



Private announcement and belief expansion: an internal perspective

Guillaume Aucher

► To cite this version:

| Guillaume Aucher. Private announcement and belief expansion: an internal perspective. Journal of Logic and Computation, 2012. hal-00856474

HAL Id: hal-00856474

<https://inria.hal.science/hal-00856474>

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.

Private Announcement and Belief Expansion: an Internal Perspective

Guillaume Aucher,
University of Luxembourg,
6 rue Coudenhove-Kalergi,
L-1359 Luxembourg
Email: guillaume.aucher@uni.lu

Abstract

AGM belief revision theory and the BMS framework of dynamic epistemic logic both deal with the formalization of belief change, the former in a single agent setting and the latter in a multi-agent setting. In this paper we study the relation between these two formalisms. To be fair, we restrict our attention to the AGM operation of expansion since the original BMS framework does not allow for belief revision. Likewise, because AGM theory follows the internal approach, instead of the original BMS framework we define and deal with its internal version. This allows us to show that the AGM operation of expansion can naturally be viewed in the multi-agent setting of the BMS framework as the operation of private announcement, which goes against the claims of [16]. In parallel, we also provide conditions under which seriality of accessibility relations is preserved during an update, in the BMS framework as well as its internal version: it is a preliminary step towards the introduction of revision mechanisms into these frameworks.

Keywords: Belief revision theory, Dynamic epistemic logic, Internal approach, Private announcement, Seriality preservation

1 Introduction

The BMS framework of dynamic epistemic logic proposed by *Baltag, Moss and Solecki* deals with the logical study of belief change in a multi-agent setting [8, 17]. There is another logic-based approach to belief change, yet in a single-agent setting, namely AGM belief revision theory proposed by *Alchourrón, Gärdenfors and Makinson* [1, 10, 11]. It seems natural to set some connections between these two formalisms. Since the original BMS framework does not allow for revision of agents' beliefs, we will focus in this paper only on AGM theory restricted to the operation of expansion. This operation consists of adding a new piece of information to the beliefs of the agent and taking all the logical consequences (regardless of whether this new piece of information is in conflict with the initial beliefs of the agent). Besides, to provide a more rigorous and fair comparison of these two formalisms, we will deal with the *internal* version of the BMS framework. Indeed, one models in the BMS framework the epistemic states of *all* the agents from an external and objective point of view, whereas in AGM theory one models the epistemic state of the agent from her point of view. In other words, the BMS framework follows the perfect external approach whereas AGM theory follows the internal approach (see [6] for details about these internal and external approaches). So, we will first give the internal version of the BMS framework and this internal version will be considered as our multi-agent generalization of AGM theory (restricted to the operation of expansion).

In a single agent setting, new information can be conveyed to the agent in a unique way and this is formalized in AGM theory by means of the operation of expansion (when the new piece of information is not in conflict with the beliefs of the agent). On the other hand, in a multi-agent setting, new information can be conveyed to a given agent via an infinite number of ways: public announcement to all the agents, private announcement to a subgroup of agents (including the agent herself), private announcement to the agent only, and so on... This infinity of events can also be formalized in the multi-agent setting of the BMS framework (which thereby provides a precise formal meaning to each of them). So, what is the most natural event (as it is defined in the BMS framework)

generalizing the operation of expansion to the multi-agent setting of the BMS framework? Given our above assumptions, we will show that the AGM operation of expansion can naturally be viewed in the multi-agent setting of (the internal version of) the BMS framework as the operation of private announcement, which goes against the claims of [16] where it is argued that it rather corresponds to *public* announcement.

As we said, AGM theory, unlike the BMS framework, is typically concerned with revision: the issue of how to incorporate conflicting information into the beliefs of the agent. The goal is then to preserve consistency of the beliefs of the agent, and so with a minimal amount of change. In the multi-agent setting of the BMS framework, this problem of preserving consistency corresponds formally to the problem of preserving seriality of accessibility relations during an update. But to solve this problem one should first determine under which conditions seriality *is* indeed preserved during a BMS update because it would indicate in which case some sort of revision should be performed (and also some conditions that this revision should fulfill). We will consider not only under which conditions the updated model is serial but also under which conditions its generated submodels are serial. We will also give similar conditions for the internal version of the BMS framework (which will turn out to be relevant for our generalization of the expansion operation to the multi-agent case). Doing so, we will introduce the notion of crazy formulas, which are formulas such that after being publicly announced at least one of the agents' beliefs become inconsistent.

The paper is organized as follows. In Section 2, we recall the BMS framework. In Section 3, we give conditions under which seriality is preserved (for the entire updated model and for its generated submodels) and we introduce the notion of crazy formula.¹ In Section 4 we give the internal version of the BMS framework, extending to a dynamic setting the internal version of epistemic logic defined in [6]. We also give the conditions of seriality preservation for this new logical framework that we derive from the previous section. Finally, in Section 5 we show using our new framework developed in the previous section that the most natural counterpart of the AGM operation of expansion in the multi-agent setting of the BMS framework is private announcement. We also explain why this formal result is intuitively reasonable.

2 The BMS framework

The formalization of belief change in the BMS framework is done in three parts. First, the perception of the initial epistemic situation by the agents is represented by an epistemic model. Second, the perception of the event occurring in this situation is represented very similarly by an event model. Third, a product update takes as argument the epistemic model and the event model and yields a new epistemic model representing how the agents perceive the resulting situation after the event has occurred. These three parts correspond to our three subsections.

2.1 Epistemic model and generated submodel

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

2.1.1 Epistemic model

An epistemic model is just a particular kind of Kripke model [9] where instead of having a single accessibility relation we have a set of accessibility relations, one for each agent.

Definition 2.1. 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 $w \in W$. $|M|$ is the cardinality of W . \triangleleft

¹This section is based on [3]

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 2.2. Let $M = (W, R, V)$ be an epistemic model and $w_a \in M$.

- We say that M is *serial* when for all $j \in G$, for all $w \in W$, $R_j(w) \neq \emptyset$.
- We say that (M, w_a) is *serial up to depth n* when for all $j \in G$, for all $w \in \left(\bigcup_{j \in G} R_j\right)^{\leq n}(w_a)$, $R_j(w) \neq \emptyset$.²

□

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

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

where p ranges over Φ and j over G . Moreover, $\varphi \vee \varphi'$ is an abbreviation for $\neg(\neg\varphi \wedge \neg\varphi')$; $\varphi \rightarrow \varphi'$ is an abbreviation for $\neg\varphi \vee \varphi'$; $\langle B_j \rangle \varphi$ is an abbreviation for $\neg B_j \neg\varphi$; $E\varphi$ is an abbreviation for $\neg U \neg\varphi$; and \perp is an abbreviation for $\neg\top$. Finally, by \mathcal{L} we denote the language \mathcal{L}^U without the universal modality U .

The degree $\deg(\varphi)$ of a formula $\varphi \in \mathcal{L}$ is defined inductively as follows: $\deg(p) = 0$, $\deg(\neg\varphi) = \deg(\varphi)$, $\deg(\varphi \wedge \psi) = \max\{\deg(\varphi), \deg(\psi)\}$, $\deg(B_j\varphi) = \deg(\varphi) + 1$, $\deg(U\varphi) = \deg(\varphi) + 1$. □

Intuitively, $B_j\varphi$ means that agent j believes that 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 2.4. Let $M = (W, R, V)$ be an epistemic model and $w \in W$. $M, w \models \varphi$ is defined inductively as follows:

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

We write $M \models \varphi$ for $M, w \models \varphi$ for all $w \in M$. If C is a class of epistemic models, we write $\models_C \varphi$ when for all $M \in C$, $M \models \varphi$. Finally, we write $\models \varphi$ when for all epistemic model M , $M \models \varphi$. □

So agent j believes φ in world w (formally $M, w \models B_j\varphi$) if φ is true in all the worlds that agent j considers possible (in world w). $M, w \models U\varphi$ 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.

²If R is a relation and $n \in \mathbb{N}$, $R^{\leq n}$ is defined by $R^{\leq n}(w) = \{v \mid \text{there is } w = w_0, \dots, w_k = v \text{ such that } w_i R w_{i+1} \text{ and } k \leq n\}$ if $n > 0$, and $R^{\leq 0}(w) = \{w\}$.

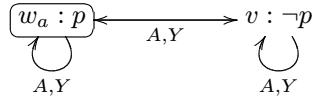


Figure 1: ‘Coin’ example.

Example 2.5. We take up more or less the coin example of [7]:

‘Ann and Yann enter a room where a quizmaster holds a coin in his hand. The quizmaster throws the coin in the air which lands in a small box on a table in front of the quizmaster. The quizmaster can see the coin but Ann and Yann cannot. The quizmaster then closes the box.’

This situation is modeled in the pointed epistemic model (M, w_a) of Figure 1. The quizmaster is not considered as an agent involved in the situation, so he is not represented in the epistemic model (he might simply be a robot or a computer program). The accessibility relations are represented by arrows indexed by A (standing for Ann) or Y (standing for Yann); p stands for ‘the coin is heads up’ and the boxed world w_a stands for the actual world. In any world, Yann and Ann consider the world where the coin is heads up and the world where the coin is tails up as being possible. So any world is accessible to any other world for A and Y . Now, thanks to the language \mathcal{L} we can express formally what is actually true in this situation. For example $M, w_a \models p \wedge (\langle B_A \rangle p \wedge \langle B_A \rangle \neg p) \wedge (\langle B_Y \rangle p \wedge \langle B_Y \rangle \neg p)$ means that the coin is heads up but both Ann and Yann do not know whether it is heads or tails up. $M, w_a \models B_Y (\langle B_A \rangle p \wedge \langle B_A \rangle \neg p) \wedge B_A (\langle B_Y \rangle p \wedge \langle B_Y \rangle \neg p)$ means that Yann believes that Ann does not know whether the coin is heads or tails up, and that Ann does so about Yann as well. \triangleleft

2.1.2 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 2.6. Let $M = (W, R, V)$ and $M' = (W', R', V')$ be two epistemic models and $w_a \in W$. Let $n \in \mathbb{N}$.

- 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* , written M_{w_a} , is the restriction of M to $\left(\bigcup_{j \in G} R_j\right)^*(w_a)$.³

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 .

\triangleleft

Proposition 2.7. Let $M = (W, R, V)$ be an epistemic model and $w_a \in W$. Then for all $w \in M_{w_a}$ and all $\varphi \in \mathcal{L}$, $M, w \models \varphi$ iff $M_{w_a}, w \models \varphi$.

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 . Finally, note that we also have that (M, w_a) is serial up to depth n if and only if M_{w_a} is serial up to depth n .

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

³if R is a relation, R^+ is defined by $R^+(w) = \{v \mid \text{there is } w = w_0, \dots, w_n = v \text{ such that } w_i R w_{i+1}\}$. R^* is defined by $R^*(w) = \{w\} \cup R^+(w)$. See [9].

event can also be described in terms of beliefs: for example, while the quizmaster announces Yann that the coin is heads up (event w_a) Ann *believes* that nothing happens (event v). This leads them to define the notion of event model whose definition is very similar to that of an epistemic model.

Definition 2.8. An *event model* M is a triple $M = (W, R, \text{Pre})$ such that

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

If $M = (W, R, \text{Pre})$ is an event model, a pair (M, w_a) where $w_a \in W$ is called a *pointed event model*. We also write $R_j = R(j)$ and $R_j(w) = \{v \in W \mid wR_j v\}$, and $w \in M$ for $w \in W$. $|M|$ is the cardinality of W . \triangleleft

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 2.9. 1. Assume that the quizmaster announces privately to Yann that the coin is heads up. This event is depicted in Figure 2. w_a stands for the event ‘the quizmaster truthfully announces that the coin is heads up’ and v stands for the event ‘nothing happens’. The boxed event corresponds to the actual event. So while the quizmaster announces to Yann that the coin is heads up (w_a), Ann believes that nothing happens (v): this explains the accessibility relation indexed by A between w_a and v . The precondition for w_a is that the coin is indeed heads up (p) while the precondition for v is any tautology (like \top) because the event where nothing happens can take place in any world.

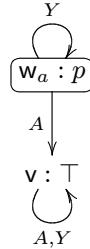


Figure 2: Private announcement of p to Yann

2. Assume that the quizmaster announces publicly that Yann believes that the coin is heads up. This event is depicted in Figure 3. There, w_a stands for ‘the quizmaster truthfully announces that Yann believes that the coin is heads up’. Because this event is correctly perceived by Ann and Yann, w_a is the only event considered possible by them. Finally, for this announcement to be made in a possible world, Yann has to believe that the coin is heads up in this world ($B_Y p$).

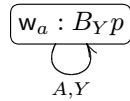


Figure 3: Public announcement of $B_Y p$

\triangleleft

2.3 Product Update

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

Definition 2.10. Let $M = (W, R, V)$ be an epistemic model and $\mathbf{M} = (\mathbf{W}, \mathbf{R}, \text{Pre})$ an event model. We define their *update product* to be the epistemic model $M \otimes \mathbf{M} = (W', R', V')$ where

1. $W' = \{(w, w) \in W \times \mathbf{W} \mid M, w \models \text{Pre}(w)\};$
2. $(v, v) \in R'_j(w, w)$ iff $v \in R_j(w)$ and $v \in \mathbf{R}_j(w)$;
3. $V'(p) = \{(w, w) \in W \otimes \mathbf{W} \mid w \in V(p)\}.$

□

Example 2.11. 1. Assume that the quizmaster privately announces to Yann that the coin is heads up. This scenario is depicted in the pointed epistemic model of Figure 4. In the resulting epistemic model, Yann knows that the coin is heads up but Ann does not know whether it is either heads or tails up and believes Yann does not know either.

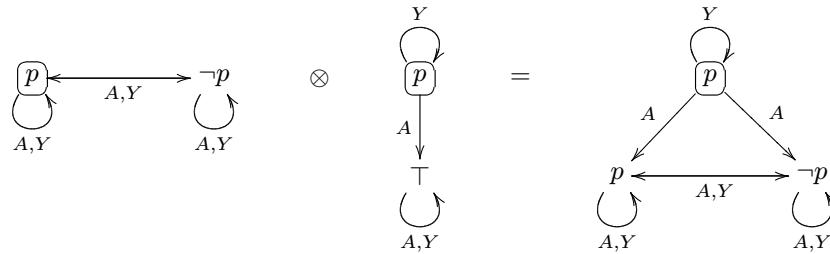


Figure 4: ‘Coin’ scenario

2. This second example shows that seriality might not be preserved during an update. Assume now that the quizmaster announces truthfully and publicly that Yann believes that the coin is heads up (formally B_{Yp}). If we update the resulting epistemic model of Figure 4 by the truthful public announcement that Yann believes that the coin is heads up (formally B_{Yp}) depicted in Figure 3 then we get the epistemic model depicted on the right of Figure 5 where Ann’s accessibility relation is not serial.

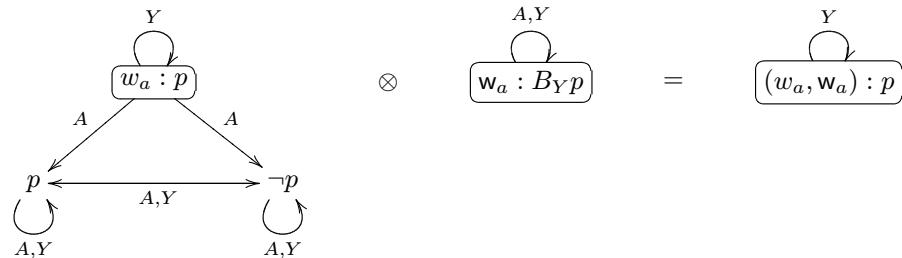


Figure 5: Failure of seriality preservation

□

3 On seriality preservation

Given a (serial) event model, we are going to give conditions that an epistemic model should fulfill so that its updated model and also the generated submodels of this updated model be serial (up to a given depth).

3.1 Seriality of the entire updated model

3.1.1 Theory

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

Proposition 3.1. *Let M be a serial event model and let M be an epistemic model.*

$$M \otimes M \text{ is defined and serial iff } M \models E \left(\bigvee_{w \in M} \text{Pre}(w) \right) \wedge U \bigwedge_{w \in M} \left(\text{Pre}(w) \rightarrow \bigwedge_{j \in G} \langle B_j \rangle \bigvee_{v \in R_j(w)} \text{Pre}(v) \right).$$

Proof. $M \models E \left(\bigvee_{w \in M} \text{Pre}(w) \right)$ clearly means that the model $M \otimes \mathbf{M}$ is defined. Now it remains to prove that $M \otimes \mathbf{M}$ is serial iff $M \models U \bigwedge_{w \in M} \left(\text{Pre}(w) \rightarrow \bigwedge_{j \in G} \langle B_j \rangle \bigvee_{v \in R_j(w)} \text{Pre}(v) \right)$.

- Assume that $M \models U \bigwedge_{w \in M} \left(\text{Pre}(w) \rightarrow \bigwedge_{j \in G} \langle B_j \rangle \bigvee_{v \in R_j(w)} \text{Pre}(v) \right)$ (*). Let $(w, \mathbf{w}) \in M \otimes \mathbf{M}$ and $j \in G$. Then $M, w \models \text{Pre}(w)$. So $M, w \models \bigwedge_{j \in G} \langle B_j \rangle \bigvee_{v \in R_j(w)} \text{Pre}(v)$ by (*). Then $M, w \models \langle B_j \rangle \bigvee_{v \in R_j(w)} \text{Pre}(v)$. So there is $v \in R_j(w)$ and $\mathbf{v} \in R_j(\mathbf{w})$ such that $M, v \models \text{Pre}(v)$. Therefore there is $(v, \mathbf{v}) \in M \otimes \mathbf{M}$ such that $(v, \mathbf{v}) \in R_j(w, \mathbf{w})$ by definition of $M \otimes \mathbf{M}$. So $M \otimes \mathbf{M}$ is serial.
- Assume that $M \not\models U \bigwedge_{w \in M} \left(\text{Pre}(w) \rightarrow \bigwedge_{j \in G} \langle B_j \rangle \bigvee_{v \in R_j(w)} \text{Pre}(v) \right)$. Then there is $w \in M$ and $\mathbf{w} \in \mathbf{M}$ such that $M, w \models \text{Pre}(w) \wedge \left(\bigvee_{j \in G} \langle B_j \rangle \bigwedge_{v \in R_j(w)} \neg \text{Pre}(v) \right)$. Then there is $j \in G$ such that $M, w \models B_j \bigwedge_{v \in R_j(w)} \neg \text{Pre}(v)$ (**). So $(w, \mathbf{w}) \in M \otimes \mathbf{M}$ but there is no $v \in R_j(w)$ and $\mathbf{v} \in R_j(\mathbf{w})$ such that $(v, \mathbf{v}) \in R_j(w, \mathbf{w})$. Indeed, otherwise we would have $M, w \models \langle B_j \rangle \bigvee_{v \in R_j(w)} \text{Pre}(v)$, which contradicts (**). So $M \otimes \mathbf{M}$ is not serial.

□

Given a (serial) event model, this proposition gives conditions that an epistemic model should fulfill so that this epistemic model updated by the event model be defined and serial. We write

$$\mathcal{S}(\mathbf{M}) = E \left(\bigvee_{w \in M} \text{Pre}(w) \right) \wedge U \bigwedge_{w \in M} \left(\text{Pre}(w) \rightarrow \bigwedge_{j \in G} \langle B_j \rangle \bigvee_{v \in R_j(w)} \text{Pre}(v) \right). E \left(\bigvee_{w \in M} \text{Pre}(w) \right) \text{ expresses that}$$

the updated model $M \otimes \mathbf{M}$ is defined and $U \bigwedge_{w \in M} \left(\text{Pre}(w) \rightarrow \bigwedge_{j \in G} \langle B_j \rangle \bigvee_{v \in R_j(w)} \text{Pre}(v) \right)$ expresses that the updated model $M \otimes \mathbf{M}$ is serial. Note that the seriality condition is evaluated only on the epistemic model M and that M does not need to be serial for the proposition to hold. However, if the event model is not serial then the product update does not yield a defined and serial updated model.

Example 3.2. Consider the event model \mathbf{M} depicted in Figure 6 corresponding to the private announcement to Y that φ holds, where φ is an epistemic formula. We then have $\mathcal{S}(\mathbf{M}) = E(\varphi \vee \top) \wedge U((\varphi \rightarrow (\langle B_A \rangle \top \wedge \langle B_Y \rangle \varphi)) \wedge (\top \rightarrow (\langle B_A \rangle \top \wedge \langle B_Y \rangle \top)))$ which is provably equivalent to $(\varphi \rightarrow \langle B_Y \rangle \varphi) \wedge \langle B_A \rangle \top \wedge \langle B_Y \rangle \top$. So, according to Proposition 3.1, for any epistemic model M , $M \otimes \mathbf{M}$ is defined and serial if and only if $M \models (\varphi \rightarrow \langle B_Y \rangle \varphi) \wedge \langle B_A \rangle \top \wedge \langle B_Y \rangle \top$. In words, given an epistemic model M , the entire updated model $M \otimes \mathbf{M}$ is serial if and only if M is already serial ($M \models \langle B_A \rangle \top \wedge \langle B_Y \rangle \top$) and in case a world of M satisfies φ then there must be another world accessible by Y from this world which satisfies φ ($M \models \varphi \rightarrow \langle B_Y \rangle \varphi$). □

From the above proposition we can easily prove the following corollary.

Corollary 3.3. *Let C be a class of epistemic models and M a serial event model.*

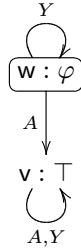


Figure 6: Private announcement of φ to agent Y

$\models_C \neg S(M)$ iff there is no epistemic model $M \in C$ such that $M \otimes M$ is defined and serial.

In other words this corollary tells us under which condition, for a given event model M , whatever epistemic model M we chose, $M \otimes M$ 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).

3.1.2 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 3.4. A *crazy formula* is a satisfiable formula $\varphi \in \mathcal{L}$ such that

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

□

Proposition 3.5. Let φ be a crazy formula and let M be the event model corresponding to the public announcement of φ . Then there is no epistemic model M such that $M \otimes M$ is defined and serial.

Proof. Thanks to Corollary 3.3 it suffices to prove that $\models \neg S(M)$ i.e. $\models E\varphi \rightarrow E\left(\varphi \wedge \bigvee_{j \in G} B_j \neg \varphi\right)$

because $S(M) = E\varphi \wedge U\left(\varphi \rightarrow \bigwedge_{j \in G} \langle B_j \rangle \varphi\right)$. Let M be an epistemic model such that $M \models E\varphi$. Let $w \in M$ such that $M, w \models \varphi$. Then by definition of a crazy formula $M, w \models \varphi \wedge \bigvee_{j \in G} B_j \neg \varphi$. Then

$M, w \models \varphi \wedge \bigvee_{j \in G} B_j \neg \varphi$ i.e. $M \models E\left(\varphi \wedge \bigvee_{j \in G} B_j \neg \varphi\right)$. □

Proposition 3.6. $\varphi = \psi \wedge B_i \neg \psi$, where $\psi \in \mathcal{L}$ is a satisfiable formula, is a crazy formula.

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

We can compare this notion of crazy formula with the notion of selfrefuting and successful formulas studied in [15]. 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.

3.2 Seriality of 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, w_a) , because this submodel captures all the information relevant to model the agents' beliefs in the new actual world (w_a, w_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 3.7. Let M be an event model, $w \in M$ and $n \in \mathbb{N}$. We define $\delta^n(w)$ inductively as follows.

- $\delta^0(w) = \text{Pre}(w) \wedge \bigwedge_{j \in G} \langle B_j \rangle \bigvee_{v \in R_j(w)} \text{Pre}(v);$
- $\delta^{n+1}(w) = \text{Pre}(w) \wedge \bigwedge_{j \in G} \langle B_j \rangle \bigvee_{v \in R_j(w)} \delta^n(v) \wedge \bigwedge_{j \in G} B_j \bigwedge_{v \in R_j(w)} (\text{Pre}(v) \rightarrow \delta^n(v)).$

□

As we are now going to show, $M, w \models \delta^n(w)$ means that $(M \otimes M, (w, w))$ is defined and serial up to depth n . In other words, it means that $(M \otimes M)_{(w, w)}$ is defined and serial up to depth n . We first need to prove two lemmas.

Lemma 3.8. Let M be an epistemic model and let M be an event model. For all $w \in M$, $w \in M$, $n \in \mathbb{N}$,

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

$M, w \models \delta^0(w)$ 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 $w = w_0 R_{j_1} w_1 R_{j_2} \dots R_{j_n} w_n = v$ such that for all $i \in \{0, \dots, n\}$, $M, w_i \models \text{Pre}(w_i)$,

$$M, v \models \bigwedge_{j \in G} \langle B_j \rangle \bigvee_{u \in R_j(v)} \text{Pre}(u)$$

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 $M, w \models \delta^{n+2}(w)$. Then $M, w \models \delta^0(w)$ because $\models \delta^{n+1}(v) \rightarrow \text{Pre}(v)$ and $\delta^0(w) = \text{Pre}(w) \wedge \bigwedge_{j \in G} \langle B_j \rangle \bigvee_{v \in R_j(w)} \text{Pre}(v)$. 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 $w = w_0 R_{j_1} w_1 R_{j_2} \dots R_{j_{n+1}} w_{n+1} = v$ such that for all $i \in \{0, \dots, n+1\}$, $M, w_i \models \text{Pre}(w_i)$.

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

Besides $w_1 \in R_{j_1}(w)$ and $M, w_1 \models \text{Pre}(w_1)$. So $M, w_1 \models \delta^{n+1}(w_1)$.

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 $w_1 = w'_1 R_{j_2} \dots R_{j_{n+1}} w'_{n+1} = v'$ such that for all i , $M, w'_i \models \text{Pre}(w'_i)$,

$$M, v' \models \bigwedge_{j \in G} \langle B_j \rangle \bigvee_{v' \in R_j(w')} \text{Pre}(v').$$

So $M, v \models \bigwedge_{j \in G} \langle B_j \rangle \bigvee_{u \in R_j(v)} \text{Pre}(u)$.

- Assume $M, w \models \delta^0(w)$ 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 $w = w_0 R_{j_1} \dots R_{j_n} w_{n+1} = v$ such that for all i , $M, w_i \models \text{Pre}(w_i)$,

$$M, v \models \bigwedge_{j \in G} \langle B_j \rangle \bigvee_{u \in R_j(v)} \text{Pre}(u).$$

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

$$\left(\bigvee_{j \in G} \langle B_j \rangle \bigvee_{v \in R_j(w)} (\text{Pre}(v) \wedge \neg \delta^{n+1}(v)) \right).$$

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

Then for all $v \in R_i(w)$ and all $v \in R_i(w)$, $M, v \models \neg \delta^{n+1}(v)$ (*).

But by assumption $M, w \models \delta^0(w)$, i.e. $M, w \models \text{Pre}(w) \wedge \bigwedge_{j \in G} \langle B_j \rangle \bigvee_{v \in R_j(w)} \text{Pre}(v)$. Then $M, w \models \langle B_i \rangle \bigvee_{v \in R_i(w)} \text{Pre}(v)$, i.e. there is $v \in R_i(w)$ and $v \in R_i(w)$ such that $M, v \models \text{Pre}(v)$ (1).

So $M, v \models \bigwedge_{j \in G} \langle B_j \rangle \bigvee_{v \in R_j(w)} \text{Pre}(v)$ (2) by assumption (take $w_1 = \dots = w_n = v$ and $w_1 = \dots = w_n = v$).

Then by (1) and (2) we get $M, v \models \delta^0(v)$.

Besides, by assumption and because wR_iv and wR_iv , for all u such that $v = v_0R_{j_1} \dots R_{j_n}u$ such that there are $v = v_0R_{j_1} \dots R_{j_n}v_n = u$ such that for all i $M, v_i \models \text{Pre}(v_i)$,

$$M, u \models \bigwedge_{j \in G} \langle B_j \rangle \bigvee_{t \in R_j(u)} \text{Pre}(t).$$

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

- Assume $M, w \models \bigvee_{j \in G} \langle B_j \rangle \bigvee_{v \in R_j(w)} (\text{Pre}(v) \wedge \neg \delta^{n+1}(v))$.

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

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

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

□

Lemma 3.9. Let M be an epistemic model and M be a serial event model. For all $w \in M$ and $w \in M$ such that $M, w \models \text{Pre}(w)$,

1. $R_j(w, w) \neq \emptyset$ for all $j \in G$ iff $M, w \models \bigwedge_{j \in G} \langle B_j \rangle \bigvee_{v \in R_j(w)} \text{Pre}(v)$;

2. If $n > 0$, $(v, v) \in \left(\bigcup_{j \in G} R_j \right)^{\leq n} (w, w)$ iff there are $w = w_0R_{j_1}w_1R_{j_2} \dots R_{j_k}w_k = v$ with $k \leq n$ and

$$w = w_0R_{j_1}w_1R_{j_2} \dots R_{j_k}w_k = v \text{ such that for all } i, M, w_i \models \text{Pre}(w_i).$$

Proof. 1. Assume $M, w \models \bigwedge_{j \in G} \langle B_j \rangle \bigvee_{v \in R_j(w)} \text{Pre}(v)$. Then for all $j \in G$, there is $v \in R_j(w)$ and $v \in R_j(w)$ such that $M, v \models \text{Pre}(v)$. Then, by definition of the product update, for all j , there is $(v, v) \in M \otimes M$ such that $(v, v) \in R_j(w, w)$. So for all $j \in G$, $R_j(w, w) \neq \emptyset$.

Assume $M, w \not\models \bigwedge_{j \in G} \langle B_j \rangle \bigvee_{v \in R_j(w)} \text{Pre}(v)$. Then there is $j \in G$ such that for all $v \in R_j(w)$ and for all

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

2. If $n > 0$, $(v, v) \in \left(\bigcup_{j \in G} R_j \right)^{\leq n} (w, w)$

iff there are j_1, \dots, j_n and $(w_1, w_1), \dots, (w_n, w_n) \in M \otimes M$ such that $(w, w)R_{j_1}(w_1, w_1)R_{j_2} \dots R_{j_n}(w_n, w_n) = (v, v)$

iff there are $w = w_0R_{j_1}w_1R_{j_2} \dots R_{j_n}w_n = v$ and $w = w_0R_{j_1}w_1R_{j_2} \dots R_{j_n}w_n = v$ such that for all i , $M, w_i \models \text{Pre}(w_i)$.

□

Theorem 3.10. Let M be an epistemic model and let \mathbf{M} be a serial event model. Let $w \in M, w \in \mathbf{M}$ and $n \in \mathbb{N}$.

$$(M \otimes \mathbf{M})_{(w,w)} \text{ is defined and serial up to depth } n \text{ iff } M, w \models \delta^n(w).$$

Proof. First, note that $(M \otimes \mathbf{M})_{(w,w)}$ is defined and serial up to depth n iff

- (w,w) is defined;

- $R_j(v,v) \neq \emptyset$ for all $(v,v) \in \left(\bigcup_{j \in G} R_j\right)^{\leq n} (w,w)$ and for all $j \in G$.

Then we get easily the expected result by Lemma 3.9 and Lemma 3.8. Indeed, (w,w) is defined and $R_j(w,w) \neq \emptyset$ for all $j \in G$ amounts to say that $M, w \models \delta^0(w)$. And if $n > 0$, $R_j(v,v) \neq \emptyset$ for all $(v,v) \in \left(\bigcup_{j \in G} R_j\right)^{\leq n} (w,w)$ and for all $j \in G$ amounts to say that 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 $w = w_0 R_{j_1} w_1 R_{j_2} \dots R_{j_n} w_n = v$ such that for all $i \in \{0, \dots, n\}$, $M, w_i \models \text{Pre}(w_i)$, $M, v \models \bigwedge_{j \in G} \langle B_j \rangle \bigvee_{u \in R_j(v)} \text{Pre}(u)$. \square

Note that because $\deg(\delta^n(w)) = n + 1$, to determine whether an epistemic model M updated by a given event model \mathbf{M} is serial up to depth n , it suffices to consider the structure of this epistemic model M only up to depth $n + 1$.

Corollary 3.11. Let M be a finite epistemic model and let \mathbf{M} be a serial event model. Let $w \in M, w \in \mathbf{M}$ and $n = |M| \cdot |\mathbf{M}| - 1$.

$$(M \otimes \mathbf{M})_{(w,w)} \text{ is defined and serial iff } M, w \models \delta^n(w).$$

The above corollary stems from the fact that if n is larger than the depth of $(M \otimes \mathbf{M})_{(w,w)}$ (which is the case if $n = |M| \cdot |\mathbf{M}| - 1$) then all the worlds accessible from (w,w) are serial. So $(M \otimes \mathbf{M})_{(w,w)}$ is indeed serial. Accordingly, this also entails that it should also be serial for any given modal depth. That is what the following proposition expresses.

Proposition 3.12. Let M be a finite epistemic model and let \mathbf{M} be a serial event model. Let $w \in M, w \in \mathbf{M}$ and $n = |M| \cdot |\mathbf{M}| - 1$.

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

Proof. The proof follows from Lemma 3.8 and the fact that for all $v \in M$ there are w_1, \dots, w_{n-1} such that $w = w_0 R_{j_1} w_1 R_{j_2} \dots R_{j_n} w_n = v$ iff there are w_1, \dots, w_{m-1} such that $w = w_0 R_{j_1} w_1 R_{j_2} \dots R_{j_m} w_m = v$. \square

Similarly, if $(M \otimes \mathbf{M})_{(w,w)}$ is serial up to a given depth n then it should also be serial up to all depths smaller than n . The following proposition proves that it is indeed the case.

Proposition 3.13. For all event models \mathbf{M} and $w \in \mathbf{M}$, if $n \geq n'$ then $\models \delta^n(w) \rightarrow \delta^{n'}(w)$.

Proof. Let \mathbf{M} be an event model and $w \in \mathbf{M}$. We prove it by induction on n . If $n = 1$ then one can easily check that the result holds. Assume it is true for a given $n \geq 1$. By induction hypothesis, for all $v \in \mathbf{M}$, $\models \delta^n(v) \rightarrow \delta^{n-1}(v)$. So

$$\begin{aligned} &\models \left(\delta^0(w) \wedge \bigwedge_{j \in G} \langle B_j \rangle \bigvee_{v \in R_j(w)} \delta^n(v) \wedge \bigwedge_{j \in G} B_j \bigwedge_{v \in R_j(w)} (\text{Pre}(v) \rightarrow \delta^n(v)) \right) \rightarrow \\ &\quad \left(\delta^0(w) \wedge \bigwedge_{j \in G} \langle B_j \rangle \bigvee_{v \in R_j(w)} \delta^{n-1}(v) \wedge \bigwedge_{j \in G} B_j \bigwedge_{v \in R_j(w)} (\text{Pre}(v) \rightarrow \delta^{n-1}(v)) \right). \end{aligned}$$

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

Example 3.14. Consider the pointed event model (M, w) of Figure 6. We list below the first $\delta^n(w)$.

- $\delta^0(w) = \varphi \wedge \langle B_A \rangle \top \wedge \langle B_Y \rangle \varphi;$
- $\delta^1(w) = \varphi \wedge \langle B_A \rangle (\langle B_A \rangle \top \wedge \langle B_Y \rangle \top) \wedge \langle B_Y \rangle (\varphi \wedge \langle B_A \rangle \top \wedge \langle B_Y \rangle \varphi) \wedge B_A (\langle B_A \rangle \top \wedge \langle B_Y \rangle \top) \wedge B_Y (\varphi \rightarrow \langle B_A \rangle \top \wedge \langle B_Y \rangle \varphi);$
- $\delta^2(w) = \varphi \wedge \langle B_Y \rangle (\varphi \wedge \langle B_A \rangle (\langle B_A \rangle \top \wedge \langle B_Y \rangle \top) \wedge \langle B_Y \rangle (\varphi \wedge \langle B_A \rangle \top \wedge \langle B_Y \rangle \varphi) \wedge B_A (\langle B_A \rangle \top \wedge \langle B_Y \rangle \top) \wedge B_Y (\varphi \rightarrow \langle B_A \rangle \top \wedge \langle B_Y \rangle \varphi)) \wedge \langle B_A \rangle (\langle B_A \rangle (\langle B_A \rangle \top \wedge \langle B_Y \rangle \top) \wedge \langle B_Y \rangle (\langle B_A \rangle \top \wedge \langle B_Y \rangle \top) \wedge B_A (\langle B_A \rangle \top \wedge \langle B_Y \rangle \top) \wedge B_Y (\langle B_A \rangle \top \wedge \langle B_Y \rangle \top)) \wedge B_Y (\varphi \rightarrow \langle B_A \rangle (\langle B_A \rangle \top \wedge \langle B_Y \rangle \top) \wedge \langle B_Y \rangle (\varphi \wedge \langle B_A \rangle \top \wedge \langle B_Y \rangle \varphi) \wedge B_A (\langle B_A \rangle \top \wedge \langle B_Y \rangle \top) \wedge B_Y (\varphi \rightarrow \langle B_A \rangle \top \wedge \langle B_Y \rangle \varphi)) \wedge B_A (\langle B_A \rangle (\langle B_A \rangle \top \wedge \langle B_Y \rangle \top) \wedge \langle B_Y \rangle (\langle B_A \rangle \top \wedge \langle B_Y \rangle \top) \wedge B_A (\langle B_A \rangle \top \wedge \langle B_Y \rangle \top) \wedge B_Y (\langle B_A \rangle \top \wedge \langle B_Y \rangle \top)).$

One can check that $\models \delta^1(w) \rightarrow \delta^0(w)$ and $\models \delta^2(w) \rightarrow \delta^1(w)$, which is in line with Proposition 3.13. These formulas can be simplified depending on the logic one considers; Proposition 5.6 in Section 5.2 shows that in the internal logic of [6], all the above $\delta^n(w)$ are actually provably equivalent to φ . \triangleleft

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 depth:

Proposition 3.15. *Let M be an epistemic model and M be a serial event model.*

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

Proof. By induction on n . \square

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 depth 0. That is actually the intuition that led to the definition of $\mathcal{S}(M)$.

Proposition 3.16. *Let M be a serial event model. Then,*

$$\models \mathcal{S}(M) \leftrightarrow E \left(\bigvee_{w \in M} \text{Pre}(w) \right) \wedge U \bigwedge_{w \in M} (\text{Pre}(w) \rightarrow \delta^0(w)).$$

$E \left(\bigvee_{w \in M} \text{Pre}(w) \right)$ expresses that the updated model is defined. The rest of the formula expresses its seriality. Note that $\delta^0(w) = \text{Pre}(w) \wedge \bigwedge_{j \in G} \langle B_j \rangle \bigvee_{v \in R_j(w)} \text{Pre}(v)$, so we have rediscovered the definition of $\mathcal{S}(M)$.

4 The internal version of the BMS framework

One can easily notice that in BMS, just as in standard epistemic logic (and also game theory), the epistemic scenarios are modeled from an external and objective point of view. Therefore the BMS framework follows the perfect external approach as described in [6]. So, it seems natural to give the internal version of the BMS framework just as we gave the internal version of epistemic logic in [6]. This new framework would model the epistemic state of *one* of the agents (called agent Y in this paper) from her point of view in a dynamic and multi-agent setting. Since the BMS framework is based on epistemic logic, it just remains to give the internal version of its dynamic part. In Section 4.1 we briefly recall the internal logic of [6]; for more details and motivations the reader should refer to [6]. In Sections 4.2 and 4.3 we complement it with dynamics.

4.1 Multi-agent possible world and internal model

In the internal approach, the modeler is one of the agents involved in the epistemic situation. The situation is modeled from the point of view of this agent, called henceforth agent Y . Both in philosophy and computer science there are formalizations of the internal point of view. Perhaps one of the dominant formalisms for this is auto-epistemic logic [13, 14]. In philosophy there are some models of full belief like the one offered by Levi [12] which is also related to ideas in auto-epistemic logic. Formally, there is a completeness theorem for a notion of full belief of the type proposed by Levi proved by Arlo-Costa in [2] (which contains some ideas about the connections between external and internal models of belief as well). Another typical example of logical formalism which follows the internal approach is AGM theory.

However, all the above formalizations of the internal approach deal with only one agent Y . The possible worlds there are supposed to represent how this agent Y perceives the surrounding world. As she is the only agent, they deal only with propositional facts about the surrounding world. Now, if we suppose that there are other agents than agent Y , a possible world for Y in that case should also deal with how the other agents perceive the surrounding world. These “multi-agent” possible worlds should then not only deal with propositional facts but also with epistemic facts. So to represent a multi-agent possible world we need to add a modal structure to our (single agent) possible worlds. We do so as follows.

Definition 4.1. A *multi-agent possible world* (M, w) is a *finite* pointed epistemic model $M = (W, R, V, w)$ generated by $w \in W$ such that R_j is serial, transitive and euclidean for all $j \in G$, and

1. $R_Y(w) = \{w\}$;
2. there is no v and $j \neq Y$ such that $w \in R_j(v)$.

□

Let us have a closer look at the definition. Condition 2 will be motivated later (after Definition 4.5), but one can actually show that any pointed epistemic model satisfying the conditions of a multi-agent possible world except condition 2 is bisimilar to a multi-agent possible world. Condition 1 ensures that in case Y is the only agent then a multi-agent possible world boils down to a possible world, as in AGM theory. Condition 1 also ensures that in case Y assumes that the situation is correctly represented by the multi-agent possible world (M, w) then for her w is the (only) actual world. In fact the other possible worlds of a multi-agent possible world are just present for technical reasons: they express the other agents’ beliefs (in world w). One could get rid of the condition that a multi-agent possible world (M, w) is generated by w but the worlds which do not belong to the submodel generated by w would have neither philosophical nor technical motivation. Finally, notice that we assume that accessibility relations are serial, transitive and euclidean. This means that the agents’ beliefs are consistent and that agents ‘know’ what they believe and disbelieve. These seem to be very natural constraints to impose on the notion of belief. Intuitively, this notion of belief corresponds for example to the kind of belief in a theorem that you have after having proved this theorem and checked the proof several times.

Example 4.2. We see in Figure 7 that a multi-agent possible world is really a generalization of a possible world. □

In the single agent case (in AGM belief revision theory), the epistemic state of the agent Y is represented by a finite set of possible worlds which expresses the fact that the agent might have some uncertainty about the situation. In a multi-agent setting, this is very similar: the epistemic state of the agent Y is represented by a (disjoint and) finite set of *multi-agent* possible worlds.

Definition 4.3. An *internal model* is a disjoint and finite union of multi-agent worlds. □

An internal model will sometimes be noted (\mathcal{M}, W_a) where W_a are the roots of its multi-agent possible worlds.

Example 4.4. In Figure 8 is depicted an example of internal model. This internal model represents how the situation described in Example 2.5 is perceived by Yann. Yann does not know whether the coin

a (single-agent) possible world:

$$[w : p, \neg q]$$

a multi-agent possible world:

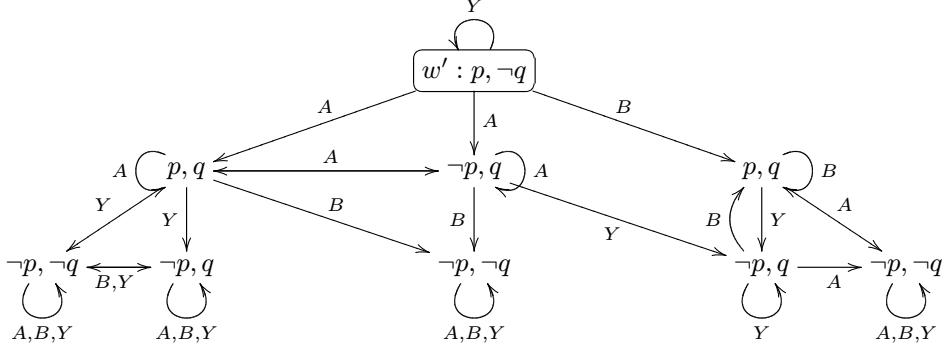


Figure 7: From possible world to multi-agent possible world

is heads or tails up (formally $\neg B_Y p \wedge \neg B_Y \neg p$). Indeed, in one multi-agent possible world (on the left) p is true at the root and in the other (on the right) p is false at the root. Yann also believes that Ann does not know whether p is true or false (formally $B_Y(\neg B_A p \wedge \neg B_A \neg p)$). Indeed, in both multi-agent possible worlds, $\neg B_A p \wedge \neg B_A \neg p$ is true at the roots. Finally, Yann believes that Ann believes that he does not know whether p is true or false (formally $B_Y B_A(\neg B_Y p \wedge \neg B_Y \neg p)$) since $B_A(\neg B_Y p \wedge \neg B_Y \neg p)$ is true at the roots of both multi-agent possible worlds. \triangleleft

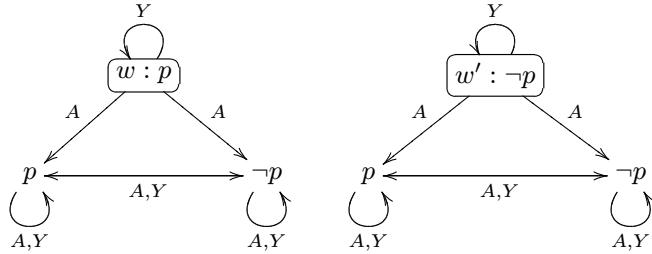


Figure 8: An internal model

The well-formed formulas of the language for the internal approach are identical to the ones of the epistemic language of Definition 2.3 (without universal modality). Its truth conditions are nevertheless a bit different and are set out below.

Definition 4.5. Let $(\mathcal{M}, \{w^1, \dots, w^n\}) = \{(M^1, w^1), \dots, (M^n, w^n)\}$ be an internal model and let $w \in \mathcal{M}$. Then $w \in M^k$ for some k , with $M^k = (W^k, R^k, V^k)$. $\mathcal{M}, w \models \varphi$ is defined inductively as follows:

- $\mathcal{M}, w \models \top$
- $\mathcal{M}, w \models p \quad \text{iff} \quad w \in V^k(p)$
- $\mathcal{M}, w \models \neg \varphi \quad \text{iff} \quad \text{not } \mathcal{M}, w \models \varphi$
- $\mathcal{M}, w \models \varphi \wedge \varphi' \quad \text{iff} \quad \mathcal{M}, w \models \varphi \text{ and } \mathcal{M}, w \models \varphi'$
- $\mathcal{M}, w \models B_Y \varphi \quad \text{iff} \quad \begin{cases} \text{for all } w^i \in W_a, \mathcal{M}, w^i \models \varphi & \text{if } w \in W_a \\ \text{for all } w' \in R_Y^k(w), \mathcal{M}, w' \models \varphi & \text{if } w \notin W_a \end{cases}$
- $\mathcal{M}, w \models B_j \varphi \quad \text{iff} \quad \text{for all } w' \in R_j^k(w), \mathcal{M}, w' \models \varphi \quad \text{if } j \neq Y$

We say that φ is *true* in (\mathcal{M}, W_a) and write $\mathcal{M}, W_a \models \varphi$ when $\mathcal{M}, w \models \varphi$ for all $w \in W_a$. \triangleleft

Note that the truth condition for the operator B_Y is defined as if there were accessibility relations indexed by Y between the roots of the multi-agent possible worlds. Therefore we could actually set an accessibility relation indexed by Y between the roots of an internal model. This would lead us to define the notion of *internal model of type 2*, which is a multi-pointed epistemic model (\mathcal{M}, W_a) whose accessibility relations are serial, transitive and euclidean and such that for all $w_a \in W_a$, called the *actual equivalence class*, $R_Y(w_a) = W_a$. One can easily show that the notions of internal model and internal model of type 2 are actually equivalent.⁴ As it turns out, the internal model of Figure 8 is an equivalent representation of the internal model of type 2 of Figure 1 (with $W_a = \{w_a, v\}$): Yann perceives correctly the situation.

Nevertheless, we stick to our current representation of internal models because it will enable us to show that generalizing AGM expansion to the multi-agent case amounts to study private announcement. We can now motivate condition 2 in light of the above truth conditions. Condition 2 warrants modularity of multi-agent possible worlds: agents j 's beliefs about agent Y 's beliefs (with $j \neq Y$) of a given multi-agent possible world stay the same regardless of the internal model which contains this multi-agent possible world. This would not be the case if condition 2 was not required because of the above truth condition of $B_Y\varphi$. This modularity plays also an important role in the generalization of AGM *revision* operation to the multi-agent case [4, 5].

4.2 Multi-agent possible event and internal event model

As we said, in the BMS framework, events are represented very similarly to epistemic situations. This formal similarity can naturally be transferred to the internal approach as well. So the way we represented agent Y 's perception of the surrounding world can easily be adapted to represent her perception of events as well. This leads us to the following definitions.

Definition 4.6. [Multi-agent possible event] A *multi-agent possible event* (M, w) is a *finite* pointed event model $M = (W, R, \text{Pre}, w)$ generated⁵ by $w \in W$ such that R_j is serial, transitive and euclidean for all $j \in G$, and

1. $R_Y(w) = \{w\}$;
2. there is no v and $j \neq Y$ such that $w \in R_j(v)$.

□

The motivations for this definition are completely similar to those for the notion of multi-agent possible world, so we do not spell them out here. Likewise, we can also define the notions of internal event model. The motivations are still very similar to that of the static case.

Definition 4.7. [Internal event model] An *internal event model* is a disjoint and finite union of multi-agent possible events. □

Example 4.8. ['Coin' example] In Figure 9, how the private announcement to Yann is perceived by Ann and Yann is depicted. As in the BMS framework, the formulas in the possible events are their preconditions and the boxed events are the events of the actual equivalence class. Because Yann perceived correctly the private announcement, his internal event model does correspond to a private announcement. On the other hand, for Ann nothing happened because she did not notice this private announcement to Yann. So her internal event model consists of a single possible event with a tautology as precondition. □

Similarly to the way we defined internal models of type 2 from internal model, we can also define *internal event models of type 2* by setting accessibility relations indexed by Y between the roots of the multi-agent possible events of an internal event model.

⁴The notion of equivalence between an internal model (\mathcal{M}, W_a) and an internal model of type 2 (\mathcal{M}', W'_a) can be defined naturally by stating that for every $w \in W_a$ there is a $w' \in W'_a$ which satisfies the same formulas as w , and vice versa.

⁵An event model M is generated from S if the restriction of M to $\{(\bigcup_{j \in G} R_j)^*(w) \mid w \in S\}$ is M itself. This definition is completely in line with Definition 2.6

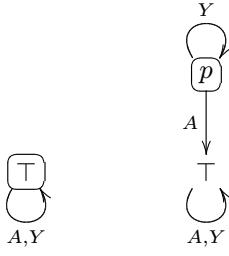


Figure 9: Internal event model for Ann (*left*) and Yann (*right*) corresponding to the private announcement to Yann that the coin is heads up

4.3 The product update

4.3.1 Definition

After (or during) the event takes place, agent Y updates her representation of the world according to her perception of this event. Formally, this amounts to define a product update between an internal model and an internal event model. That is what we are going to do now.

Definition 4.9. [Product update] Let $(\mathcal{M}, W_a) = \{(M^1, w^1); \dots; (M^n, w^n)\}$ be an internal model (with $M^i = (W^i, R^i, V^i)$). Let $(\mathcal{A}, A_a) = \{(\mathbf{M}^1, \mathbf{w}^1); \dots; (\mathbf{M}^m, \mathbf{w}^m)\}$ be an internal event model (with $\mathbf{M}^k = (\mathbf{W}^k, \mathbf{R}^k, \text{Pre}^k)$).

If $\mathcal{M}, w^i \not\models \text{Pre}^k(\mathbf{w}^k)$ then $M^i \otimes_{int} \mathbf{M}^k$ is not defined. Otherwise, $M^i \otimes_{int} \mathbf{M}^k$ is defined and it is the submodel of $M = (W, R, V)$ generated by (w^i, \mathbf{w}^k) , where

- $W = \{(w, \mathbf{w}) \mid w \in W^i, \mathbf{w} \in \mathbf{M}^k, \mathcal{M}, w \models \text{Pre}^k(\mathbf{w})\};$
- $(w', \mathbf{w}') \in R_j(w, \mathbf{w})$ iff $w' \in R_j^i(w)$ and $\mathbf{w}' \in \mathbf{R}_j^k(\mathbf{w})$;
- $(w, \mathbf{w}) \in V(p)$ iff $w \in V^i(p)$ for all $p \in \Phi$.

Then the *updated model* $(\mathcal{M} \otimes_{int} \mathcal{A})$ is defined as follows. If for all $(M^i, w^i) \in \mathcal{M}$ and $(\mathbf{M}^k, \mathbf{w}^k) \in \mathcal{A}$, $M^i \otimes_{int} \mathbf{M}^k$ is not defined then $\mathcal{M} \otimes_{int} \mathcal{A}$ is not defined. Otherwise,

$$\mathcal{M} \otimes_{int} \mathcal{A} = \{(M^i \otimes_{int} \mathbf{M}^k, (w^i, \mathbf{w}^k)) \mid (M^i, w^i) \in \mathcal{M}, (\mathbf{M}^k, \mathbf{w}^k) \in \mathcal{A} \text{ and } M^i \otimes_{int} \mathbf{M}^k \text{ is defined}\}$$

The updated model is written $(\mathcal{M} \otimes_{int} \mathcal{A}, W_a \otimes_{int} A_a)$. \triangleleft

Note that in the definition of W , the preconditions $\text{Pre}(\mathbf{w}^k)$ are evaluated in \mathcal{M} and not in M^i . Indeed, in case $\text{Pre}^k(\mathbf{w}^k)$ is a formula of the form $B_Y \varphi$, then we need the whole internal model to evaluate it (see the truth conditions for $B_Y \varphi$ in Definition 4.5). That is why we need to redefine the product update completely. However, note that our definition is completely in line with the BMS one. In fact, one can easily show that given an internal model and an internal event model, updating with the BMS product update the corresponding internal (event) models of type 2 yields an internal model of type 2 which is equivalent (in the sense of footnote 4) to the internal model obtained with our product update.

Example 4.10. ['Coin' example] In Figure 10 is depicted how the 'coin' scenario depicted in Figure 4 is perceived by Yann. Because Yann perceived correctly the scenario, this representation is equivalent to the one of Figure 4. Importantly, note that the updated model is the multi-agent possible world of the original internal model which satisfies the (private) announcement that p ('the coin is heads up'). \triangleleft

Because the update mechanisms are essentially the same as in the BMS product update, it also entails that, as in BMS, the updated models are not necessarily serial. So the updated model is not necessarily an internal model. We are now going to study under which conditions it is an internal model.

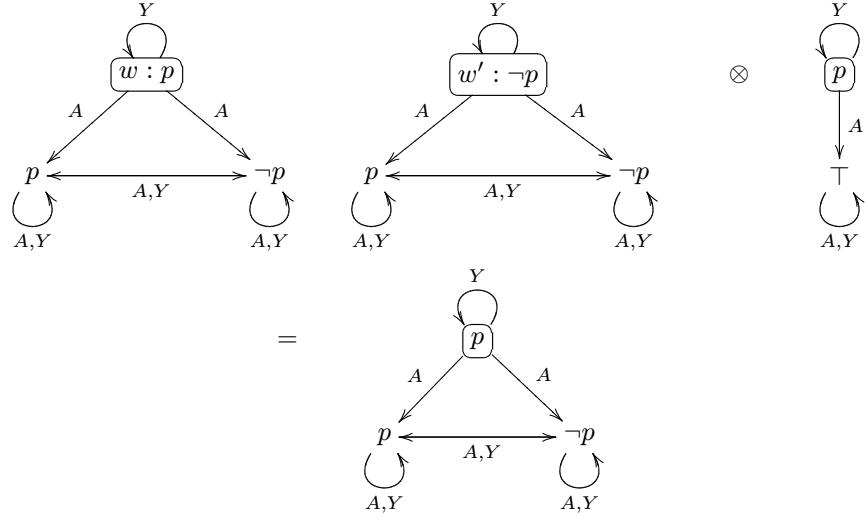


Figure 10: Yann's perception of the ‘coin’ scenario

4.3.2 Under which conditions updated models are internal models?

The seriality condition is a bit involved. This is because we do not use exactly the BMS product but employ a slightly different one, which is nevertheless based on the BMS one.

Proposition 4.11. Let $(\mathcal{M}, W_a) = \{(M^1, w^1); \dots; (M^n, w^n)\}$ be an internal model (with $M^i = (W^i, R^i, V^i)$). Let $(\mathcal{A}, M_a) = \{(M^1, w^1); \dots; (M^m, w^m)\}$ be an internal event model (with $M^k = (W^k, R^k, Pre^k)$). Assume that $(\mathcal{M} \otimes_{int} \mathcal{A}, W_a \otimes_{int} A_a)$ is defined. Then,

$$(\mathcal{M} \otimes_{int} \mathcal{A}, W_a \otimes_{int} A_a) \text{ is an internal model iff for all } w^i, \mathcal{M}, w^i \models \bigwedge_{(A^k, w^k) \in \mathcal{A}} (Pre^k(w^k) \rightarrow \delta^d(w^k))$$

where $d = \max\{|M^i| \cdot |M^k| \mid i \in \{1, \dots, n\}, k \in \{1, \dots, m\}\}$ and δ is defined in Definition 3.7.

Proof. First we prove a lemma.

Lemma 4.12. Let $d_{i,k} = |M^i| \cdot |M^k| - 1$. Then $\mathcal{M}, w^i \models \delta^{d_{i,k}}(w^k)$ iff $M^i \otimes_{int} M^k$ is defined and serial.

Proof. Because the update product $M^i \otimes_{int} M^k$ is not the same as the BMS update product, we cannot apply directly Corollary 3.11.

- Assume $\mathcal{M}, w^i \models \delta^{d_{i,k}}(w^k)$. Then $\mathcal{M}, w^i \models Pre^k(w^k)$, so $M^i \otimes_{int} M^k$ is defined. Besides, note that $M^i \otimes_{int} M^k$ is serial iff (E)
 1. for all $j \in G$, $R_j(w^i, w^k) \neq \emptyset$
 2. for all $j \neq Y$, for all $v^j \in R_j(w^i)$ and $v^j \in R_j(w^k)$ such that $\mathcal{M}, v^j \models Pre^k(v^j)$, the submodel of $M^i(v^j) \otimes M^k(v^j)$ generated by (v^j, v^j) is serial, where $M^i(v^j)$ is the submodel of M^i generated by v^j and $M^k(v^j)$ is the submodel of M^k generated by v^j .

$M^i(v^j) \otimes M^k(v^j)$ is the usual BMS update product. Indeed, for all $\varphi \in \mathcal{L}$, all $u^j \in M^i(v^j)$, $M^i(v^j), u^j \models \varphi$ iff $\mathcal{M}, u^j \models \varphi$. So for the worlds of $M^i(v^j)$, the update product \otimes_{int} is the same as the BMS update product \otimes . This will allow us to apply Corollary 3.11.

1. $\mathcal{M}, w^i \models \delta^1(w^k)$ because $d_{i,k} \geq 1$ and Proposition 3.13. So $\mathcal{M}, w^i \models \bigwedge_{j \in G} \langle B_j \rangle \bigvee_{v \in R_j(w^k)} Pre^k(v)$.

So $R_j(w^i, w^k) \neq \emptyset$ for all $j \in G$.

2. Let $v^j \in R_j(w^i)$ and $\mathbf{v}^j \in \mathbf{R}_j(\mathbf{w}^k)$ such that $\mathcal{M}, v^j \models \text{Pre}(v^j)$, i.e. $M^i, v^j \models \text{Pre}^k(v^j)$. Then $M^i, v^j \models \delta^{d_{i,k}-1}(v^j)$ by definition of $\delta^{d_{i,k}}(\mathbf{w}^k)$. Besides, $|M^i(v^j)| \leq |M^i| - 1$ and $|\mathbf{M}^k(v^j)| \leq |\mathbf{M}^k| - 1$. So $n = |M^i(v^j)| \cdot |\mathbf{M}^k(v^j)| \leq d_{i,k} - 1$.

So $M^i, v^j \models \delta^n(v^j)$ by Proposition 3.13.

So the submodel of $M^i(v^j) \otimes \mathbf{M}^k(v^j)$ generated by (v^j, \mathbf{v}^j) is defined and serial by Corollary 3.11.

- We prove by induction on $m \leq d_{i,k}$ that if $\mathcal{M}, w^i \models \neg\delta^m(\mathbf{w}^k)$ then $M^i \otimes_{int} \mathbf{M}^k$ is either not defined or not serial.

m=0 $\delta^0(\mathbf{w}^k) = \text{Pre}^k(\mathbf{w}^k)$. So $\mathcal{M}, w^i \models \neg\text{Pre}^k(\mathbf{w}^k)$. So $M^i \otimes_{int} \mathbf{M}^k$ is not defined.

m+1 $\mathcal{M}, w_i \models \neg\delta^m(\mathbf{w}^k)$ iff $\mathcal{M}, w_i \models \neg\text{Pre}^k(\mathbf{w}^k)$ or $\mathcal{M}, w^i \models \bigvee_{j \in G} B_j \wedge_{v \in \mathbf{R}_j(\mathbf{w}^k)} \neg\delta^{m-1}(v)$ or

$$\mathcal{M}, w^i \models \bigvee_{j \in G} \langle B_j \rangle \bigvee_{v \in \mathbf{R}_j(\mathbf{w}^k)} (\text{Pre}(v) \wedge \neg\delta^{m-1}(v)).$$

1. If $\mathcal{M}, w^i \models \neg\text{Pre}^k(\mathbf{w}^k)$ then $M^i \otimes_{int} \mathbf{M}^k$ is not defined.

2. If $\mathcal{M}, w^i \models \bigvee_{j \in G} B_j \wedge_{v \in \mathbf{R}_j(\mathbf{w}^k)} \neg\delta^{m-1}(v)$ then there is $j \in G$ such that $\mathcal{M}, w^i \models B_j \wedge_{v \in \mathbf{R}_j(\mathbf{w}^k)} \neg\delta^{m-1}(v)$

(a) If $j = Y$ then $\mathcal{M}, w^i \models \neg\delta^{m-1}(\mathbf{w}^k)$. So $M^i \otimes \mathbf{M}^k$ is not defined or not serial by Induction Hypothesis.

(b) If $j \neq Y$ then for all $v^j \in R_j(w^i)$, all $\mathbf{v}^j \in \mathbf{R}_j(\mathbf{w}^k)$, $M^i, v^j \models \neg\delta^{m-1}(v^j)$. So $M^i(v^j) \otimes \mathbf{M}^k(\mathbf{v}^j)$ is not defined or not serial by Corollary 3.11. So, because M^i and \mathbf{M}^k are serial, $M^i \otimes_{int} \mathbf{M}^k$ is not serial by observation (E).

3. If $\mathcal{M}, w^i \models \bigvee_{j \in G} \langle B_j \rangle \bigvee_{v \in \mathbf{R}_j(\mathbf{w}^k)} (\text{Pre}(v) \wedge \neg\delta^{m-1}(v))$ then there is $j \in G$ and $v^j \in R_j(w^i)$

and $\mathbf{v}^j \in \mathbf{R}_j(\mathbf{w}^k)$ such that $\mathcal{M}, v^j \models \text{Pre}^k(v^j) \wedge \neg\delta^{m-1}(v^j)$.

(a) If $j = Y$ then $\mathcal{M}, w^i \models \neg\delta^{m-1}(\mathbf{w}^k)$. So the result holds by Induction Hypothesis.

(b) If $j \neq Y$ then by the same reasoning as in (2)(b), we get that $M^i \otimes_{int} \mathbf{M}^k$ is not defined or not serial.

□

- Assume for all w^i that $\mathcal{M}, w^i \models \bigwedge_{(\mathbf{M}^k, \mathbf{w}^k) \in \mathcal{A}} (\text{Pre}^k(\mathbf{w}^k) \rightarrow \delta^d(\mathbf{w}^k))$ (*). Let $(M^i, w^i) \in \mathcal{M}$ such that there is $(\mathbf{M}^k, \mathbf{w}^k) \in \mathcal{A}$ such that $\mathcal{M}, w^i \models \text{Pre}^k(\mathbf{w}^k)$. Then $\mathcal{M}, w^i \models \delta^d(\mathbf{w}^k)$ by (*).

But $\models \delta^d(\mathbf{w}^k) \rightarrow \delta^{d_{i,k}}(\mathbf{w}^k)$ because $d_{i,k} \leq d$ and because of Proposition 3.13.

So $\mathcal{M}, w^i \models \delta^{d_{i,k}}(\mathbf{w}^k)$. Then $M^i \otimes_{int} \mathbf{M}^k$ is serial by Lemma 4.12, and so for all $(M^i, w^i) \in \mathcal{M}$ and $(\mathbf{M}^k, \mathbf{w}^k) \in \mathcal{A}$ such that $M^i \otimes_{int} \mathbf{M}^k$ is defined.

So finally, $\mathcal{M} \otimes_{int} \mathcal{A}$ is an internal model.

- Assume that there is w^i and $(\mathbf{M}^k, \mathbf{w}^k) \in \mathcal{A}$ such that $\mathcal{M}, w^i \models \text{Pre}^k(\mathbf{w}^k) \wedge \neg\delta^d(\mathbf{w}^k)$.
 $d \geq d_{i,k} = |M^i| \cdot |\mathbf{M}^k| - 1$ by assumption. So $\mathcal{M}, w^i \models \neg\delta^{d_{i,k}}(\mathbf{w}^k)$ by Proposition 3.12. So $M^i \otimes_{int} \mathbf{M}^k$ is defined but not serial by Lemma 4.12. So $\mathcal{M} \otimes_{int} \mathcal{A}$ is not an internal model.

□

5 Application: generalizing expansion to a multi-agent setting

5.1 AGM theory and expansion

Belief revision theory was developed before the beginning of dynamic epistemic logic. It really started with Alchourrón, Gärdenfors and Makinson's seminal paper [1]. In AGM belief revision theory, the epistemic state of the agent is often represented by a belief set. A *belief set* K is a set of propositional formulas that is closed under logical consequence. These propositional formulas represent the beliefs of the agent. AGM distinguishes three types of belief change: expansion, revision and contraction. The

expansion of K with a propositional formula φ , written $K + \varphi$, consists of adding φ to K and taking all the logical consequences. Note that this might yield inconsistency. The revision of K with φ , written $K * \varphi$, consists of adding φ to K , but in order that the resulting set be consistent, some formulas are removed from K . Finally, the contraction of K with φ , written here $K \doteq \varphi$, consists in removing φ from K , but in order that the resulting set be consistent, some other formulas are also removed. Of course there are some connections between these operations. From a contraction operation, one can define a revision operation thanks to the Levi identity:

$$K * \varphi = (K \doteq \neg\varphi) + \varphi.$$

And from a revision operation, one can define a contraction operation thanks to the Harper identity:

$$K \doteq \varphi = K \cap (K * \neg\varphi).$$

In this paper, we will focus on the expansion operation. We will show how this operation can be generalized to a multi-agent setting. For a generalization of the revision operation to the multi-agent case in the same line, see [5].

In this section, we assume that the set of propositional letters Φ is finite, and that all the formulas belong to the propositional language \mathcal{L}_0 defined over Φ . Let $Cn(\cdot)$ be the classical consequence operation, i.e. for a set of propositional formulas Σ , $Cn(\Sigma) = \{\chi \mid \Sigma \vdash \chi\}$. We can now define formally a belief set.

Definition 5.1. [Belief set] A *belief set* K is a set of propositional formulas in \mathcal{L}_0 such that $Cn(K) = K$. We denote by K_\perp the unique inconsistent belief set consisting of all propositional formulas. \triangleleft

Classically, in AGM theory, we start by proposing rationality postulates that belief change operations must fulfill. These postulates make precise our intuitions about these operations and what we mean by rational change. Below are the rationality postulates for the expansion operation + proposed by Gärdenfors [10].

K+1 $K + \varphi$ is a belief set

K+2 $\varphi \in K + \varphi$

K+3 $K \subseteq K + \varphi$

K+4 If $\varphi \in K$ then $K = K + \varphi$

K+5 $K + \varphi$ is the smallest set satisfying $K + 1 - K + 4$.

$K + 1$ tells us that the expansion operation + is a function from pairs of belief set and formula to belief sets. This entails that we can iterate the expansion operation. $K + 2$ tells us that when the agent expands her belief set by φ then as a result φ is one of her beliefs. All the other postulates refer to some kind of minimal change. $K + 3$ tells us that when the agent expands by φ she does not throw away any of her former beliefs. $K + 4$ tells us that if the agent already believes φ then expanding by φ should not change her beliefs: the change made to add φ to the belief set is minimal.

The following (representation) theorem tells us that these postulates actually determine a *unique* expansion operation on belief sets.

Theorem 5.2 ([10]). *A function + satisfies $K + 1 - K + 5$ iff for each belief set K and formula φ , $K + \varphi = Cn(K \cup \{\varphi\})$.*

So from now on, we define the expansion operation + by $K + \varphi = Cn(K \cup \{\varphi\})$.

So far our approach to expansion was syntactically driven. Now we are going to give a semantical approach to expansion and set some links between these two approaches. We use the possible world semantics. First we consider the set \mathcal{W} consisting of all the (logically) possible worlds. A possible world w can be viewed as an interpretation, i.e. a function from Φ to $\{\top, \perp\}$ which specifies which propositional letters (such as ‘it is raining’) are true in this world w . For a propositional formula χ , we

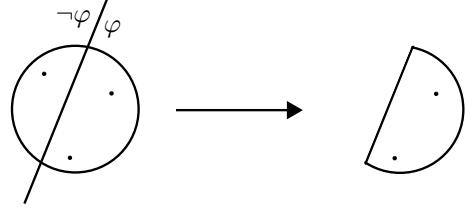


Figure 11: AGM expansion by φ

write $w \models \chi$ when χ is true at w in the usual sense⁶. Then a formula χ is true in a set W of possible worlds, written $W \models \chi$, if and only if for all $w \in W$, $w \models \chi$. Besides, because Φ is finite, \mathcal{W} is also finite. We can then represent the agent's epistemic state by a subset W of \mathcal{W} (which is consequently finite as well). Intuitively, W is the smallest set of possible worlds in which the agent believes that the actual world is located.

There is actually a very close correspondence between belief sets and sets of possible worlds.

Definition 5.3. Let W be a finite set of possible worlds. We define the belief set K_W associated to W by $K_W = \{\chi \mid W \models \chi\}$.

Let K be a belief set. We define the set of possible worlds W_K associated to K by $W_K = \{w \mid w \models \chi \text{ for all } \chi \in K\}$. Then,

$$W \models \chi \text{ iff } \chi \in K_W, \text{ and } \chi \in K \text{ iff } W_K \models \chi.$$

□

Now we can define the semantic counterpart of the expansion operation defined previously.

Definition 5.4. [(Semantic) expansion] Let W be a finite set of possible worlds and φ of formula. The *expansion* of W by φ , written $W + \varphi$, is defined as follows.

$$W + \varphi = \{w \in W \mid w \models \varphi\}.$$

□

This semantic counterpart of the expansion is described graphically in Figure 11. The initial model W is on the left of the arrow and the expanded model $W + \varphi$ is on the right of the arrow. The dots represent possible worlds and the straight line separates the worlds satisfying φ from the worlds satisfying $\neg\varphi$.

Finally, we show that these two definitions of expansion, syntactic and semantic, are in fact equivalent.

Theorem 5.5. For all belief sets K and all finite set of possible worlds W ,

$$\chi \in K + \varphi \text{ iff } W_K + \varphi \models \chi, \text{ and } W + \varphi \models \chi \text{ iff } \chi \in K_W + \varphi.$$

Proof. $\chi \in K + \varphi$

$$\text{iff } \chi \in Cn(K \cup \{\varphi\})$$

$$\text{iff } K \cup \{\varphi\} \vdash \chi$$

$$\text{iff for all } w \text{ such that } w \models \psi \text{ for all } \psi \in K \cup \{\varphi\}, w \models \chi$$

$$\text{iff for all } w \in W_K, \text{ if } w \models \varphi \text{ then } w \models \chi$$

$$\text{iff for all } w \in W_K + \varphi, w \models \chi$$

$$\text{iff } W_K + \varphi \models \chi.$$

$$\chi \in K_W + \varphi$$

$$\text{iff } \chi \in Cn(K_W \cup \{\varphi\})$$

$$\text{iff } K_W \cup \{\varphi\} \vdash \chi$$

$$\text{iff for all } w \text{ such that } w \models K_W \cup \{\varphi\}, w \models \chi$$

⁶ $w \models \chi$ is defined inductively by: $w \models p$ iff $w(p) = \top$; $w \models \neg\chi$ iff not $w \models \chi$; and $w \models \chi \wedge \chi'$ iff $w \models \chi$ and $w \models \chi'$.

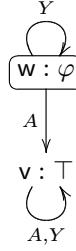


Figure 12: Private announcement of φ to agent Y

iff for all w such that $w \models K_W$ and $w \models \varphi$, $w \models \chi$
 iff for all $w \in W + \varphi$, $w \models \chi$
 iff $W + \varphi \models \chi$. □

This ends our account about expansion. Now we are going to show that the AGM operation of expansion can naturally be viewed in a multi-agent setting of the BMS framework as the operation of private announcement.

5.2 Expansion and private announcement

Assume we are now in a multi-agent setting and we follow the internal approach. Let us have a closer look at the case of private announcement as it is modeled in the BMS framework. The event model of a private announcement of $\varphi \in \mathcal{L}$ to agent Y is depicted in Figure 12.

Proposition 5.6. *Let $(\mathcal{A}, \{w\}) = (\{w, v\}, R, \text{Pre}, \{w\})$ be the internal event model of a private announcement of $\varphi \in \mathcal{L}$ to Y . Then for all $n \in \mathbb{N}$, $\models_{\text{Ext}} \delta^n(v) \leftrightarrow \top$ and $\models_{\text{Int}} \delta^n(w) \leftrightarrow \varphi$.*

Proof. We prove both results by induction on n . Clearly, the result holds for $n = 0$. Assume it is true for n .

Then $\delta^{n+1}(w) = \delta^0(w) \wedge \langle B_Y \rangle \delta^n(w) \wedge \langle B_A \rangle \delta^n(v) \wedge B_Y(\text{Pre}(w) \rightarrow \delta^n(w)) \wedge B_A(\text{Pre}(v) \rightarrow \delta^n(v))$ by Definition 3.7. So $\models_{\text{Int}} \delta^{n+1}(w) \leftrightarrow \varphi \wedge \langle B_Y \rangle \varphi \wedge \langle B_A \rangle \top \wedge B_Y(\varphi \rightarrow \varphi) \wedge B_A(\varphi \rightarrow \top)$ by induction hypothesis. Then $\models_{\text{Int}} \delta^{n+1}(w) \leftrightarrow \varphi \wedge \langle B_Y \rangle \varphi$. So $\models_{\text{Int}} \delta^{n+1}(w) \leftrightarrow \varphi$ by axiom \top .

Besides, $\delta^{n+1}(v) = \delta^0(v) \wedge \langle B_Y \rangle \delta^n(v) \wedge \langle B_A \rangle \delta^n(v) \wedge B_Y(\text{Pre}(v) \rightarrow \delta^n(v)) \wedge B_A(\text{Pre}(v) \rightarrow \delta^n(v))$ by Definition 3.7. So $\models_{\text{Ext}} \delta^{n+1}(v) \leftrightarrow \top \wedge \langle B_Y \rangle \top \wedge \langle B_A \rangle \top \wedge B_Y(\top \rightarrow \top) \wedge B_A(\top \rightarrow \top)$, i.e. $\models_{\text{Ext}} \delta^{n+1}(v) \leftrightarrow \top$. □

Corollary 5.7. *Let (\mathcal{M}, W_a) be an internal model of type 1 and let $(\mathcal{A}, \{w\})$ be the internal event model of the private announcement of $\varphi \in \mathcal{L}$ to Y . Then $(\mathcal{M} \otimes_{\text{int}} \mathcal{A}, W_a \otimes_{\text{int}} \{w\})$ is defined and is an internal model of type 1 iff there is $w_a \in W_a$ such that $\mathcal{M}, w_a \models \varphi$.*

Proof. Assume that $\mathcal{M} \otimes_{\text{int}} \mathcal{A}$ is defined. We know that for all $n \in \mathbb{N}$, $\models_{\text{Int}} (\text{Pre}(w) \rightarrow \delta^n(w)) \leftrightarrow (\text{Pre}(w) \rightarrow \text{Pre}(w))$ by Proposition 5.6. So for all $w \in W_a$, $\mathcal{M}, w \models (\text{Pre}(w) \rightarrow \delta^n(w))$. This means that $\mathcal{M} \otimes_{\text{int}} \mathcal{A}$ is an internal model by Proposition 4.11. So, if $\mathcal{M} \otimes_{\text{int}} \mathcal{A}$ is defined then it is an internal model. But $\mathcal{M} \otimes_{\text{int}} \mathcal{A}$ is defined iff there is $w_a \in W_a$ such that $\mathcal{M}, w_a \models \varphi$. So we get the result. □

This corollary is obtained thanks to our study on seriality preservation for generated submodels made in Section 3.2. It tells us that as soon as the updated model by a private announcement is defined then it must be an internal model, in other words it is serial. This might seem strange at first sight but the following crucial theorem provides a good explanation for that.

Theorem 5.8. *Let $(\mathcal{M}, W_a) = \{(M^1, w^1), \dots, (M^n, w^n)\}$ be an internal model and $(\mathcal{A}, \{w\})$ be the internal event model of a private announcement of $\varphi \in \mathcal{L}$ to Y . Then,*

$$\mathcal{M} \otimes_{\text{int}} \mathcal{A} = \{(M^i, w^i) \mid \mathcal{M}, w^i \models \varphi\}$$

Proof. It suffices to apply the definition of \otimes_{int} . □

This theorem is very important: it bridges the gap between AGM belief revision theory and dynamic epistemic logic as viewed by BMS. Indeed, one can note that the definition of $\mathcal{M} \otimes_{int} \mathcal{A}$ is very similar to the semantic definition of expansion in Definition 5.4. On the one hand, the expansion of a set of possible worlds by a propositional formula φ consists in the worlds that satisfy φ . On the other hand, the updated model of an internal model by a private announcement of an epistemic formula φ consists in the *multi-agent* possible worlds that satisfy φ . This similarity is depicted in Figure 13 where the triangles represent multi-agent possible worlds. Note that this theorem also explains why an updated model is an internal model as soon as it is defined: it is because the updated model is a submodel of the original internal model.

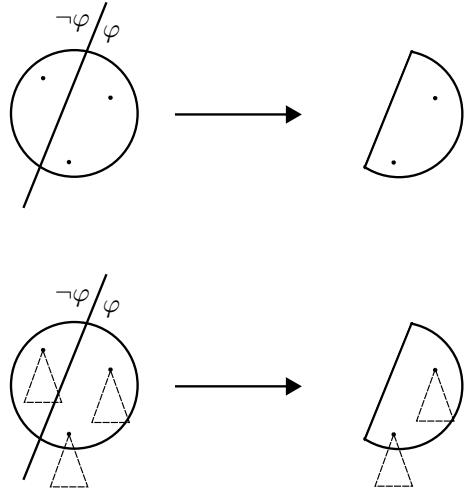


Figure 13: AGM expansion by φ (*above*) and BMS update by a private announcement of φ (*below*)

So, informally, the BMS update by a private announcement can be viewed as a ‘multi-agent’ AGM expansion. In other words, AGM expansion can be viewed as a particular case of a BMS update by a private announcement in which Y is the only agent. This means that private announcement is the BMS generalization to the multi-agent case of AGM expansion. This goes against van Ditmarsch, van der Hoek and Kooi’s claim that *public* announcement can be viewed as a belief expansion [16]. However, their comparison is rather syntactical and in that respect they only consider special kinds of formulas to represent belief sets and (public) announcements, namely ‘positive’ formulas. Here, our results hold independently of any particular chosen language because we compare only the semantics of expansion and private announcement. We believe this semantic correspondence to be deeper than any syntactic one because the languages in the single agent case and the multi-agent case are anyway quite different and so do not allow for a straightforward comparison. Moreover, the fact that private announcement can be viewed as a generalization of expansion in a multi-agent setting is not accidental. Indeed, an important property of private announcement is that not only the actual world does not change but also the agents’ beliefs do not change (except of course for agent Y ’s beliefs). For example, suppose you (Y) believe p , and agent j believes p (and perhaps even that p is common belief of Y and j). When a third external agent privately tells you that $\neg p$ then afterwards j still believes p and you still believe that j believes p (and that j believes that p is common belief). This static aspect of private announcements is similar to the static aspect of AGM belief revision in a single-agent case: in both cases the world does not change but only agent Y ’s beliefs about the world change.

Example 5.9. [‘Coin’ example] Let us take up the ‘coin example’ and let us consider Yann’s internal point of view. Yann’s internal model representing the initial situation is recalled in Figure 14. Then, according to Theorem 5.8, the resulting situation after the private announcement to Yann that the coin is heads up (p) is the internal model composed only of the multi-agent possible world (M, w) on the left of Figure 14. We can check that this result is correct in the scenario depicted in Figure 10. \triangleleft

Remark 5.10. In the internal approach, just as in the external approach of the BMS framework, private announcement of Moore formulas of the form $\varphi \wedge \neg B_Y \varphi$ do not trivialize the internal product update, because these formulas can be internally satisfiable. After the private announcement, the

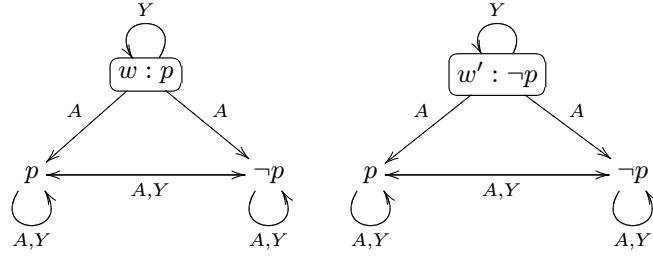


Figure 14: Yann’s initial internal model $\{(M, w), (M', w')\}$

Moore formula is just not believed by agent Y , because she then believes φ (see [15] for more details). Our dynamic approach somehow refines the standard ‘static’ analysis of Moore formulas by the previous internal versions of epistemic logic. On the other hand, private announcement of crazy formulas of the form $\varphi \wedge B_Y \neg \varphi$ trivialize the internal product update, because these formula are not internally satisfiable. To avoid it, one would need to introduce some sort of revision mechanisms into our product update. \triangleleft

6 Conclusion

Our contribution in this paper is threefold. First, we proposed a dynamic epistemic logic for the internal approach (the internal version of the BMS framework), extending naturally to a dynamic setting the internal version of epistemic logic [6]. This framework can also be viewed as a generalization of AGM theory (restricted to the expansion operation) to the multi-agent setting and it provides a logical framework for the internal approach in a *dynamic* setting. Second, we provided conditions under which seriality is preserved during an update (for the entire updated model and for its generated submodels), for the BMS framework as well as its internal version. Third, we showed using this new logical framework that the AGM operation of expansion can naturally be viewed in a multi-agent setting as the BMS operation of private announcement.

So in a sense we have bridged the gap between AGM and BMS. Nevertheless, we focused here only on the generalization to the multi-agent setting of the operation of expansion and we did not consider the operation of revision. The generalization of the AGM operation of revision to the multi-agent setting is dealt with in [5]. The idea there is simply from a semantic point of view to replace possible worlds by multi-agent possible worlds, and from a syntactic point of view to replace the propositional language by an epistemic language: AGM representation theorems can then easily be lifted to the multi-agent case. However and as a consequence of this paper, this generalization would correspond to a kind of revision agent Y performs when she receives a conflicting information via a *private* announcement. It remains to propose revision mechanisms for all the other kinds of events which can be formalized by means of an (internal) event model: public announcement, private announcement to a subgroup, and so on... Our results on the conditions for preserving seriality are a first step in that direction. Indeed, they indicate in which case a revision is necessary and which conditions the revised model should satisfy.

References

- [1] Carlos Alchourrón, Peter Gärdenfors, and David Makinson. On the logic of theory change: Partial meet contraction and revision functions. *Journal of Symbolic Logic*, 50(2):510–530, 1985.
- [2] Horacio Arlo-Costa. Qualitative and probabilistic models of full belief. In S. Buss, P. Hajek, and P. Pudlak, editors, *Logic Colloquium’98*, Lecture Notes on Logic 13, 1999.
- [3] Guillaume Aucher. Consistency preservation and crazy formulas in BMS. In Steffen Hölldobler, Carsten Lutz, and Heinrich Wansing, editors, *JELIA*, volume 5293 of *Lectures Notes in Computer Science*, pages 21–33. Springer, 2008.

- [4] Guillaume Aucher. Internal models and private multi-agent belief revision. In Muller Padgham, Parkes and Parsons, editors, *Proceedings of Autonomous Agents and Multi-agent Systems (AAMAS 2008)*, pages 721–727, Estoril, Portugal, 2008.
- [5] Guillaume Aucher. Generalizing AGM to a multi-agent setting. *Logic Journal of the IGPL*, 2010. DOI: 10.1093/jigpal/jzp037.
- [6] Guillaume Aucher. An internal version of epistemic logic. *Studia Logica*, 1:1–22, 2010.
- [7] Alexandru Baltag and Larry Moss. Logic for epistemic programs. *Synthese*, 139(2):165–224, 2004.
- [8] Alexandru Baltag, Larry Moss, and Slawomir Solecki. The logic of common knowledge, public announcement, and private suspicions. In I. Gilboa, editor, *Proceedings of the 7th conference on theoretical aspects of rationality and knowledge (TARK98)*, pages 43–56, 1998.
- [9] Patrick Blackburn, Maarten de Rijke, and Yde Venema. *Modal Logic*, volume 53 of *Cambridge Tracts in Computer Science*. Cambridge University Press, 2001.
- [10] Peter Gärdenfors. *Knowledge in Flux (Modeling the Dynamics of Epistemic States)*. Bradford/MIT Press, Cambridge, Massachusetts, 1988.
- [11] Peter Gärdenfors and Hans Rott. *Handbook of Logic in Artificial Intelligence and Logic Programming*, volume Volume 4, Epistemic and temporal reasoning, chapter Belief Revision, pages 35–132. Clarendon Press, Oxford, 1995.
- [12] Isaac Levi. *The covenant of reason: rationality and the commitments of thought*, chapter The Logic of Full Belief, pages 40–69. Cambridge University Press, 1997.
- [13] R.C.Moore. Possible-world semantics for autoepistemic logic. In *Proceedings of the Non-Monotonic Reasoning Workshop*, pages 344–354, New Paltz NY, 1984.
- [14] R.C.Moore. *Logic and Representation*. CSLI Lecture Notes, 1995.
- [15] Hans van Ditmarsch and Barteld Kooi. The secret of my success. *Synthese*, 151:201–232, 2006.
- [16] Hans van Ditmarsch, Wiebe van der Hoek, and Barteld Kooi. Public announcements and belief expansion. In *Advances in Modal Logic*, pages 335–346, 2004.
- [17] Hans van Ditmarsch, Wiebe van der Hoek, and Barteld Kooi. *Dynamic Epistemic Logic*, volume 337 of *Synthese library*. Springer, 2007.