

A Combined System for Update Logic and Belief Revision

Guillaume Aucher

Department of Computer Science
University of Otago
PO Box 56 DUNEDIN 9015
New Zealand
`aucher@atlas.otago.ac.nz`

Abstract. In this paper we propose a logical system combining the update logic of A. Baltag, L. Moss and S. Solecki (to which we will refer to by the generic term BMS, [BMS04]) with the belief revision theory as conceived by C. Alchourrón, P. Gärdenfors and D. Mackinson (that we will call the AGM theory, [GardRott95]) viewed from the point of view of W. Spohn ([Spohn90], [Spohn88]). We also give a proof system and a comparison with the AGM postulates.

Introduction and Motivation : Update logic is a modal logic trying to model epistemic situations involving several agents, and changes that can occur in these situations due to incoming information or more generally incoming action. Belief revision theory typically deals with changes (revisions) that a database representing a belief state of a unique agent must undergo after adding conflicting information to the database. Roughly speaking, these two theories thus deal with the same kind of phenomenon. However, there are some dissimilarities. On the one hand, belief revision theory is not a logic and it deals with a single agent, unlike update logic. On the other hand, belief revision theory deals with revision (and expansion) of information unlike update logic which deals only with expansion of information. Far from being in contradiction, it seems then that these theories have a lot to give each other. So it makes sense to look for a way in which they can be merged.

In Sect.1, we will set out the BMS theory and the AGM theory viewed from the point of view of W. Spohn. In Sect.2 we will propose a system combining these two theories. In Sect.3, we will give an axiomatization of it with a soundness and completeness proof. In Sect.4, we will show that it fulfills the 8 AGM postulates.

1 Update Logic and Belief Revision Theory.

1.1 Update Logic.

In this section we set out the core of update logic as viewed by BMS. We split this account into three parts: 1. static part, 2. dynamic part ('dynamic' because

we deal with actions) and 3. update mechanism. Throughout this exposition and this paper we follow a simple example called the ‘coin’ example taken from [BMS04]. This is the following:

“A and B enter a large room containing a remote-control mechanical coin flipper. One presses the button, and the coin spins through the air, landing in a small box on a table. The box closes. The two people are much too far to see the coin. The coin actually heads up.”

1. **Static Part.** We classically represent the above (static) situation s by the ‘epistemic model’ depicted in Fig.1.




Fig. 1. BMS model for the ‘coin’ example

The tokens w and v represent *possible worlds*. The double border around w means that it is the *actual world*. In this world, the coin is heads up. This last point is rendered formally by assigning the propositional letter H to w , which stands for ‘the coin is Heads up’. Similarly, in the possible world v the coin is tails up. this is rendered formally by assigning the propositional letter T to v , which stands for ‘the coin is Tails up’. This assignment of propositional letters to worlds is rendered formally by what we call a *valuation*: see definition below.

The *accessibility relation* $w \rightarrow_A v$ intuitively means that while A is in world w where the coin is heads up, he still considers possible that he is in world v where the coin is tails up (because he does not know whether the coin is heads or tails up). More generally, we set an accessibility relation $w \rightarrow_j v$ when ‘on the basis of agent j ’s information in world w , the world v is a possible world’.

This epistemic representation of a particular situation is caught by the following general definition:

Definition: We call epistemic model M a tuple $M = (W, \rightarrow_j, V, w_0)$ where W is a set of possible worlds, \rightarrow_j are finitely many accessibility relations indexed by the agents j , V is a valuation function which assigns a set of possible worlds to each propositional letter, and w_0 is the actual world. \diamond

We can then ‘say things’ about specific epistemic models (modeling specific situations) by introducing a language whose one of the components is a knowledge operator K_j defined like that:

$M, w \models K_j \phi$ iff for all v such that $w \rightarrow_j v$, $M, v \models \phi$.

Intuitively $M, w \models K_j \phi$ means ‘in world w , j Knows that ϕ ’. We can then check with this definition that in our example, the epistemic model of Fig.1 captures what we want (e.g. the sentence ‘in the actual world, A does not know whether the coin is Heads or Tails up’ is rendered by the formula $M, w \models \neg K_A H \wedge \neg K_A T$).

See [FHMV95] for an extensive account of what is just outlined here.

2. **Dynamic Part.** Now we consider the following epistemic action a : ‘A cheats and learns that the coin is Heads up, B suspecting anything about it’. We use the term “epistemic” (in “epistemic action”) in the sense that the action doesn’t change facts in the world. We represent how this action is perceived by the agents (just as we represented above how a situation is perceived by the agents) by the action model depicted in Fig.2.




Fig. 2. BMS action model for the action ‘A cheats’

The token σ represents the *simple action* ‘A looks at the coin and observes that the coin is heads up’. A double border around σ means that it is the *actual action*. For this action to be carried out in a particular possible world, the coin needs to be Heads up in this possible world. That’s the intuitive meaning of the *precondition* H in the action model. The token τ represents the simple action ‘nothing happens’. This action can be carried out in any possible world, hence its precondition is the tautology *True*, which is true in any possible world.

The *accessibility relation* $\sigma \rightarrow_B \tau$ intuitively means that ‘while A looks at the coin and observes that it is heads up (σ), for B nothing actually happens (τ)’. More generally, we set an accessibility relation $\sigma \rightarrow_j \tau$ when the following condition is fulfilled: ‘if σ occurs then in j ’s view τ is one of the action that might have happened’.

This epistemic representation of a particular action is caught by the following general definition:

Definition: We call an action model Σ a tuple $\Sigma = (\Sigma, \rightarrow_j, Pre, \sigma_0)$ where Σ is a set of simple action tokens, \rightarrow_j are finitely many accessibility relations indexed by the agents j , Pre is a function which assigns preconditions to each action token, and σ_0 is the actual action. \diamond

3. Update Mechanism. Now, in reality the agents update their beliefs according to these two pieces of information: action a and situation s . This gives rise to a new situation $s \times a$. This actual update is rendered formally by the following mathematical update product:

Definition Let $M = (W, \rightarrow_j, V, w_0)$ be an epistemic model and $\Sigma = (\Sigma, \rightarrow_j, V, \sigma_0)$ an action structure. We define their update product to be the epistemic model $M \otimes \Sigma = (W \otimes \Sigma, \rightarrow'_j, V', w'_0)$ where

1. $W \otimes \Sigma = \{(w, \sigma) \in W \times \Sigma; w \in V(Pre(\sigma))\}$.
2. $(w, \sigma) \rightarrow'_j (v, \tau)$ iff $w \rightarrow_j v$ and $\sigma \rightarrow_j \tau$.
3. $V'(p) = \{(w, \sigma) \in W \otimes \Sigma; w \in V(p)\}$.
4. $w'_0 = (w_0, \sigma_0)$. \diamond

Intuitive interpretation: 1. The possible worlds that we consider after the update are all the ones resulting from the performance of one of the actions in one of the worlds, under the assumption that the action can ‘possibly’ take place in the corresponding world (assumption expressed by the function Pre).

2. The components of our action model are ‘simple’ actions (in the sense of BMS, see [BMS04] for more precision). It allows us to state that the accessibility (or uncertainty) relations for the epistemic model and the epistemic action model are independent from one another. This independence allows us to ‘multiply’ these uncertainties to compute the new accessibility (or uncertainty) relation.

3. The definition of the valuation exemplifies the fact that our actions do not change facts. (That is why we call them *epistemic* actions, as already said above.)

4. Finally, we naturally assume that the actual action can ‘possibly’ take place in the actual world.

Let us get back to our ‘coin’ example. The update product of Fig.1 and Fig.2 yields the model depicted in Fig.3. This model presents some flaws and will be discussed in the rest of the paper.

We have set out the core of update logic as viewed by BMS. Yet, bear in mind that in [BMS04] a genuine logical system is built out of it, that we do not expound here.




Fig. 3. BMS model corresponding to the situation after the action ‘A cheats’

1.2 Belief Revision Theory: W.Spohn’s Approach.

In this section, we set out a simplified account of W.Spohn’s approach to belief revision theory as conceived by AGM (see [GardRott95]).

Generally speaking, belief revision theory deals with changes that must undergo a database representing a belief state of an agent after adding to the database information. (Note that it deals only with the notion of belief and not with the one of knowledge like in update logic.)

The format of the database can take two main different forms: syntactic and semantic. The former consists of a belief set K that consists of *propositional* formulas (also called sentences, representing the facts accepted in the belief state) and that is closed under logical consequences. The latter consists of a set W of possible worlds (representing the narrowest set of possible worlds in which the individual believes that the actual world is located). It can be shown that these two representations are actually equivalent.

The type of change for a state of belief which interests us most is *revision* (the other classical ones are *expansion* and *contraction*). It consists of adding to the belief set K a new sentence ϕ that is typically inconsistent with K . In order that the resulting belief set $K * \phi$ be consistent, some of the old sentences in ϕ are deleted. Now two basic questions come up to mind:

1. What general conditions this revised belief set $K * \phi$ must fulfill in order that the revision process be the closest possible to one performed by rational agents ? This is the concern of the 8 AGM postulates that can be found in [GardRott95].

2. What sentences should be actually deleted from the belief set in order to form the new belief set $K * \phi$? In the literature, there are several explicit procedures that compute the new belief set $K * \phi$ after a revision. We focus on the one proposed by W.Spohn based on a possible world semantics ([Spohn90], [Spohn88]). His approach satisfies moreover the 8 AGM postulates.

Definition: An ordinal conditional function is a function κ from a given set W of possible worlds into the class of ordinals such that some possible worlds are assigned the smallest ordinal 0. \diamond

Intuitively, κ represents a plausibility grading of the possible worlds: the worlds that are assigned the smallest ordinals are the most plausible, according to the beliefs of the individual. Then,

Definition: We define $\kappa(\phi)$ as $\kappa(\phi) := \min\{\kappa(w); w \in \phi\}$.

We say that a formula ϕ is *believed* (with degree of firmness α) when $\kappa^{-1}(0) \subseteq \{w; w \in \phi\}$ (resp. and $\kappa(\neg\phi) = \alpha$).

The belief set K associated with the ordinal conditional function κ is the set of all propositions believed in κ . \diamond

Now assume the sentence ϕ is announced and the agent believes it with a degree of firmness α . We can then define the resulting ordinal conditional function $\kappa * (\phi, \alpha)$ representing the new state of belief:

Definition: Let ϕ be a proposition such that $\{w; w \in \phi\} \neq \emptyset$. We define the ordinal conditional function $\kappa * (\phi, \alpha)$ by:

$$\kappa * (\phi, \alpha)(w) = \begin{cases} \kappa(w) - \kappa(\phi) & \text{if } w \in \phi \\ \alpha + \kappa(w) - \kappa(\phi^c) & \text{if } w \in \phi^c. \end{cases} \diamond$$

Note that in this new belief state, ϕ is believed with firmness α . Finally,

Proposition: If we define $K * \phi$ as the belief set associated with $\kappa * (\phi, \alpha)$, the revision function $*$ thus defined satisfies the 8 AGM postulates. \diamond

So we have set out update logic and belief revision theory as viewed by W. Spohn. Now we are going to propose a system combining these two theories and see what insights it provides us regarding information change. As in the BMS exposition, we split our account in three parts: 1. Static part 2. Dynamic part 3. Update mechanism (inspired from W. Spohn's theory).

2 A Combined System

2.1 The Static Part

Definition. Just as in the BMS system, we want to represent how a static situation is perceived by the agents from the point of view of their beliefs *and* knowledge. That is to say, we want to represent what the agents know and believe about the actual world and also about what the other agents know and believe in general. We do that thanks to what we call a belief epistemic model.

From now on and in the rest of the paper, Max is an arbitrary fixed natural number different from 0.

Definition 1. A belief epistemic model (be-model) $M = (W, \{\sim_j; j \in G\}, \{\kappa_j; j \in G\}, V, w_0)$ is a tuple where:

1. W is a set of possible worlds.
2. w_0 is the possible world corresponding to the actual world.
3. \sim_j is an equivalence relation defined on W for each agent j .
4. κ_j is an operator, ranging from 0 to Max , defined on the set of possible worlds.
5. V is a valuation.
6. G is a set of agents.

Intuitive Interpretation. Points 1,2,5,6 are clear (see Sect.1.1). It remains to give intuitive interpretations for points 3 and 4.

3. The equivalence relation \sim_j intuitively models the notion of knowledge. Its intuitive interpretation is:

$w \sim_j v$ iff agent j 's knowledge in w and v is the same.

Note that this implies that j cannot distinguish world w from v (otherwise she would not have the same knowledge in w and v) and that her information is the same in w and v . This also implies that \sim_j is an equivalence relation, as mentioned in the definition.

4. The plausibility assignment κ_j intuitively models the notion of belief. Among the worlds j cannot distinguish (the worlds where her knowledge is the same), there are worlds that j might consider more plausible than others. This is expressed by the plausibility grading κ_j : the more plausible a world is for the agent j , the closer its plausibility value is to 0 (this is of course completely similar to W. Spohn's approach set out in Sect.1.2). A maximal degree of plausibility Max (originally needed for technical reasons: see Sect.3.1) is introduced and we assume that beyond a certain degree of plausibility (Max), the agent can not distinguish two different worlds of different plausibility.

Remark 1. Note that one could argue that this plausibility assignment should be dependent on the world w in which the agent dwells. This is wrong. Indeed, in a particular world w , j bases her plausibility assignment only on information she has in w . Yet this information is the same in any world indistinguishable from w , as noted in point **3**. So the assignment will be the same for any world indistinguishable from w : that is why we consider a 'global' plausibility assignment.

Static Language and Example. We can easily define a language \mathcal{L}_{St} for be-models (St for *Static*).

Definition 2. The syntax of the language \mathcal{L}_{St} is defined by,

$$\phi := p \mid \neg\phi \mid \phi \wedge \psi \mid K_j\phi \mid B_j^n\phi \text{ where } n \in \mathbb{N}$$

Its semantics is defined by,

$$M, w \models K_j\phi \text{ iff for all } v \text{ st } w \sim_j v, M, v \models \phi.$$

$$M, w \models B_j^n\phi \text{ iff for all } v \text{ st } w \sim_j v \text{ and } \kappa_j(v) \leq n, M, v \models \phi.$$

The intuitive meaning of $M, w \models B_j^n\phi$ is that ‘in world w , j believes with plausibility (a degree of) at most n that ϕ is true’. The definition of B_j^n is taken from [vDL03]. K_j is the usual knowledge operator (see [FHMV95] or Sect.1.1).

Example 1. Let us get back to the ‘coin’ example introduced in Sect.1.1. We model the initial situation in our system by the be-model depicted in Fig.4. The same situation is modeled in the BMS system by the model depicted in Fig.1.




Fig. 4. be-model for the ‘coin’ example

The belief epistemic model depicted in Fig.5 corresponds to the situation resulting from the action ‘A cheats and learns that the coin is heads up’ occurring in the initial situation. We assume *moreover* that in this action ‘B suspects that A cheats’, unlike the BMS framework. The labeling of the worlds will become clear in Sect.2.3; for the time being, just ignore it. Relations in the model are equivalence relations; again, H is for ‘Heads’, T is for ‘Tails’ and the double bordered world corresponds to the actual world.

In this model and in the actual world (w, σ) , A *knows* that the coin is heads up (formally: $K_A H$) and B *believes* that A doesn’t know whether the coin is Heads or Tails (formally: $B_B^0(\neg K_A H \wedge \neg K_A T)$).

The corresponding model of BMS is depicted in Fig.6 (or Fig.3.). In this model, The nuance of concepts (belief B and knowledge K) is not displayed because we use the same crude accessibility relation \rightarrow_j for both the notions of knowledge and belief. So, when we read what is true in the actual world (w, σ) , we have *personally* to introduce this nuance of concept (belief B and knowledge K) because it is not displayed in the formalism itself. This is of course a flaw of the BMS system.




Fig. 5. be-model corresponding to the situation after the action ‘A cheats’




Fig. 6. BMS model corresponding to the situation after the action ‘A cheats’

2.2 The Dynamic Part

Again, just as in BMS, we want to represent how an action is perceived by several agents from the point of view of their beliefs and knowledge. That is to say, we want to do the same thing as in the last section, but with an action instead of a static situation. We do that thanks to what we call a belief epistemic action model.

Definition 3. A belief epistemic action model (be-action-model) Σ is a tuple $(\Sigma, \{\sim_j; j \in G\}, \{\kappa_j^*; j \in G\}, Pre, \sigma_0)$ such that:

1. Σ is a set of possible actions.
2. σ_0 is the possible action corresponding to the actual action.
3. \sim_j is an equivalence relation on Σ indexed by the set of agents.
4. κ_j^* is an operator indexed by the set of agents, ranging from 0 to Max.
5. Pre is a function from the set of simple actions to the formulas of \mathcal{L}_{St} .
6. G is a set of agents.

The intuitive interpretation is very similar to the one spelled out for the notion of a belief epistemic model. So we refer the reader to the previous section for a correct interpretation of the definition: the term ‘world’ just has to be replaced by the term ‘action’. The only differences concern the absence of a valuation and the introduction of the function Pre . Intuitively, $Pre(\sigma)$ is a necessary condition for the action σ to be performed in a particular world (see Sect.1.1).

Here again, \sim_j and κ_j , modeling respectively the notions of knowledge and belief, are refinements of the crude epistemic relation \rightarrow_j of BMS.

Example 2. We reconsider the example of cheating (see Sect.1.1): ‘A cheats and learns that the coin is heads up, B suspecting that A cheats’. We propose the belief epistemic action model depicted in Fig.7 for this action (where the double border corresponds to the actual action and relations are equivalence relations). In this be-action-model, while ‘A looks at the coin and observes H’ (action σ), B *believes* ‘nothing happens’ (action μ) but nevertheless considers plausible (with plausibility α and β) that A looked at the coin (actions σ and τ respectively), because she suspects that A has cheated.

The corresponding action model of BMS is depicted in Fig.8 (or Fig.2). Contrary to this one, we add one other possible action, that B may consider possible : ‘A looks at the coin and observes tail’ (depicted as τ). Indeed, in our framework, B suspects A of having cheated but doesn’t know whether in that case A has observed Heads or Tails. Hence we have to consider two possible actions for cheating: ‘A looks at the coin and observes H’ (action σ) and ‘A looks at the coin and observes T’ (action τ). The third action μ represents the action where ‘nothing happens’.




Fig. 7. be-action-model for the action ‘A cheats’




Fig. 8. BMS action model for the action ‘A cheats’

Example 3. We consider the action consisting of the public announcement ‘A knows whether the coin is Heads or Tails’ (formally: $K_A H \vee K_A T$). There is no essential difference between our model (depicted in Fig.9) and the BMS model (depicted in Fig.10) but we need their introduction for the purpose of the following section.

$$\boxed{\boxed{\rho : K_A H \vee K_A T}} \\ \kappa_A(\rho) = \kappa_B(\rho) = 0.$$

Fig. 9. be-action-model for the action ‘public announcement that A knows whether the coin is Heads or Tails’




Fig. 10. BMS action model for the action ‘public announcement that A knows whether the coin is Head or Tail’

2.3 The Update Mechanism

We now define the mathematical update product. It is supposed to render the actual update performed by the agents which follows their apprehension of the action. The apprehension (or perception) of the action corresponds to what we modeled in the last section by a be-action-model. Note that this process (apprehension + update) may be done simultaneously in reality, but in our formalism we clearly separate it.

Definition 4. Given a belief epistemic model $M = (W, \sim_j, \kappa_j, V, w_0)$ and a belief epistemic action model $\Sigma = (\Sigma, \sim_j, \kappa_j^*, Pre, \sigma_0)$ we define their update product to be the belief epistemic model $M \otimes \Sigma = (W \otimes \Sigma, \sim'_j, \kappa'_j, V', w'_0)$, where:

1. $W \otimes \Sigma = \{(w, \sigma) \in W \times \Sigma; w \in V(Pre(\sigma))\}$.
2. $(w, \sigma) \sim'_j (v, \tau)$ iff $w \sim_j v$ and $\sigma \sim_j \tau$
3. $\kappa'_j(w, \sigma) = Cut_{Max}(\kappa_j^*(\sigma) + \kappa_j(w) - \kappa_j^w(\phi))$

where $\phi = Pre(\sigma)$, $\kappa_j^w(\phi) = \min\{\kappa_j(v); v \in V(\phi) \text{ and } v \sim_j w\}$ and

$$Cut_{Max}(x) = \begin{cases} x & \text{if } 0 \leq x \leq Max. \\ Max & \text{if } x > Max \end{cases}$$

4. $V'(p) = \{(w, \sigma) \in W \times \Sigma; w \in V(p)\}$
5. $w'_0 = (w_0, \sigma_0)$.

Intuitive interpretation. Points 1,2,4,5 have exactly the same interpretation as in BMS (see Sect.1.1). It remains to motivate the key point concerning the update of plausibility, which is inspired from W. Spohn’s ordinal conditional function $\kappa * (\phi, \alpha)$ (see Sect.1.2). We will give two justifications. The first one is ‘intuitive’ and the second one is related to probability theory.

First of all, $\kappa_j^*(\sigma) + \kappa_j(w) - \kappa_j^w(\phi)$ is the core of the update. Cut_{Max} is just a minor technical device so that the new plausibility assignment fits into the set

$\{0, \dots, Max\}$. However it has also an intuitive import if we refer to the assumption motivating the introduction of *Max* in Sect.1.1.1.

1. Now the first justification. We are interested in an update performing a genuine belief revision. So, our rational intuition should guide us in order to determine the correct plausibility update. In that respect, it seems intuitively clear that,

‘If you believe an action has taken place, then after the update you should believe what is *then* (after the update) true in the worlds where the action has taken place.’

Or more generally and precisely,

‘In a current world w , if you believe with plausibility $\kappa_j^*(\sigma)$ that an action σ has taken place, then after the update you should believe with plausibility $\kappa_j^*(\sigma)$ what is *then* true in the worlds where the action has taken place and that you cannot distinguish from your current world w .’

So we would be tempted to assign roughly to the worlds accessible from w where the action σ has taken place the plausibility $\kappa_j^*(\sigma)$. Yet doing so, we would lose part of the overtones and information present in the former model amongst the worlds where the action σ has taken place (and that are accessible from w). So we add $\kappa_j(w) - \kappa_j^w(\phi)$ to $\kappa_j^*(\sigma)$ in order to keep track of and incorporate this former information (glance at the definition of $\kappa_j^w(\phi)$).

2. Now another justification. W. Spohn showed in [Spohn90] that we can draw a precise and rigorous parallel between probability theory and his plausibility theory. More precisely, he showed that sum, multiplication, and division of probabilities can be replaced respectively by the minimum, addition, and subtraction of plausibilities. This will be of interest for us: we will jump from plausibility to probability, then use probability results to get what we want and finally jump back to plausibility by translating our probabilistic outcome.

For our purpose, note that we can perfectly replace in this justification plausibility of worlds $\kappa_j(w)$ (and actions $\kappa_j(\sigma)$) by probability of worlds $P_j(w)$ (and actions $P_j(\sigma)$ respectively). Now, we want to determine $\kappa_j(w, \sigma)$. That is to say, in a probabilistic setting, we want to determine $P_j(w, \sigma) = P_j(W \cap A)$ where W stands for ‘we were in world w ’ and A for ‘action σ occurred in w ’. Probability theory tells us that

$$P_j(W \cap A) = P_j(A) \cdot P_j(W|A).$$

Clearly,

$$P_j(A) = P_j(\sigma).$$

So it remains to determine $P_j(W|A)$, that is to say the probability that we were in world w *given* the extra assumption that action σ occurred in this world. We reasonably claim

$$P_j(W|A) = \frac{P_j(w)}{\sum\{P_j(v); w \sim_j v \text{ and } v \in V(Pre(\sigma))\}}.$$

That is to say, we *conditionalize* the probability of w for j ($P_j(w)$) to the worlds where the action σ has taken place and that may correspond for j to the actual world w ($\{v; w \sim_j v \text{ and } v \in V(Pre(\sigma))\}$). That is how it would be done in classical probability theory. The intuition behind it is that we now possess the extra piece of information that σ occurred in w , so the worlds indistinguishable from w where the action σ did *not* occur do not play a role anymore for the determination of the probability of w : we can then get rid of them and conditionalize on the relevant worlds.

Finally we get:

$$P_j(w, \sigma) = \frac{P_j(\sigma) \cdot P_j(w)}{\sum\{P_j(v); w \sim_j v \text{ and } v \in V(Pre(\sigma))\}}.$$

Now with the translation from probability to plausibility proved in [Spohn90] ('sum \rightarrow minimum', 'multiplication \rightarrow addition', 'division \rightarrow subtraction'), we get the expected outcome:

$$\kappa'_j(w, \sigma) = \kappa_j^*(\sigma) + \kappa_j(w) - \kappa_j^w(\phi)$$

Remark 2. The first justification of the plausibility assignment stresses the priority of the plausibility assignment of the be-action-model upon the plausibility assignment of the be-model. So in a sense, it stresses also the priority of new information upon former information, just as in belief revision theory. We think this is how the plausibility assignment of action should be interpreted.

Remark 3. We have implicitly assumed in these two justifications that the plausibility of the actions are independent from the worlds in which they are performed. However, this is wrong for some cases. Indeed, for example consider a vague announcement of a formula ϕ that the agent j cannot distinguish from ϕ' because she is not sure whether she heard it correctly. In a \sim_j -equivalence class where the agent j knows more formulas that logically imply ϕ than ϕ' , j will find the announcement ϕ more plausible than ϕ' because she will have more actual evidence at her disposal to think so. However, in another \sim_j -equivalence class of the same model where the agent j knows more formulas that logically imply ϕ' than ϕ , j will find the announcement of ϕ less plausible than ϕ' because she will have less actual evidence that would prompt her to think so.

Example 4. In this example, we are going to see the added value of a combined system.

Yet, first, let us briefly give an example of update. If we update the be-model depicted in Fig.4 by the be-action-model depicted in Fig.7, we get the be-model depicted in Fig.5. That is what we want. (Note that correlatively, in the BMS system, the update of the model depicted in Fig.1 by the action model depicted in Fig.8 (or Fig.2) would yield the model depicted in Fig.6 (or Fig.3).)

Now in the actual world of the model depicted in Fig.5 (respectively Fig.6 for the BMS system), B *believes* that A doesn't know whether the coin is Heads or Tails. Yet, A actually knows whether the coin is Heads or Tails. B's belief is consequently wrong. So, what happens if we update these models by a public announcement that 'A knows whether the coin is Heads or Tails'? Indeed this public announcement would contradict B's beliefs and then B would have to *revise* her beliefs.

This public announcement is depicted in Fig.9 (respectively Fig.10 for the BMS system). This update yields the be-model depicted in Fig.11 (respectively Fig.12 for the BMS system). In this be-model (Fig.11), B now believes that A knows whether the coin is Heads or Tails, so B did revise her beliefs. On the other hand, in the BMS model (Fig.12) B now believes everything. In a sense we could say that the announcement drives her 'crazy' because it contradicts her beliefs.

So, we see in this example that the BMS system does not perform belief revision, unlike our system. This is of course a flaw of the BMS system.




Fig. 11. be-model corresponding to the situation after the announcement 'A knows whether the coin is Head or Tail'




Fig. 12. BMS model corresponding to the situation after A has cheated and after the announcement 'A knows whether the coin is Head or Tail'

Generally speaking, it seems that a system allowing for misperception (like cheating) *has* to incorporate a revision of belief feature.

2.4 The Full Language $\mathcal{L}(\Sigma_{sg})$

We extend the static language defined in Sect.2.1. in order to incorporate the dynamic feature. This dynamic feature will be displayed in the full language by the programs and the modality $[\pi]$.

Let $\Sigma_{sg} = (\Sigma, \sim_j, \kappa_j^*)$ be a fixed action signature (see [BMS04]).

Definition 5. *The syntax of the language $\mathcal{L}(\Sigma_{sg})$ is defined by,*

- Sentences $\phi := True \mid p \mid \neg\phi \mid \phi \wedge \psi \mid K_j\phi \mid B_j^n\phi \mid [\pi]\phi$
where $n \in \mathbb{N}$
- Programs $\pi := \sigma\psi \mid \pi + \rho \mid \pi.\rho$ (see [BMS04])

Its semantics is completely similar to the BMS one (see [BMS04]), except for the operator B_j^n (see Sect.2.1).

Intuitively, $M, w \models [\pi]\phi$ says that ‘in world w , after the action corresponding to the program π has been performed, ϕ will hold’.

Our system is a semantically driven logical system. So naturally, we can try to find an axiomatization and a completeness proof for it. That is the concern of the next section.

3 Logic of Combined Update and Revision

3.1 The Full Proof System

We set out in this section the core of the proof system AX w.r.t. the semantics of $\mathcal{L}(\Sigma_{sg})$. The core of the sub-proof system AX’ w.r.t. the semantics of \mathcal{L}_{St} (see Sect.2.1) is labeled with *. (For a full version of the proof system, see the web-site <http://www.ilc.uva.nl/Publications/reportlist.php?Series=MoL>.)

1. $\vdash [\sigma_i\psi]p \leftrightarrow (\psi_i \rightarrow p)$
2. $\vdash [\sigma_i\psi]\neg\chi \leftrightarrow (\psi_i \rightarrow \neg[\sigma_i\psi]\chi)$
3. $\vdash [\sigma_i\psi]\phi \wedge \chi \leftrightarrow ([\sigma_i\psi]\phi \wedge [\sigma_i\psi]\chi)$
4. $\vdash [\sigma_i\psi]K_j\phi \leftrightarrow \{\psi_i \rightarrow \bigwedge\{K_j[\sigma_k\psi]\phi; \sigma_k \sim_j \sigma_i\}\}$
5. $\vdash [\sigma_i\psi]B_j^n\phi \leftrightarrow (\psi_i \rightarrow \bigwedge\{B_j^{p-1}\neg\psi_k \wedge \neg B_j^p\neg\psi_k \rightarrow B_j^{n+p-\kappa_j^*(Pre(\sigma_k))}[\sigma_k\psi]\phi; \sigma_k \sim_j \sigma_i \text{ and } p \in \{0..Max\}\})$ where $n < Max$.
6. $\vdash [\pi.\rho]\phi \leftrightarrow [\pi][\rho]\phi$
7. $\vdash [\pi + \rho]\phi \leftrightarrow [\pi] \wedge [\rho]\phi$
8. $*\vdash K_j\phi \rightarrow \phi$
9. $*\vdash B_j^n\phi \rightarrow K_j B_j^n\phi$ for all $n \in \mathbb{N}$
10. $*\vdash \neg B_j^n\phi \rightarrow K_j \neg B_j^n\phi$ for all $n \in \mathbb{N}$
11. $*\vdash B_j^n\phi \rightarrow B_j^{n'}\phi$ for all $n \geq n'$

12. $*\vdash K_j\phi \leftrightarrow B_j^n\phi$ for all $n \geq Max$.

Axiom 12 is somewhat problematic. Indeed, we ‘jump’ from the notion of belief with highest plausibility to the notion of knowledge: this is somewhat mysterious ! Note that we could avoid that by allowing an infinite number of degrees of belief (and then get rid of Max), but then in axiom 5 we would get an infinite conjunction in the second term.

Moreover, it seems unfortunately impossible to give an intuitive import to axiom 5.

3.2 Completeness and Soundness Proofs

An exhaustive completeness and soundness proof can be found on the website <http://www.illc.uva.nl/Publications/reportlist.php?Series=MoL> . We only provide a sketch of the proof here.

Soundness Proof. We only give the soundness proof of axiom 5. The soundness of the other axioms can be easily checked or are spelled out in [BMS04].

First, note that in any be-model M ,

Fact: $\kappa_j^w(\phi) = l \Leftrightarrow M, w \models B_j^{l-1}\neg\phi \wedge \neg B_j^l\neg\phi$.

Now, if we spell out the definition of $M, w \models [\sigma_i\psi]B_j^n\phi$, we get at a certain point to the expression $Cut_{Max}(\kappa_j^*(\sigma_j) + \kappa_j(v) - \kappa_j^w(\psi_j)) \leq n$. Yet, $n < Max$, so this expression is equivalent to $\kappa_j^*(\sigma_j) + \kappa_j(v) - \kappa_j^w(\psi_j) \leq n$ which is again equivalent to $\kappa_j(v) \leq n + \kappa_j^w(\psi_j) - \kappa_j^*(\sigma_j)$. That is how the positions of $[\sigma_i\psi]$ and B_j are swapped in axiom 5. Finally the value of $\kappa_j^w(\psi_j)$ is determined by the Fact.

Completeness Proof. First we show:

Theorem 1. *AX' is a sound and strongly complete axiomatization with respect to the semantics of \mathcal{L}_{St} .*

Proof. We use the following canonical model:

$M^c = (W^c, \sim_j, \kappa_j, V)$ where,

- $W^c = \{w_W; W \text{ maximal AX'-consistent set}\}$.
- $\sim_j = \{(w_V, w_W); V/K_j \subseteq W\}$ where $V/K_j = \{\phi; K_j\phi \in V\}$.
- $\kappa_j(w_W) = \min\{n; W/B_j^n \subseteq W\}$.
- $w_W \in V(p)$ iff $p \in W$.

Lemma 1. *For all $\phi \in \mathcal{L}(\Sigma_{sg})$, there is $\phi_{St} \in \mathcal{L}_{St}$ such that $\vdash \phi \leftrightarrow \phi_{St}$*

Proof. We prove it by successive inductions (see the web-site mentioned above). We use in great extent the ‘reduction’ axioms 1 to 5 : they all ‘push through’ the epistemic operators and connectives, except for the basic case 1 where $[\sigma_i\psi]$ disappears.

Theorem 2. *AX is strongly complete with respect to the semantics of $\mathcal{L}(\Sigma_{sg})$.*

Proof. Thanks to lemma 1 and the soundness of our logic, the completeness proof with respect to the semantics of the full language boils down to the completeness proof with respect to the static language. This last point has been shown in theorem 1, so we have the expected result.

4 Comparison with the AGM Postulates

To check whether the AGM postulates are fulfilled, we first need to define, relatively to a world w and for an agent j , the belief set, the expanded belief set and the revised belief set. We will deal with propositional language as in the AGM theory (see Sect.1.2). The type of be-action-model we naturally consider for the update is a public announcement of a propositional formula ϕ , depicted in Fig.13.

$$\boxed{\boxed{\sigma : \phi}}$$

$$\kappa_j(\sigma) = 0 \text{ for all } j.$$

Fig. 13. be-action-model for the action ‘public announcement of the propositional formula ϕ ’

Definition 6. *For each world w , we define*

- the belief set $K^w = \{\phi \in \mathcal{L}; M, w \models B_j^0\phi\}$,
- the revision of the belief set K^w by ϕ , $K^w * \phi = \{\psi \in \mathcal{L}; M, w \models [\sigma, \phi]B_j^0\psi\}$,
- the expansion of the belief set K^w by ϕ , $K^w + \phi = \{\psi \in \mathcal{L}; M, w \models B_j^0[\sigma, \phi](\phi \rightarrow \psi)\}$,


where \mathcal{L} is the propositional language.

Theorem 3. *(If $M, w \models \neg K_j\neg\phi$ then) $*$ defined by $K^w * \phi$ satisfies the 8 AGM postulates.*

Proof. The proof is standard and can be found on the web-site <http://www.ilc.uva.nl/Publications/reportlist.php?Series=MoL>.

Note 1. In the theorem, the assumption within brackets is a natural one.

Remark 4. If we consider the epistemic language (i.e. with knowledge) instead of the propositional language \mathcal{L} for the formation of belief sets, then some AGM postulates are not fulfilled. This failure is due to the fact that the epistemic formulas satisfiable in any world may change after an update with a public announcement (phenomenon called ‘persistence’). Yet, assume we slightly change the be-action-model depicted in Fig.13 and replace it with the one depicted in Fig.14.

$$\kappa_j^*(\sigma) = 0 \text{ and } \kappa_j^*(\tau) = \alpha_j \text{ for all agents } j.$$


The diagram shows two boxes representing belief sets. The left box contains the text $\sigma : \phi$ and the right box contains $\tau : \neg\phi$. A horizontal line connects the two boxes, with a small 'j' positioned above the line, indicating an action performed by agent j.

Fig. 14.

Then the 8 AGM postulates are satisfied for the epistemic language. (For a proof see the web-site mentioned above.) Indeed, the epistemic formulas satisfiable in any world are the same after an update with this type of be-action-model.

Moreover, after the update, ϕ is believed by all the agents j with *firmness* α_j in world w (see Sect.1.2 for a definition of the notion of firmness).

General Conclusion : We have set out a logical system merging update logic and belief revision theory. This system satisfies the AGM postulates. However, extending the BMS framework with an accessibility relation for the notion of belief and updating it by ‘multiplication’ does not perform a *revision* of belief but rather an *expansion* of beliefs.

Nevertheless, our system presents some limitations. First, as mentioned in remark 3, we assume that the plausibility of an action does not depend on the world in which it is performed. Unfortunately, we cannot make small variations to our current system in order to avoid this assumption. This then shrinks the set of actions we can consider in our system. Second, the relationship between the notions of knowledge and belief is not properly rendered as it is suggested in axiom 12, and the theorem $\neg K_j\phi \rightarrow K_j\neg K_j\phi$ is inadmissible in many types of situations (although not in our example).

Concerning the second point, we believe one can tackle epistemic issues and describe (actual) epistemic situations accurately only with an epistemic formalism which renders properly all the nuances and overtones within and between these notions. This paper is a first step towards it: we enriched and refined our epistemic formalism by introducing the notion of belief explicitly. Further and better refinement is the concern of ongoing research.

Acknowledgement : I thank my master’s thesis supervisors Johan van Benthem and Hans van Ditmarsch for their support during the process of my master’s thesis. I also want to thank Andreas Herzig, Gary Cleveland and three anonymous referees for their comments about this paper.

References

- [BMS04] A. Baltag, L.S. Moss, and S. Solecki. Logic for epistemic program. In *Synthese* Volume 139, Issue, March 2004. Pages: 165 - 224.
- [FHMV95] R. Fagin, J.Y. Halpern, Y. Moses, and M.Y. Vardi. *Reasoning about knowledge*. MIT Press, Cambridge MA, 1995.
- [GardRott95] P. Gardenfors and H. Rott, 1995, 'Belief Revision', in D. M. Gabbay, C. J. Hogger and J. A. Robinson, eds., *Handbook of Logic in Artificial Intelligence and Logic Programming 4*, Oxford University Press, Oxford 1995.
- [Spohn90] W. Spohn. A general non-probability theory of inductive reasoning. In R. D. Schachter, T. S. Levitt, L. N. Kanal, and J. F. Lemmer, editors, *Uncertainty in artificial intelligence 4*, pages 149-155. North-Holland, Amsterdam, 1990.
- [Spohn88] W. Spohn. Ordinal conditional functions: A dynamic theory of epistemic states. In W. L. Harper and B. Skyrms, editors, *Causation in Decision, Belief Change, and Statistics*, vol.2, pages 105-134. Reidel, Dordrecht, 1988.
- [vDL03] H.P. van Ditmarsch and W.A. Labuschagne. *A multimodal language for revising defeasible beliefs*. In E. Álvarez and R. Bosch and L. Villamil, editors, *Proceedings of the 12th International Congress of Logic, Methodology, and Philosophy of Science (LMPS)*, pages 140-141, Oviedo University Press, 2003.