



Consistency preservation and crazy formulas in BMS

Guillaume Aucher

► **To cite this version:**

Guillaume Aucher. Consistency preservation and crazy formulas in BMS. Logics in artificial intelligence, 11th european conference (JELIA), Sep 2008, Dresden, Germany. 2008. <inria-00556148>

HAL Id: inria-00556148

<https://hal.inria.fr/inria-00556148>

Submitted on 1 Sep 2013

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Consistency Preservation and Crazy Formulas in BMS^{*}

Guillaume Aucher

University of Otago (NZ) - Université Paul Sabatier (F)
IRIT, 118 route de Narbonne, 31062 Toulouse cedex 9 (F)
aucher@irit.fr

Abstract. We provide conditions under which seriality is preserved during an update in the BMS framework. We consider not only whether the entire updated model is serial but also whether its generated submodels are serial. We also introduce the notion of crazy formulas which are formulas such that after being publicly announced at least one of the agents' beliefs become inconsistent.

1 Introduction

Providing formalisms which allow agents to reason adequately about belief and belief change is an important goal in artificial intelligence so that artificial agents can act autonomously and rationally in a given environment. An obvious requirement for these agents should be that their beliefs always remain consistent whatever happens.

One way to represent formally the agents' beliefs about a given (static) situation is by means of an epistemic model. Expressing that the agents' beliefs are consistent amounts to assume that the accessibility relations of the epistemic model are serial.

The next step is to introduce events and to model their effects on the agents' beliefs. An influential formalism has been proposed in dynamic epistemic logic by Baltag, Moss and Solecki (to which we will refer by the term BMS [3]) that deals with this issue. Their idea is to represent how an event occurring in this situation is perceived by the agents by means of an event model and then to define a formal update mechanism that specifies how the agents update their beliefs according to this event model and the original epistemic model. This yields a new epistemic model corresponding to the resulting situation. However, as it turns out, it is quite possible formally that this new epistemic model is not serial, even if the original epistemic model and the event model were serial. For example, if agent A believes $\neg\phi$ then after a public announcement of ϕ agent A 's accessibility relation is not serial anymore. Specifying formally under which conditions this happens in general (and not only for public announcements) has not been studied so far. In this paper we tackle this issue and also introduce some

* I thank my supervisors Hans van Ditmarsch and Andreas Herzig for useful comments on this paper.

formulas, called ‘crazy formulas’, such that some of the agents’ beliefs always become inconsistent after they are publicly announced: they become ‘crazy’.

The paper is organized as follows. In Section 2, we recall the BMS system together with the definitions of seriality and generated (sub)model. In Section 3, we provide conditions under which the entire updated epistemic model is serial and introduce what we call ‘crazy formulas’. In Section 4, we investigate under which conditions a generated submodel of the entire updated epistemic model is serial. Finally, we conclude in Section 5.

2 The BMS System

In this paper, Φ is a set of propositional letters and G is a finite set of agents.

2.1 Epistemic Model and Generated Submodel

Epistemic model. An epistemic model is just a particular kind of Kripke model [5] where instead of having a single accessibility relation we have a set of accessibility relations, one for each agent.

Definition 1. An epistemic model M is a triple $M = (W, R, V)$ such that

- W is a non-empty set of possible worlds;
- $R : G \rightarrow 2^{W \times W}$ assigns an accessibility relation to each agent;
- $V : \Phi \rightarrow 2^W$ assigns a set of possible worlds to each propositional letter and is called a valuation.

If $M = (W, R, V)$ is an epistemic model, a pair (M, w_a) with $w_a \in W$ is called a pointed epistemic model. We also write $R_j = R(j)$ and $R_j(w) = \{w' \in W \mid wR_jw'\}$, and $w \in M$ for $w \in W$.

Intuitively, a pointed epistemic model (M, w_a) represents from an external point of view how the actual world w_a is perceived by the agents G . The possible worlds W are the relevant worlds needed to define such a representation and the valuation V specifies which propositional facts (such as ‘it is raining’) are true in these worlds. Finally the accessibility relations R_j model the notion of belief. We set $w' \in R_j(w)$ in case the world w' is compatible with agent j ’s belief in world w . We can then define the notion of seriality for epistemic models.

Definition 2. Let $M = (W, R, V)$ be an epistemic model. We say that M is serial when for all $j \in G$, R_j satisfies the following condition:

Seriality: for all $w \in W$, $R_j(w) \neq \emptyset$.

Intuitively, an epistemic model which is not serial means that in the corresponding situation there is an agent j whose beliefs are inconsistent. More precisely, it means that it is not common belief that the agents’ beliefs are consistent.

Now inspiring ourselves from modal logic, we can define a language for epistemic models. The modal operator is just replaced by a ‘belief’ operator, one for each agent.

Definition 3. The language \mathcal{L}^U is defined as follows:

$$\mathcal{L}^U : \phi ::= \top \mid p \mid \neg\phi \mid \phi \wedge \phi \mid B_j\phi \mid U\phi$$

where p ranges over Φ and j over G . Moreover, $\phi \vee \phi'$ is an abbreviation for $\neg(\neg\phi \wedge \neg\phi')$; $\phi \rightarrow \phi'$ is an abbreviation for $\neg\phi \vee \phi'$; $\hat{B}_j\phi$ is an abbreviation for $\neg B_j\neg\phi$; $O\phi$ is an abbreviation for $\neg U\neg\phi$; and \perp is an abbreviation for $\neg\top$.

Finally, by \mathcal{L} we denote the language \mathcal{L}^U without the universal modality U .

Intuitively, $B_j\phi$ means that agent j believes that the formula ϕ is true. U is the universal modality which is introduced here only for technical reasons in order to express the seriality preservation conditions. Now we can give a genuine meaning to the formulas of this language by defining truth conditions for these formulas on the class of epistemic models.

Definition 4. Let $M = (W, R, V)$ be an epistemic model and $w \in W$. $M, w \models \phi$ is defined inductively as follows:

$$\begin{aligned} M, w &\models \top \\ M, w &\models p && \text{iff } w \in V(p) \\ M, w &\models \neg\phi && \text{iff not } M, w \models \phi \\ M, w &\models \phi \wedge \phi' && \text{iff } M, w \models \phi \text{ and } M, w \models \phi' \\ M, w &\models B_j\phi && \text{iff for all } v \in R_j(w), M, v \models \phi \\ M, w &\models U\phi && \text{iff for all } v \in W, M, v \models \phi \end{aligned}$$

We write $M \models \phi$ for $M, w \models \phi$ for all $w \in M$. If \mathcal{C} is a class of epistemic models, we write $\models_{\mathcal{C}} \phi$ when for all $M \in \mathcal{C}$, $M \models \phi$. Finally, we write $\models \phi$ when for all epistemic model M , $M \models \phi$.

So agent j believes ϕ in world w (formally $M, w \models B_j\phi$) if ϕ is true in all the worlds that agent j considers possible (in world w). $M, w \models U\phi$ expresses that ϕ is valid in the model M . The universal modality is thus a stronger notion than the common belief modality often used in epistemic logic.

Example 1. We take up the coin example of [3]. Assume there are two agents Ann and Bob who are in a room where there is a coin in a box. The coin is actually heads up but the box is closed. So both of them do not know whether the coin is heads or tails up. Now assume that Bob cheats and looks at the coin, Ann suspecting nothing about it. This resulting situation is modeled in the pointed epistemic model (M, w_a) of Figure 1. The accessibility relations are represented by arrows indexed by A (standing for Ann) or B (standing for Bob); p stands for ‘the coin is heads up’ and the boxed world w_a stands for the actual world. Now thanks to the language \mathcal{L} we can express what is true in this situation. For example, Bob correctly believes that the coin is heads up: $M, w_a \models p \wedge B_B p$; while Ann believes that he does not know whether the coin is heads or tails up: $M, w_a \models B_A(\neg B_B p \wedge \neg B_B \neg p)$.




Fig. 1. The coin example

Generated submodel. An epistemic model might contain some information that is not relevant to model a given situation. We now define the notion of generated submodel that discards this useless information.

Definition 5. Let $M = (W, R, V)$ and $M' = (W', R', V')$ be two epistemic models and $w_a \in W$.

- We say that M' is a submodel of M if $W' \subseteq W$; for all $j \in G$, $R'_j = R_j \cap (W' \times W')$ and for all $p \in \Phi$, $V'(p) = V(p) \cap W'$. We also say that M' is the restriction of M to W' .
- The submodel of M generated by w_a is the restriction of M to $(\bigcup_{j \in G} R_j)^*(w_a)^1$.

In case the submodel of M generated by w_a is M itself, we say that M is generated by w_a and that w_a is the root of M .

Proposition 1. Let $M = (W, R, V)$ be an epistemic model and M' a submodel of M generated by some $w_a \in W$. Then for all $w \in M'$ and all $\phi \in \mathcal{L}$, $M, w \models \phi$ iff $M', w \models \phi$.

This proposition entails that in a pointed epistemic model (M, w_a) where w_a stands for the actual world, the part of the model M that is really relevant for us to model the corresponding situation is the submodel of M generated by w_a .

2.2 Event Model

Epistemic models are used to model how the agents perceive the actual world in terms of beliefs about the world and about the other agents' beliefs. The insight of the BMS approach is that one can describe how an event is perceived by the agents in a very similar way. Indeed, the agents' perception of an event can also be described in terms of beliefs: for example, while Bob looks at the coin and sees that it is heads up (event a_a) Ann *believes* that nothing happens (event b). This leads them to define the notion of event model whose definition is very similar to that of an epistemic model.

¹ If R is a relation, R^+ is defined by $R^+(w) = \{v \mid \text{there is } w_1, \dots, w_n = v \text{ such that } w_i R w_{i+1}\}$. R^* is defined by $R^*(w) = \{w\} \cup R^+(w)$. See [5].




Fig. 2. Public announcement of B_{BP}

Definition 6. An event model A is a triple $A = (E, R, Pre)$ such that

- E is a finite and non-empty set of possible events;
- $R : G \rightarrow 2^{E \times E}$ assigns an accessibility relation to each agent;
- $Pre : E \rightarrow \mathcal{L}$ assigns an epistemic formula to each possible event.

If $A = (E, R, Pre)$ is an event model, a pair (A, a_a) where $a_a \in E$ is called a pointed event model. We also write $R_j = R(j)$ and $R_j(a) = \{b \in E \mid aR_j b\}$, and $a \in A$ for $a \in E$.

The main difference with the definition of an epistemic model is that we no longer have a valuation V but instead a function Pre . This function is supposed to specify under which condition an event can physically take place in a possible world.

Example 2. Assume that an external agent announces publicly that Bob believes that the coin is heads up (formally B_{BP}). This event is depicted in Figure 2. There, a_a stands for ‘the external agent truthfully announces that Bob believes that the coin is heads up’. Because this event is correctly perceived by Ann and Bob, a_a is the only event considered possible by them. Finally, for this truthful announcement to be made in a possible world, Bob has indeed to believe that the coin is heads up in this world (B_{BP}).

2.3 Product Update

Now, in reality after (or during) this event e , the agents update their beliefs by taking into account these two pieces of information: the event e and the initial situation s . This gives rise to a new situation $s \times e$. This actual update is rendered formally by the following mathematical update product between a pointed epistemic model and a pointed event model.

Definition 7. Let $M = (W, R, V, w_a)$ be a pointed epistemic model and $A = (E, R, Pre, a_a)$ a pointed event model such that $M, w_a \models Pre(a_a)$. We define their update product to be the pointed epistemic model $M \otimes A = (W \otimes E, R', V', w'_a)$ where

1. $W \otimes E = \{(w, a) \in W \times E \mid M, w \models Pre(a)\}$;
2. $(v, b) \in R'_j(w, a)$ iff $v \in R_j(w)$ and $b \in R_j(a)$;
3. $V'(p) = \{(w, a) \in W \otimes E \mid w \in V(p)\}$;
4. $w'_a = (w_a, a_a)$.




Fig. 3. Failure of seriality preservation

Example 3. This example shows that seriality might not be preserved during an update. If we update the epistemic model depicted in Figure 1 by the truthful public announcement that Bob believes that the coin is heads up (formally $B_B p$) depicted in Figure 2 then we get the epistemic model depicted on the right of Figure 3 where Ann’s accessibility relation is not serial.

3 Seriality Preservation for the Entire BMS Product

3.1 Theory

First of all, for a given epistemic model M and a given event model A , we say that the update product $M \otimes A$ is *defined* if there is $w \in M$ and $a \in A$ such that $M, w \models \text{Pre}(a)$. We introduce this definition because seriality of updated models makes sense only for defined updated models.

Proposition 2. *Let A be a serial event model and let M be an epistemic model.*

$M \otimes A$ is defined and serial iff

$$M \models O \left(\bigvee_{a \in A} \text{Pre}(a) \right) \wedge U \bigwedge_{a \in A} \left(\text{Pre}(a) \rightarrow \bigwedge_{j \in G} \hat{B}_j \bigvee_{b \in R_j(a)} \text{Pre}(b) \right).$$

Proof. $M \models O \left(\bigvee_{a \in A} \text{Pre}(a) \right)$ clearly means that the model $M \otimes A$ is defined.

Now it remains to prove that $M \otimes A$ is serial iff

$$M \models U \bigwedge_{a \in A} \left(\text{Pre}(a) \rightarrow \bigwedge_{j \in G} \hat{B}_j \bigvee_{b \in R_j(a)} \text{Pre}(b) \right).$$

- Assume that $M \models U \bigwedge_{a \in A} \left(\text{Pre}(a) \rightarrow \bigwedge_{j \in G} \hat{B}_j \bigvee_{b \in R_j(a)} \text{Pre}(b) \right)$ (*). Let $(w, a) \in M \otimes A$ and $j \in G$. Then $M, w \models \text{Pre}(a)$. So $M, w \models \bigwedge_{j \in G} \hat{B}_j \bigvee_{b \in R_j(a)} \text{Pre}(b)$ by (*). Then $M, w \models \hat{B}_j \bigvee_{b \in R_j(a)} \text{Pre}(b)$. So there is $v \in R_j(w)$ and $b \in R_j(a)$ such that $M, v \models \text{Pre}(b)$. Then there is $(v, b) \in M \otimes A$ such that $(v, b) \in R_j(w, a)$ by definition of $M \otimes A$. So $M \otimes A$ is serial.

- Assume that $M \not\models U \bigwedge_{a \in A} \left(\text{Pre}(a) \rightarrow \bigwedge_{j \in G} \hat{B}_j \bigvee_{b \in R_j(a)} \text{Pre}(b) \right)$. Then there is $w \in M$ and $a \in A$ such that $M, w \models \text{Pre}(a) \wedge \left(\bigvee_{j \in G} B_j \bigwedge_{b \in R_j(a)} \neg \text{Pre}(b) \right)$. Then there is $j \in G$ such that $M, w \models B_j \bigwedge_{b \in R_j(a)} \neg \text{Pre}(b)$ (**). So $(w, a) \in M \otimes A$ but there is no $v \in R_j(w)$ and $b \in R_j(a)$ such that $(v, b) \in R_j(w, a)$. Indeed, otherwise we would have $M, w \models \hat{B}_j \bigvee_{b \in R_j(a)} \text{Pre}(b)$, which contradicts (**). So $M \otimes A$ is not serial.

We write $\mathcal{S}(A) = O \left(\bigvee_{a \in A} \text{Pre}(a) \right) \wedge U \bigwedge_{a \in A} \left(\text{Pre}(a) \rightarrow \bigwedge_{j \in G} \hat{B}_j \bigvee_{b \in R_j(a)} \text{Pre}(b) \right)$. $O \left(\bigvee_{a \in A} \text{Pre}(a) \right)$ expresses that the updated model $M \otimes A$ is defined.

$U \bigwedge_{a \in A} \left(\text{Pre}(a) \rightarrow \bigwedge_{j \in G} \hat{B}_j \bigvee_{b \in R_j(a)} \text{Pre}(b) \right)$ expresses that the updated model $M \otimes A$ is serial. Note that the seriality conditions bear only on M and that M does not need to be serial for the proposition to hold. But of course if M is serial then the proposition still holds. However, if the event model is not serial then one cannot get a serial updated model. From this proposition we can easily prove the following corollary.

Corollary 1. *Let \mathcal{C} be a class of epistemic models and A a serial event model.*

$$\models_{\mathcal{C}} \neg \mathcal{S}(A)$$

iff there is no epistemic model $M \in \mathcal{C}$ such that $M \otimes A$ is defined and serial.

In other words this corollary tells us under which condition, for a given event model A , whatever epistemic model M we chose, $M \otimes A$ will not be defined or not serial. If this condition is fulfilled that would mean intuitively that in any epistemic situation, if the event (corresponding to this event model) is performed, then afterwards in any case (some of) the agents' beliefs are inconsistent. This is of course counter intuitive and we should then avoid such kinds of event (models).

3.2 Crazy Formulas

We are going to give an example of a class of epistemic formulas such that after they are publicly announced some of the agents' beliefs become inconsistent.

Definition 8. *A crazy formula is a satisfiable formula $\phi \in \mathcal{L}$ such that*

$$\models \phi \rightarrow \bigvee_{j \in G} B_j \neg \phi.$$

Proposition 3. *Let ϕ be a crazy formula and let A be the event model corresponding to the public announcement of ϕ . Then there is no epistemic model M such that $M \otimes A$ is defined and serial.*

Proof. Thanks to Corollary 1 it suffices to prove that $\models \neg \mathcal{S}(A)$ i.e. $\models O\phi \rightarrow O\left(\phi \wedge \bigvee_{j \in G} B_j \neg \phi\right)$ because $\mathcal{S}(A) = O\phi \wedge U\left(\phi \rightarrow \bigwedge_{j \in G} \hat{B}_j \phi\right)$. Let M be an epistemic model such that $M \models O\phi$. Let $w \in M$ such that $M, w \models \phi$. Then by definition of a crazy formula $M, w \models \phi \wedge \bigvee_{j \in G} B_j \neg \phi$. Then $M, w \models \phi \wedge \bigvee_{j \in G} B_j \neg \phi$ i.e. $M \models O\left(\phi \wedge \bigvee_{j \in G} B_j \neg \phi\right)$.

Proposition 4. $\phi = \psi \wedge B_i \neg \psi$, where $\psi \in \mathcal{L}$ is a satisfiable formula, is a crazy formula.

Proof. One can easily show that $\models \phi \rightarrow B_i \neg \phi$.

We can compare this notion of crazy formula with the notion of selfrefuting and successful formulas studied in [7]. Selfrefuting formulas are formulas that are no longer true after they are publicly announced. An example of such formulas is Moore's sentence $p \wedge \neg B_j p$: if it is announced then p becomes common belief and in particular $B_j p$ becomes true. Here our formulas are a bit different: after they are publicly announced some of the agents' beliefs become inconsistent. On the other hand, successful formulas are formulas which are always true after being publicly announced. One can show that crazy formulas are not successful.

4 Seriality Preservation for Generated Submodels

One should note that it is quite possible that an updated model consists of several disjoint epistemic models. But in practice, as we said in Section 2.1, the epistemic model we are really interested in is the submodel of the entire updated model generated by the actual world (w_a, a_a) . So, more generally, we would like to know under which conditions a particular generated submodel of the entire updated model is serial. That is what we are going to investigate now.

Definition 9. Let A be an event model, $a \in A$ and $n \in \mathbb{N}$. We define $\delta^n(a)$ inductively as follows.

- $\delta^0(a) = \text{Pre}(a)$;
- $\delta^{n+1}(a) = \delta^0(a) \wedge \bigwedge_{j \in G} \hat{B}_j \bigvee_{b \in R_j(a)} \delta^n(b) \wedge \bigwedge_{j \in G} B_j \bigwedge_{b \in R_j(a)} (\text{Pre}(b) \rightarrow \delta^n(b))$.

Intuitively, $M, w \models \delta^n(a)$ means that the submodel of $M \otimes A$ generated by (w, a) is defined and serial up to modal depth n . This interpretation is endorsed by the following two lemmas which will be used to prove the main proposition.

Lemma 1. Let M be an epistemic model and let A be an event model. For all $w \in M$, $a \in A$, $n \in \mathbb{N}$,

$$M, w \models \delta^{n+1}(a) \text{ iff}$$

$M, w \models \delta^1(a)$ and for all $v \in M$ such that $w = w_0 R_{j_1} w_1 R_{j_2} \dots R_{j_n} w_n = v$ such that there are $a = a_0 R_{j_1} a_1 R_{j_2} \dots R_{j_n} a_n = b$ such that for all $i \in \{0, \dots, n\}$, $M, w_i \models \text{Pre}(a_i)$,

$$M, v \models \bigwedge_{j \in G} \hat{B}_j \bigvee_{c \in R_j(b)} \text{Pre}(c)$$

Proof. We prove it by induction on n . The case $n = 0$ is clear. We prove the induction step. Assume the property is true for n .

- Assume $M, w \models \delta^{n+2}(a)$. Then $M, w \models \delta^1(a)$ because $\delta^{n+1}(b) \rightarrow \text{Pre}(b)$ and $\delta^1(a) = \text{Pre}(a) \wedge \bigwedge_{j \in G} \hat{B}_j \bigvee_{b \in R_j(a)} \text{Pre}(b)$. Let $v \in M$ such that $w = w_0 R_{j_1} w_1 R_{j_2} \dots R_{j_{n+1}} w_{n+1} = v$ and such that there are $a = a_0 R_{j_1} a_1 R_{j_2} \dots R_{j_{n+1}} a_{n+1} = b$ such that for all $i \in \{0, \dots, n+1\}$, $M, w_i \models \text{Pre}(a_i)$.

By assumption, $M, w \models \bigwedge_{j \in G} B_j \bigwedge_{b \in R_j(a)} (\text{Pre}(b) \rightarrow \delta^{n+1}(b))$. So $M, w_1 \models \bigwedge_{b \in R_{j_1}(a)} (\text{Pre}(b) \rightarrow \delta^{n+1}(b))$. Besides $a_1 \in R_{j_1}(a)$ and $M, w_1 \models \text{Pre}(a_1)$. So $M, w_1 \models \delta^{n+1}(a_1)$.

Then, by induction hypothesis, for all v' such that $w_1 = w'_1 R_{j_2} \dots R_{j_{n+1}} w'_{n+1} = v'$ such that there are $a_1 = a'_1 R_{j_2} \dots R_{j_{n+1}} a'_{n+1} = a'$ such that for all i , $M, w'_i \models \text{Pre}(a'_i)$,

$$M, v' \models \bigwedge_{j \in G} \hat{B}_j \bigvee_{b' \in R_j(a')} \text{Pre}(b').$$

So $M, v \models \bigwedge_{j \in G} \hat{B}_j \bigvee_{c \in R_j(b)} \text{Pre}(c)$

- Assume $M, w \models \delta^1(a)$ and assume that for all $v \in M$ such that $w = w_0 R_{j_1} \dots R_{j_n} w_{n+1} = v$ such that there are $a = a_0 R_{j_1} \dots R_{j_n} a_{n+1} = b$ such that for all i , $M, w_i \models \text{Pre}(a_i)$,

$$M, v \models \bigwedge_{j \in G} \hat{B}_j \bigvee_{c \in R_j(b)} \text{Pre}(c).$$

Now, assume $M, w \not\models \delta^{n+2}(a)$. Then $M, w \models \neg \text{Pre}(a) \vee \left(\bigvee_{j \in G} B_j \bigwedge_{b \in R_j(a)} \neg \delta^{n+1}(b) \right)$

$\vee \left(\bigvee_{j \in G} \hat{B}_j \bigvee_{b \in R_j(a)} (\text{Pre}(b) \wedge \neg \delta^{n+1}(b)) \right)$.

- $M, w \models \neg \text{Pre}(a)$ is impossible by assumption.
- Assume $M, w \models \bigvee_{j \in G} B_j \bigwedge_{b \in R_j(a)} \neg \delta^{n+1}(b)$. Then for some $i \in G$, $M, w \models B_i \bigwedge_{b \in R_i(a)} \neg \delta^{n+1}(b)$. Then for all $v \in R_i(w)$ and all $b \in R_i(a)$, $M, v \models \neg \delta^{n+1}(b)$ (*).

But by assumption $M, w \models \delta^1(a)$, i.e. $M, w \models \text{Pre}(a) \wedge \bigwedge_{j \in G} \hat{B}_j \bigvee_{b \in R_j(a)} \text{Pre}(b)$. Then $M, w \models \hat{B}_i \bigvee_{b \in R_i(a)} \text{Pre}(b)$, i.e. there is $v \in R_i(w)$ and $b \in R_i(a)$ such that $M, v \models \text{Pre}(b)$ (1).

So $M, v \models \bigwedge_{j \in G} \hat{B}_j \bigvee_{b \in R_j(a)} \text{Pre}(b)$ (2) by assumption (take $w_1 = \dots = w_n = v$ and $a_1 = \dots = a_n = b$).

Then by (1) and (2) we get $M, v \models \delta^1(b)$.

Besides, by assumption and because $wR_i v$ and $aR_i b$, for all u such that $v = v_0 R_{j_1} \dots R_{j_n} u$ such that there are $b = b_0 R_{j_1} \dots R_{j_n} b_n = c$ such that for all i $M, v_i \models \text{Pre}(b_i)$

$$M, u \models \bigwedge_{j \in G} \hat{B}_j \bigvee_{d \in R_j(c)} \text{Pre}(d).$$

So $M, v \models \delta^{n+1}(b)$ by induction hypothesis. This is impossible by (*).

- Assume $M, w \models \bigvee_{j \in G} \hat{B}_j \bigvee_{b \in R_j(a)} (\text{Pre}(b) \wedge \neg \delta^{n+1}(b))$.

Then there is $i \in G$, $v \in R_i(w)$ and $b \in R_i(a)$ such that $M, v \models \text{Pre}(b) \wedge \neg \delta^{n+1}(b)$.

By the same argument as above we get to a contradiction.

So finally $M, w \models \delta^{n+2}(a)$.

Lemma 2. *Let M be a finite epistemic model and A be a finite serial event model. Let $n = |M| \cdot |A|^2$. For all $w \in M$ and $a \in A$ such that $M, w \models \text{Pre}(a)$,*

1. $R_j(w, a) \neq \emptyset$ for all $j \in G$ iff $M, w \models \bigwedge_{j \in G} \hat{B}_j \bigvee_{b \in R_j(a)} \text{Pre}(b)$;

2. $(v, b) \in \left(\bigcup_{j \in G} R_j \right)^+ (w, a)$ iff there are $w = w_0 R_{j_1} w_1 R_{j_2} \dots R_{j_n} w_{n-1} = v$ and $a = a_0 R_{j_1} a_1 R_{j_2} \dots R_{j_n} a_{n-1} = b$ such that for all i , $M, w_i \models \text{Pre}(a_i)$.

Proof. 1. Assume $M, w \models \bigwedge_{j \in G} \hat{B}_j \bigvee_{b \in R_j(a)} \text{Pre}(b)$. Then for all $j \in G$, there is

$v \in R_j(w)$ and $b \in R_j(a)$ such that $M, v \models \text{Pre}(b)$. Then, by definition of the product update, for all j , there is $(v, b) \in M \otimes A$ such that $(v, b) \in R_j(w, a)$. So for all $j \in G$, $R_j(w, a) \neq \emptyset$.

Assume $M, w \not\models \bigwedge_{j \in G} \hat{B}_j \bigvee_{b \in R_j(a)} \text{Pre}(b)$. Then there is $j \in G$ such that for

all $v \in R_j(w)$ and for all $b \in R_j(a)$, $M, v \not\models \text{Pre}(b)$. Then, by definition of the product update, there is no $(v, b) \in M \otimes A$ such that $(v, b) \in R_j(w, a)$. So $R_j(w, a) = \emptyset$ for some $j \in G$.

2. $M \otimes A$ is of cardinality at most n due to our hypothesis that $n = |M| \cdot |A|$. So

every world $(v, b) \in M \otimes A$ such that $(v, b) \in \left(\bigcup_{j \in G} R_j \right)^+ (w, a)$ is accessible from (w, a) in at most $n - 1$ steps. So,

² $|M|$ (resp. $|A|$) is the number of possible worlds (resp. events) of M (resp. A).

$$(v, b) \in \left(\bigcup_{j \in G} R_j \right)^+ (w, a) \text{ iff}$$

there are j_1, \dots, j_{n-1} and $(w_1, a_1), \dots, (w_{n-1}, a_{n-1}) \in M \otimes A$ such that $(w, a)R_{j_1}(w_1, a_1)R_{j_2} \dots R_{j_{n-1}}(w_{n-1}, a_{n-1}) = (v, b)$ iff

there are $w = w_0R_{j_1}w_1R_{j_2} \dots R_{j_{n-1}}w_{n-1} = v$ and $a = a_0R_{j_1}a_1R_{j_2} \dots R_{j_{n-1}}a_{n-1} = b$ such that for all i , $M, w_i \models \text{Pre}(a_i)$.

Proposition 5. *Let M be a finite epistemic model and let A be a finite serial event model. Let $w \in M$, $a \in A$ and $n = |M| \cdot |A|$.*

The submodel of $M \otimes A$ generated by (w, a) is defined and serial iff $M, w \models \delta^n(a)$.

Proof. First, note that the submodel of $M \otimes A$ generated by (w, a) is defined and serial iff

- (w, a) is defined;
- $R_j(w, a) \neq \emptyset$ for all $j \in G$;
- $R_j(v, b) \neq \emptyset$ for all $(v, b) \in \left(\bigcup_{j \in G} R_j \right)^+ (w, a)$ and for all $j \in G$.

Then we get easily the expected result by Lemma 2 and Lemma 1. Indeed, (w, a) is defined and $R_j(w, a) \neq \emptyset$ for all $j \in G$ amounts to say that $M, w \models \delta^1(a)$. And

$R_j(v, b) \neq \emptyset$ for all $(v, b) \in \left(\bigcup_{j \in G} R_j \right)^+ (w, a)$ and for all $j \in G$ amounts to say that for all $v \in M$ such that $w = w_0R_{j_1}w_1R_{j_2} \dots R_{j_n}w_n = v$ such that there are $a = a_0R_{j_1}a_1R_{j_2} \dots R_{j_n}a_n = b$ such that for all $i \in \{0, \dots, n\}$, $M, w_i \models \text{Pre}(a_i)$, $M, v \models \bigwedge_{j \in G} \hat{B}_j \bigvee_{c \in R_j(b)} \text{Pre}(c)$.

This proposition is coherent with our interpretation of $M, w \models \delta^n(a)$. As we said, intuitively, $M, w \models \delta^n(a)$ means that the submodel of $M \otimes A$ generated by (w, a) is (defined and) serial up to modal depth n . So, if n is larger than the modal depth of the submodel $M \otimes A$ generated by (w, a) (which is the case if $n = |M| \cdot |A|$) then all the worlds accessible from (w, a) are serial. So this generated submodel is indeed serial. Accordingly, this also entails that it should be serial for any given modal depth. That is what the following property expresses.

Proposition 6. *Let M be a finite epistemic model and let A be a finite and serial event model. Let $w \in M$, $a \in A$ and $n = |M| \cdot |A|$.*

If $M, w \models \delta^n(a)$ then for all $m \geq n$, $M, w \models \delta^m(a)$.

Proof. The proof follows from Lemma 1 and the fact that for all $v \in M$ there are w_1, \dots, w_{n-1} such that $w = w_0R_{j_1}w_1R_{j_2} \dots R_{j_n}w_n = v$ iff there are w_1, \dots, w_{m-1} such that $w = w_0R_{j_1}w_1R_{j_2} \dots R_{j_m}w_m = v$.

Similarly, if a submodel of $M \otimes A$ generated by (w, a) is serial up to a given modal depths d then it should also be serial up to all modal depth smaller than d . The following proposition proves that it is indeed the case.

Proposition 7. *For all event models A and $a \in A$, if $n \geq n'$ then $\models \delta^n(a) \rightarrow \delta^{n'}(a)$.*

Proof. Let A be an event model and $a \in A$. We prove it by induction on n . If $n = 0$ or $n = 1$ then the result trivially holds. Assume it is true for a given $n \geq 1$. Assume $\models \delta^{n+1}(a)$, i.e. $\models \delta^0(a) \wedge \bigwedge_{j \in G} \hat{B}_j \bigvee_{b \in R_j(a)} \delta^n(b) \wedge \bigwedge_{j \in G} B_j \bigwedge_{b \in R_j(a)} (Pre(b) \rightarrow \delta^n(b))$.

By induction hypothesis, for all $b \in A$, $\models \delta^n(b) \rightarrow \delta^{n-1}(b)$. So

$$\begin{aligned} & \models \left(\delta^0(a) \wedge \bigwedge_{j \in G} \hat{B}_j \bigvee_{b \in R_j(a)} \delta^n(b) \wedge \bigwedge_{j \in G} B_j \bigwedge_{b \in R_j(a)} (Pre(b) \rightarrow \delta^n(b)) \right) \rightarrow \\ & \left(\delta^0(a) \wedge \bigwedge_{j \in G} \hat{B}_j \bigvee_{b \in R_j(a)} \delta^{n-1}(b) \wedge \bigwedge_{j \in G} B_j \bigwedge_{b \in R_j(a)} (Pre(b) \rightarrow \delta^{n-1}(b)) \right). \end{aligned}$$

i.e. $\models \delta^{n+1}(a) \rightarrow \delta^n(a)$. So for all $n' \leq n + 1$, $\models \delta^{n+1}(a) \rightarrow \delta^{n'}(a)$ by induction hypothesis.

Finally, we can strike some relationship between the seriality conditions for the entire updated model and for the generated submodels of the entire updated model. Indeed, if the entire updated model is serial then all its generated submodels should be serial up to any modal depth:

Proposition 8. *Let M be an epistemic model and A be a serial event model.*

$$\text{if } M \models \mathcal{S}(A) \text{ then for all } n \geq 0 \text{ } M \models \bigwedge_{a \in A} (Pre(a) \rightarrow \delta^n(a)).$$

Proof. By induction on n .

Besides, one can notice that the entire updated model is serial if and only if all its generated submodels are serial. But in fact, because we consider *all* the generated submodels, it suffices that these generated submodels be serial only up to modal depth 1. That is actually the intuition that led to the definition of $\mathcal{S}(A)$.

Proposition 9. *Let M be an epistemic model and let A be a serial event model. Then,*

$$M \models \mathcal{S}(A) \leftrightarrow O \left(\bigvee_{a \in A} Pre(a) \right) \wedge U \bigwedge_{a \in A} (Pre(a) \rightarrow \delta^1(a)).$$

$O \left(\bigvee_{a \in A} Pre(a) \right)$ expresses that the updated model is defined. The rest of the formula expresses its seriality. Note that $\delta^1(a) = Pre(a) \wedge \bigwedge_{j \in G} \hat{B}_j \bigvee_{b \in R_j(a)} Pre(b)$, so we have rediscovered the definition of $\mathcal{S}(A)$.

5 Conclusion

We have given conditions under which an entire updated model and its generated submodels are serial. We also introduced the notion of crazy formula which are formulas such that after being publicly announced at least one of the agents' beliefs become inconsistent. We could wonder whether other properties are also preserved during an update. It has been shown that most of the relevant ones like reflexivity, transitivity and euclidicity are preserved [3].

The fact that seriality is not preserved in the BMS system means that it should be enriched with some sort of revision mechanisms so that seriality is restored. For example, in Example 3 Ann should revise her beliefs about Bob after the public announcement. To do this, most of the existing approaches resort to a richer framework by introducing plausibility [1,6] or probability [2,4] but none of them tackles the issue directly in its original form by manipulating accessibility relations. This paper is a preliminary step in that direction.

References

1. Aucher, G.: A combined system for update logic and belief revision. In: Barley, M., Kasabov, N. (eds.) *PRIMA 2004. LNCS (LNAI)*, vol. 3371, pp. 1–17. Springer, Heidelberg (2005)
2. Aucher, G.: Interpreting an action from what we perceive and what we expect. *Journal of Applied Non-classical Logics* 17(1), 9–38 (2007)
3. Baltag, A., Moss, L.: Logic for epistemic program. *Synthese* 139(2), 165–224 (2004)
4. Baltag, A., Smets, S.: Probabilistic dynamic belief revision. In: van Benthem, J., Ju, S., Veltman, F. (eds.) *A Meeting of the Minds: Proceedings of the Workshop on Logic, Rationality and Interaction*, London. Computing Series, vol. 8, College Publications (2007)
5. Blackburn, P., de Rijke, M., Venema, Y.: *Modal Logic*. Cambridge Tracts in Computer Science, vol. 53. Cambridge University Press, Cambridge (2001)
6. van Ditmarsch, H.: Prolegomena to dynamic logic for belief revision. *Synthese* 147, 229–275 (2005)
7. van Ditmarsch, H., Kooi, B.: The secret of my success. *Synthese* 151, 201–232 (2006)