(i)$. In applications this can be useful, for instance when an agent discovers that a particular piece of information did not make any sense at all, for instance due to a sensor-malfunction.

5.3 Characteristics of the information selection function

One of the reasons for introducing the selection function in our system, was to prevent our agents forming inconsistent beliefs. We will call a selection function which does this job, a proper selection function.

Definition 5.10 (Proper selection functions)

An information selection function $\text{Sel} : \mathcal{M} \times \mathcal{S} \rightarrow \mathcal{A} \rightarrow [\mathcal{L}_0 \times \mathcal{SC}] \rightarrow [\mathcal{L}_0]$ is called proper if for every proper Kripke world $(M, s) = ((S, \pi, R, B, (M_o, s_o), H, r, c, \text{Sel}), s)$ and agent i holds

$$(M_o, s_o) \not\models B_i \perp \Rightarrow (M, s) \not\models B_i \perp.$$

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 given by $\text{Sel}_i(M, s)(i)(H(i)) = []$ is proper, but by using this selection function our agents will never update their knowledge and beliefs. This leads to the definition of eager selection functions. We call a selection function which lets the agents use every piece of information as long as that does not lead to a contradiction an eager selection function.

Definition 5.11 (Eager selection functions)

A proper selection function $\text{Sel} : \mathcal{M} \times \mathcal{S} \rightarrow \mathcal{A} \rightarrow [\mathcal{L}_0 \times \mathcal{SC}] \rightarrow [\mathcal{L}_0]$ is called eager, if for every Kripke world $(M, s) = ((S, \pi, R, B, (M_o, s_o), H, r, c, \text{Sel}), s)$ and agent i holds

$$((S_o, \pi_o, R_o, B_o, (M_o, s_o), H, r, c, \text{Sel}), s_o)[\text{Filter}(M, s)(\text{Sel}_n)(H)] \not\models \mathbf{B}_i \perp$$

\Rightarrow

$$\text{Sel}(M, s)(i)(H(i)) = \text{Sel}_n(M, s)(i)(H(i))$$

where $\text{Sel}_n(M, s)(i)[(\phi_0, t_0), \dots, (\phi_k, t_k)] = [\phi_0, \dots, \phi_k]$ (the naive selection function)

Because the definition of eager selection functions does not say anything about which information should be selected when an agent receives contradicting information, demanding that a selection function is eager still does not ensure that the selection function works in an intuitive way. But because further restrictions would imply arbitrary choices to a certain extent, and because we can prove a lot of nice properties of our system by only demanding that the selection function is eager, we do not suggest further restrictions.

5.4 Properties of our system

In this section we give properties of our system of belief update, they serve to strengthen the opinion that our system behaves in an intuitive way.

Proposition 5.2 (Properties of our system of belief update)

For every proper Kripke model $(M, s) = ((S, \pi, R, B, (M_o, s_o), H, r, c, \text{Sel}), s)$, that contains an eager selection function, $\phi \in \mathcal{L}_0$, $\psi \in \mathcal{L}$, $t \in \mathcal{SC}$ the following holds:

$$\begin{aligned} (M, s) \models \neg \mathbf{B}_i \neg \phi &\quad \rightarrow \quad [\text{do}_i(\text{p_n_i}(\phi, t))] \mathbf{B}_i \phi \\ (M, s) \models \mathbf{B}_i \phi \wedge \neg \mathbf{B}_i \perp &\quad \rightarrow \quad [\text{do}_i(\text{p_n_i}(\neg \phi, t))] \neg \mathbf{B}_i \perp \\ (M, s) \models \mathbf{B}_i \psi &\quad \rightarrow \quad [\text{do}_i((\text{p_n_i}(\phi, t); (\text{f_j}(\phi, t)))] \mathbf{B}_i \psi \end{aligned}$$

The first clause states that if agent a does not believe ϕ , before receiving ϕ , it will believe ϕ afterwards. The second clause states that if agent a does believe ϕ , receiving $\neg \phi$ will not lead to inconsistent beliefs. The third clause expresses that forgetting after processing information does not lead to a change in the beliefs of an agent.

For a selection function, that compromises that the trust of agent i in source I, exceeds the trust of the agent in source II, we have:

$$\begin{aligned} (M, s) &\models \neg B_i \neg \phi \wedge \neg B_i \phi \rightarrow [\text{do}_i((\text{p_n_i}(\phi, I); (\text{p_n_i}(\neg \phi, II)))] B_i \phi \\ (M, s) &\models K_j(\neg B_i \neg \phi \wedge \neg B_i \phi \rightarrow [\text{do}_i((\text{p_n_i}(\phi, I); (\text{p_n_i}(\neg \phi, II)))] B_i \phi) \end{aligned}$$

The first clause states that agents give priority to information from more trustworthy sources. The second clause is an example of the fact that agents are aware of each other selection functions. This can or cannot be desired for a specific implementation of a multi-agent system. If this property is unwanted the notion of Kripke models should be changed such that different worlds can contain different selection functions. In this way one can model uncertainty about other agents' selection functions.

6 Critique on the KARO-model

For our model for belief update we have been inspired by critique on the KARO-model. In this section we treat the main points of critique and show that our model is not sensitive to them.

6.1 Unwanted side effects to the beliefs of other agents

[Gerbrandy, 1998] shows that in the KARO-model an update of the beliefs of agent a leads to changes in the beliefs of other agents. This is unwanted because we assume that other agents are not aware of the update of agent a .

It follows directly from our definition of proper belief expansion functions that in our model the knowledge and beliefs of other agents remain fixed when an agent updates its beliefs.

6.2 Priority to incoming information

[Halpern and Friedman, 1998] claim that the principle of incoming information, which states that observing ϕ leads to believing ϕ , is to be doubted. They argue that this 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 in stead 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 is the case, they state that "we need to describe how to combine a given strength of belief in the observation with the strengths of beliefs in the original epistemic state".

This is exactly what we have done in our model of belief update. 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 (proper) beliefs, it will decide by means of its selection function which of the sources is more trustworthy.

More or less the same point of critique is made by [Dragoni *et al.*, 1994]. They argue that "in a multi-agent environment, where information comes from a variety of [...] sources with different degrees of reliability [...] the chronological sequence of the informative acts has nothing to do with their credibility or importance".

7 Conclusion

In this paper we have given a model of belief update which describes how the knowledge and beliefs of agents should 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 by an agent*. 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 of that they 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 have shown that our model of belief update behaves in an intuitive way and that it is not sensitive to the critique on the KARO-model.

We have limited ourselves to the updates of the knowledge and beliefs of a single agent with a propositional formula. A suggestion for further research would be to extend the model to updates of knowledge and beliefs of groups of agents with modal formulas.

The knowledge and belief models that we defined in 2.3 are not closed under expansion. This is easy to see; there is a knowledge and belief model that entails that agent a knows that agent b does not believe ϕ . Expansion of the knowledge and beliefs of agent b with ϕ leads to false knowledge of a . We have not been able to define restrictions on the accessibility relations of our knowledge and belief models, that yield the greatest subset that is closed under expansion. This could be a topic of further research.

References

- [Baltag *et al.*] Alexandru Baltag, Lawrence S. Moss and Slawomir Solecki (to appear), *The Logic of Public Announcements and Common Knowledge*, <http://www.math.indiana.edu/home/moss/articles.html>.
- [Dragoni *et al.*, 1994] Aldo Franco Dragoni and Paolo Giorgini, *Belief Revision Through the Belief Function Formalism in a Multi-Agent Environment*, volume 1193 of *Lecture Notes in Computer Science (subseries LNAI)*, Springer Verlag, 1994.
- [Gerbrandy, 1998] Jelle Gerbrandy, *Bisimulations on Planet Kripke*, Ph.D. Thesis, Universiteit van Amsterdam, 1998.
- [Goldblatt, 1992] R. Goldblatt, *Logics of Time and Computation*, volume 7 of *CSLI Lecture Notes*, Stanford, 1992.
- [Halpern and Friedman, 1998] Joseph Y. Halpern and Nir Friedman, *Belief Revision: A Critique*, International Conference on Knowledge Representation and Reasoning, 421-431, 1996. <http://simon.cs.cornell.edu/Info/People/halpern/>, 1998.
- [Hoek *et al.*, 1999] W. van der Hoek, B. van Linder, and J.-J. Ch. Meyer, *An Integrated Modal Approach to Rational Agents in Foundations of Rational Agency*, M.Woolridge and A.Rao (eds.), 133-167, 1999
- [Kraus and Lehmann, 1988] S.Kraus, D. Lehman, Knowledge, belief and time, *Theoretical Computer Science*, 58:155-174, 1988.
- [Linder *et al.*, 1995] B. van Linder, W. van der Hoek, and J.-J. Ch. Meyer, *Actions That Make You Change Your Mind in Knowledge and Belief in Philosophy and Artificial Intelligence*, Armin Laux, Heinrich Wansing (eds.), 104-146, 1995