A Combined System for Update Logic and Belief Revision

Guillaume Aucher
aucher_guillaume@yahoo.fr

December 5, 2003

Aknowledgements

I want to thank very much my supervisors, Johan van Benthem and Hans van Ditmarsch for their insightful comments. The size of their e-mails always impressed me, and they were really involved in my work. Thanks! I also want to thank all my friends in Amsterdam and my family who, all, in some way or another contributed to achieve this thesis.

Contents

General Introduction

Roughly speaking, in this thesis we will propose a logical system merging update logic as conceived by A.Baltag, L.Moss, S.Solecki (BMS)[BMS03] on the one hand; and belief revision theory as conceived by C.Alchourron, P.Gardenfors and D.Mackinson (AGM)([GardRott95])(viewed from the point of view of W.Spohn ([Spohn90],[Spohn88])) on the other hand. Before tackling the topic, we need to set out some general assumptions about the type of phenomenon that we intend to study thanks to these theories. It will also indirectly provide us a framework for our future work, and give an idea of the topic of this thesis (and these theories).

We assume that any situation s involving several agents can be rendered from the point of view of the agents' knowledge and beliefs of the situation by a mathematical model M. We assume that this association is correct, in the sense that every intuitive judgement concerning s corresponds to a formal assertion concerning M. Now in the situation s, an action a may occur. We also assume that this action a can be *correctly* (see above) rendered from the point of view of the agents' knowledge and beliefs by a mathematical model Σ. Now in reality the agents update their knowledge and beliefs according to these two pieces of information: action a and situation s, giving rise to a new actual situation $s \times a$. We assume again that we can render this update mechanism by a mathematical update $\otimes$ such that, as above, $M \otimes \Sigma$ corresponds *correctly* (see above) to $s \times a$. Moreover we assume that the update mechanism concerning the agents' beliefs be the closest possible to a belief revision (conceived by AGM). Note that in reality, once the agents receive the new information carried by the action a, update their knowledge and beliefs. This double process may be done simultaneously in reality by the agents. Yet we carefully separate it in our approach by introducing Σ because these are two conceptually different things: apprehension of the new information (corresponding to Σ); and update (corresponding to the update $\otimes$), on the basis of this apprehension.

By the very nature of the BMS and AGM theories (see chapter 1), merging them seems one of the best ways to give concrete form to these assumptions. Yet the resulting system should be a genuine *logical* system and we must keep that in mind. So, first we will set out these theories. Second, we will propose a merge system of these theories. Third we will propose a proof system for this system (with the introduction of a special canonical model in the completeness proof). Finally, we will compare our system with other similar systems, and also the AGM theory.

Chapter 1

The Ingredients: Update Logic and Belief Revision Theory

Introduction: As stated in the general introduction, our main goal is to merge update logic with belief revision theory. Our first task will be to present an account of these theories. In the section 1.1, we will present update logic as conceived by BMS. In the section 1.2, we will present belief revision theory as conceived by AGM and highlight W.Spohn's approach.

1.1 Update Logic

Introduction: In this chapter we will set out what is generally called update logic. First we will present a rough overview of a simple subcase called the logic of public information, whose aim will be to give the main underlying ideas of the topic (see[vB03]). Second we will present a broader set up which enables us to deal with more complex information (not only public) conceived by A.Baltag, L.Moss and S.Solecki in [BMS03].

1.1.1 The Simple Case: Dynamic Epistemic Logic

Update logic is a formalism trying to model epistemic situations involving several agents, and changes that can occur to these situations after incoming information or more generally incoming action. Update logic is composed of two main features. First, we model thanks to an epistemic modal model the knowledge a group of agents possesses: this is the static part. Second, we have to update this model after for example a public announcement of a formula to all the agents: this public announcement and correlative update constitute the dynamic part. Note that these epistemic actions can be much more complex than simple public announcement, including hiding information for some of the agents or cheating,etc...This complexity is dealt with in the next section by introducing the notion of a simple action structure. In this first account, we will concentrate only on public announcement to get an intuition of the main underlying ideas that occur in update logic. We will start by giving a concrete example where update logic can be used, to better understand what is going on. Then we will present a sketchy formalization of the phenomenon called the logic of public information update (We will base our exposition on [vB03] and [BMS03])

Example: Muddy children

We have two children, A and B, both dirty. A can see B but not himself, and B can see A but not herself. Let p be the proposition stating that A is dirty, and q be the proposition stating that B is not dirty. We represent this situation by the following epistemic modal model where relations are equivalence relations (for a precise definition of a modal model, see [BdRV01]; and for its epistemic interpretation, see [FHMV95]):

$$\begin{tikzcd} \boxed{\boxed{s : p, q}} \arrow[r, dash, "A"] \arrow[d, dash, "B"] & \boxed{t : \neg p, q} \arrow[d, dash, "B"] \\ \boxed{u : p, \neg q} \arrow[r, dash, "A"'] & \boxed{v : \neg p, \neg q} \end{tikzcd}$$

States s, t, u, v intuitively represent possible worlds, a proposition (for example p) satisfiable at one of these states intuitively means that in the possible world corresponding to this state, the intuitive interpretation of p (p is dirty) is true. From now on, we will name equivalently possible worlds and states. The links between states labelled by agents (A or B) intuitively express a notion of indistinguishability for the agent at stake between two possible worlds. For example, the link between s and t labelled by A intuitively means that A can not distinguish the possible world s from t and vice versa. Indeed, A can not see himself, so he can not distinguish between a world where he is dirty and one where he is not dirty. However, he can distinguish between worlds where B is dirty or not because he can see B. With this intuitive interpretation we are brought to assume that our relations between states are equivalence relations.

Now, suppose that their father comes and announces that at least one is dirty. Then we update the model by (explanations follow):

$$\begin{tikzcd} \boxed{\boxed{s : p, q}} \arrow[r, dash, "A"] \arrow[d, dash, "B"] & \boxed{t : \neg p, q} \\ \boxed{u : p, \neg q} & \end{tikzcd}$$

What we actually do is suppressing the worlds where the content of the announcement is not fulfilled. In our case this is the world where $\neg p$ and $\neg q$ are true. This suppression is what we call the update. We then get the model depicted. As a result of the announcement, both A and B do know that at least one of them is dirty. We can read this from the model.

Now suppose there is a second (and final) announcement that says that neither knows they are dirty (an announcement can express facts about the situation as well as epistemic facts about the knowledge held by the agents). We then update similarly the model by suppressing the worlds which do not satisfy the content of the announcement, or equivalently by keeping the worlds which do satisfy the announcement. This update process thus yields the following model:

$\boxed{\boxed{s : p, q}}$

By interpreting this model, we get that A and B both know that they are dirty, which seems to contradict the content of the announcement. Actually, we see here at work one of the main features of the update process: a proposition is not necessarily true after being announced. That is what we technically call "self-persistence" (see [vB03]) and this problem arises for epistemic formulas (unlike propositional formulas). One must not confuse the announcement and the update induced by this announcement, which might cancel some of the information encoded in the announcement (like in our example).

Logic of Public Announcement

Now we roughly present the main ingredients of a logic called the logic of public information, which formalizes these ideas and combines epistemic and dynamic logic (see [vB03]).

We have seen that a public announcement of a proposition φ changes the current epistemic model M with actual world w as follows:

Eliminate all worlds which currently do not satisfy φ.

So the universe we consider is composed of states which are epistemic models and we have got several relations from one state (corresponding to an epistemic action model) to another labelled with any epistemic formula A. We strike them when the final state is the epistemic model of the original state updated with the epistemic formula A. These relations (for each formula) are partial functions since there is always at most a single updated model possible after an announcement.

Our language is then composed of assertions of the form $[A!]\varphi$ which intuitively means that after a truthful announcement of A, φ holds. The logic of this system merges epistemic with dynamic logic. There is a complete and decidable axiomatization with key axioms:

$< A! > p \leftrightarrow A\&p$ for atomic facts p
$< A! > \neg\varphi \leftrightarrow A\&\neg < A! > \varphi$
$< A! > \varphi \vee \psi \leftrightarrow < A! > \varphi \vee < A! > \psi$
$< A! > K_i\varphi \leftrightarrow A\&K_i(A \rightarrow < A! > \varphi)$
(or equivalently: $[A!]K_j\varphi \leftrightarrow A \rightarrow K_j(A \rightarrow [A!]\varphi)$)

We can also incorporate common knowledge.

What is important in this survey is that we view epistemic updates (after a public announcement) as relations between epistemic models. In the language, we then refer to updates through modalities ($[A!]$). These two features will also be present in the more general approach of the next chapter that we now set out.

1.1.2 The General Case: BMS theory

We are going to present a framework which deals with more general actions than the simple case of public announcement of the last chapter.

A Simple Intuitive Example

First we are going to consider a simple example that we will name from now on the "coin example". This is the following:

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

We classically represent this scenario by the following epistemic modal model:

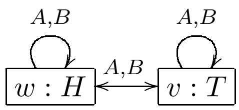

where w is the actual world.

Note that here epistemic relations are not equivalence relations but simply binary relations. A link $w \rightarrow_A v$ intuitively means that dwelling in the possible world w, A considers the possible world v possible (or has access to it).

Now suppose that A cheats and secretly opens the box. The resulting situation is depicted as follow:

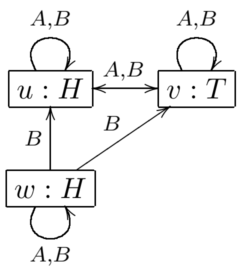

In the actual world w, A *knows* that the coin is Heads up but B *believes* that A doesn't know whether the coin is Heads or Tails. Note that in the BMS approach, knowledge and beliefs are somehow intertwined as we can clearly see in this example. The distinction will be clearly made in our approach (see chapter 2).

What we want to model in order to have a complete and faithful rendering of the process is the way the agents epistemically perceive the action, and how we update our model (the first diagram in our example) thanks to this piece of information, giving rise to a new model (the second diagram in our example). To come back to our general introduction, we would like to determine Σ and $\otimes$. Σ is called the simple action structure and $\otimes$ the update product.

The Core of BMS

We are going to give straightaway the definition of a simple action structure with its intended interpretation, followed by the definition of the update product, with its intended interpretation.

<u>simple action structure</u>

Definition 1.1.1 A simple action structure is a tuple $\Sigma = (\Sigma, \rightarrow_j, Pre)$ where:

- Σ is a set of simple actions.
- $\rightarrow_j$ is a family of relations on Σ indexed by the set of agents.
- $Pre : \Sigma \longrightarrow \Phi$ is a function from Σ to the collection of all epistemic propositions.

$\triangleleft$

Σ intuitively represents the possible actions, just as states intuitively represent possible worlds for modal models. The accessibility relations $\rightarrow_j$ express the uncertainty the agents have about the action taking place. All this is completely similar to modal logic. The introduction of Pre is necessary to ensure that an action σ can be performed in a certain world w: the world at stake w must fulfill this precondition $Pre(\sigma)$ (see below the update product).

To come back to our example of cheating, the corresponding simple action structure is:

$$\begin{tikzcd} \boxed{\sigma : H} \arrow[loop above, "A"] \arrow[r, "B"] & \boxed{\tau : True} \arrow[loop above, "{A,B}"] \end{tikzcd}$$

σ represents the action of cheating with precondition that the coin must be Heads up (which is as expected because the coin is actually heads up). τ represents the action where nothing happens, with precondition $True$ because it can be performed everywhere. The accessibility says that B *believes* nothing happens whereas A *knows* that he has cheated. We still see here the blend between the notions of knowledge and belief.

Now we can precisely study the update product of a simple action structure with a simple action structure.

Update product

Definition 1.1.2 Let $M = (W, \rightarrow_j, V)$ be an epistemic modal model and $\Sigma = (\Sigma, \rightarrow_j, V)$ a simple action structure. We define their update product to be the epistemic model
$M \otimes \Sigma = (W \otimes \Sigma, \rightarrow'_j, V')$, where

- $W \otimes \Sigma = \{(w, \sigma) \in W \times \Sigma; w \in V(Pre(\sigma))\}$.
- $(w, \sigma) \rightarrow'_j (v, \tau)$ iff $w \rightarrow_j v$ and $\sigma \rightarrow_j \tau$.
- $V'(p) = \{(w, \sigma) \in W \otimes \Sigma; w \in V(p)\}$

$\triangleleft$

The worlds that we consider after the update are all the possible worlds resulting from the performance of one of the actions in one of the worlds, under the assumption that the action can "possibly" take place in the corresponding world (expressed by Pre).

The product arrows $\rightarrow'_j$ on the output frame represent agent j's epistemic uncertainty about the output state. The intuition is that for the particularly simple kinds of actions that we take as "simple actions", the uncertainty regarding the action is assumed to be independent of the uncertainty regarding the current state. This independence allows us to modify these two uncertainties in order to compute the uncertainty regarding the output state: if whenever the input state is w, agent j thinks the input might be some other state w', and if whenever the current action happening is σ, agent j thinks the current action might be some other action σ', then whenever the output state (w, σ) is reached, agent j thinks the alternative output state (w, σ') might have been reached. Moreover, these are the only output states that j considers possible.

The definition of the valuation intuitively means that our actions, if successful, do not change the facts. Thus we are concerned with purely epistemic actions (that is why we call it epistemic action structure: it renders the *epistemic* structure of an action perceived by several agents).

To come back to our earlier example, we can then easily check that we get from the update product of the epistemic model and the epistemic action model of cheating the expected epistemic model.

The additional material necessary to build a genuine logical system

What we just set out forms the essence of the BMS approach. Now we need some more technical apparatus in order to build a genuine logical system. We will just give roughly the necessary definitions to achieve it, with a few comments.

Definition 1.1.3 An epistemic action model is a pair (Σ, Γ) consisting of a simple action structure Σ and a set Γ of designated simple actions. $\triangleleft$

Each of the simple actions $\gamma \in \Gamma$ can be thought as being a possible "deterministic resolution" of the non-deterministic action a.

Definition 1.1.4 An epistemic relation between epistemic models $M = (W, \rightarrow_j, V)$ and $M' = (W', \rightarrow'_j, V')$ is a relation between the sets W and W'. We write $r : M \longrightarrow M'$ for this.

An update $\mathtt{r}$ is a pair of operations on the set of epistemic models,
$\mathtt{r} = (M \mapsto M(\mathtt{r}), M \longrightarrow \mathtt{r}_M)$
where for each epistemic model M,$\mathtt{r}_M : M \longrightarrow M(\mathtt{r})$ is an epistemic relation. We call $M \longrightarrow M(\mathtt{r})$ the update map, and $M \mapsto \mathtt{r}_M$ the update relation. $\lhd$

This definition exemplifies what we said in section 1.1.1 about the way we represent an update, namely by a relation between models.

Now, from an epistemic action model, we can get an update.

Definition 1.1.5 Let (Σ, Γ) be an epistemic action model. We define an update that we also denote (Σ, Γ) as follows:

1. $M(\Sigma, \Gamma) = M \otimes \Sigma$.
2. $w(\Sigma, \Gamma)_M(w', \sigma)$ iff $w = w'$ and $\sigma \in \Gamma$

We call this the update induced by (Σ, Γ) $\lhd$

Finally the language is composed of two sorts of objects: the programs (denoted π) and the sentences (such as $[\pi]\varphi, K_j\varphi$). The interpretation of programs are updates. So the interpretation of $[\pi]\varphi$ is:

$M, w \models [\pi]\varphi$ iff for all v such that $w[\![\pi]\!]_M v$,$M([\![\pi]\!]), v \models \varphi$.

Again, this exemplifies what we have said in section 1.1.1 about the fact that we refer to updates through modalities.

Final remark of section 1.1: The BMS system will be the logical system on which we will base our further investigations. Our combined system will be very close to it, so we need to stress its importance.

1.2 The Belief Revision Theory

Introduction: Now, we will present the second ingredient of our merge system: the belief revision theory. We will particularly focus on the AGM theory (see [GardRott95]) and stress the AGM postulates. We will also set out the very important (for our account) W.Spohn's theory of ordinal conditional function as a "constructive" way (just as A.Grove models) to look at belief revision theory (see [Spohn90],[Spohn88]).

1.2.1 The AGM Theory

Belief revision deals with changes that must undergo a database representing a belief state (of an agent for example) after adding to (or subtracting from) the database information. This phenomenon is involved under different forms in many fields.

Generally speaking, some fundamental questions about this phenomenon appear almost immediately. First, what database format will be used? This will be the motivation of the first paragraph: "representing belief states". Second, how are the choices made concerning what to retract (if necessary) from the database, and what conditions must fulfill the outcome after a change? This will be the concern of the remaining paragraphs and the following section about W.Spohn's theory. In answering these questions, a small number of basic rationality postulates, or as we might equally call them, integrity constraints, can be considered as operative and will often be used in the following:

1. The belief in a database should be kept consistent whenever possible.
2. If the beliefs in a database logically entail a sentence A, then A should be included in the database.
3. The amount of the information lost in a belief change should be kept to a minimal.
4. In so far as some beliefs are considered more important or entrenched than others, one should retract the least important ones.

There are actually different modes of performing belief changes. We will adopt the following, called logic constrained, which takes the integrity constraints as constraints for the very process of belief change (as already announced) and operate directly on the current logically closed database without taking into account the old inputs that enabled to form the current belief state. This method of theory change is most eminently instantiated by C.Alchourron-P.Gardenfors-D.Markinson (AGM), which will form the focus of this account.

In the AGM approach, one can distinguish three main types of belief changes:

Expansion A new sentence is added to a belief system K(the database) regardless of the consequences of the larger set so formed. The belief system that results from expanding K by a sentence A together with the logical consequences will be denoted $K + A$ and consists of $\{B : K \cup \{A\} \vdash B\}$.

Revision A new sentence that is (typically) inconsistent with a belief system K is added, but in order that the resulting belief system be consistent - satisfies integrity constraint 1 - some of the old sentences in K are deleted. The result of revising K by a sentence A will be denoted $K * A$.

Contraction A sentence in K is retracted without adding any new facts. In order that the resulting system satisfies integrity constraint 2, some other sentences from K must be given up. The result of contracting K with respect to the sentence A will be denoted $K \doteq A$.

We will focus in this paper only on the revision and expansion operations.

Note that in these three cases, the format of representation remains the same after a change.

The only belief change that poses problems are revision and contraction: there are several outcomes possible and we have to choose one of them (i.e. there is indeterminacy). Indeed, an approach based solely on logic is not enough to determine it.

Hence, we will have to propose constructions of the revision process (that is a revision function *): this will be the A.Grove semantics and W.Spohn's theory. However, we have also to propose first postulates (regulations) for such constructions (that is rationality criteria for revisions and contractions of the databases) that the database K and its revision $(K * A)$ must verify: that will be the AGM postulates.

Representing Belief States

The most common representation of belief states in computational context and the one adopted in this account are sentential or propositional in the sense that the elements constituting the database are coded as formulas representing sentences. The underlying logic chosen includes classical propositional logic, that is monotonic and compact and that validates the cut rule and the deduction theorem. If H is a set of sentences we use the notation $Cn(H)$ for all consequences of H.

Another question that must be answered when representing a state of belief is whether the justifications for the beliefs should be part of the representation or not. The affirmative answer leads to the foundation theory and the negative one to the coherence one. We shall consider in this account the second theory. According to the coherence theory, the main objectives are to maintain consistency in the revised epistemic state that guarantee sufficient overall coherence (which is in accord with our integrity constraints).

Besides the syntactic representation mentioned above, we will see that we can also use a semantic one. We will then consider two main representations: belief set(for the syntactic part) and possible world models (for the semantic part).

Belief set. A belief set is a set K of sentences that is closed under logical consequences (i.e. satisfies integrity constraint 2). The interpretation of such a set is that it contains all the sentences that are accepted in the represented belief state.

Possible worlds models. An obvious objection to using sets of sentences as representations of belief state is that the objects of belief are normally not sentences but rather the contents of sentences, that is propositions. Thereby, the basic (and classical) semantic ideas connecting sentences with propositions is then that a sentence expresses a given proposition iff it is true in exactly those possible worlds that constitute the set of worlds representing the proposition. Then we can represent a belief set by a set W of possible worlds. The epistemic interpretation of W is that it is the narrowest set of possible worlds in which the individual in the represented belief state believes that the actual world is located.

Moreover, there is a very close correspondence between belief sets and possible world models. For any set W of possible worlds, we can define a corresponding belief set K_W as the set of those sentences that are true in all worlds in W. Conversely, for any belief set K, we can define a corresponding possible worlds model W_K by identifying the possible worlds in W_K with the maximal consistent extensions of K. Then we say that a sentence A is true in such an extension W iff $A \in w$.

The AGM Postulates

We will offer a basic presentation of the postulates without entering into details. They can easily be rediscovered by examining them carefully. However, we wish to mention that they have been determined in order to respect as much as possible the four integrity constraints of the beginning and be in accord with the intuition.

$(K*1)$ **(closure)** for any sentence A and any belief set K, $K*A$ is a belief set. (recall introduction)

$(K*2)$ **(success)** $A \in K*A$.

$(K*3)$ **(expansion 1)** $K*A \subseteq K+A$.

$(K*4)$ **(expansion 2)** if $\neg A \notin K$ then $K+A \subseteq K*A$.

$(K*6)$ **(extensionality)** $K*A = K_\perp$ only if $\vdash \neg A$ ($K_\perp$ is the inconsistent logic)

$(K*6)$ **(extensionality)** if $\vdash A \leftrightarrow B$ then $K*A$=$K*B$.

The postulates $(K*1)$-$(K*6)$ are elementary requirements that connect K, A and $K*A$. The final two conditions concern composite belief revisions of the form $K*A \wedge B$:

$(K*7)$ **(conjunction 1)** $K*A \wedge B \subseteq (K*A)+B$.

$(K*8)$ **(conjunction 2, rational monotony)** if $\neg B \notin K*A$ then $(K*A)+B \subseteq K*A \wedge B$.

We next turn to consequences of these 8 postulates.

$(K*8r)$ **(disjunction)** $K*A \vee B \subseteq Cn(K*A \cup K*B)$.

$(K*7c)$ **(cut)** if $B \in K*A$ then $K*A \wedge B \subseteq K*A$.

$(K*8c)$ **(cautious monotony)** if $B \in K*A$ then $K*A \subseteq K*A \wedge B$.

As already mentioned in the introduction, the postulates $(K*1)$-$(K*8)$ do not uniquely characterize the revision $K*A$ in terms of only K and A, and logical properties are not sufficient here.

Grove Models

As mentioned in the introduction, logic alone cannot tell us how to revise a given database (e.g. a belief set). We need some additional information, a selection mechanism (i.e. a revision function), in order to be able to decide rationally which sentences to give up and which to keep. In this chapter we will only investigate the A.Grove method. In the next chapter, we will present W.Spohn's theory.

We consider a language L of propositional logic. We naturally (classically) define an interpretation I of L, a model of a sentence, and let $Mod(A)$ (and $Mod(H)$) denote the set of all models of A (and H respectively) and $\mathcal{I}$ the set of all interpretations.

We will define an ordering (depending on a belief set K) on the set of all interpretations and then use this ordering to decide which interpretations should constitute models of $K * A$, (and thus indirectly determine $K * A$ in this way). Such an ordering should reflect the fact that some models of A are closer models to K than others.

A.Grove defines

Definition 1.2.1 $M(K) = \cup\{K \perp B; B \in K \text{ and } \nvdash B\}$ where $K \perp B$ is the set of all maximal subsets of K that do not entail B. $\triangleleft$

We then get the following proposition:

Proposition 1.2.2 *There is a one to one correspondence between the elements of $M(K)$ and $\mathcal{I} - Mod(K)$.*

Using the fact that there is an ordering of the elements of $M(K)$, we can define an ordering of the elements of $\mathcal{I} - Mod(K)$ that we extend to the whole set $\mathcal{I}$. We then notice that in order to respect $(K * 3)$ and $(K * 4)$, the models of K must be the smallest elements in $\mathcal{I}$. Technically we define:

Definition 1.2.3 A relation $\leq_K$ over $\mathcal{I}$ depending on K is said persistent if:

1. If $I \in Mod(K)$ then $I \leq_K J$ for all interpretations J.

2. If $I \in Mod(K)$ and $J \notin Mod(K)$, then $I \leq_K J$.

$\triangleleft$

As a consequence we obtain the following theorem:

Theorem 1.2.4 *A revision function * satisfies $(K * 1)$ to $(K * 8)$ iff there exists a persistent total pre-ordering $\leq_K$ such that $Mod(K * A) = Min(Mod(A), \leq_K)$.*

Another model-theoretic approach has been proposed by Grove which uses a system of "spheres" centered on the set of possible worlds where K is satisfiable (i.e. $Mod(K)$).

Definition 1.2.5 A system of spheres is a collection of subsets of interpretations which satisfies the following conditions:

(S1) S is totally ordered by $\subseteq$; that is, if S and S' are in S, then either $S \subseteq S'$ or $S' \subseteq S$.

(S2) $Mod(K)$ is the $\subseteq$-minimum element of S.

(S3) The set of all interpretations is in S.

(S4) If A is a sentence and there is a sphere S in S intersecting $Mod(A)$, then there is a smallest sphere in S intersecting $Mod(A)$.

$\lhd$

Note that we can actually draw a correspondence between $\leq_K$ and a system of spheres.

Next, let $\mathcal{S}_A$ denote the smallest sphere in S intersecting $Mod(A)$ (which exists according to (S4)). Grove then proves the following representation result:

Theorem 1.2.6 *A revision function * satisfies* $(K * 1)$ *to* $(K * 8)$ *iff there exists a system S of spheres such that* $Mod(K * A) = Mod(A) \cap \mathcal{S}_A$.

Note that the condition 3 of integrity constraint is respected here because $\mathcal{S}_A$ is the set of the closest worlds to $Mod(K)$ in which A holds. We understand here the purpose of (S4).

We can see that these two theorems somehow link the syntax (belief set) and semantic (possible worlds models) approaches.

1.2.2 W.Spohn's Approach to Belief Revision Theory

Grove models can be seen as a step towards expressing a degree of plausibility of possible worlds, in the sense that the system of spheres implicitly assign a ranking among possible worlds. A clear attempt to achieve this has been proposed by W.Spohn in his theory of ordinal conditional functions, apart the many other approaches.

We must however note that W.Spohn's approach considers only a single agent , unlike our future multi-agent approach (see chapter 2) based on W.Spohn's single one.

We now dive into the matter. (Note that our account is a simplified version of the real theory.)

Definition 1.2.7 An ordinal conditional function is a function κ from a given set W of possible worlds into the class of ordinals such that some possible worlds are assigned the smallest ordinal 0.

$\lhd$

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

The plausibility ranking of possible worlds can be extended to a ranking of propositions viewed as a set of worlds (the worlds where it is satisfiable):

Definition 1.2.8 Let A be a proposition, which in our setting corresponds to a set of worlds. The plausibility of the proposition A is defined by

$\kappa(A) = min\{\kappa(w); w \in A\}$ $\lhd$

This definition entails that the plausibility ranking of propositions has the following properties:

Proposition 1.2.9 *For all propositions A, either $\kappa(A) = 0$ or $\kappa(\neg A) = 0$ or both.*
For all non empty propositions A and B, $\kappa(A \cup B) = min\{\kappa(A), \kappa(B)\}$

We then say that a formula A is believed when $\kappa^{-1}(0) \subseteq A$. Besides this crude notion of belief, we can talk of greater or lesser plausibility or firmness of belief. We then say that A is believed more firmly than B iff $\kappa(\neg A) > \kappa(\neg B) > 0$, that is if the most plausible worlds outside A are less plausible than the most plausible worlds outside B (note that this condition implies that A and B are believed by the last proposition).

We can even extrapolate this definition to any kind of propositions even if they are not believed:

Definition 1.2.10 Let A and B be propositions. A is more plausible than B iff $\kappa(A) < \kappa(B)$ or $\kappa(\neg B) < \kappa(\neg A)$. $\triangleleft$

Now we turn our attention to how W.Spohn deals with epistemic changes such as revision. A key point is that the epistemic inputs are not only propositions but propositions together with a degree of plausibility. We shall note this input as a pair (A, α) where A is the proposition and α the degree of plausibility with which the formula should be believed by the agent after his revision of beliefs. We call this "degree of plausibility" also "degree of firmness", to be coherent with what has just been said about plausibility of propositions.

To do so, W.Spohn introduces the auxiliary concept of A-part of κ:

Definition 1.2.11 Let κ be an ordinal conditional function. The A-part of κ is the function $\kappa(. \mid A)$ defined on A for which for all $w \in A$, $\kappa(w \mid A) = -\kappa(A) + \kappa(w)$ $\triangleleft$

(For technical reasons, W.Spohn here makes use of the notion of left sided subtraction of ordinals, which is defined as follows: if a and b are two ordinals with $a \leq b$, then $-a + b$ is the uniquely determined ordinal c such that $a + c = b$).

One might say that the A-part of κ is the restriction of κ to A shifted to 0, that is, in such a way that $\kappa(A \mid A) = 0$.

With the aid of this concept, we can now define the ordinal conditional function $\kappa * (A, \alpha)$ (called the (A, α)-conditionalisation of κ) representing the new state of belief:

Definition 1.2.12 Let A be a proposition such that $A \neq \emptyset$

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

$\triangleleft$

Thus the (A,α)-conditionalisation of κ is the union of the A-part of κ and the $\neg A$-part of κ shifted up by α degrees of plausibility. It follows from the definition that $\kappa * (A, \alpha)(A) = 0$ and $\kappa * (A, \alpha)(\neg A) = \alpha$. Hence A is believed in $\kappa * (A, \alpha)$ with firmness α.

It conforms to the intuitive requirement that getting informed about A does not change the epistemic state restricted to A or $\neg A$, i.e. the plausibility grading within A or A^c. In other words, the (A, α)-conditionalisation of κ leaves the A-part as well as the A^c-part of κ

unchanged; they are only shifted in relation to one another. Thereby, we give meaning to the relative distances of possible worlds.

Now, if $\neg A$ is accepted in κ, that is, if $\kappa(\neg A) = 0 < \kappa(A)$, then the process of forming the (A, α)-conditionalisation (where $\alpha > 0$) is a generalization of a revision of belief sets. To make this more precise we need the following natural definitions.

Definition 1.2.13 The belief set K *associated with* the ordinal conditional function κ is the set of all propositions believed in κ $\lhd$

Definition 1.2.14 $K * A$ is the belief set associated with $\kappa * (A, \alpha)$ where $\alpha > 0$ $\lhd$

Then finally, we use these newly defined sets in the following proposition.

Proposition 1.2.15 *The revision function * defined on belief sets by the last two definitions satisfies the AGM postulates* $(K * 1) - (K * 8)$

Hence we indeed can see (A, α)-conditionalisation as a generalisation of belief revision. Moreover, when $\kappa(A) = 0$, the (A, α)-conditionalisation of κ corresponds to an expansion.

Final remark of section 1.2: W.Spohn's theory is very interesting for our purpose. It will turn out that it will play a central role for our purpose sketched in the general introduction.

Conclusion of chapter 1: We clearly see that the BMS approach is very similar to the approach sketched through the assumptions of the general introduction. The only difference is that it deals only with knowledge, and that the notions of knowledge and beliefs are somehow intertwined and identified. Moreover AGM deals explicitly with update of beliefs, namely with the notion of revision. These two approaches seem somehow complementary for our concerns. So, it makes sense to look for a way in which these two systems can be merged.

Chapter 2

A Combined System

Introduction: Now we will merge the two systems. One of the obvious point that needs to be done is to introduce the double notion of belief and knowledge; this will provide more overtones and precision. This will be the concern of the section 2.1. Afterwards, we will need to redefine our notions of update and action structure in order to incorporate the belief feature (see section 2.2 and 2.3). Concerning the belief update, it should be close to an AGM belief revision as said in the general introduction. W.Spohn's theory will turn out to be inspiring for that purpose.

2.1 The Static Part

In this chapter we will set out how we mathematically model a social situation involving knowledge and beliefs (see general introduction) via the notion of belief epistemic model. We will see that the proposed modelisation is actually a kind of extension of the basic case of knowledge.

2.1.1 The Notion of Belief Epistemic Model

We start with a simple and intuitive example from [vDL03].

Consider a situation where an operator j monitors a control system that consists of a light and a fan. Both can be on and off. Neither is known. The light is normally on. Both light and fan are in fact off.

There are thus four possible worlds corresponding to the actual world: light and fan on, light on and fan off,etc... Each of them is indistinguishable from the others for the operator. We model it this way using an epistemic model:

$$\begin{tikzcd} \boxed{s : p, q} \arrow[r, dash, "j"] \arrow[d, dash, "j"] & \boxed{t : p, \neg q} \arrow[d, dash, "j"] \\ \boxed{u : \neg p, q} \arrow[r, dash, "j"'] & \boxed{\boxed{v : \neg p, \neg q}} \end{tikzcd}$$

s,t,u,v are the possible worlds of the model. p and q are propositional variables whose intuitive meaning is respectively "light is on" and "fan is on". Now we can add a belief structure to our model (thanks to the information provided about the actual situation). To do so we introduce a plausibility grading: each world has a certain plausibility, is more or less likely to correspond to the actual world according to the operator. We grade it with a natural number ranging from 0 to M (where M is an arbitrary fixed natural number): the lower the grading is, the more the world is considered plausible.

So in our example we set $\kappa_j(s) = \kappa_j(t) = 0 < \kappa_j(u) = \kappa_j(v)$. $\kappa_j(x)$ represents the plausibility grading of the world x for j.

Indeed the light is normally on, so naturally the operator considers more plausible s and t than u and v because in these worlds s and t , the light is on.

We can now give the general definition:

Definition 2.1.1 A belief epistemic model $M = (W, \{\sim_j; j \in G\}, \{\kappa_j; j \in G\}, V, w_0)$ is a tuple where:

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

$\lhd$

- The interpretation of points 1,2,5,6 is clear.
- For point 3, we strike a link $\sim_j$ between two possible worlds whenever the agent j has no information which enables him to distinguish between these two worlds.
- The assignment of κ_j is carried out by considering individually each $\sim_j$ equivalence class. Indeed , a $\sim_j$-equivalence class corresponds to indistinguishable worlds for the agent j. Among them, the agent has a plausibility preference, and that is what we model by κ_j. So, for each $\sim_j$-equivalence class, we assign a plausibility grading. However, because the $\sim_j$ equivalence classes form a partition of W, it amounts to assigning a plausibility to every world. Nevertheless, we must stress the fact that this assignment is made relative to each equivalence class. A plausibility value for a world and an agent j makes sense only when it is considered relatively to the corresponding $\sim_j$-equivalence class, even if at first sight, by looking at the definition, we are tempted to think that the assignment is made individually and absolutely for each world.

  The grading goes from 0 to M where M is an arbitrary fixed natural number. The more a world is plausible for the agent j, the closer its plausibility value is to 0, and the less plausible a world is the closer it is to M. The intuitive import of introducing a natural number M is that beyond a certain degree of plausibility the agent cannot distinguish two different worlds of different plausibility. This is coherent and that is what often happens in reality. It is also introduced for compactness reason as we will see in section 3.3.2.

Important remarks:

- One would be tempted to think that this plausibility assignment for any equivalence should be made relative to a particular world of this equivalence class. That is, according to which world the agent dwells in, his plausibility assignment is different, because the information available in this world is different from the information available in another world of the epistemic class. Yet, this is a spurious reasoning. Indeed, the information available in a world for an agent j is what he knows and what he considers possible and not the formulas satisfiable in this world. Yet, in any world of the same $\sim_j$-equivalence class, what j knows and considers possible is the same. Hence, dwelling in any world of the same equivalence class the agent j will base his plausibility grading on the same information. So, the plausibility grading will be identical in every world of the $\sim_j$-equivalence class. That is what we wanted to show.

  To put it differently, intuitively if I cannot distinguish between two worlds, then I should believe the same things in these two worlds.

- This use of an equivalence relation $\sim_j$ and a plausibility grading κ_j replaces the former relation used in [BMS03]. Indeed formally, the classical arrow of [BMS03] $v \rightarrow_j w$ is replaced in our setting by an equivalence relation $v \sim_j w$ and by the plausibility grading κ_j where $\kappa_j(v) > \kappa_j(w)$. So, the two notions of knowledge and belief which were somehow intertwined in [BMS03] are clearly separated by use of the plausibility grading. Indeed, for [BMS03], the notion of belief was a particular type of knowledge and was formally identified with the notion of knowledge in their system. Here, the plausibility grading provides us with more richness and overtones concerning these two notions of knowledge and belief, as we will explicitly see in the forthcoming study of the language (see section 3.1 and 2.1.3).

- Obviously, this plausibility grading corresponds to the plausibility grading of W.Spohn's theory (see section 1.2.2). And for a particular situation, each $\sim_j$-equivalence class in our approach corresponds to the set of possible worlds for the agent j in W.Spohn's approach. However, we don't assume that for each $\sim_j$-equivalence class there is a world of plausibility 0 like in W.Spohn's theory, because as we will see in section 2.3 (update mechanism) it is quite possible that after an update, no world of the resulting equivalence class has the plausibility 0 (whereas it was the case before the update). Nevertheless, without this assumption, the treatment and underlying ideas are exactly the same.

2.1.2 Examples

We reconsider the examples of the previous chapter:

- Muddy children.

  The corresponding belief epistemic model for the initial situation is

$$\begin{tikzcd} \boxed{\boxed{s : p, q}} \arrow[r, dash, "A"] \arrow[d, dash, "B"] & \boxed{t : p, \neg q} \arrow[d, dash, "B"] \\ \boxed{u : \neg p, q} \arrow[r, dash, "A"'] & \boxed{v : \neg p, \neg q} \end{tikzcd}$$

  with $\kappa_A(s) = \kappa_A(t) = \kappa_A(u) = \kappa_A(v) = 0$, and the same for B.

- coin example (see section 1.1.2).

  The corresponding belief epistemic model for the initial situation is :

$$\begin{tikzcd} \boxed{\boxed{s : H}} \arrow[r, dash, "A{,}B"] & \boxed{t : T} \end{tikzcd}$$

  with $\kappa_A(s) = \kappa_A(t) = \kappa_B(s) = \kappa_B(t) = 0$

  The model resulting from a private announcement to A (suspected by B) that the coin is lying heads up is the following:

$$\begin{tikzcd} \boxed{(s) : H} \arrow[r, dash, "A{,}B"] \arrow[d, dash, "A"] & \boxed{(t) : T} \arrow[dl, dash, "A"] \\ \boxed{\boxed{(s) : H}} & \end{tikzcd}$$

  with $\kappa_B(u) = 0$ for all u, and $\kappa_A(s) = 1, \kappa_A(s) = \kappa_A(t) = 0$

2.1.3 The Language for the Static Part

We can easily define a language $\mathcal{L}_{St}$ for belief epistemic models (St for static):

$\mathcal{L}_{St}\ \varphi := p \mid \neg\varphi \mid \varphi \wedge \psi \mid K_j\varphi \mid B_j^k\varphi$ where $k \in \mathbb{Z}$

Its semantics is the following:

Let M be a belief epistemic model and w a world in M. We only give the semantics for κ_j and B_j^k, the rest is classical.

$M, w \models K_j\varphi$ iff for all v st $w \sim_j v$, $M, v \models \varphi$.
$M, w \models B_j^k\varphi$ iff for all v st $w \sim_j v$ and $\kappa_j(v) \leq k$, $M, v \models \varphi$.

The intuitive meaning of $M, w \models B_j^k \varphi$ is that in a certain world w, j believes with plausibility (a degree of) at most k that φ is true.

K_j is the usual knowledge operator.

2.2 The Dynamic Part

2.2.1 The Notion of Belief Epistemic Action Model

In this section we will define the notion of epistemic action model. This notion is intended to model how the agents perceive a particular action taking place from the point of view of their knowledge and beliefs (see general introduction). The way the action is perceived by the agents is, for example, which representation of the actual action (that we call possible action) they consider to be the most plausible, what they know about what the other agents believe (about the actual action taking place), etc... This is of course relevant for our study because it will determine the resulting epistemic and belief situation after the performance of the actual action, which we will model by an updated belief epistemic action model (see next section). Thus our epistemic action has to be designed in order to catch the essential elements that pertain to the way the agents epistemically perceive the action.

First an intuitive example. Consider again the coin example. The action that we will consider is a completely private truthful announcement to A that the coin is lying Heads up. The corresponding belief epistemic action model is:

$$\boxed{\boxed{\sigma : H}} \overset{B}{\rule[0.5ex]{3em}{0.4pt}} \boxed{\tau : True}$$

with $\Sigma = \{\sigma, \tau\}$, $\sigma \sim_j \tau$, $Pre(\sigma) = H$, $Pre(\tau) = True$, $\kappa_B^*(\tau) = 0$, $\kappa_B^*(\sigma) = \alpha$.

In this actual action, there are two possible representations: nothing happens (for B), or there is a truthful announcement of H (for A) (the actual action). So we model it by respectively introducing the "simple" actions τ and σ. The actual action is σ and is denoted by a double bordered line. Both have "preconditions" which express the fact that they can not be performed in any possible world. Indeed, for σ to be performed in any possible world, the coin must be H because it is a *truthful* announcement of H. On the contrary τ can be performed in every world because nothing happens, so the precondition is a propositional tautology (which is satisfied in any world). Now B has no objective information which would enable her to distinguish σ from τ because σ is by definition a private action and thus an action that she cannot perceive. So, there is a $\sim_B$ link between σ and τ. On the contrary, A has at his disposal actual evidence that enables him to distinguish σ from τ: he is aware of the private announcement to him. So there is no $\sim_A$ link between σ and τ. Now B has clearly a preference between σ and τ: she believes the actual action is τ. We model it by introducing a plausibility assignment κ_B^*.

Now B can suspect more or less that there has been a private announcement to A and we can perfectly express in our setting this degree of suspicion thanks to α. Indeed the higher α is, the less B suspects A of having cheated; or to put it differently, the lower α is, the more B suspects A of having cheated. This is a natural consequence of our definition of the degree of plausibility: the higher value is assigned to a world or an action, the less this world or action is considered plausible.

Note that this action model can also correspond to other types of actions than just a private announcement, like for example cheating (A lifted the cup and checked the outcome H).

Now the general definition:

Definition 2.2.1 A belief epistemic action model Σ is a tuple $(\Sigma, \sim_j, \kappa_j^*, Pre, \sigma_0)$ such that:

1. Σ is a set of simple actions.
2. $\sim_j$ is an equivalence relation indexed by the set of agents on Σ.
3. κ_j^* is a function from the set of simple actions to the set of natural numbers $\{0, ..., M\}$ indexed by the set of agents, such that for each $\sim_j$-equivalence class, (at least) one of the actions is assigned the plausibility 0.
4. Pre is a function from the set of simple actions to the formulas of $\mathcal{L}_{St}$.
5. w_0 is the possible world corresponding to the actual world.

$\lhd$

2.2.2 Intended Interpretation and Examples

1. As already sketched in the last example Σ is the set of the possible determinations of the actual actions. We call them "simple" actions because we assume that we can not decompose the actual structure of a simple action σ into (sub) actions temporally smaller, whose succession would form the original action σ and that would still be relevant for our study.
2. We set $\sigma \sim_j \tau$ (with σ and τ simple actions) when j has no information which would enable him to distinguish σ from τ.
3. κ_j^* expresses the plausibility preference that the agent has among actions that he can not objectively distinguish.
4. Pre assigns to each simple action a precondition that a world w must fulfill in order for this σ to be performed in this world w.
5. σ_0 denotes the actual action.

In the following, for convenience we will denote a belief epistemic action model simply an action model.

We see in this definition that the belief epistemic action model is very similar to the corresponding BMS one. It is normal because as already stressed in the static part (see section 1.1.2) the main difference of our system with BMS is that we replace the crude notion of knowledge (represented in BMS by arrows $\rightarrow_j$) by the double notion of knowledge and belief (represented in our system by $\sim_j$ and κ_j^*).

Some other examples are offered in order to better understand this notion of belief epistemic action model.

Simple Examples

- Public announcement of φ to all the agents:

  $$\boxed{\boxed{\sigma : \varphi}}$$

  The intuition behind this model is clear.

- Announcement of φ with degree of firmness α_j to all the agents j:

  $$\boxed{\sigma : \varphi} \overset{j}{\text{——}} \boxed{\boxed{\tau : \neg\varphi}}$$

  with $\kappa_j^*(\sigma) = 0$ and $\kappa_j^*(\tau) = \alpha_j$

  Here j accepts the information φ with "degree of firmness" $\kappa_j^*(\tau) = \alpha_j$. This degree of firmness corresponds to the degree of firmness introduced by W.Spohn (see section 1.2.2). Note that we must not confound the notion of degree of plausibility and of degree of firmness. The first is assigned directly to a world (or action) and bears on this world (or action). the second is an additional and optional piece of information which bears on the degree of plausibility. It is *indirectly* assigned via a different (and linked) world or set of worlds (see section 1.1.2). Precisely and intuitively, it expresses at which degree of plausibility , the degree of plausibility of a world (or an action) is assigned. It can thus be viewed as a "meta"-degree of plausibility which reveals an extra introspection of the agent. However, we can not talk of degree of firmness in any case (see section 1.2.2). It just provides more overtones and precision to the first notion of plausibility. Classically, we talk of degree of firmness when the degree of plausibility of the world at stake is 0 (Anyway, that is why we talked of firmness of beliefs only in section 1.2.2).

  We will see (see section 4.3) that the update (see section 2.3) of a belief epistemic action model with this action model corresponds to a belief revision for agent j of his belief with the proposition φ (with a degree of plausibility α_j) exactly as in W.Spohn's work.

- In the muddy children example, A lies to B by telling him that she is not dirty:

$$\boxed{\boxed{\sigma : q \wedge K_A q}} \overset{B}{\text{———}} \boxed{\tau : \neg q \wedge K_A \neg q}$$

with $\kappa^*_B(\tau) = 0$ and $\kappa^*_B(\sigma) = \alpha$

σ represents the simple action of A lying by telling $\neg q$ (this is the actual action). This presupposes that q is true and A knows it ($K_A q$)because otherwise it would not be a lie. τ represents the announcement that $\neg q$ is the case. Concerning τ, for the announcement to be made presupposes of course that $\neg q$ is the case but also that A knows that $\neg q$ is true because in this simple action τ, A is supposed to be faithful. B has no information which enables her to distinguish whether A is lying or not, that is the actual action is σ or τ. So the $\sim_B$ link between σ and τ. Now B believes the actual action is τ, so $\kappa^*_B(\tau) = 0 < \kappa^*_B(\sigma)$. B can more or less believe what A told her and it is possible to express it in our model thanks to the number α. Indeed the higher α is, the more B believes what B told her; or to put it differently, the closer α is to $0 = \kappa_B(\sigma)$, the less B believes what A told her. Indeed if the two degrees of plausibility are close, B will somehow consider the two simple actions (lie and truthful announcement) equally plausible; whereas if there is a gap between them, one of the action will be much more plausible than the other. These considerations are similar to those we had in the cheating example and are very close to the notion of degree of firmness spelled out in the last example.

A more complex example

This example involves several $\sim_j$-equivalence classes for each agent j. It is the following. Consider again the "coin" situation.

The action:"Given H or T under the cup, with A and B ignorant about it, each suspects the other of having "cheated" (i.e. lifted the cup and checked the outcome). Actually nobody cheated and H is under the cup. A now asks B: 'Do you know whether it is H or T'"

The action model leading up to that question is the following:

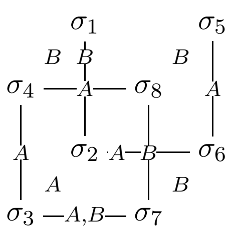

It has 8 possible actions:

1. σ_1: A and B learn that H. $Pre(\sigma_1) = H$.

2. σ_2: A learns that H and B not. $Pre(\sigma_2) = H$.

3. σ_3(the actual action): neither A nor B learn that H. $Pre(\sigma_3) = H$.

4. σ_4: B learns that H and A not. $Pre(\sigma_4) = H$.

5. σ_5: A and B both learn that T. $Pre(\sigma_5) = T$.

6. σ_6: A learns that T and B not. $Pre(\sigma_6) = T$.

7. σ_7: Neither A nor B learn that T. $Pre(\sigma_7) = T$.

8. σ_8: B learns that T and A not. $Pre(\sigma_8) = T$.

There are clearly for A and B three equivalence classes.

- For A: $E_1^A = \{\sigma_1, \sigma_2\}, E_2^A = \{\sigma_5, \sigma_6\}, E_3^A = \{\sigma_3, \sigma_4\sigma_7, \sigma_8\}$
- For B: $E_1^B = \{\sigma_1, \sigma_4\}, E_2^B = \{\sigma_5, \sigma_8\}, E_3^B = \{\sigma_2, \sigma_3, \sigma_6, \sigma_7\}$

Now, IF A learns that H or learns that T, he considers it more likely that B does not (i.e. prefers the actions where B remains ignorant).That is

In $E_1^A : \kappa_A^*(\sigma_1) > \kappa_A^*(\sigma_2) = 0$.
In $E_2^A : \kappa_A^*(\sigma_5) > \kappa_A^*(\sigma_6) = 0$.

Similarly for B:

In $E_1^B : \kappa_B^*(\sigma_1) > \kappa_B^*(\sigma_4) = 0$.
In $E_2^B : \kappa_B^*(\sigma_5) > \kappa_B^*(\sigma_8) = 0$.

Dually, IF A doesn't learn H or T, he considers it more likely that B does (i.e. prefers the actions where B learns H or T). That is,

In $E_3^A : \kappa_A^*(\sigma_3) > \kappa_A^*(\sigma_4) = 0, \kappa_A^*(\sigma_7) > \kappa_A^*(\sigma_8) = 0$

Similarly for B:

In $E_3^B : \kappa_B^*(\sigma_3) > \kappa_B(\sigma_2) = 0, \kappa_B^*(\sigma_7) > \kappa_B^*(\sigma_6) = 0$.

2.2.3 Technical Apparatus

These definitions will be necessary for the next chapter and the chapter about completeness. They are just necessary technical devices and don't bring any new conceptual insights.

Definition 2.2.2 An action signature is a structure $\Sigma_{sg} = (\Sigma, \sim_j, \kappa_j^*, (\sigma_1, ..., \sigma_n))$ where:

1. Σ is a set composed of the simple actions $\{\sigma_1, ..., \sigma_n\}$ (which is a list without repetition)

2. For each agent j, $\sim_j$ is an equivalence relation on Σ.

3. For each agent j, κ_j^* is a function from the simple actions $\{\sigma_1, ..., \sigma_n\}$ to the class of epistemic formulas.

$\triangleleft$

We can get from an action signature a full belief epistemic action model thanks to this definition:

Definition 2.2.3 Let Σ_{sg} be an action signature, σ a simple action of Σ, and $\psi = \psi_1, ..., \psi_n$ a list of epistemic propositions. A signature-based action model, based on the signature Σ_{sg}, σ and ψ, that we note $(\Sigma_{sg}, \sigma)(\psi_1, ..., \psi_n)$, is the belief epistemic action model $(\Sigma, \sigma, \sim_j, \kappa_j^*, Pre)$ defined by:

1. $\Sigma, \sim_j, \kappa_j^*$ are the ones given by Σ_{sg}.
2. the actual action is σ.
3. for each $j \in \{1, ..., n\}, Pre(\sigma_j) = \psi_j$.

$\triangleleft$

2.3 The Update Mechanism

We have now set the two main components. First, the belief epistemic model M as a formal counterpart (see general introduction) of the way an actual situation s is perceived by the agents from the point of view of their beliefs and knowledge. Second, the belief epistemic action model Σ as a formal counterpart (see intro) of the way an actual situation a is perceived by the agents from the point of view of their beliefs and knowledge. Now in reality the agents update their knowledge and beliefs according to these two pieces of information: action a and situation s, giving rise to a new actual situation $s \times a$. We would like to render this update formally with a mathematical model $\otimes$ such that $M \otimes \Sigma$ would be a new belief epistemic model which would be a formal counterpart of $s \times a$ (see general introduction).

So, first we will define the update product of a belief epistemic model with a belief epistemic action model (corresponding to the $\otimes$ above). Second, we will give its intended interpretation. Third, we will present some examples of update. Fourth, we will define some mathematical objects necessary to get a full logical system out of this semantically driven approach. As part of it and as the last chapter, we will define a full language incorporating dynamic features, and more precisely this update product. This full language will be an extension of the static language defined in section 2.1.3.

2.3.1 The Notion of Update Product

Given a belief epistemic model $\mathbf{M} = (M, \sim_j, V, \kappa_j, w_0)$ and a belief epistemic action model $\Sigma = (\Sigma, \sim_j, Pre, \kappa_j^*, \sigma_0)$ we define their update product to be the epistemic action model

$M \otimes \Sigma = (W \otimes \Sigma, w_o', \sim_j', V', \kappa_j')$

given by the following: the new states are pairs of old states w and simple action σ

$W \otimes \Sigma = \{(w, \sigma) \in W \times \Sigma; w \in V(Pre(\sigma))\}$.

The new accessibility relation is taken to be the products of the corresponding accessibility relations in the two frames: i.e., for (w, σ), $(w', \sigma') \in W \times \Sigma$, we put

$(w, \sigma) \sim'_j (w', \sigma')$ iff $w \sim_j w'$ and $\sigma \sim_j \sigma'$

The new valuation map is essentially given by the old valuation:

$V'(p) = \{(w, \sigma) \in W \times \Sigma; w \in V(p)\}$ and

$w'_0 = (w_0, \sigma_0)$.

Note that this state necessarily exists because it corresponds to what actually happened in the real world. The fact that $w_0 \in V(Pre(\sigma_0))$ is a meta(and necessary)-assumption of our system. Eventually, we define κ'_j as follows:

$\kappa'_j((w, \tau)) = Cut_M(\kappa_j(w) + \kappa^*_j(\sigma) - \kappa^w_j(\varphi))$

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

$$Cut_M(x) = \begin{cases} x & \text{if } 0 \leqslant x \leqslant M. \\ M & \text{if } x > M \end{cases}$$

2.3.2 Intended Interpretation

$W \otimes \Sigma$: A new possible world (w, σ) in the resulting model corresponds to the resulting situation of performing the action corresponding to σ in the world corresponding to w. However, the action σ can be performed in w only if the precondition $Pre(\sigma)$ of the action σ is satisfied in this world: actions cannot be performed in an arbitrary world, they presuppose some "material" preconditions in the world. For example, if the action is a truthful announcement of Head (cf coin example), then this action can be carried out only in worlds where H holds because it is a *truthful* announcement. (Note that it is completely similar to the BMS system.)

$\sim'_j$: Two new worlds are indistinguishable for j if the worlds they come from were already indistinguishable before the action and if the actions which have been performed in these worlds are indistinguishable for j.

Thus uncertainty among states can only come from existing uncertainty via indistinguishable actions. This is natural and that is one would expect. (Note that this is again completely similar to the BMS system.)

V': We essentially take the same valuation as the one of the input model. If a world w "survives" an action, then the same fact p holds at the output world (w, σ) as at the input world w. This means that our actions, when performed, do not change the facts.

This condition can of course be relaxed in various ways, to allow for fact changing actions. However, in this paper we are primarily concerned with purely *belief epistemic* action, such as the examples in the last chapter. (Note that it is still completely similar to the BMS system.)

w'_0: This state necessarily exists because it corresponds to what actually happened in the real world. The fact that $w_0 \in V(Pre(\sigma_0))$ is a meta(and necessary)-assumption of our system.

κ'_j: Here we update the plausibility assignment thanks to the new information provided by the action model Σ.

$\kappa_j(t) + \kappa^*_j(\sigma) - \kappa^w_j(\varphi)$ (*) is the core of the update. Cut_M is a technical device in order that the new plausibility assignment fits in the range of the plausibility scale of the new belief epistemic model (i.e. fits in the set $\{0, .., M\}$). The role of this function is also intuitively correct because as already said in section 2.1.1 (about the maximum value M), after a certain stage of plausibility value (denoted here M), the agents cannot distinguish between two different plausibility values. So it is as if the plausibility values beyond this stage M for the agent were the same as the value M of this stage. This thereby gives an intuitive motivation to this function Cut_M, besides its technical role.

Now we will motivate and interpret why (*), the "formal" new plausibility of (w, σ) (Cut_M corresponding to the "actual" plausibility) has this shape.

(*)= [1]+[2]+[3] ,where

[1]=plausibility for j that t is the actual world.

[2]=plausibility for j that σ is the actual action taking place in w.

[3]=relativization of the plausibility [1] to the relevant worlds: the worlds where the action σ will take place and that may correspond for j to the actual situation w.

We assume that the plausibilities of the actions are independent from the particular world in which they are performed (see remarks below)

1 $=\kappa_j(w)$ clearly.

2 $=\kappa^*_j(\sigma)$. The plausibility of the actions are independent between each other, there is no relative plausibility between two actions, these plausibilities are absolute and assigned once by the agent j. Moreover, these plausibilities are also independent from the world they can take place, as just said. So they are doubly independent: from the current world and from the other actions. This entails that in any world s where the action σ is taking place, j will consider that τ can be performed with plausibility $\kappa^*_j(\tau)$.

3 $=-\kappa_j^w(\varphi) = min\{\kappa_j(u); u \sim_j w$ and $u \in V(Pre(\sigma))\}$. We relativize the ordinal assignment to the relevant worlds, that is the worlds where σ will take place and that may correspond for j to the actual situation w. Indeed the former ordinal assignment κ_j made sense only if we considered all the worlds that may correspond for j to the actual situation w. However, formally, now the action σ is taking place, so j only has to consider and restrict his attention to the set of worlds where the action σ will take place (and that correspond to the ones that possibly represent the actual situation w he dwells in) because the other ones do not play a role anymore. So j must relativize his plausibility ordering to this set by rescaling the ordinal assignment, in order to start again and deal with the action σ with a self defined plausibility ordering and abstract from the former one.

Important remarks:

- We could motivate our definition of plausibility update somewhat differently. Indeed, it is intuitively clear that,

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

  Or more generally and precisely,

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

  So we would be tempted to assign roughly to the worlds accessible from w where the action σ has taken place the plausibility $\kappa_j^*(\tau)$. Yet doing so, we would lose part of the information and overtones present in the former model among the worlds where the action σ has taken place (and that are accessible from w). So we add to $\kappa_j^*(\sigma)$ the expression $\kappa_j(w) - \kappa_j^w(\varphi)$ in order to incorporate this former information. This represents, as already explained, the plausibility for j that his current world is w ($\kappa_j(w)$) relativized to the world where the action σ will take place ($-\kappa_j^w(\varphi)$). Doing that, we will still be sure that there is a (v, σ) such that $\kappa_j(v, \sigma) = \kappa_j^*(\sigma)$ and $\kappa_j^*(\sigma)$ is the minimum plausibility among the worlds where the action σ has taken place. That is what we wanted (cf (*)). Moreover, we also keep track of the former plausibility structure by this addition of $\kappa_j(w) - \kappa_j^w(\varphi)$.

  these thoughts are very well exemplified by the last example of the section 2.3.3

- These two interpretations (this one just set and the one above) are somehow equivalent. Indeed, they both stress the priority of the plausibility assignment of the action model upon the plausibility assignment of the model. For the first interpretation, this priority is expressed in the interpretation of the subtraction of $\kappa_j^w(\varphi)$ (namely via the notion of relativization, which indirectly gives priority to the action at stake). Moreover, these action plausibilities of the action structure are a guideline for the plausibility update process. That is according to their relative distribution and repartition that the agents globally revise (update) their beliefs. This last thought will be exemplified by our last example.

- Of course there are several ways to update plausibilities depending on the background motivations. In this work (as already mentioned in the introduction), we are interested in one performing a genuine belief revision. Thus we want to be close to an ideal belief revision. That's what explains our intuitive choices in the first remark and in the main comments

  We will also see in section 4.1 that what makes our plausibility update one of belief revision is the subtraction of $\kappa_j^w(\varphi)$.

- We see in the first remark that the plausibility update can be done separately from the rest of the update.

- This plausibility update process is essentially local: it considers a particular world and a particular action only, abstracting the other possible actions. All this thanks to the subtraction of $\kappa_j^w(\varphi)$.

- The formula (*) has to be put in parallel with the (φ, α)-conditionalisation of W.Spohn's theory (see section 1.2.2). They are very similar. We will see that this similarity allows for belief revision as conceived by W.Spohn in section 1.2.2. Indeed it is quite possible to perform a revision of beliefs as viewed by W.Spohn for each agent, satisfying moreover the 8 AGM postulates by updating, in the way we just described, a model with a particular action model

- Last but not least, we have assumed that the plausibility of the actions are independent from the worlds in which they are performed. However, this is wrong. Indeed, for example consider a hazy announcement of a formula φ that the agent j cannot distinguish from φ'. In a $\sim_j$-equivalence class where the agent j knows more formulas that logically imply φ than φ', j will find the announcement φ more plausible than φ' because he will have more actual evidence at his disposal to think so. However, in another $\sim_j$-equivalence class of the same model where the agent j knows more formulas that logically imply φ' than φ, j will find the announcement of φ less plausible than φ' because he will have less actual evidence that would prompt him to think so. This extra assumption is a limitation of our system but an attempt to avoid it would render our system much more complicated. It just restricts the set of actions we consider to actions whose epistemic and belief feature doesn't change according to which world they are performed.

2.3.3 Examples

- *Update with a public announcement.*

  1. Let us come back to the muddy children example. According to section 2.1.2, the initial model has this shape:

     $$\begin{tikzcd} \boxed{\boxed{s : p, q}} \arrow[r, dash, "A"] \arrow[d, dash, "B"] & \boxed{t : p, \neg q} \arrow[d, dash, "B"] \\ \boxed{u : \neg p, q} \arrow[r, dash, "A"'] & \boxed{v : \neg p, \neg q} \end{tikzcd}$$

     with $\kappa_A(s) = \kappa_A(t) = \kappa_A(u) = \kappa_A(v) = 0$, and the same for B.

  2. A public announcement by the father telling that at least one of them is dirty corresponds to this action model:

     $$\boxed{\boxed{\sigma : p \vee q}}$$

     with $\kappa^*_A(\sigma) = \kappa^*_B(\sigma) = 0$.

  3. Now our update product yields:

     $$\begin{tikzcd} \boxed{\boxed{(s, \sigma) : p, q}} \arrow[r, dash, "A"] \arrow[d, dash, "B"] & \boxed{(t, \sigma) : \neg p, q} \\ \boxed{(u, \sigma) : \neg p, q} & \end{tikzcd}$$

     This is the expected model (see section 1.1.1).

- *Update with a private announcement.*

  1. Consider the "coin" example (see section 1.1.2)

     $$\begin{tikzcd} \boxed{\boxed{w : H}} \arrow[r, dash, "A{,}B"] & \boxed{v : T} \end{tikzcd}$$
     with $\kappa_A(w) = \kappa_A(v) = \kappa_B(w) = \kappa_B(v) = 0$

2. A private announcement to A that the coin is lying Heads up corresponds to the following action model:

$$\begin{tikzcd} \boxed{\sigma : H} \arrow[r, dash, "B"] & \tau : True \end{tikzcd}$$

with $\kappa^*_B(\tau) = 0$, $\kappa^*_B(\sigma) = \alpha$.

3. The update product yields:

$$\begin{tikzcd} (w,\tau) : H \arrow[r, dash, "{A,B}"] \arrow[d, dash, "B"] & (v,\tau) : T \arrow[dl, dash, "B"] \\ \boxed{(w,\sigma) : H} & \end{tikzcd}$$

with $\kappa_A(u) = 0$ for all u, and $\kappa_B(w,\sigma) = \alpha, \kappa_B(w,\tau) = \kappa_B(v,\tau) = 0$

In the actual world, A knows that the coin is H, B believes that A does not know whether the coin is H or T and B does not know whether it is H or T. This is the expected intuitive outcome.

- *Update with a lie*

1. Consider again the muddy children example:

$$\begin{tikzcd} \boxed{s : p, q} \arrow[r, dash, "A"] \arrow[d, dash, "B"] & t : p, \neg q \arrow[d, dash, "B"] \\ u : \neg p, q \arrow[r, dash, "A"'] & v : \neg p, \neg q \end{tikzcd}$$

with $\kappa_A(s) = \kappa_A(t) = \kappa_A(u) = \kappa_A(v) = 0$, and the same for B.

2. A lie by A to B by telling her that she is not dirty corresponds to the following action model:

$$\begin{tikzcd} \boxed{\sigma : q \wedge K_A q} \arrow[r, dash, "B"] & \tau : \neg q \wedge K_A \neg q \end{tikzcd}$$

with $\kappa^*_B(\tau) = 0$ and $\kappa^*_B(\sigma) = \alpha$

3. The update product yields:

$$\begin{tikzcd} (s,\sigma) : p, q \arrow[r, dash, "A"] \arrow[d, dash, "B"] & (t,\sigma) : p, \neg q \arrow[d, dash, "B"] \\ (u,\tau) : \neg p, q \arrow[r, dash, "A"'] & (v,\tau) : \neg p, \neg q \end{tikzcd}$$

with $\kappa_B(u,\tau) = \kappa_A(v,\tau) = 0$, and $\kappa_B(s,\sigma) = \kappa_B(t,\sigma) = \alpha$ the same for B.

There is no change for A in his beliefs and knowledge. In the actual world, B believes she is not dirty. That is what we expected.

- *Announcement of a formula with a degree of firmness*

  1. Consider the following belief epistemic model with a single agent a.:

     $M = (W = W_1 \sqcup W_2, w_0, \sim_a, \kappa_a, V)$ such that:
     - $W_1 \cap W_2 = \emptyset$.
     - $\sim_a = W \times W$.
     - $\kappa_a(p) = 0$, $\kappa_a(\neg p) = 2$ for some propositional variable p
     - $V(p) = W_1, V(\neg p) = W_2$.

     p is more believed than $\neg p$ by the agent in the actual world.
  2. An announcement of $\neg p$ with a degree of firmness α to all (actually only one) the agents:

     $$\boxed{\sigma : p} \overset{a}{\text{———}} \boxed{\boxed{\tau : \neg p}}$$

     with $\kappa^*_a(\tau) = 0$ and $\kappa^*_a(\sigma) = \alpha$
  3. The update product yields the model $M' = (W, w_0, \sim_0, \kappa_0, V)$ where $\kappa_0(p) = 1$ and $\kappa_0(\neg p) = 0$.

     Hence p is now less believed than $\neg p$ by the agent in the actual world. This exemplifies what we have just said in the remarks, namely that the plausibility assignment of the action model has priority over the former plausibility assignment of the model (thanks to the relativization term). Indeed, a believes the announcement is τ, and after the update, a believes what is true in the worlds where τ has taken place.

     We will see in section 4.3 that this action model corresponds to a revision of beliefs (satisfying the 8 AGM postulates).

2.3.4 Technical Apparatus

These definitions are almost equal to the ones of the BMS theory (see section 1.1.2).

First some general definitions.

Definition 2.3.1 A belief epistemic relation between belief epistemic models M and M' is a relation between the sets W and W'. We write $r : M \longrightarrow M'$ for this.

An update $\mathtt{r}$ is a pair of operations on the set of belief epistemic model,

$\mathtt{r} = (M \longrightarrow M(\mathtt{r}), M \mapsto \mathtt{r}_M)$,

where for each belief epistemic model M, $\mathtt{r}_M : M \longrightarrow M(\mathtt{r})$ is a belief epistemic relation. We call $M \longrightarrow M(\mathtt{r})$ the update map, and $M \mapsto \mathtt{r}_M$ the update relation.

$\lhd$

The collection of updates is closed in various ways:

Definition 2.3.2 Let $\mathtt{r}$ and $\mathtt{s}$ be belief epistemic updates.

- Their composition is again a belief epistemic update, where $M(\mathtt{r.s}) = M(\mathtt{r})(\mathtt{s})$ and $(\mathtt{r.s})_M = \mathtt{r}_M.\mathtt{s}_{M(r)}$ (we are writing relational composition from left to right).
- Their sum $\mathtt{r+s}$ is again a belief epistemic update, where $M(\mathtt{r+s}) = M(\mathtt{r}) + M(\mathtt{s})$, and $(\mathtt{r+s})_M = \mathtt{r}_M + \mathtt{s}_M$. (we are writing $+$ for the disjoint union operation on set, extended in the obvious way to relations).

$\lhd$

Now a link with what has been done: we can get updates induced naturally by belief epistemic action models.

Definition 2.3.3 Let Σ be a belief epistemic action model. We define an update which we also denote Σ (corresponding to the $\mathtt{r}$ of the definition of update) as follows:

1. $M(\Sigma) = M \otimes \Sigma$.
2. For all $w \in M, (v, \sigma) \in M(\Sigma)$, $w\Sigma_M(v, \sigma)$ iff w=v and σ is the actual action of Σ

We call this the update induced by Σ. $\lhd$

2.3.5 The Full Language

Now that we have at our disposal the material concerning the dynamic part and the update part of our system, we can extend the static language $\mathcal{L}_{St}$ defined in section 2.1.3 in order to cover and express this dynamic feature in our very language. We must, however, mention the fact that we do not consider in our study the notion of common knowledge and common belief.

First we have to say that our logical system will deal with only a fixed and unique type of action. So our reasoning will be based on a single type of action. This type of action will be determined by our choice of an action signature Σ_{sg}. Let n be the number of actions in Σ_{sg}. First we define the syntax of the full language $\mathcal{L}(\Sigma_{sg})$.

Definition 2.3.4 The syntax of the language $\mathcal{L}(\Sigma_{sg})$ is defined by,

- Sentences $\varphi := True \mid p \mid \neg\varphi \mid \varphi \wedge \psi \mid K_j\varphi \mid B_j^k\varphi \mid [\pi]\varphi$
  where $k \in \mathbb{Z}$

- Programs $\pi := \sigma_i\psi_1, ..., \psi_n \mid \pi + \rho \mid \pi.\rho$.

$\triangleleft$

We note that we have two sorts of syntactic objects: sentences and programs. We call programs of the form $\sigma, \psi_1, ..., \psi_n$ simple programs. Par abus de notation, we will also denote this simple program σ, ψ without any ambiguity.

Now the semantics of $\mathcal{L}(\Sigma_{sg})$. The semantics of programs are updates:

- $[\![\sigma_i, \psi_1, ..., \psi_n]\!] = (\Sigma_{sg}, \sigma_i)(\psi_1, ..., \psi_n)$ (recall that the left term of the equality rather denotes an update than a belief epistemic action model, see definition 2.3.3).
- $[\![\pi, \rho]\!] = [\![\pi]\!].[\![\rho]\!]$.
- $[\![\pi + \rho]\!] = [\![\pi]\!] + [\![\rho]\!]$.

The operations + and . were defined in definition 2.3.2.

Remark:

- We see here the usefulness of introducing this notion of action signature: it enables us, once fixed a type of action, to consider different displays of this type of action (by changing the precondition). that is why we will call our logic like in BMS theory, for example, logic of public announcement, logic of private announcement...

Now we can give the semantics of the sentences.

Let $M = (W, w_0, \sim_j, \kappa_j^*, V)$ be a belief epistemic model.

- $M, w \models p$ iff $w \in V(p)$
- $M, w \models \neg\varphi$ iff $M, w \not\models \varphi$.
- $M, w \models \varphi \wedge \psi$ iff $M, w \models \varphi$ and $M, w \models \psi$.
- $M, w \models K_j\varphi$ iff for all v such that $w \sim_j v$ $M, v \models \varphi$.
- $M, w \models B_j^k\varphi$ iff for all v such that $w \sim_j v$ and $\kappa_j(v) \leq k$, $M, v \models \varphi$.
- $M, w \models [\pi]\varphi$ iff for all v such that $w[\![\pi]\!]_M v$, $M([\![\pi]\!]), v \models \varphi$.

Remarks:

- The semantics of $[\pi]\varphi$ motivates the way we defined an update in definition 2.3.1. Indeed we need two data. The first concerning the updated model $M(\llbracket\pi\rrbracket)$. The second concerning the relation between the original model and the updated model: this is the role of the update map $\llbracket\pi\rrbracket_M$. these two data form the core of our definition of an update.

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

- Finally, in the syntax of B_j^k, k is in $\mathbb{Z}$. However, the values of κ_j are always included between 0 and $M > 0$, so it seems useless to assign such a big range for the degree of plausibility k. Actually, this extension is necessary because as we will see in the next chapter, in the axiom 9, it is quite possible formally that k goes beyond the range of κ_j (more precisely because $k = m + l - \kappa_j^*(\sigma_k)$)

Conclusion of chapter 2: We have got now a combined system which seems to fulfill our expectations. However, we don't know yet if the belief update is close to a belief revision. This will be the concern of the last chapter which will deal with several comparisons, notably the AGM postulates. Beforehand, we need to complete our system by determining a proof system, because our approach has been semantically driven so far. This is the concern of the following chapter.

Chapter 3

Logic of Combined Update and Revision

Introduction: In this chapter, we will present a proof system AX for our semantically driven system exposed so far, together with a soundness and completeness proof. Concerning the completeness proof, we will first define a sub-proof system AX' for which we prove completeness with respect to the semantics of $\mathcal{L}_{St}$. We will then prove that every formula of the proof system AX is provably equivalent in AX to a formula of AX'. This will enable us to get the completeness proof for the full proof system AX with respect to the semantics of $\mathcal{L}(\Sigma_{sg})$.

3.1 The Full Proof System

We set out in this chapter the proof system AX for $\mathcal{L}(\Sigma_{sg})$. The sub-proof system AX' corresponding to $\mathcal{L}_{St}$ (see section 2.1.3) is labelled with * :

1. * All propositional tautologies.

2. $\vdash [\pi](\varphi \to \psi) \to ([\pi]\varphi \to [\pi]\psi)$ $[\pi]$-distribution.
3. * $\vdash B_j^m(\varphi \to \psi) \to (B_j^m\varphi \to B_j^m\psi)$ B_j^m-distribution.
4. * $\vdash K_j(\varphi \to \psi) \to (K_j\varphi \to K_j\psi)$

5. $\vdash [\sigma_i, \psi]p \leftrightarrow (\psi_i \to p)$
6. $\vdash [\sigma_i, \psi]\neg\chi \leftrightarrow (\psi_i \to \neg[\sigma_i, \psi]\chi)$
7. $\vdash [\sigma_i, \psi]\varphi \wedge \chi \leftrightarrow ([\sigma_i, \psi]\varphi \wedge [\sigma_i, \psi]\chi)$
8. $\vdash [\sigma_i, \psi]K_j\varphi \leftrightarrow \{\psi_i \to \bigwedge\{K_j[\sigma_k, \psi]\varphi; \sigma_k \sim_j \sigma_i\}\}$

9. $\vdash [\sigma_i, \psi]B_j^m\varphi \leftrightarrow \{\psi_i \rightarrow \bigwedge\{B_j^{l-1}\neg\psi_k \wedge \neg B_j^l\neg\psi_k \rightarrow B_j^{m+l-\kappa_j^*(\sigma_k)}[\sigma_k, \psi]\varphi; \sigma_k \sim_j \sigma_i$ and $l \in \{0..M\}\}\}$ where $m < M$

10. $\vdash [\pi][\rho]\varphi \leftrightarrow [\pi.\rho]\varphi$

11. $\vdash [\pi + \rho]\varphi \leftrightarrow [\pi]\varphi \wedge [\rho]\varphi$

12. * $\vdash K_j\varphi \rightarrow \varphi$

13. * $\vdash B_j^k\varphi \rightarrow K_jB_j^m\varphi$ for all $m \in \mathbb{Z}$

14. * $\vdash \neg B_j^m\varphi \rightarrow K_j\neg B_j^m\varphi$ for all $m \in \mathbb{Z}$

15. * $\vdash B_j^m\varphi \rightarrow B_j^{m'}\varphi$ for all $m \geqslant m'$

16. * $\vdash K_j\varphi \leftrightarrow B_j^m\varphi$ for all $m \geqslant M$

17. * From $\vdash \varphi$ and $\vdash \varphi \rightarrow \psi$ infer $\vdash \psi$

18. * From $\vdash \varphi$ infer $\vdash B_j^m\varphi$ $[B_j]$-generalization

19. * From $\vdash \varphi$ infer $\vdash K_j\varphi$ $[K_j]$-generalization

20. From $\vdash \varphi$ infer $\vdash [\pi]\varphi$ $[\pi]$-generalization

Remark: The axioms of our proof system AX concerning the interaction between the notions of belief and knowledge are intuitively appealing. Indeed, axiom 13 says informally that "If you believe something, you know that you believe it". Axiom 14 says that "if you do not believe something, you know that you do not believe it". Axiom 15 says that "if you know something, you also believe it". In a sense, this is a sign whereby our choice of function κ_j as independent from the current world (see section 2.1.1) was intuitively correct. Moreover, axiom 8 is exactly the same as in the BMS system, which is not surprising, given that our approach is based on the one of BMS. Axiom 16 is only here for technical reasons (see section 3.2).

3.2 Completeness Proof for the Static Part

In this chapter, we are going to prove strong completeness of AX' with respect to the semantics of the language $\mathcal{L}_{St}$. We will need the following three lemmas:

Lemma 3.2.1 *Every AX'-consistent set $\Gamma \subseteq \mathcal{L}_{St}$ can be extended to a maximal AX'-consistent set w.r.t. $\mathcal{L}_{St}$. In addition, if Γ is a maximal AX'-consistent set, then it satisfies the following properties:*

1. *For every formula $\varphi \in \mathcal{L}_{St}$, exactly one of φ and $\neg\varphi$ is in Γ.*

2. $\varphi \wedge \psi \in \Gamma$ *iff* $\varphi \in \Gamma$ *and* $\psi \in \Gamma$.

3. *if* φ *and* $\varphi \to \psi$ *are both in* Γ, *then* ψ *is in* Γ.

4. *if* $\psi \in AX'$ *then* $\psi \in \Gamma$.

Proof: It is a classical proof that we don't give. $\square$

Lemma 3.2.2 *For all maximal consistent sets* V, W, *for all* $k \in \mathbb{Z}$,
$V/_{K_j} \subseteq W \to (V/_{B_j^k} \subseteq W \leftrightarrow W/_{B_j^k} \subseteq W)$

Proof: $V/_{K_j} \subseteq W \Leftrightarrow$ for all formula $\varphi, K_j\varphi \in V \to \varphi \in W$ (*)

We are going to show that under the assumption (*),
$B_j^k\varphi \in V \Leftrightarrow B_j^k\varphi \in W$ for all k, all formulas φ

- Assume $B_j^k\varphi \in V$.

  We know that $B_j^k\varphi \to K_jB_j^k\varphi \in V$ because it is an axiom (see lemma 3.2.1).
  So $K_jB_j^k\varphi \in V$.
  Then by (*), $B_j^k\varphi \in W$

- Assume $B_j^k\varphi \in W$.

  Assume $B_j^k\varphi \notin V$. Then $\neg B_j^k\varphi \in V$ by lemma. But $\neg B_j^k\varphi \to K_j\neg B_j^k\varphi \in V$.
  So $K_j\neg B_j^k\varphi \in V$.
  Then $\neg B_j^k\varphi \in W$ by (*)
  i.e. $B_j^k\varphi \notin W$ (by lemma 3.2.1). This is impossible.

  So $B_j^k\varphi \in V$.

We have proved $B_j^k\varphi \in V \Leftrightarrow B_j^k\varphi \in W$ for all k, all formulas φ.

Now, if $V/_{K_j} \subseteq W$
then $B_j^k\varphi \in V \Leftrightarrow B_j^k\varphi \in W$ for all φ
then $(B_j^k\varphi \in V \to \varphi \in W) \Leftrightarrow (B_j^k\varphi \in W \to \varphi \in W)$ for all formulas φ
then $V/_{B_j^k} \subseteq W \leftrightarrow W/_{B_j^k} \subseteq W$

So, $V/_{K_j} \subseteq W \to (V/_{B_j^k} \subseteq W \leftrightarrow W/_{B_j^k} \subseteq W)$ $\square$

Lemma 3.2.3 $V/_{B_j^k} \subseteq V/_{B_j^{k'}}$ *for all* $k \geq k'$.

Proof:
Let $\varphi \in V/_{B_j^k}$.
Then $B_j^k\varphi \in V$
but $B_j^k\varphi \rightarrow B_j^{k'}\varphi \in AX'$
so $B_j^k\varphi \rightarrow B_j^{k'}\varphi \in V$ by lemma 3.2.1
then $B_j^{k'}\varphi \in V$ by lemma 3.2.1 again
i.e. $\varphi \in V/_{B_j^{k'}}$ $\square$

Theorem 3.2.4 *AX' is a sound and strongly complete axiomatization w.r.t. the semantics of* $\mathcal{L}_{St}$.

Proof: Soundness can easily be checked. We only prove completeness.
To prove completeness, we must show that

Every AX'-consistent set Γ of formulas is satisfiable on some belief epistemic model M^c (*).

We prove it using a general technique that works for a wide variety of modal logics. We construct a special structure M^c which is a belief epistemic model called the canonical structure for AX'. M^c has a state w_V corresponding to every maximal AX'-consistent set V. Then we show:

$M^c, w_V \models \varphi$ iff $\varphi \in V$ (**)

That is, we show that a formula is true at a state w_V exactly if it is one of the formulas in V. Note that (**) suffices to prove (*), for by lemma 3.2.1, if Γ is AX'-consistent, then Γ is contained in some maximal AX'-consistent set V. From (**) it follows that $M^c, w_V \models \Gamma$, and so Γ is satisfiable in M^c. Therefore, Γ is satisfiable w.r.t. the semantics of $\mathcal{L}_{St}$, as desired.

We thus define M^c like that :

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

- $W^c = \{w_V; V$ maximal AX'-consistent set $\}$
- $\sim_j = \{(w_V, w_W); V/_{K_j} \subseteq W\}$ used before where $V/_{K_j} = \{\varphi; K_j\varphi \in V\}$
- $\kappa_j(w_W) = min\{k; W/_{B_j^k} \subseteq W\}$
- $w_W \in V(p)$ iff $p \in W$

1. First, we are going to show:

   $M^c, w_W \models \varphi$ iff $\varphi \in V$

We prove it by induction on φ.

The propositional case is straightforward. The knowledge case is classical. We only prove the belief case.

- Assume $\varphi = B_j^k\psi \in V$.
  i.e. $\psi \in V/_{B_j^k}$

  For all maximal consistent sets (m.c.s.) W such that $V/_{K_j} \subseteq W$, $V/_{B_j^k} \subseteq W \to \psi \in W$ because $\psi \in V_{B_j^k}$.
  i.e. For all W such that $V/_{K_j} \subseteq W$, and $V/_{B_j^k} \subseteq W$, $\psi \in W$.
  i.e. For all W such that $V/_{K_j} \subseteq W$, and $W/_{B_j^k} \subseteq W$, $\psi \in W$ by lemma 3.2.1.
  i.e. For all W such that $w_V \sim_j w_W$ and $\kappa_j(w_W) \leq k$, $\psi \in W$ by definition.
  i.e. For all $w_V \sim_j w_W$ and $\kappa_j(w_W) \leq k$, $M^c, w_W \models \psi$ by induction hypothesis.
  i.e. $M^c, w_V \models B_j^k\psi$.
  i.e. $M^c, w_V \models \varphi$.
- Assume $M^c, w_W \models \varphi = B_j^k\psi$

  $V_{B_j^k} \cup \{\neg\psi\}$ is not AX'-consistent.

  For suppose otherwise.
  There is then a m.c.s W such that $V/_{B_j^k} \cup \{\neg\psi\} \subseteq W$

  If $k \geq M$, then $V/_{B_j^k} = V/_{K_j}$ because $B_j^k\varphi \leftrightarrow K_j\varphi \in V$ (1).

  If $k < M$, then $V/_{B_j^M} \subseteq V/_{B_j^k}$ by lemma 3.2.3
  i.e. $V/_{K_j} \subseteq V/_{B_j^k}$ because of (1)

  In both cases, $V/_{K_j} \subseteq V/_{B_j^k}$

  Moreover, $V/_{B_j^k} \subseteq W$
  So, $V/_{K_j} \subseteq W$ and $V/_{B_j^k} \subseteq W$
  then $V/_{K_j} \subseteq W$ and $W/_{B_j^k} \subseteq W$ by lemma 3.2.2
  i.e. $w_V \sim_j w_W$ and $\kappa_j(w_W) \leq k$

  By the induction hypothesis, we have $M^c, w_W \models \neg\psi$ because $\neg\psi \in W$

  So, $M^c, w_V \models \neg B_j^k\psi$, which is impossible.

Since $V/_{B_j^k} \cup \{\neg\psi\}$ is not AX'-consistent, there must be some finite subset, say $\{\varphi_1, .., \varphi_n, \psi\}$ which is not AX consistent. Thus by propositional reasoning, we have:

$\vdash \varphi_1 \to (\varphi_2 \to (... \to (\varphi_n \to \psi)))$

By the axiom of B_j^k-generalisation we have
$\vdash B_j^k(\varphi_1 \to (\varphi_2 \to (... \to (\varphi_n \to \varphi))))$

By the axiom of B_j^k-distributivity and modus ponens,

$\vdash B_j^k\varphi_1 \to (B_j^k\varphi_2 \to ... \to (B_j^k\varphi_n \to B_j^k\psi))$.

So by lemma 3.2.1,
$B_j^k\varphi_1 \to (B_j^k\varphi_2 \to ... \to (B_j^k\varphi_n \to B_j^k\psi)) \in V$.
Then again by the lemma 3.2.1 because $B_j^k\varphi_i \in V$ for all i, we get
$B_j^k\psi \in V$, i.e. $\varphi \in V$

2. Now we have to prove that M^c is well-defined and that M^c is a belief epistemic model.

   - To prove definedness, we have to show that κ_j is well defined. It is the case because
     - $\kappa_j(w_W)$ exists for all w_W.

       Indeed $V/_{K_j} \subseteq V$ by axiom 12
       and $V/_{K_j} = V/_{B_j^M}$ by axiom 16
       So $\{k; V/_{B_j^k} \subseteq V\} \neq \emptyset$
       Then the minimum exists.

     - $\kappa_j(w_W) \leq k \Rightarrow$ for all $k' \geq k$ $\kappa_j(w_W) \leq k'$

       i.e. $V/_{B_j^k} \subseteq V \Rightarrow$ for all $k' \geq k$ $V_{B_j^{k'}} \subseteq V$
       This is true by lemma 3.2.3

     - $0 \leq \kappa_j(w_W) \leq M$ for all w_W by the first point.

   - The proof that M^c is a belief epistemic model amounts to prove that $\sim_j$ is an equivalence relation. This is a classical proof that we don't give, it uses the axioms 12-13-14(and 16).

     This ends the proof. $\square$

3.3 Completeness and Soundness Proofs for the Dynamic Part

3.3.1 Completeness

We are going to prove that every formula of the proof system AX is provably equivalent in AX to a formula of AX'. We will proceed step by step with successive lemmas.

Lemma 3.3.1 *If $\vdash \varphi \leftrightarrow \psi$ and $\vdash \varphi' \leftrightarrow \psi'$ then $\neg\varphi, \varphi \wedge \psi, B_j^m\varphi, [\pi]\varphi, K_j\varphi$ are respectively equivalent to $\neg\varphi', \varphi' \wedge \psi', B_j^m\varphi', [\pi]\varphi', K_j\varphi'$*

Proof: The cases $\neg\varphi, \varphi\wedge\psi$ are obvious (it only uses propositional tautologies). $B_j^m\varphi, [\pi]\varphi$ and $K_j\varphi$ are very similar. We will only treat $[\pi]\varphi$.

1. $\vdash \varphi \leftrightarrow \psi$ hypothesis.
2. $\vdash \varphi \rightarrow \psi$ propositional tautology and 1.
3. $\vdash [\pi](\varphi \rightarrow \psi)$ $[\pi]$-generalization.
4. $[\pi](\varphi \rightarrow \psi) \rightarrow ([\pi]\varphi \rightarrow [\pi]\psi)$ $[\pi]$-distribution.
5. $[\pi]\varphi \rightarrow [\pi]\psi$ modus ponens 3,4.

We reason similarly to prove $\vdash [\pi]\psi \rightarrow [\pi]\varphi$.

Now, with $\vdash p \rightarrow (q \rightarrow p \wedge q)$, we get $\vdash [\pi]\varphi \leftrightarrow [\pi]\psi$. $\square$

Now thanks to this lemma, we get

Proposition 3.3.2 *If $\vdash \psi \leftrightarrow \psi'$ and ψ is a subformula of φ, then $\vdash \varphi \leftrightarrow \varphi'$ where φ' is the formula obtained from φ by substituting ψ by ψ'*

Proof: By induction on φ.

- $\varphi = p$, then the only possibility for ψ is $\psi = \varphi$. So, obviously, $\vdash \varphi \leftrightarrow \varphi'$.
- induction step: it uses lemma 3.3.1.

Lemma 3.3.3 *For all $\varphi \in \mathcal{L}_{St}$ and all simple program $\sigma_i, \psi_1, .., \psi_n$, there is $\varphi_{St} \in \mathcal{L}_{St}$ such that*
$\vdash [\sigma_i, \psi_1, .., \psi_n]\varphi \leftrightarrow \varphi_{St}$.

Proof: By induction on φ

- $\varphi = p$. We then get the result by the axiom 5, by taking $\varphi' = \psi_i \rightarrow p$.
- $\varphi = \neg\chi$.
  Then $[\sigma_i, \psi_1, .., \psi_n]\neg\chi \leftrightarrow (\psi_i \rightarrow \neg[\sigma_i, \psi]\chi)$ by axiom 6.
  Yet, by the induction hypothesis, there is $\chi' \in \mathcal{L}_{St}$ such that $\vdash [\sigma_i, \psi]\chi \leftrightarrow \chi'$.
  So $\vdash [\sigma_i, \psi_1, .., \psi_n]\varphi \leftrightarrow (\psi_i \rightarrow \neg\chi')$ by proposition 3.3.2, and $\varphi' := (\psi_i \rightarrow \neg\chi') \in \mathcal{L}_{St}$.

- The other cases are dealt with similarly, using the fact that the axioms all reduce $[\sigma_i, \psi_1, .., \psi_n]\varphi$ to something with basic action modalities only over shorter static formulas. $\square$

Lemma 3.3.4 *For all $\varphi \in \mathcal{L}_{St}$ and all action programs π, there is $\varphi_{St} \in \mathcal{L}_{St}$ such that $\vdash [\pi]\varphi \leftrightarrow \varphi_{St}$*

Proof: By induction on the action program length.

- If π is a simple program, then we get the result by lemma 3.3.3.
- If $\pi = \tau + \rho$, $\vdash [\pi]\varphi \leftrightarrow [\tau]\varphi \wedge [\rho]\varphi$ by axiom 11.

  Yet by induction hypothesis, $\vdash [\tau]\varphi \leftrightarrow \varphi_1$ where $\varphi_1 \in \mathcal{L}_{St}$
  and $\vdash [\rho]\varphi \leftrightarrow \varphi_2$ where $\varphi_2 \in \mathcal{L}_{St}$.

  So, $\vdash [\pi]\varphi \leftrightarrow \varphi_1 \wedge \varphi_2$, thanks to proposition 3.3.2, and $\varphi_1 \wedge \varphi_2 \in \mathcal{L}_{St}$.
  Thereby, we get the expected result.
- If $\pi = \tau.\rho$,
  then $\vdash [\pi]\varphi \leftrightarrow [\tau][\rho]\varphi$ by axiom 10.
  Yet, by the induction hypothesis, $\vdash [\rho]\varphi \leftrightarrow \varphi_1$ with $\varphi_1 \in \mathcal{L}_{St}$.
  Then $\vdash [\pi]\varphi \leftrightarrow [\tau]\varphi_1$ thanks to proposition 3.3.2.
  However, again by the induction hypothesis, $\vdash \varphi_1 \leftrightarrow \varphi_2$ where $\varphi_2 \in \mathcal{L}_{St}$.

  So $\vdash [\pi]\varphi \leftrightarrow \varphi_2$. That is what we wanted. $\square$

Proposition 3.3.5 *For all $\varphi \in \mathcal{L}(\Sigma_{sg})$, there is $\varphi_{St} \in \mathcal{L}_{St}$ such that $\vdash \varphi \leftrightarrow \varphi_{St}$*

Proof: By the induction on the number of occurrence of action programs

- If there is no occurrence of action programs, then the result is clear.
- Let $\varphi \in \mathcal{L}(\Sigma_{sg})$ composed of $n > 0$ occurrences of action programs.

  Let ψ be a subformula of φ of the form $\psi = [\pi]\psi'$ where $\psi' \in \mathcal{L}_{St}$. Because $n > 0$, it is quite possible to find one (just pick an innermost occurrence of action formula of this shape). Now, by lemma 3.3.4, $\vdash \psi \leftrightarrow \psi_{St}$ where $\psi_{St} \in \mathcal{L}_{St}$. So by proposition 3.3.2, $\vdash \varphi \leftrightarrow \varphi'$ where φ' is obtained from φ by substituting ψ by ψ_{St}. Then clearly, in φ' there are less than n occurrences of action programs because the occurrence of $[\pi]$ has been suppressed.

  We can then apply the induction hypothesis, and there is $\varphi'' \in \mathcal{L}_{St}$ such that $\vdash \varphi' \leftrightarrow \varphi''$.

Thus, $\vdash \varphi \leftrightarrow \varphi''$ and $\varphi'' \in \mathcal{L}_{St}$. That is what we wanted. $\square$

Finally, we get the expected outcome:

Theorem 3.3.6 *AX is strongly complete with respect to the semantics of* $\mathcal{L}(\Sigma_{sg})$

Proof: Assume $T \models \varphi$ where $T \subseteq \mathcal{L}(\Sigma)$ and $\varphi \in \mathcal{L}(\Sigma)$.
For all $\psi \in T$, $\vdash \psi \leftrightarrow \psi_{St}$ where $\psi_{St} \in \mathcal{L}_{St}$ by proposition 3.3.5.

So, $\models T$ iff $\models T_{St}$ where $T_{St} = \{\psi_{St}; \psi \in T\}$.
Moreover, $\models \varphi$ iff $\models \varphi_{St}$.

So, $T \models \varphi$ iff $T_{St} \models \varphi_{St}$.

Now, by strong completeness of AX' with respect to the semantics of $\mathcal{L}_{St}$ (see theorem 3.2.4),
$T_{St} \vdash_{AX'} \varphi_{St}$.

However, because AX is an extension of AX',
$T_{St} \vdash \varphi_{St}$

So, $T \vdash \varphi_{St}$.
Then $T \vdash \varphi$ because $\vdash \varphi_{St} \leftrightarrow \varphi$. That is what we wanted. $\square$

3.3.2 Soundness

Theorem 3.3.7 *AX is sound w.r.t. the semantics of* $\mathcal{L}(\Sigma_{sg})$.

Proof:We only prove axioms 8 and 9, the others can be done without great difficulty.

We first prove the soundness of the axiom 9.

First note that in a belief epistemic model M,

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

It is indeed coherent because for all $l \neq l'$,
$\vdash \neg\{(B_j^{l-1}\neg\varphi \wedge \neg B_j^l\neg\varphi) \wedge (B_j^{l'-1}\neg\varphi \wedge \neg B_j^{l'}\neg\varphi)\}$

Now, let $k < M$

$M, w \models [\sigma_i, \psi] B_j^k \varphi$

iff, if $M, w \models \psi_i$ then $M(\Sigma), (w, \sigma_i) \models B_j^k \varphi$, where $M(\Sigma)$ is the updated model.

iff, if $M, w \models \psi_i$ then,
for all $(v, \sigma_j) \sim_j (w, \sigma_i)$ and $\kappa_j(v, \sigma_j) \leq k$, $M(\Sigma), (v, \sigma_j) \models \varphi$.

iff, if $M, w \models \psi_i$ then,
for all σ_j such that $\sigma_i \sim_j \sigma_j$,
for all v such that $w \sim_j v$ and $Cut_M(\kappa_j(v) - \kappa_j^w(\psi_j) + \kappa_j^*(\sigma_j)) \leq k$ and $M, v \models \psi_j$,
$M(\Sigma), (v, \sigma_j) \models \varphi$

iff, if $M, w \models \psi_i$ then,
for all σ_j such that $\sigma_i \sim_j \sigma_j$,
for all v such that $w \sim_j v$ and $\kappa_j(v) - \kappa_j^w(\psi_j) + \kappa_j^*(\sigma_j) \leq k$ then,
if $M, v \models \psi_j$ then $M(\Sigma), (v, \sigma_j) \models \varphi$

because if $k < M$, then $Cut_M(x) \leq k \Leftrightarrow x \leq k$.

iff, if $M, w \models \psi_i$ then,
for all σ_j such that $\sigma_i \sim_j \sigma_j$,
for all v such that $w \sim_j v$ and $\kappa_j(v) - \kappa_j^w(\psi_j) + \kappa_j^*(\sigma_j) \leq k$ then,
$M, v \models [\sigma_j, \psi]\varphi$

iff, if $M, w \models \psi_i$ then,
for all σ_j such that $\sigma_i \sim_j \sigma_j$, for all $l \in \{0, .., M\}$,
for all v such that $w \sim_j v$ and $\kappa_j^w(\psi_j) = l$ and $\kappa_j(v) \leq l - \kappa_j^*(\sigma_j) + k$ then,
$M, v \models [\sigma_j, \psi]\varphi$

iff, if $M, w \models \psi_i$ then,
for all σ_j such that $\sigma_i \sim_j \sigma_j$, for all $l \in \{0, .., M\}$,
if $\kappa_j^w(\psi_j) = l$ then
for all v such that $w \sim_j v$ and $\kappa_j(v) \leq l - \kappa_j^*(\sigma_j) + k$ then,
$M, v \models [\sigma_j, \psi]\varphi$

iff, if $M, w \models \psi_i$ then,
for all σ_j such that $\sigma_i \sim_j \sigma_j$, for all $l \in \{0, .., M\}$,
if $M, w \models B_j^{l-1} \neg\psi_j \wedge \neg B_j^l \neg\psi_j$ then
for all v such that $w \sim_j v$ and $\kappa_j(v) \leq l - \kappa_j^*(\sigma_j) + k$ then,
$M, v \models [\sigma_j, \psi]\varphi$

iff, if $M, w \models \psi_i$ then,
for all σ_j such that $\sigma_i \sim_j \sigma_j$, for all $l \in \{0, .., M\}$,
if $M, w \models B_j^{l-1} \neg\psi_j \wedge \neg B_j^l \neg\psi_j$ then $M, w \models B_j^{k+l-\kappa_j^*(\sigma_j)} [\sigma_j, \psi]\varphi$

iff $M, w \models \psi_i \rightarrow \wedge\{B_j^{l-1}\neg\psi_j \wedge \neg B_j^l \neg\psi_j \rightarrow B_j^{k+l-\kappa_j^*(\sigma_j)}[\sigma_i, \psi]\varphi; \sigma_j \sim_j \sigma_i \text{ and } l \in \{0, .., M\}\}$

Now we prove the soundness of the axiom 8.

Let M be a belief epistemic model.

$M, w \models [\sigma_i, \psi]K_j\varphi$ then $M(\Sigma), (w\sigma_i) \models K_j\varphi$

iff, if $M, w \models \psi_i$ then for all (v, σ_k) such that $(w, \sigma_i) \sim_j (v, \sigma_k)$,
$M(\Sigma), (v, \sigma_k) \models \varphi$

iff, if $M, w \models \psi_i$ then for all v such that $w \sim_j v$ and all σ_k such that $\sigma_i \sim_j \sigma_k$ and $M, v \models \psi_k$,
$M(\Sigma), (v, \sigma_k) \models \varphi$

iff, if $M, w \models \psi_i$ then for all σ_k such that $\sigma_i \sim_j \sigma_k$, for all v such that $w \sim_j v$,
if $M, v \models \psi_k$ then $M(\Sigma), (v, \sigma_k) \models \varphi$

iff, if $M, w \models \psi_i$ then for all σ_k such that $\sigma_i \sim_j \sigma_k$, for all v such that $w \sim_j v$,
if $M, v \models [\sigma_k, \psi]\varphi$

iff, if $M, w \models \psi_i$ then for all σ_k such that $\sigma_i \sim_j \sigma_k$,
if $M, v \models K_j[\sigma_k, \psi]\varphi$

iff, if $M, w \models \psi_i \rightarrow \wedge\{K_j[\sigma_k, \psi]\varphi; \text{ for all } \sigma_k \sim_j \sigma_i\}$

That is what we wanted. $\square$

We notice in the proof of soundness of axiom 9 without the introduction of this maximum of plausibility value M, we would get an infinite conjunction. Then we would be obliged somehow to introduce an infinite conjunction in our language. Doing so, our logic would not been compact. Compactness is a nice property that we do not want to lose. That is one of the motivations for introducing a maximum of plausibility value M (see section 2.1.1).

Conclusion: Now we have got a full logical system. We can then compare it with other systems. This is what we do in the following chapter.

Chapter 4

Comparisons with other Similar Systems

4.1 Comparison with the BMS System

Introduction: In this section, we will compare the BMS system ([BMS03]) with our system S. In order that the comparison be the fairest possible, we will modify the current BMS system to incorporate explicitly a belief component. This will be the first part of this section and the resulting system will be very much in line with the current BMS approach; then we will compare these two systems. The major point of this comparison will be that BMS does not do real belief revision. Finally, we will present a modification of our system that embeds the BMS one. This will thus enable us to isolate in our system the formal core of belief revision.

4.1.1 The BMS System with both Belief and Knowledge

We roughly define a BMS model, a BMS action model and a BMS product update. The rest of the technical apparatus is almost the same (update relation, action signature,...).

Definition 4.1.1 A BMS model M is a tuple $M = (W, \sim_j, \rightarrow_j, V, w_0)$ where:

- $\sim_j$ is an equivalence relation on W for each agent j.
- $\rightarrow_j$ is a binary relation on W for each agent j such that $\rightarrow_j \subseteq \sim_j$ and,
  (C_m) : if $w \rightarrow v$ and $w \sim_j w'$ then $w' \rightarrow v$.

$\lhd$

Note that, in the definition, the condition (C_m) is intuitively correct and is needed later for technical reasons (see appendix C1).

Definition 4.1.2 A BMS action model Σ is a tuple $\Sigma = (\Sigma, \sim_j^a, \rightarrow_j^a, Pre, \sigma_0)$ where

- $\sim_j^a$ is an equivalence relation on Σ.

- $\rightarrow_j^a$ is a binary relation such that $\rightarrow_j^a \subseteq \sim_j^a$ and,

  (C_a): if $\sigma \rightarrow_j^a \tau$ then for all σ' st $\sigma' \sim_j^a \sigma$, $\sigma' \rightarrow_j^a \tau$.

$\lhd$

Note that in the definition, the condition (C_a) is intuitively correct and is again needed later for technical reasons (see appendix C1).

Definition 4.1.3 Given a BMS model $M = (W, \sim_j, \rightarrow_j, V, w_0)$ and a BMS action model $\Sigma = (\Sigma, \sim_j^a, \rightarrow_j^a, Pre, \sigma_0)$, we define their update product to be the BMS model $M \otimes \Sigma = (W \otimes \Sigma, \sim_j', \rightarrow_j', V', w_0')$ where:

- $(w, \sigma) \rightarrow_j' (v, \tau)$ iff $w \rightarrow_j v$ and $\sigma \rightarrow_j^a \tau$.
- $(w, \sigma) \sim_j' (v, \tau)$ iff $w \sim_j v$ and $\sigma \sim_j^a \tau$.

The definition for the other components being as in the classical version. $\lhd$

These definitions are clearly in line with the former BMS approach.

From now on, we fix a BMS action signature Σ_{BMS}. We also need to define the language $\mathcal{L}(\Sigma_{BMS})$:

$\varphi := p \mid \neg\varphi \mid \varphi \wedge \psi \mid K_j\varphi \mid B_j\varphi \mid [\sigma, \psi_1, .., \psi_n]\varphi$

We now define its semantics:

$M, w \models K_j\varphi$ iff for all v such that $w \sim_j v$, $M, v \models \varphi$.
$M, w \models B_j\varphi$ iff for all v such that $w \rightarrow_j v$, $M, v \models \varphi$.

4.1.2 A Different Update of Beliefs

Intuitively, the BMS system expands the agents' beliefs but keeps the former beliefs true. Indeed, in the update product, an $\rightarrow_j'$ arrow will be possible only if there was already an $\rightarrow_j$ arrow; so the set of $\rightarrow_j'$ accessible worlds after an update *seems* to shrink, enhancing at the same time the set of beliefs. The belief update seems then close to an expansion operation. On the contrary, still intuitively, in our system, it is quite possible to modify the structural shape of the beliefs, and especially get rid of former beliefs. Indeed, thanks to $-\kappa_j^w(\varphi)$, we can very well get after an update a world with plausibility 0 which was not of plausibility 0 beforehand (unlike BMS as just said). This can dramatically change the shape of beliefs and cancel some former beliefs. The update seems then close to a revision.

We are going to make this more precise in this section. We first start by giving some technical results whose proofs can be found in the appendix C2. Then we will give their interpretation and use them for our analysis.

Technical Results

Corollary 4.1.4 *For all $\sigma_1, ..., \sigma_k$ that are $\sim_j$ accessible from σ and all propositional formula φ, $\models_{BMS} K_j\varphi \rightarrow [\sigma, \varphi_1, .., \varphi_n]K_j\varphi$.*

Corollary 4.1.5 *For all $\sigma_1, ..., \sigma_k$ that are $\sim_j$ accessible from σ and all propositional formula φ, $\models_S K_j\varphi \rightarrow [\sigma, \varphi_1, ..., \varphi_n]K_j\varphi$*

Corollary 4.1.6 *Let $\sigma_1, ..., \sigma_k$ be the actions $\rightarrow_j$ accessible from σ and $\psi = Pre(\sigma)$ and φ a propositional formula. Then,*
$B_j\varphi \rightarrow [\sigma, \varphi_1, ..., \varphi_n]B_j\varphi$

Proposition 4.1.7 *for all consistent sets of propositional formulae $\{\varphi_1, ..., \varphi_n, \varphi\}$ and all simple action σ,*
$\not\models_S B_j\varphi \rightarrow [\sigma, \varphi_1, ..., \varphi_n]B_j\varphi$.

Corollary 4.1.8 *Let $\sigma_1, ..., \sigma_k$ be the actions that are $\rightarrow_j$ accessible from σ,*
$\models_{BMS} B_j\neg\varphi \rightarrow [\sigma, \varphi, .., \varphi, \varphi_{k+1}, .., \varphi_n]B_j \perp$.

Corollary 4.1.9 *Let $\sigma_1, .., \sigma_k$ be the actions $\sim_j$ accessible from σ and of plausibility 0.*
$\models_S \neg K_j\neg\varphi \wedge B_j\neg\varphi \rightarrow [\sigma, \varphi, .., \varphi, \varphi_{k+1}, .., \varphi_n]\neg_j \perp$.

Corollary 4.1.10 *Let $\sigma_1, .., \sigma_k$ be the actions that are $\sim_j$ accessible from σ and of plausibility 0,*
$\models_S K_j\neg\varphi \rightarrow [\sigma, \varphi, .., \varphi, \varphi_{k+1}, .., \varphi_n]B_j \perp$.

Proposition 4.1.11 *For all belief epistemic formula φ (then without action term) and all belief epistemic model M,*
$F(M), w \models_{BMS} \varphi$ iff $M, w \models_S \varphi$

Proposition 4.1.12 *For all $\varphi_1, .., \varphi_n, \varphi$ epistemic formulae,*
If $M, w \models_S [\sigma, \varphi_1, .., \varphi_n]B_j\varphi$ then $F(M), w \models_{BMS} [\sigma, \varphi_1, .., \varphi_n]B_j\varphi$

Proposition 4.1.13 *For all model M of S and all epistemic formula φ,*
$M, w \models_S [\sigma, \varphi_1, .., \varphi_n]K_j\varphi$ iff $F(M), w \models [\sigma, \varphi_1, .., \varphi_n]K_j\varphi$

Interpretation of the Technical Results

Corollary 4.1.4 and 4.1.5 tell us that in BMS and S the agents keep their propositional knowledge after the update. This means that the agents can only increase their propositional knowledge as the situation is evolving. Corollary 4.1.6 tells us that what we just said about propositional knowledge can also be said for propositional belief in BMS; but this is not the case in S as proposition 4.1.7 tells us.

So BMS and S behave identically concerning the update of propositional knowledge, but differ on the update of propositional belief. BMS "expands" its propositional belief and knowledge during an update, whereas S only "expands" its propositional knowledge. Now, we can interpret $B_j\neg\varphi$ as "j believes that only an action with precondition $\neg\varphi$ can occur". In that case if j believes some action(s) have occurred but whose precondition(s)is (are) φ, then this incoming new information contradicts his beliefs. So, either j revises his beliefs or j does not. The first case can be expressed formally by $\neg B_j \perp$ (i.e. his new beliefs are consistent), and the second case by $B_j \perp$(i.e. his new beliefs are inconsistent). Indeed, after a revision of beliefs, the resulting set of beliefs is always consistent in our approach by assumption, and this is the hallmark of belief revision as conceived by AGM. Corollary 4.1.8 then tells us that in BMS, in case j encounters new information which contradicts his beliefs, he does not revise his beliefs. On the contrary, corollary 4.1.9 tells us that in S, if new information contradicts j's beliefs but still j considers the information possible, he will revise his beliefs. We see here a difference between the nature of the belief update in S and BMS. In S, the belief update seems to be similar to a belief revision (see section 4.3); whereas in BMS, we are sure that it is not the case by the above.

On the other hand, both behave similarly when it comes to updating knowledge (corollary 4.1.10 and proposition 4.1.13) but this is normal: the epistemic structure is the same after any of these two updates.

Finally, proposition 4.1.12 and proposition 4.1.13 tell us that j will believe more epistemic formulas after a BMS update than after a S update, that is the belief set after an update is always bigger in BMS than in S. This displays the fact that new unexpected worlds of plausibility 0 can be created after an update due to the subtraction of $\kappa_j^w(\varphi)$. This phenomenon does not occur in BMS, which is why more epistemic formulas are believed.

To sum up, both BMS and S update knowledge similarly but differ with regard to the updating of beliefs. S seems to follow a belief revision (see section 4.3) whereas BMS clearly does not. BMS update seems rather to be similar to an "expansion" as shown in the propositional case.

4.1.3 The Formal Reason of this Difference Concerning the Update of Beliefs

We modify our system by only changing the plausibility update.

We set: $\kappa_j(w,\sigma) = Cut_M(\kappa_j(w) + \kappa_j^*(\sigma))$

instead of: $\kappa_j(w,\sigma) = Cut_M(\kappa_j(w) + \kappa_j^*(\sigma) - \kappa_j^w(\varphi))$

This yields a new system S'.

We then have the

Theorem 4.1.14 *Let M be a belief epistemic action model, $w \in M$ and $\varphi \in \mathcal{L}(\Sigma_{BMS})$. then there exists some translation function T from $\mathcal{L}(\Sigma_{BMS})$ to S'-formulas, and a surjective application F from the set of belief epistemic model to BMS model such that:*

$M, w \models_{S'} T(\varphi)$ *iff* $F(M), w \models_{BMS} \varphi$

Proof: see appendix C1.

Note: S'-formulas are formulas in $\mathcal{L}(\Sigma_{S'})$ for some suitable belief epistemic action signature $\Sigma_{S'}$ (see appendix C1).

This theorem tells us that any BMS-formula satisfiable in a BMS model is also satisfiable in a S'-model, modulo a translation. That is what we intuitively wanted.

The difference between S' and our system only consists in the suppression of the term $-\kappa_j^w(\varphi)$ in the plausibility update. Now S' embeds BMS as just shown. This implies that the differences observed between our system and BMS are due to this term. Thereby, the term $-\kappa_j^w(\varphi)$ is the hallmark of belief *revision* as intuitively sensed in the introduction of this section.

Conclusion: Though our approach is based on the BMS one, we clearly saw an important difference concerning the update of beliefs, namely that our update is closer to a genuine belief revision, thanks to the subtraction of $-\kappa_j^w(\varphi)$ as shown in section 4.1.3. So our system and the BMS one are just two different ways to look at how we update beliefs.

Anyway, as we said in the remarks of section 2.3.2, our choice of plausibility update could have been different: it just depends on our background motivations and intentions.

4.2 Comparison with H.v.Ditmarsch and W.Labuschagne's System

Introduction: In this short section we will present a comparison between our system and H.v.Ditmarsch and W.Labuschagne system which can be found in [vDL03]. We do not sum up their work but directly give the comparison. We will indicate with a label (*) the notions introduced in [vDL03].

Concerning the static part, belief epistemic models are completely equivalent to doxastic epistemic models(*). Indeed, we can prove that every formula satisfiable in a belief epistemic model is also satisfiable in a doxastic epistemic model modulo a translation, and vice versa. One of the differences is that in [vDL03], the language allows common knowledge, unlike our language (but we can easily incorporate it). Another difference is that we use a maximal plausibility value M, at the difference of [vDL03].

Concerning the dynamic part, there are many more difference. Our system allows a single type of actions fixed by the action signature, whereas in the other approach we deal at the same time with several actions: revision, test, learning. It has, of course, advantages and shortcomings in both cases. Now, let us have a look at the different types of actions present in [vDL03].

1. The test operation $(?\varphi)$(*) has no counterpart in our system.

2. The revision operator ($*_G\varphi$) has a corresponding action signature in our system, namely the second example in section 2.2.2. However the plausibility update is slightly different.

3. The union operator ($L_G\alpha$) intuitively means that the group G learns the occurrence of an action α. So, all the agents learning a formula φ ($L_N?\varphi$) has a natural counterpart in our system, namely a public announcement of φ, and this can be proved formally. However for the case of a group of agents G different from the whole group of agent N, things are more complicated and we can not reduce one type of action to another.

Conclusion: So, we can not reduce one system to the other. Indeed our system allows more flexibility concerning the action we can consider than the one exposed in [vDL03]. But we can only deal with one action at the same time in our logical system, unlike H.v.Ditmarsch and W.Labuschagne's system.

4.3 Comparison with the AGM Postulates

In this section we want to check whether the AGM postulates are fulfilled after an update. So, as a preliminary step, we need to define belief sets, and decide what type of action signature we consider for the update. For the first point we have to choose the language we consider for the formation of belief sets. This will divide our analysis in three parts: 1)propositional language, 2)epistemic language, 3)epistemic-dynamic language. For the second point, the most natural and "purest" action we would be prompted to consider would be a public announcement. We will however see that this creates problems for the verification of the AGM postulates in the case of the epistemic language. These problems can be solved if we consider a more refined type of action signature than a simple public announcement.

In the following we consider a belief epistemic model $M = (W, \sim_j, \kappa_j, V, w_0)$, $w \in W$ and an agent j.

4.3.1 Propositional Language for Formation of Belief Bets

Let $\mathcal{L}_1$ denotes the propositional language.

The action signature Σ_{sg1} that we consider for our logic $\mathcal{L}(\Sigma_{sg1})$ is the one of public announcement with single simple action $A!$ (that we have so far usually denoted σ)

Let $A \in \mathcal{L}_1$

Definition 4.3.1 We define:

$K^w = \{\varphi \in \mathcal{L}_1; M, w \models B^0_j\varphi\}$
$K^w * A = \{\varphi \in \mathcal{L}_1; M, w \models [A!, A]B^0_j\varphi\}$
$K^w + A = \{\varphi \in \mathcal{L}_1; M, w \models B^0_j[A!, A](A \to \varphi)\}$ $\triangleleft$

K^w is the (propositional) belief set of the agent j dwelling in world w.

$K^w * A$ is the revised belief set for the agent j dwelling in world w after the announcement of A.

$K^w + A$ is the expansion of the belief set K^w by A after the announcement of A.

We then have the crucial:

Theorem 4.3.2 *In the $\mathcal{L}(\Sigma_{sg1})$ framework, the revision function * defined by $K^w * A$ satisfies the 8 AGM postulates.*

Proof: see appendix E2.

So the propositional case works out. Let us now turn our attention to the more problematic epistemic case.

4.3.2 Epistemic Language for the Formation of Belief Sets

First, we consider the paradigmatic case of public announcement for performing belief revision, as we did in the propositional case. The definition of belief sets are exactly the same as above, except that we replace $\mathcal{L}_1$ by the epistemic language $\mathcal{L}_2$ (epistemic language *without* belief formulas: see concluding remarks). So for example, $K^w = \{\varphi \in \mathcal{L}_2; M, w \models B_j\varphi\}$.

Unfortunately,

Theorem 4.3.3 *In the $\mathcal{L}(\Sigma_{sg1})$ framework, the revision function * defined by $K^w * A$ satisfies only the postulates $(K * 1), (K * 5)$ and $(K * 6)$.*

Proof: see appendix E1.

As is spelled out in the appendix, the postulates are not all satisfied because the formulas satisfiable at any world may change after an update. Indeed, by updating a public announcement, we get rid of worlds; so some $\sim_j$ links disappear, changing the epistemic structure of the model. This is a classical phenomenon called "relativization" and "self persistence" (see [vB03] and appendix).

However, we can cope with this problem by slightly modifying the action structure in order to keep the whole epistemic structure of the model after the update. We then consider for our logical system the following action signature Σ_{sg2}:

$$\boxed{\boxed{A!}} \overset{j}{\text{———}} \boxed{\tau}$$

with $\kappa_j^*(A!) = 0$ and $\kappa_j^*(\tau) = \alpha_j$ for all agents j.

Note that this corresponds to the second example of section 2.2.2.

We keep the same definition as before for the belief sets (with $A \in \mathcal{L}_2$):

Definition 4.3.4 We define:

$K^w = \{\varphi \in \mathcal{L}_2; M, w \models B_j^0\varphi\}$

$K^w * A = \{\varphi \in \mathcal{L}_2; M, w \models [A!, A, \neg A] B_j^0 \varphi\}$
$K^w + A = \{\varphi \in \mathcal{L}_2; M, w \models B_j^0 [A!, A, \neg A](A \rightarrow \varphi)\}$ $\triangleleft$

We then get the very appealing

Theorem 4.3.5 *In the $\mathcal{L}(\Sigma_{sg2})$ framework, the revision function * defined by $K^w * A$ satisfies the 8 AGM postulates. Moreover, after an update, A is believed by all the agents j with firmness α_j at world w (see section 1.2.2)*

Proof: see appendix E3.

Now we can turn our attention to the last case.

4.3.3 Epistemic Dynamic Language for the Formation of Belief sets

We shall not go through an extensive study of this case because we can know beforehand that the AGM postulates will not be satisfied, thanks to the so-called "Gardenfors paradox".

This paradox, discovered by AGM itself, says that adding counterfactuals to belief states, encoding revisions that one would make if given the antecedent, leads to inconsistency.

So in our case, as transposition of counterfactuals, revision with formulas of the form $\psi \rightarrow [A!, A]\varphi$ (which are epistemic dynamic formulas) will invalidate the AGM postulates.

Final remarks: Note that for the formation of belief sets we ruled out belief formulas. It seems counter-intuitive because the agents could in reality have beliefs about what the other agents believe. For example, it is possible that j believes that i believes φ. However, following this example, we will face a problem if we get $\neg\varphi$ as incoming public information. Indeed i will have to revise his beliefs, but how does j know how he will do that? j's beliefs concerning i's beliefs will be hardly determinable, at least with logical considerations, except in that j will believe that i believes $\neg\varphi$. This difficulty prompted us to avoid dealing with other agents' beliefs, but a natural improvement of our work would of course consist in incorporating such considerations.

Conclusion: So we have seen that in almost all the cases (except the last one), we can always find a particular action signature which performs belief revision as conceived by AGM. In a sense, this shows that our update of beliefs is close to a belief revision. Indeed, the action signatures considered in this section are basic ones. So, they are the most liable to check whether or not our system performs belief revision because they are close to a "pure" incoming information.

Chapter 5

General Conclusion

So we have achieved our goal sketched in the introduction. Indeed, we have built a logical system involving both notions of beliefs and knowledge, whose update of agent's beliefs is close to a (belief) revision (at least in basic cases, see section 4.3). However, it presents some limitations.

First, as mentioned in the remarks of section 2.3.2, the plausibility of an action does not depend on the world in which they are performed. This is of course counter-intuitive, as explained in section 2.3.2. However, trying to remedy this by making small variations to our current system seems difficult. Indeed, the plausibilities of the possible actions in the action signature are fixed once and for all (see section 2.3.5), whatever their preconditions are. However, these preconditions can vary because they are assigned in the language (unlike the plausibilities of the actions as remarked above).

Secondly, another limitation is that, as mentioned in section 4.2, we can only deal in our logical system with a single type of action (unlike in the approach of [vDL03]) which is fixed in the logic by the action signature Σ_{sg}.

Finally, as already mentioned in section 4.3, the access for an agent to the other agents' beliefs is unclear, especially after an update. Indeed, in reality, what does an agent know about how the other agents update (and especially revise) their beliefs? These considerations are not present in the AGM theory of belief revision because they consider only a single agent, unlike our multi-agent perspective. So, a formal and logical treatment of this type of issue remains to be done.

In this appendix we will present some material which has been removed from the rest of the thesis for better readability, or whose relevance was not important enough. We will first introduce the notion of bisimulation.Second, we will make a comparison of our system with the probabilistic one as exposed in [vBJoLLI03]. Third, we will give the proofs missing in the comparison with the BMS system (see section 4.1). finally, we will do the same regarding the AGM theory (see section 4.3).

Appendix A

Bisimulation

Bisimulation is a useful notion in modal logic. It generally holds that if two structures are bisimilar, then they are behaviorally indistinguishable. In the case of belief epistemic model, behaviorally indistinguishable means satisfying the same formulas. A well known result in modal logic is that if two pointed models are bisimilar, then they satisfy the same formulas (for details, see [BdRV01]). in this section we show that such a result holds for our logic as well.

Definition A.0.6 Let $M = (W, \sim_j, \kappa_j, V, w_0)$ and $M' = (W', \sim'_j, \kappa_j, V', w'_0)$ be two belief epistemic models. A non empty binary relation $Z \subseteq W \times W'$ is called a bisimulation between M and M' (notation $Z : M \rightleftarrows M'$) if the following conditions are satisfied:

1. If wZw' then w and w' satisfy the same proposition letters.
2. if[wZw' and ($w \sim_j v$ and $\kappa_j(v) \leqslant m$)] then [there is v' such that vZv', $w' \sim_j v'$ and $\kappa_j(v') \leqslant m$] (the forth condition).
3. If[wZw' and ($w' \sim_j v'$ and $\kappa_j(v') \leqslant m$)] then [there is v such that vZv', $w \sim_j v$ and $\kappa_j(v) \leqslant m$] (the back condition)

When Z is a bisimulation linking two states w in M and w' in M' we say that w and w' are bisimilar and we write $M, w \leftrightarrows M', w'$.

$\lhd$

Theorem A.0.7 *Let M, M' be belief epistemic models. If $M, w \leftrightarrows M', w'$ then for all formula φ, $M, w \models \varphi$ iff $M', w' \models \varphi$.*

Proof: by induction on φ.

- The case where φ is a propositional letter follows from clause 1, and the one where φ is $\perp$ is immediate.
- The boolean cases are immediate from the induction hypothesis.

- As for formulas of the form $B_j^m\varphi$, assume $M, w \models B_j^m\varphi$.
  Then for all v st $w \sim_j v$ and $\kappa_j(v) \leqslant m$, $m, v \models \varphi$.

  Let v' st $w' \sim_j v'$ and $\kappa_j(v') \leqslant m$. Then by back condition, there is v st vZv' and $w \sim_j v$ and $\kappa_j(v) \leqslant m$.

  So by assumption, $M, v \models \varphi$.

  But vZv', so by induction hypothesis $M', v' \models \varphi$. And this for all v' st $w' \sim_j v'$ and $\kappa_j(v') \leqslant m$.

  Thus $M', w' \models B_j^m\varphi$.

  The reciprocity is very similar.

- For formulas of the form $[\pi]\varphi$, we know that they are provably equivalent to formulas without action terms and this case has just been studied. This ends this proof.

$\square$

Appendix B

Comparison with the Probabilistic Approach

In the following, S refers to our system and P refers to the probabilistic system.

First we will present an overview of the probabilistic approach. Then we will compare systematically the two approaches as regard to their dynamic part, static part and update product.

B.1 Overview of the Probabilistic Approach (P)

The overall structure of P is completely similar to our system S: P is composed of a static part (the probabilistic epistemic model) and a dynamic part (action model). Moreover the epistemic structure and update is exactly the same. So we focus on the probabilistic part of the system.

First, for the static part, the epistemic action model falls into equivalence classes

$D_{j,w} = \{v; w \sim_j v\}$

Over these, probability functions are given, satisfying the usual conditions:

$P_{j,w}$ is defined on the probabilistic space $D_{j,w}$.

And we assume that $P_{j,w}$ assigns probabilities $P_{j,w}(v)$ to single worlds v. So the probability of a formula φ at state w is then defined by $P_{j,w}(\varphi) = \sum\{P_{j,w}(v); v \in D_{j,w} \text{ and } M, v \models \varphi\}$.

Secondly, for the dynamic part, we assign probabilities to actions depending on which world this action is performed:

Probabilities $P_{j,w}(a)$ for individual actions at state w in the current world.

Finally, the probabilistic update is defined at the world (v, b) in $D_{j,(w,a)} = \{(v, b); (v, b) \sim_j (w, a)\}$ by:

$$P_{j,(w,a)}(v, b) = \frac{P_{j,w}(v).P_{j,v}(b)}{\sum\{P_{j,w}(u).P_{j,u}(c);(u,c)\in D_{j,(w,a)}\}}$$

B.2 Dynamic Part

Both action structures in P and S have the same shape. They are both composed of states representing simple actions and are both composed of epistemic links between simple actions which are equivalence relations. So the epistemic structure is similar. What differs is the belief part and the probabilistic part. Nevertheless, both assign to actions among the same $\sim_j$-equivalence class, a degree of plausibility in S and a probability in P.

But in P, the same simple action may have different probabilities at different states of a process. That is to say, the probability of performance of an action in a given world for a given agent depends on this current world. It is unlike S where the plausibility of an action b, $\kappa_j^*(b)$ doesn't depend on the current world the agent dwells in. The actions considered in S thus form a subset of the ones of P because they must satisfy the extra condition that their possible occurrence is independent from the current world.

This last feature makes impossible to translate a belief action model to a probabilistic action model , and hence more generally we can not embed the probabilistic approach in the belief approach and vice versa.

Finally, the assignment is made differently in S and P: in S it is achieved locally and in P it is achieved globally. Indeed, in S, among an $\sim_j$-equivalence class, the agent j considers each action individually (it is hence a local process) and assigns on the basis of the nature of this single action and its scale, a plausibility value. So the belief value is assigned locally (individually) to each simple action. On the contrary, in P, the assignment is made globally. Indeed, in a $\sim_j$-equivalence class in a certain world w, the agent j assigns at the same time a probability to each action. And this assignment for a simple action depends on the other simple actions of the same $\sim_j$-equivalence class. That is why we say that it is a global assignment, because the assignment is made relatively to the other actions. This is exemplified by the extra condition that $\sum_{b\in D_{j,a}} P_{j,v}(b) = 1$ (where $D_{j,a}$ is the set of the actions of the equivalence class containing a), because we see here that these probabilities are somehow linked by an extra assumption.

We must however moderate our analysis because in S, there is though as well a global condition, namely that one of the simple action in an equivalence class should be assigned the plausibility 0. However this doesn't really spoil the two different spirits of plausibility and probability assignment in S and P. We will indeed further see that this duality is present throughout the analysis.

B.3 Static Part

B.3.1 General Remarks

Like for the actions models, both type of models in S and P have the same shape. They are composed of states representing possible worlds and are both composed of epistemic links between worlds which are equivalence relations. So the epistemic structure is similar. What differs is the probability part and the belief part. Nevertheless both assign to possible worlds among the same $\sim_j$-equivalence class a degree of plausibility in S and a probability in P.

B.3.2 The Assignment of Probability and Plausibility: Locality versus Globality

Once again the assignment is made differently in S and P: in S it is made locally and in P it is made globally. Indeed, in S the plausibility ordering is assigned *individually* to each accessible world without taking into account the other worlds, whereas in P the individual assignment formally depends on the other worlds. In S, to assign a plausibility to a world t, we just have to consider on the one hand this world, and on the other hand our ordinal scale; we abstract the other worlds linked to s (this is very similar to the update of the plausibility ordering, see below). In P, the assignment is done simultaneously (not individually like in S) and instantly to all the worlds accessible from the given world; this process then compares each accessible world to the others unlike S. The assignment in P is a global process in contrast to the local assignment in S. In P this phenomenon is exemplified by the fact that the probability assignments at a given world s must fulfill $\sum_{w \in D_{j,v}} P_{i,w}(v) = 1$: we see that they are somehow globally linked via this relation.

however a similarity is that in both S and P, the plausibility of a world doesn't depend on the current world the agent j dwells in and is the same over a $\sim_j$ equivalence class.

B.3.3 An Attempt to Strike Relationships between both Type of Models: Aristocracy versus Democracy

It is easy from a belief epistemic model to get a probabilistic epistemic model (but not the converse). Indeed, a natural translation would be:

Definition B.3.1 Let $T : S \longrightarrow P$ be the operator defined from the set of belief epistemic model $\mathcal{B}$ to the set of probabilistic epistemic model by:

For all belief epistemic model $M = (W, \sim_j, \kappa_j, V)$, $T(M) = (W, \sim_j, P_{j,w}, V)$ where

$$P_{j,w}(v) = \frac{(M - \kappa_j(v)) + 1}{\sum\{(M - \kappa_j(v)) + 1; v \sim_j w\}} (*)$$

$\lhd$

The adjunction of 1 in (*) prevents against the case where $M - \kappa_j(v) = 0$ for all $t \sim_j s$(in that case the denominator would be equal to 0). Clearly, $\sum_{v \in D_{j,w}} P_{j,w}(v) = 1$, and $P_{i,w}$ is a probability function defined on $D_{j,w} = \{v; w \sim_j v\}$. So T is well-defined . We then have the following lemma which makes our translation sound:

Lemma B.3.2 $P_{i,w}(v) < P_{j,w}(v') \Leftrightarrow \kappa_j(v) > \kappa_j(v')$

Indeed, if an agent considers a world more *plausible* than another, he will also consider it more *probable*, and reciprocally.

We then pose:

Definition B.3.3 Let M_P be a probabilistic epistemic model and M a belief epistemic model,

$M_P \models P_j(\varphi) > P_j(\psi)$ iff $\sum\{P_{j,w}(v); w \sim_j v$ and $M_P, w \models \varphi\} > \sum\{P_{j,w}(v); w \sim_j v$ and $M_P, v \models \psi\}$

$M, w \models \kappa_j(\varphi) > \kappa_j(\psi)$ iff $\kappa_j^w(\varphi) = min\{\kappa_j(v); w \sim_j v$ and $M, v \models \varphi\} > \kappa_j^w(\psi)$.

$\lhd$

$M_P, w \models P_j(\varphi) > P_j(\varphi)$ intuitively means that in world w, j considers formula φ more probable than ψ. $M, w \models K_j(\varphi) > K_j(\psi)$ intuitively means that in world w, j considers formula *psi* more plausible than φ. The last case corresponds exactly to W.Spohn's interpretation (see section 1.2.2).

So we would expect that the ranking of formulas in P after the translation T corresponds to the ranking of formulas in S .Unfortunately,

Proposition B.3.4 *There is a belief epistemic model M and a world $w \in M$ and formulas φ and ψ such that*

$T(M), w \models P_j(\varphi) > P_j\psi) \not\Leftrightarrow M, w \models \kappa_j(\varphi) < \kappa_j(\psi)$

Proof:Consider the belief epistemic model $M = (\{w_0, w_1, w_2\} = W, \sim_j, \kappa_j, V)$ where

- the maximum of plausibility is $Max = 3$.
- $\sim_j = W \times W$,
- $\kappa_j(w_0) = 0$ and $\kappa_j(w_1) = \kappa_j(w_2) = 2$
- $V(p) = \{w_0\}$

Then $M, w_0 \models \kappa_j(p) < \kappa_j(\neg p)$
but in $T(M)$, $P_{j,w_0} = 1/2$, and $V(p) = \{w_0\}$
So, $T(M), w_0 \not\models P_j(p) > P_j(\neg p)$ $\Box$

This result is not due to the translation T (see lemma B3.2) but rather to the way we define the plausibility of a formula. Indeed, we state that the formula φ is the *minimum* of the degree of plausibility of states where φ holds and that are accessible via w. Doing so we introduce a compound of layers whose each of them is composed of states accessible from w and that have the same degree of beliefs. Moreover by stating that we take the minimum of them, we give a priority ordering over the different layers: a layer corresponding to a certain plausibility "cancels" the effects of a layer of lower degree. Our definition forgets the global plausibility structure of worlds accessible from w and focus locally on the worlds accessible from w of higher plausibility.

This feature goes somehow in line with what we just said about the difference between local and global aspects of assignment in P and S inasmuch we see that the local aspect of the plausibility assignment plays a more central role when it comes to determining the plausibility of a formula.

To sum up, a plausibility assignment implicitly gives priority of worlds upon others in our setting, whereas a plausibility assignment doesn't do so. A metaphor will surely express better what we want to say.

The latter (P) is more "democratic", the former is rather "aristocratic". Indeed, in S if we identify worlds with peoples, the plausibility grading gives some people a kind of power (the grading) which has priority upon other's power, just like some people have a priority power when it comes to decide something in society:that's why we call it aristocratic. On the contrary, P is democratic because if we again identify worlds with peoples, peoples hold a part of power which is assigned by the probability assignment but don't have priority upon other when it comes to decide something: the more power a people has, the higher the probability the current world is. Thus power is assigned democratically to some people in P, contrary to the aristocratic assignment in S.

B.4 The Update Product

The update product is completely identical for the epistemic structure of the models. The only difference concerns the belief and probabilistic part.

These are the beliefs and probabilistic updates for S and P:

S: $\kappa_j(w, a) = \kappa_j(w) - \kappa_j^w(Pre(a)) + \kappa_j^*(a)$

P:

$$P_{j,(w,a)}(v, b) = \frac{P_{j,w}(v).P_{j,v}(b)}{\sum\{P_{j,w}(u).P_{j,u}(c); (u, c) \in D_{j,(w,a)}\}}$$

$P_{j,w}(v).P_{j,v}(b)$ clearly corresponds to $\kappa_j^*(b) + \kappa_j(v)$. Note that this replacement of the sum, multiplication and division of probabilities by respectively the minimum, addition and subtraction of plausibility is not an accident and has been studied in depth by W.Spohn in [Spohn90].

So, $-\kappa_j^w(Pre(a))$ would correspond to the division by $\sum\{P_{j,w}(u).P_{j,u}(c); (u, c) \in D_{j,(w,a)}\}$. Though they are formally identical, they are conceptually different. The former bears on a single action b whereas the latter bears on all the actions $\sim_j$-accessible from a. In P we just need a rescaling so that $P_{j,(w,a)}$ is indeed a probability function. So we have to divide by $\sum\{P_{j,w}(u).P_{j,u}(c); (u, c) \in D_{j,(w,a)}\}$. But in S, we don't need such a global rescaling. The global rescaling is already implicitly present in the assignment of the $\kappa_j^*(a)$ for any action a (see section 2.3 ,update product). $-\kappa_j^w(Pre(a))$ is rather the hallmark of revision (see comparison with BMS, section 4.1), even if it is originally a local rescaling (see section 2.3.2).

This local rescaling goes again in line with the fact already stressed that the plausibility assignment is something local in S, and hence that its update must also be dealt locally. That's what actually happens in the update where we not only focus on the particular world at stake, but also on the particular action taking place in this world and abstract the other accessible ones (just like we abstract the other accessible worlds when we assign a plausibility to a particular world). Doing so we really stress the local feature. We don't really care about the global shape of the final outcome which is determined mainly by the epistemic action model (see the example). It is of course different from the probabilistic case where the global feature is highlighted by rescaling to *all* the actions accessible from the one taking place in a particular world.

Appendix C

Proofs missing in the "Comparison with the BMS System"

C.1 Proof of theorem 4.1.14

Definition C.1.1 We define by induction the translation from a $\mathcal{L}(\Sigma_{BMS})$-formula *without* action term to a S'-formula:

T1: $T(p) = p$.

T2: $T(\varphi \wedge \psi) = T(\varphi) \wedge T(\psi)$.

T3: $T(\neg\varphi) = \neg T(\varphi)$.

T4: $T(K_j\varphi) = K_j T(\varphi)$.

T5: $T(B_j\varphi) = B_j^0 T(\varphi)$.

$\lhd$

The translation from a $\mathcal{L}(\Sigma_{BMS})$-formula containing action terms to a S'-formula needs to define a function from the BMS action model Σ_{BMS} to the S' action model $\Sigma_{S'}$.

Definition C.1.2 Let $\Sigma_{BMS} = (\Sigma, \sigma_0, \sim_j, \to_j^a, Pre)$ be a BMS action model. We define the belief epistemic action model $\Sigma_{S'} = (\Sigma, \sigma_0, \sim_j, \kappa_j^*, Pre')$ by:

- $$\kappa_j^*(\sigma) = \begin{cases} 0 & \text{if there is } \tau \text{ such that } \tau \to_j^a \sigma \\ M & \text{otherwise} \end{cases}$$

- $Pre'(\sigma) = T(Pre(\sigma))$

$\lhd$

We can now set the inductive clause for action-formula:

Definition C.1.3 **T6:** $T([\sigma, \psi_1, ...\psi_n]\varphi) = [\sigma, T(\psi_1), ..., T(\psi_n)]T(\varphi)$.

$\triangleleft$

So we have got a full translation from a $\mathcal{L}(\Sigma_{BMS})$-formula to a $\mathcal{L}(\Sigma_{S'})$-formula.
Thanks to (C_a), we get the fact:

Fact:$\tau \rightarrow_j^a \sigma$ iff $\kappa_j^*(\sigma) = 0$ and $\tau \sim_j \sigma$
Finally, we define F

Definition C.1.4 Let $M = (W, w_0, \sim_j, \kappa_j, V)$ be a S'-model. We define the BMS-model $F(M) = (W, w_0, \sim_j, \rightarrow_j, V)$ by
$w \rightarrow_j v$ iff $\kappa_j(v) = 0$ and $w \sim_j v$. $\triangleleft$

Note that F is surjective thanks to (C_m).

We finally get the:

Theorem C.1.5 *Let $\varphi \in \mathcal{L}(\Sigma_{BMS})$ and M a S' model. Then,*
$M, w \models T(\varphi)$ iff $F(M), w \models \varphi$.

Proof: By induction on φ.

T1 to T4 are obvious. We now prove T5:

$$
\begin{array}{lcl}
M, w \models T(B_j\varphi) & \Leftrightarrow & M, w \models B_j^0 T(\varphi) \\
& \Leftrightarrow & \text{for all v such that } w \sim_j v \text{ and } \kappa_j(v) = 0, M, t \models T(\varphi). \\
& \Leftrightarrow & \text{for all v such that } w \sim_j v \text{ and } \kappa_j(v) = 0 F(M), v \models \varphi \text{ by induction hypothesis.} \\
& \Leftrightarrow & \text{for all v such that } w \rightarrow_j v, F(M), v \models \varphi \text{ by definition??} \\
& \Leftrightarrow & F(M), s \models B_j\varphi
\end{array}
$$

To prove T6, we need the:

Lemma C.1.6 $F(M \otimes \Sigma_{S'}) = F(M) \otimes \Sigma_{BMS}$

Proof: The only point to check is the belief relation. Indeed, the fact that both members of the equality have the same epistemic structure comes from T1-T5 and the definition of F. For all $(w, \sigma), (v, \tau) \in F(M) \otimes \Sigma_{BMS}$,

$$
\begin{array}{lcl}
(w, \sigma) \rightarrow_j (v, \tau) & \Leftrightarrow & w \rightarrow_j v \text{ and } \sigma \rightarrow_j^a \tau \\
& \Leftrightarrow & (\kappa_j(w) = 0 \text{ and } w \sim_j v) \text{ and } (\kappa_j^*(\tau) = 0 \text{ and } \sigma \sim_j \tau) \text{ by definition?? and Fact} \\
& \Leftrightarrow & (w, \sigma) \sim_j (t, \tau) \text{ and } \kappa_j(v, \tau) = 0 \text{ by definition of the plausibility update.}
\end{array}
$$

That is what was expected. $\square$

We can then easily prove T6. $\square$

C.2 Proof of the other Technical Results

the propositional case.

Proposition C.2.1 *For all $\sigma_1, ..., \sigma_k$ that are $\sim_j$ accessible from σ and all propositional formula φ,*

$\models_{BMS} [\sigma, \varphi_1, ..., \varphi_n]K_j\varphi \leftrightarrow \psi \rightarrow \bigwedge\{K_j(\varphi_i \rightarrow \varphi); i \in \{1..k\}\}$

Proof: First assume $M, w \models [\sigma, \varphi_1, .., \varphi_n]K_j\varphi$.

Assume $M, w \models \psi$
then $M(\Sigma), (w, \sigma) \models K_j\varphi$ ($M(\Sigma)$ is the updated model)
For all $i \in \{1, .., k\}$
let v st $w \sim_j v$
if $M, v \models \varphi_i$, then $(w, \sigma) \sim_j (v, \sigma_i)$ because $w \sim_j v$ and $\sigma \sim_j \sigma_i$
then $M(\Sigma), (v, \sigma_i) \models \varphi$ because $M(\Sigma), (w, \sigma) \models K_j\varphi$.
then $M, v \models \varphi$ because φ is a propositional formula.
so $M, v \models \varphi_i \rightarrow \varphi$ and so for all v st $v \sim_j w$.
then $M, w \models K_j(\varphi_i \rightarrow \varphi)$ and so for all $i \in \{1, .., k\}$
So $M, w \models \wedge\{K_j(\varphi_i \rightarrow \varphi); i \in \{1, .., k\}\}$
Then $M, w \models \psi \rightarrow \wedge\{K_j(\varphi \rightarrow \varphi); i \in \{1, .., k\}\}$

Assume $M, w \models \psi \rightarrow \wedge\{K_j(\varphi_i \rightarrow \varphi); i \in \{1, .., k\}\}$

Assume $M, w \models \psi$.
For all (v, τ) st $(w, \sigma) \sim_j (v, \tau)$,
$\tau \in \{\sigma_1, .., \sigma_k\}$, so $M, v \models \varphi_i$ for some $i \in \{1, .., k\}$
then $M, v \models \varphi$ because $w \sim_j v$ and $M, w \models K_j(\varphi_i \rightarrow \varphi)$,
then $M(\Sigma), (v, \sigma) \models \varphi$ because φ is a propositional formula
so $M(\Sigma), (w, \sigma) \models K_j\varphi$
then $M, w \models [\sigma, \varphi_1, .., \varphi_n]K_j\varphi$ $\square$

Proposition C.2.2 *For all $\sigma_1, ..., \sigma_k$ that are $\sim_j$ accessible from σ and all propositional formula φ,*

$\models [\sigma, \varphi_1, .., \varphi_n]K_j\varphi \leftrightarrow \{\psi \rightarrow \bigwedge\{K_j(\varphi_i \rightarrow \varphi); i \in \{1..k\}\}\}$

Proof: Assume $M, w \models [\sigma, \varphi_1, .., \varphi_n]K_j\varphi$.

If $M, w \models \psi$
then $M(\Sigma), (w, \sigma) \models K_j\varphi$.

For all $i \in \{1,..,k\}$

let v st $v \sim_j w$

if $M, v \models \varphi_i$, then $(w, \sigma) \sim_j (v\sigma_i)$

then $M(\Sigma), (v, \sigma_i) \models \varphi$

then $M, v \models \varphi$ because φ is a propositional formula.

so $M, v \models \varphi_i \rightarrow \varphi$

so $M, w \models K_j(\varphi \rightarrow \varphi)$ and so for all i

so $M, w \models \wedge\{K_j(\varphi \rightarrow \varphi); i \in \{1,..,k\}\}$

so $M, w \models \psi \rightarrow \wedge\{K_j(\varphi_i \rightarrow \varphi); i \in \{1,..,k\}\}$

Assume $M, w \models \psi \rightarrow \wedge\{K_j(\varphi_i \rightarrow \varphi); i \in \{1,..,k\}\}$

assume $M, w \models \psi$

for all (v, τ) st $(w, \sigma) \sim_j (v, \tau)$,

$\tau \in \{\sigma_1, .., \sigma_k\}$, so $M, v \models \varphi_i$ for some $i \in \{1,..,k\}$

then $M, v \models \varphi$ because $w \sim_j v$ and $M, w \models K_j(\varphi_i \rightarrow \varphi)$

$M(\Sigma), (v, \tau) \models \varphi$ because φ is a propositional formula

so $M(\Sigma), (w, \sigma) \models K_j\varphi$

then $M, w \models [\sigma, \varphi_1, .., \varphi_n]K_j\varphi$ $\square$

Corollary C.2.3 *For all $\sigma_1, ..., \sigma_k$ that are $\sim_j$ accessible from σ and all propositional formula φ, $\models_S K_j\varphi \rightarrow [\sigma, \varphi_1, ..., \varphi_n]K_j\varphi$*

Proposition C.2.4 *Let $\sigma_1, ..., \sigma_k$ be the actions $\rightarrow_j$ accessible from σ and $\psi = Pre(\sigma)$ and φ a propositional formula. Then,*

$\models_{BMS} \{\psi \rightarrow \bigwedge\{B_j(\varphi_i \rightarrow \varphi); i \in \{1..k\}\}\} \leftrightarrow [\sigma, \varphi_1, .., \varphi_k, .., \varphi_n]B_j\varphi.$

i.e.$\models_{BMS} \{\psi \rightarrow B_j(\varphi_1 \vee .. \vee \varphi_k \rightarrow \varphi)\} \leftrightarrow [\sigma, \varphi_1, .., \varphi_k, .., \varphi_n]B_j\varphi$

Proof: Assume $M, w \models \varphi \rightarrow B_j(\varphi_1 \vee .. \vee \varphi_k \rightarrow \varphi)$

Assume $M, w \models \psi$

Then $M, w \models B_j(\varphi_1 \vee .. \vee \varphi_k \rightarrow \varphi)$

for all (v, τ) st $(w, \sigma) \rightarrow_j (v, \tau)$,

$w \rightarrow_j v$ and $M, v \models \varphi_1 \vee .. \vee \varphi_k$ because $M, v \models Pre(\sigma)$

then $M, v \models \varphi$ because $M, w \models B_j(\varphi_1 \vee .. \vee \rightarrow \varphi)$

then $M(\Sigma), (v, \tau) \models \varphi$ because φ is a propositional formula

and then $M, w \models [\sigma, \varphi_1, .., \varphi_n]B_j\varphi$

Assume $M, w \models [\sigma, \varphi_1, .., \varphi_n]B_j\varphi$

assume $M, w \models \psi$

then $M(\Sigma), (w, \sigma) \models B_j\varphi$

then for all (v, τ) st $(w, \sigma) \rightarrow_j (v, \tau)$, $M(\Sigma), (v, \tau) \models \varphi$

then for all (v, τ) st $(w, \sigma) \rightarrow_j (v, \tau)$, $M, v \models \varphi$ because φ is a propositional formula

then for all v st $w \rightarrow_j v$ and $M, v \models \varphi_1 \vee .. \vee \varphi_k$ $M, v \models \varphi$
i.e. $M, v \models B_j(\varphi_1 \vee .. \vee \varphi_k \rightarrow \varphi)$
So $M, w \models \psi \rightarrow B_j(\varphi_1 \vee .. \vee \varphi_k \rightarrow \varphi)$ $\square$

Corollary C.2.5 *Let $\sigma_1, ..., \sigma_k$ be the actions $\rightarrow_j$ accessible from σ and $\psi = Pre(\sigma)$ and φ a propositional formula. Then,*
$B_j\varphi \rightarrow [\sigma, \varphi_1, ..., \varphi_n]B_j\varphi$

Proposition C.2.6 *for all consistent set of propositional formulae $\{\varphi_1, ..., \varphi_n, \varphi\}$ and all simple action σ,*
$\not\models_S B_j\varphi \rightarrow [\sigma, \varphi_1, ..., \varphi_n]B_j\varphi$.

Proof: Let $\sigma_1, .., \sigma_k$ be the actions $\sim_j$ accessible from σ and of plausibility 0, $\{\varphi_1, .., \varphi_n, \varphi\}$ consistent set of propositional formulas and σ action such that $Pre(\sigma) = \psi$

Let M be a model st there is w st
$M, w \models \psi, M, w \models B_j\varphi$
and such that there is v st $w \sim_j v$ $M, v \models \neg\varphi \wedge \varphi_i$ for some i, and $\kappa_j(v) = \kappa_j^w(\varphi_i) \neq 0$
then $M(\Sigma), (v, \sigma_i) \models \neg\varphi$ and $\kappa_j(v, \sigma_i) = 0$
so $M(\Sigma), (w, \sigma) \models \neg B_j\varphi$
so $M, w \models_S [\sigma, \varphi_1, .., \varphi_n]\neg B_j\varphi$

Now assume $\models_S B_j\varphi \rightarrow [\sigma, \varphi_1, .., \varphi_n]B_j\varphi$
then $M, w \models B_j\varphi \rightarrow [\sigma, \varphi_1, .., \varphi_n]B_j\varphi$
but $M, w \models B_j\varphi$,
so $M, w \models [\sigma, \varphi_1, .., \varphi_n]B_j\varphi$
thereby $M, w \models [\sigma, \varphi_1, .., \varphi_n]B_j\varphi \wedge \neg B_j\varphi$
i.e. $M, w \models [\sigma, \varphi_1, .., \varphi_n] \perp$
so $M, w \models \neg\varphi$. which is impossible. $\square$

Is the belief update a genuine revision?

Proposition C.2.7 *Let $\sigma_1, ..., \sigma_k$ be the actions that are $\rightarrow_j$ accessible from σ, $\models_{BMS}$ $B_j\neg\varphi_1 \wedge ... \wedge B_j\neg\varphi_k \rightarrow [\sigma, \varphi_1, ..., \varphi_k, \varphi_{k+1}, ..., \varphi_n]B_j \perp$*

Proof: Assume $M, w \models B_j\neg\varphi_1 \wedge .. \wedge B_j\neg\varphi_k$ (*), then

assume$M, w \models Pre(\sigma)$
assume there is $(v, \sigma_i) \in M(\Sigma)$ st $(w, \sigma) \rightarrow_j (v, \sigma_i)$
then $w \rightarrow_j v$ and $\sigma \rightarrow_j \sigma_i$ and $M, v \models \sigma_i$
then $M, v \models \neg\varphi_1 \wedge .. \wedge \neg\varphi_k$ because (*)
and $\sigma_i \in \{\sigma_1, .., \sigma_k\}$, so $M, v \models \varphi_i$ for some $i \in \{1, .., k\}$
this is impossible
so $M(\Sigma), (w, \sigma) \models B_j \perp$.
Thus $M, w \models [\sigma, \varphi_1, .., \varphi_n]B_j \perp$ $\square$

Corollary C.2.8 *Let $\sigma_1, ..., \sigma_k$ be the actions that are $\rightarrow_j$ accessible from σ,*
$\models_{BMS} B_j \neg\varphi \rightarrow [\sigma, \varphi, .., \varphi, \varphi_{k+1}, .., \varphi_n] B_j \perp$.

Proposition C.2.9 *Let $\sigma_1, ..., \sigma_k$ be the actions that are $\sim_j$ accessible from σ and of plausibility at most m,*
$\models_S (\neg K_j \neg\varphi_1 \wedge B_j^m \varphi_1 \vee .. \vee (\neg K_j \neg \varphi_k \wedge B_j^m \neg\varphi_k) \rightarrow [\sigma, \varphi_1, .., \varphi_k, \varphi_{k+1}, .., \varphi_n] \neg B_j^m \perp$.

Proof: Let $i \in \{1, .., k\}$.

Assume $M, w \models Pre(\sigma)$
because there is v_i st $v_i \sim_j w$ and $M, v_i \models \varphi_i$
let v_i st $\kappa_j(v_i) = min\{\kappa_j(v_i); w \sim_j v_i$ and $M, v_i \models \varphi_i\}$
then $M(\Sigma), (w, \sigma) \models \neg B_j^m \perp$
then $M, w \models [\sigma, \varphi_1, .., \varphi_k, \varphi_{k+1}, .., \varphi_n] \neg B_j^m \perp$ □

Corollary C.2.10 *Let $\sigma_1, .., \sigma_k$ be the actions $\sim_j$ accessible from σ and of plausibility 0.*
$\models_S \neg K_j \neg\varphi \wedge B_j \neg\varphi \rightarrow [\sigma, \varphi, .., \varphi, \varphi_{k+1}, .., \varphi_n] \neg_j \perp$.

Corollary C.2.11 *Let $\sigma_1, .., \sigma_k$ be the actions that are $\sim_j$ accessible from σ and of plausibility 0,*
$\models_S K_j \neg\varphi \rightarrow [\sigma, \varphi, .., \varphi, \varphi_{k+1}, .., \varphi_n] B_j \perp$.

Proposition C.2.12 *For all formula belief epistemic formula φ (then without action term) and all belief epistemic model M,*
$F(M), w \models_{BMS} \varphi$ iff $M, w \models_S \varphi$

Proposition C.2.13 *For all $\varphi_1, .., \varphi_n, \varphi$ epistemic formulae,*
If $M, w \models_S [\sigma, \varphi_1, .., \varphi_n] B_j \varphi$ then $F(M), w \models_{BMS} [\sigma, \varphi_1, .., \varphi_n] B_j \varphi$

Proof: assume $[\sigma, \varphi_1, .., \varphi_n] B_j \varphi$

assume $F(M), w \models \psi$ where $\psi = Pre(\sigma)$
then $M, w \models \psi$ because ψ is an epistemic formula
then $M(\Sigma), (w, \sigma) \models B_j \varphi$
then for all (v, τ) st $(w, \sigma) \sim_j (v, \tau)$ and $\kappa_j(v, \tau) = 0$, $M(\Sigma), (v, \tau) \models \varphi$
But the epistemic structure of $M(\Sigma)$ is the same as the one of $F(M)(\Sigma')$, and φ is an epistemic formula, so
for all (v, τ) st $(w, \sigma) \sim_j (v, \tau)$ and $\kappa_j(v, \tau) = 0$, $F(M)(\Sigma'), (v, \tau) \models \varphi$
then for all (v, τ) st $(w, \sigma) \rightarrow_j (v, \tau)$, $F(M)(\Sigma'), (v, \tau) \models \varphi$ because if $(w, \sigma) \rightarrow_j (v, \tau)$ then $(w, \sigma) \sim_j (v, \tau)$ and $\kappa_j(v, \tau) = 0$
then $F(M)(\Sigma'), (w, \sigma) \models B_j \varphi$
so $F(M), w \models [\sigma, \varphi_1, .., \varphi_n] B_j \varphi$ □

Proposition C.2.14 *For all model M of S and all epistemic formula φ,*
$M, w \models_S [\sigma, \varphi_1, .., \varphi_n] K_j \varphi$ iff $F(M), w \models [\sigma, \varphi_1, .., \varphi_n] K_j \varphi$

Proof: Because the epistemic structure of $M(\Sigma)$ is the same as $F(M)(\Sigma')$.

Appendix D

Proofs missing in the section "Comparison with the AGM postulates"

D.1 Proof of theorem 4.3.3

In this proof we will indicate as extra assumption during the calculus of the postulate unsatisfied the reasons why they are not fulfilled.

Now we are going to check if the 8 postulates are verified:

$(K * 1)$: yes.

$(K * 2)$: no.

$$\begin{aligned} A \in K^w * A &\Leftrightarrow M, w \models [A!, A]B_j A \\ &\Leftrightarrow \text{if } M, w \models A \text{ then } M(A!), (w, A!) \models B_j A, \text{ where} M(A!) \text{ is the model } M \text{ updated with } A! \\ &\Leftrightarrow M(A!), (w, A!) \models B_j A \end{aligned}$$

This last equivalence holds if for all $v \in M$, $M, v \models A \Leftrightarrow M(A!), (v, A!) \models A$ (see issue 0 in the next section).

$(K * 3) - (K * 4)$: no.

$$\begin{aligned} \neg A \notin K^w &\Leftrightarrow M, w \models \neg B_j \neg A \\ &\Leftrightarrow \kappa_j^w(A) = 0.(*) \end{aligned}$$

$$\begin{aligned} \varphi \in K^w * A &\Leftrightarrow M, w \models [A!, A]B_j \varphi \\ &\Leftrightarrow \text{if } M, w \models A \text{ then } M(A!), (w, A!) \models B_j \varphi. \\ &\Leftrightarrow M(A!), (w, A!) \models B_j \varphi. \end{aligned}$$

$$
\begin{aligned}
\varphi \in K^w +' A &\Leftrightarrow M, w \models B_j[A!, A](A \to \varphi)\\
&\Leftrightarrow \text{ for all v st } \kappa_j^w(v) = 0, M, v \models [A!, A]A \to \varphi\\
&\Leftrightarrow \text{ for all v st } \kappa_j^w(v) = 0, \text{ if } M, v \models A \text{ then } M(A!), (v, A!) \models A \to \varphi\\
&\Leftrightarrow \text{ for all v st } \kappa_j^w(v) = 0 \text{ and } M, v \models A, M(A!), (v, A!) \models A \to \varphi\\
&\Leftrightarrow \text{ for all v st } \kappa_j^w(v) = 0 \text{ and } M, v \models A, M(A!), (v, A!) \models \varphi.\\
&\quad\ \text{if } M, v \models A \Leftrightarrow (A!), (v, A!) \models A.\\
&\Leftrightarrow \text{ for all } (v, A!) \text{ st } \kappa_j^{(w,A!)}(v, A!) = 0, M(A!), (v, A!) \models \varphi, \text{ because of (*).}\\
&\quad\ \text{if } M, v \models A \Leftrightarrow M(A!), (v, A!) \models A.\\
&\Leftrightarrow M(A!), (w, A!) \models B_j\varphi\\
&\quad\ \text{if } M, v \models A \Leftrightarrow M(A!), (v, A!) \models A.\\
&\Leftrightarrow \varphi \in K^w * A \text{ if } M, v \models A \Leftrightarrow M(A!), (v, A!) \models A \text{ (see issue 0).}
\end{aligned}
$$

$(K * 5)$: Clearly, if $\vdash \neg A$ then $K^w * A = K_\perp$ the inconsistent logic because of soundness.

Now, $M, w \models A$ by assumption, so there is v st $(w, A!) \sim_j (v, A!)$ and $\kappa_j^{(w,A!)}(v, A!) = 0$ because $\{\kappa_j^{(w,A!)}(v, A!); v \sim_j w\} \neq \emptyset$. But $M, w \models [A!, A]B_j\varphi \Leftrightarrow M(A!), (w, A!) \models B_j\varphi$. So, $M(A!), (w, A!) \models \neg B_j\varphi$. Then $\perp \notin K^w * A$. And then $K^w * A \neq K_\perp$. So, if $\neg \vdash \neg A$ then $K^w * A \neq K_\perp$; i.e. if $K^w * A = K_\perp$ then $\vdash \neg A$.

Finally, $\vdash \neg A \Leftrightarrow K * A = K_\perp$.

$(K * 6)$: yes, clearly.

$(K * 7) - (K * 8)$: No.

$$
\begin{aligned}
\neg B \notin K^w * A &\Leftrightarrow M, w \models \neg[A!, A]B_j\neg B\\
&\Leftrightarrow M(A!), (w, A!) \models \neg B_j \neg B\\
&\Leftrightarrow \kappa_j^{(w,A!)}(B) = 0(*)
\end{aligned}
$$

$$
\begin{aligned}
\varphi \in K^w * A \wedge B &\Leftrightarrow M, w \models [A \wedge B!, A \wedge B]B_j\varphi\\
&\Leftrightarrow \text{ if } M, w \models A \wedge B \text{ then } M(A \wedge B), (w, A \wedge B!) \models B_j\varphi\\
&\Leftrightarrow M(A \wedge B), (w, A \wedge B!) \models B_j\varphi\\
&\Leftrightarrow \text{ for all } (v, A \wedge B!) \text{ st } \kappa_j^{(w,A\wedge B!)}(v, A \wedge B!) = 0, M(A \wedge B!) \models \varphi.
\end{aligned}
$$

$\varphi \in K^w * A +' B$

$\Leftrightarrow$ $M, w \models [A!, A]B_j[B!, B](B \to \varphi)$

$\Leftrightarrow$ $M(A!), (w, A!) \models B_j[B!, B](B \to \varphi)$

$\Leftrightarrow$ for all $(v, A!)$, st $\kappa_j^{(w,A!)}(v, A!) = 0, M(A!), (v, A!) \models [B!, B](B \to \varphi)$

$\Leftrightarrow$ for all $(v, A!)$, st $\kappa_j^{(w,A!)}(v, A!) = 0,$
if $M(A!), (v, A!) \models B$ then $M(A!), (B!), ((v, A!), B!) \models B \to \varphi$

$\Leftrightarrow$ for all $(v, A!)$, st $M(A!), (v, A!) \models B$ and $\kappa_j^{(w,A!)}(v, A!) = 0,$
$M(A!), (B!), ((v, A!), B!) \models B \to \varphi$ because of (*).

$\Leftrightarrow$ for all v st $M, v \models A \wedge B$ and $\kappa_j^{(w,A!),B!)}((v, A!), B!) = 0,$
$M(A!)(B!)((v, A!), B!) \models B \to \varphi$
if $M, v \models B \Leftrightarrow M(A!), (v, A!) \models B$ (see issue 1 in the next section)

$\Leftrightarrow$ for all v st $M, v \models A \wedge B$ and $\kappa_j^{(w,A\wedge B!)}((v, A \wedge B!) = 0,$
$M(A!)(B!)((v, A!), B!) \models B \to \varphi$
if $M, v \models B \Leftrightarrow M(A!), (v, A!) \models B$ (see issue 1)

$\Leftrightarrow$ for all v st $M, v \models A \wedge B$ and $\kappa_j^{(w,A\wedge B!)}((v, A \wedge B!) = 0,$
$M, v \models B \to \varphi$
1)if $M, v \models B \Leftrightarrow M(A!), (v, A!) \models B$ (see issue 1)
2)and $M(A!), (v, A!) \models B \to \varphi \Leftrightarrow M(A!)(B!), ((v, A!), B!) \models B \to \varphi$ (see issue 1)
3)and $M, v \models B \to \varphi \Leftrightarrow M(A!), (v, A!) \models B \to \varphi$ (see issue 1).

$\Leftrightarrow$ for all v st $M, v \models A \wedge B$ and $\kappa_j^{(w,A\wedge B!)}((v, A \wedge B!) = 0, M, v \models \varphi$
under the same assumptions 1)2)3).

$\Leftrightarrow$ for all $(v, A \wedge B!)$st $\kappa_j^{(w,A\wedge B!)}(v, A \wedge B!) = 0, M(A \wedge B!), (v, A \wedge B!) \models \varphi$
under the assumptions 1)2)3) and
4)$M, v \models \varphi \Leftrightarrow M, (A \wedge B!), (v, A \wedge B!) \models \varphi$ (see issue 1)

$\Leftrightarrow$ $\varphi \in K^w * (A \wedge B)$ under the assumption 1)2)3)4).

D.2 Explanation of the Failure of some Postulates in the Theorem 4.3.3

In these proofs we remark that they fail for almost always the same kind of problems. Indeed, by updating a public announcement,we get rid of worlds, so epistemic link disappear. Hence formulas true at a given world are different before and after an announcement.

More precisely, we have to cope with these 2 issues:

issue 0: Which epistemic formulas imply their self relativization?

issue 1: Which forms of epistemic assertion remain true at a world whenever other worlds are eliminated from the model ?(persistence)

Note that (see [vB03]) Issue 0 can be solved by the

Fact: in each model, every public announcement has a persistent equivalent.

Moreover issue 1 can be solved by the

theorem: The epistemic formulas without common knowledge that are preserved under submodels are precisely those definable using literals $p, \neg p$, conjunction, disjunction, and K_i-operators.

D.3 Proof of theorem 4.3.2

The proof is completely similar to theorem 4.3.3, but here we do not encounter the issues 0 and 1 because we deal with propositional formulas (see section 4.2). So all the postulates are verified.

D.4 Proof of theorem 4.3.5

The proof is completely similar to theorem 4.3.3 but as in theorem 4.3.2, issues 0 and 1 are not problematic anymore because the epistemic structure is preserved after the update. The rest of the proof (concerning firmness) is obvious.

Bibliography

[BMS03] A. Baltag, L.S. Moss,and S. Solecki. The logic of public announcements, common knowledge and private suspicions. Manuscript, originally presented at TARK 98, 2003.

[BdRV01] P. Blackburn, M. de Rijke,and Y. Venema. *Modal Logic*. Cambridge University Press, Cambridge, 2001. Cambridge Tracts in Theoretical Computer Science 53.

[FHMV95] R. Fagin, J.Y. Halpern, Y. Moses, and M.Y. Vardi. *Reasoning about knowledge.* MIT Press, Cambridge MA, 1995.

[GardRott95] P. Gardenfors and H. Rott, 1995, 'Belief Revision', in D. M. Gabbay, C. J. Hogger and J. A. Robinson, eds., *Handbook of Logic in Artificial Intelligence and Logic Programming* 4, Oxford University Press, Oxford 1995.

[Spohn90] W. Spohn. A general non-probability theory of inductive reasoning. In R. D. Schachter, T. S. Levitt, L. N. Kanal, and J. F. Lemmer, editors, *Uncertainty in artificial intelligence 4*, pages149-15. Norht-Holland, Amsterdam, 1990.

[Spohn88] W. Spohn. Ordinal conditional functiuns: A dynamic theory of epistemic states. In W. L. Harper and B. Skyrms,editors, *Causation in Decision, Belief Change, and Statistics*, vol.2, pages 105-134. reidel, Dordrecht, 1988.

[vB03] J. van Benthem. *'One is a Lonely Number':on the logic of communication.* Logic Colloquium 2002.

[vBJoLLI03] J. van Benthem. Conditional probability meets update logic. In *Journal of Logic, Language and Information* Volume 12, Issue 4, Fall 2003. Pages: 409 - 421.

[vDL03] H.P. van Ditmarsch and W.A. Labuschagne, Dynamic Doxastic Logic for Defeasible Belief Revision , work in progress (version July 2003)