

Temporal interaction of information and belief

Giacomo Bonanno*

Department of Economics,
University of California,
Davis, CA 95616-8578, USA
e-mail: gfbonanno@ucdavis.edu

December 2005

Abstract

The temporal updating of an agent's beliefs in response to a flow of information is modeled in a simple modal logic that, for every date t , contains a normal belief operator B_t and a non-normal information operator I_t which is analogous to the 'only knowing' operator discussed in the computer science literature. Soundness and completeness of the logic are proved and the relationship between the proposed logic and the AGM theory of belief revision is discussed.

1 Introduction

Belief revision is a central topic in several fields. In game theory, belief revision is the main building block of two widely used solution concepts for dynamic (or extensive) games, namely perfect Bayesian equilibrium (see, for example, Battigalli [2], Bonanno [5] and Fudenberg and Tirole [9]) and sequential equilibrium (Kreps and Wilson [15]). The idea behind these solution concepts is that, during the play of the game, a player should revise his beliefs by using Bayes' rule "as long as possible". Thus if an information set has been reached that had positive prior probability, then beliefs at that information set are obtained by using Bayes' rule (with the information being represented by the set of nodes in the information set under consideration). If an information set is reached that had zero prior probability, then new beliefs are formed more or less arbitrarily, but from that point onwards these new beliefs must be used in conjunction with Bayes' rule, unless further information is received that is inconsistent with those revised beliefs. In computer science the theory of belief revision pioneered by Alchourron *et al* [1] (known as the AGM theory) has been studied extensively

*A first draft of this paper was presented at the Workshop on Belief Change in Rational Agents: Perspectives from Artificial Intelligence, Philosophy and Economics, Dagstuhl (Germany), August 2005 (<http://drops.dagstuhl.de/portals/05321/>).

(for an overview see Gärdenfors [10]). While in game theory beliefs are typically represented by a probability distribution over a set of states and belief revision is modeled in terms of Bayes' rule, in the AGM theory beliefs are modeled syntactically as sets of formulas in a given language (called belief sets). Information is thought of as a formula in this language and belief revision is modeled as an operation that transforms a belief set into a new belief set that incorporates the information. Within the AGM tradition the issue of iterated belief revision has recently received considerable attention (see, for example, Nayak *et al* [18]).

In this paper we propose a simple modal logic for *iterated* belief revision, extending the two-period framework of Bonanno [6]. For every date $t \in \mathbb{N}$ (where \mathbb{N} is the set of natural numbers) we postulate a belief operator B_t and an information operator I_t . The interaction of information and belief over time is captured by several axioms. We start with three simple axioms and show that the corresponding logic is sound and complete with respect to the class of Kripke structures that satisfy the iterative version of the qualitative content of Bayes' rule. The addition of further axioms yields a logic which is sound and complete with respect to the class of structures that capture the basic postulates of the AGM theory. Finally, we show that a strengthening of this logic corresponds to the class of structures where any belief revision history can be rationalized in terms of a plausibility ordering of the set of states.

In the next section we begin with an example that illustrates the formal logic and structures analyzed in the remainder of the paper.

2 A motivating example

A doctor examines a patient who reports an outbreak of skin rashes. The patient claims not to have made any changes in his diet and gives the doctor a list of medications that he has been taking for a number of years. Based on her experience, the doctor narrows down the possible causes to four: bacterial infection (B), viral infection (V), allergic reaction to food (F) and allergic reaction to medication (M). An initial assessment of the case (**Time 0**) leads the doctor to believe that it is an infection. Since she knows of no effective treatment for a viral infection, she prescribes antibiotics. A few days later (**Time 1**) the patient reports that there has been no change in his symptoms. The doctor treats this report as information that it is not a case of bacterial infection and becomes convinced that it is a viral infection. She informs the patient that, unfortunately, there are no drugs that would be effective against a viral infection. The patient requests a blood test. A positive result would confirm the presence of an infection (which, however, could be either bacterial or viral), while a negative result would rule out an infection. The doctor yields to the patient's request and a few days later (**Time 2**) the lab reports a negative result to the blood test. Based on this information, the doctor reaches the conclusion that the patient must have developed a sensitivity to one of the drugs and advises the patient to stop taking all his medications. A few days later (**Time 3**) the patient reports, once again, that there has been no change in his symptoms. The doctor then

concludes that it must be an allergic reaction to food and instructs the patient to keep a detailed food diary.

In this example the doctor’s beliefs evolve over time in response to information. The history of the doctor’s beliefs can be represented using a sequence of Kripke structures (Kripke [16]): the set of states or possible worlds is $\Omega = \{B, V, M, F\}$ (where B denotes bacterial infection, V viral infection, M medication allergy and F food allergy) and at every date t the doctor’s beliefs can be represented by a binary relation \mathcal{B}_t on Ω as shown in Figure 1. In all the figures we represent a binary relation $R \subseteq \Omega \times \Omega$ as follows: (1) if a rounded rectangle encloses a set of states then, for any two states ω and ω' in that rectangle, $\omega R \omega'$ and (2) if there is an arrow from a state ω to a rounded rectangle, then for any state ω' in that rectangle, $\omega R \omega'$. For example, in Figure 1 we have that $\mathcal{B}_0 = \{(B, B), (B, V), (V, B), (V, V), (M, B), (M, V), (F, B), (F, V)\}$, $\mathcal{B}_1 = \{(B, V), (V, V), (M, V), (F, V)\}$, etc.




Figure 1
The evolution of the doctor’s beliefs over time

Syntactically, let B_t be the belief operator at time t , so that the interpretation of $B_t\phi$ is “at time t the individual believes that ϕ ”. If ω is a state and ϕ a formula, we denote by $\omega \models \phi$ the fact that ϕ is true at state ω . The truth of the formula $B_t\phi$ at state ω is then determined as usual: $\omega \models B_t\phi$ if and only, for every ω' such that $\omega \mathcal{B}_t \omega'$, $\omega' \models \phi$, that is, if ϕ is true at every state that, at date t , the individual considers possible at state ω . For example, if ϕ is the proposition “the patient has an infection” then, in our example, it is true at -

and only at - states B and V. If the true state is M then at that state it is false that the patient has an infection and yet it is true that at date 0 the doctor believes that the patient has an infection ($M \models \neg\phi$ and $M \models B_0\phi$).

Changes in the doctor’s beliefs are brought about by the receipt of new information. It would be desirable to represent the flow of information over time in the same way in which we represent beliefs, namely by means of a sequence of binary relations \mathcal{I}_t for every date t . We propose to model information in a way which is reminiscent of the notion of “only knowing” (Levesque [17]). Intuitively, we interpret “I am informed that ϕ ” as equivalent to “all I am told is ϕ ”. We can capture this interpretation of information by means of the following validation rule. Let I_t be the time t information operator, so that the interpretation of $I_t\phi$ is “at time t the individual is informed that ϕ ”. Then we set $\omega \models I_t\phi$ if and only if two conditions hold: (1) for every ω' such that $\omega\mathcal{I}_t\omega'$, $\omega' \models \phi$, and (2) for every $\omega' \in \Omega$, if $\omega' \models \phi$ then $\omega\mathcal{I}_t\omega'$. That is, $I_t\phi$ is true at state ω if the set of states reachable from ω by means of the relation \mathcal{I}_t coincides with the truth set of ϕ .¹ With this interpretation in mind, we can represent the flow of information received by the doctor by means of the sequence of binary relations depicted in Figure 2.




Figure 2
The flow of information

The doctor’s beliefs at time $t + 1$ are the result of the interaction between her beliefs at time t and the information received at time $t + 1$. In our example, at time 0 the doctor believes that the patient has a (bacterial or viral) infection and at time 1 she is informed that the patient does not have a bacterial infection. Combining the two, the doctor modifies her beliefs and comes to the new belief that the patient has a viral infection. In a probabilistic setting, this new belief is what would be required by Bayes’ rule. The interaction of old beliefs and

¹Thus our information operator plays the role of the conjunction of the two operators K and O proposed by Levesque [17] where $K\phi$ is interpreted as “the individual knows that ϕ ”, while $O\phi$ is interpreted as “the individual *only* knows that ϕ ”.

new information in our example is shown in Figure 3, which puts together the structures of Figures 1 and 2.




Figure 3
The interaction of beliefs and information over time

It is worth noting that, while the first piece of information received by the doctor (“it is not a bacterial infection”) does not cause surprise – since it is compatible with the doctor’s initial belief (it could be a viral infection) – later pieces of information do cause surprise, since they contradict the doctor’s beliefs.

In the next section we propose a logic that is sound and complete with respect to the class of Kripke structures that satisfy the qualitative version of Bayes’ rule (the structure illustrated in Figure 3 belongs to this class). We then consider a strengthening of that logic, obtained by adding three hypotheses: that beliefs are consistent, that information is non-contradictory and that information is always believed (even when it contradicts current beliefs; the structure of Figure 3 satisfies all these hypotheses). We show that the individual’s belief revision over time satisfies the basic postulates of the AGM belief revision theory (see Alchourron *et al* [1]). Finally, we show that when information becomes more refined over time (so that every new piece of information implies earlier information, as is the case in the structure of Figure 3) the individual’s belief revision can be rationalized in terms of a plausibility ordering over the set of states: at every date the individual considers possible all and only those

states that are most plausible among the ones that are compatible with the new information received.

3 A logic for iterated belief revision

Let \mathbb{N} be the set of natural numbers. We consider a modal propositional logic based on the following operators: a belief operator B_t and an information operator I_t , for every date $t \in \mathbb{N}$, and a “global” operator² A . The intended interpretation is as follows:

- $I_t\phi$ at time t all the individual is informed of is that ϕ
- $B_t\phi$ at time t (after revising his earlier beliefs in light of the information just received) the individual believes that ϕ
- $A\phi$ it is globally true that ϕ .

The formal language is built in the usual way from a countable set S of atomic propositions, the connectives \neg (for “not”) and \vee (for “or”) and the modal operators.³ Thus the set Φ of formulas is defined inductively as follows: $q \in \Phi$ for every atomic proposition $q \in S$, and if $\phi, \psi \in \Phi$ then all of the following belong to Φ : $\neg\phi$, $\phi \vee \psi$, $A\phi$ and, for every $t \in \mathbb{N}$, $B_t\phi$ and $I_t\phi$.

We denote by \mathfrak{L} the logic determined by the following axioms and rules of inference.

AXIOMS:

1. All propositional tautologies.
2. For every $t \in \mathbb{N}$, axiom K for B_t and A (note the absence of an analogous axiom for I_t):

$$\begin{aligned} \text{K}_B. \quad & B_t\phi \wedge B_t(\phi \rightarrow \psi) \rightarrow B_t\psi \\ \text{K}_A. \quad & A\phi \wedge A(\phi \rightarrow \psi) \rightarrow A\psi \end{aligned}$$

3. S5 axioms for A :

$$\begin{aligned} \text{T}_A. \quad & A\phi \rightarrow \phi \\ 5_A. \quad & \neg A\phi \rightarrow A\neg A\phi \end{aligned}$$

4. For every $t \in \mathbb{N}$, inclusion axiom for B_t (note the absence of an analogous axiom for I_t):

$$\text{Incl.} \quad A\phi \rightarrow B_t\phi$$

5. Axioms to capture the non-standard semantics for I :

²For a thorough discussion of the global (or universal) modality see Goranko and Passy [11].

³See, for example, Blackburn *et al* [3]. The connectives \wedge (for “and”), \rightarrow (for “if ... then ...”) and \leftrightarrow (for “if and only if”) are defined as usual: $\phi \wedge \psi = \neg(\neg\phi \vee \neg\psi)$, $\phi \rightarrow \psi = \neg\phi \vee \psi$ and $\phi \leftrightarrow \psi = (\phi \rightarrow \psi) \wedge (\psi \rightarrow \phi)$.

$$\begin{aligned} \text{I}_1. \quad & (I_t\phi \wedge I_t\psi) \rightarrow A(\phi \leftrightarrow \psi) \\ \text{I}_2. \quad & A(\phi \leftrightarrow \psi) \rightarrow (I_t\phi \leftrightarrow I_t\psi) \end{aligned}$$

RULES OF INFERENCE:

1. Modus Ponens: $\frac{\phi, \phi \rightarrow \psi}{\psi}$ (MP)
2. Necessitation for A : $\frac{\phi}{A\phi}$ (Nec_A)

On the semantic side, a *frame* is a collection $\langle \Omega, \{\mathcal{B}_t\}_{t \in \mathbb{N}}, \{\mathcal{I}_t\}_{t \in \mathbb{N}} \rangle$ where Ω is a set of *states* and, for every $t \in \mathbb{N}$, \mathcal{B}_t and \mathcal{I}_t are binary relations on Ω , whose interpretation is as follows:

- $\alpha \mathcal{B}_t \beta$ at time t and state α the individual considers state β possible
- $\alpha \mathcal{I}_t \beta$ at state α , state β is compatible with the information received at time t .

Let $\mathcal{B}_t(\omega) = \{\omega' \in \Omega : \omega \mathcal{B}_t \omega'\}$ denote the set of states that, at date t , the individual considers possible at state ω . Define $\mathcal{I}_t(\omega)$ similarly.

The connection between syntax and semantics is given by the notion of model. Given a frame $\langle \Omega, \{\mathcal{B}_t\}_{t \in \mathbb{N}}, \{\mathcal{I}_t\}_{t \in \mathbb{N}} \rangle$, a *model* is obtained by adding a *valuation* $V : S \rightarrow 2^\Omega$ (where 2^Ω denotes the set of subsets of Ω , usually called *events*) which associates with every atomic proposition $p \in S$ the set of states at which p is true.⁴ The truth of an arbitrary formula at a state is then defined inductively as follows ($\omega \models \phi$ denotes that formula ϕ is true at state ω ; $\|\phi\|$ is the truth set of ϕ , that is, $\|\phi\| = \{\omega \in \Omega : \omega \models \phi\}$):

- if q is an atomic proposition, $\omega \models q$ if and only if $\omega \in V(q)$,
- $\omega \models \neg\phi$ if and only if $\omega \not\models \phi$,
- $\omega \models \phi \vee \psi$ if and only if either $\omega \models \phi$ or $\omega \models \psi$ (or both),
- $\omega \models B_t\phi$ if and only if $\mathcal{B}_t(\omega) \subseteq \|\phi\|$,
- $\omega \models I_t\phi$ if and only if $\mathcal{I}_t(\omega) = \|\phi\|$,
- $\omega \models A\phi$ if and only if $\|\phi\| = \Omega$.

Remark 1 *Note that, while the truth condition for $B_t\phi$ is the standard one, the truth condition of $I_t\phi$ is unusual in that the requirement is $\mathcal{I}_t(\omega) = \|\phi\|$ rather than merely $\mathcal{I}_t(\omega) \subseteq \|\phi\|$. This is what establishes the similarity between our information operator and the “only knowing” operator discussed in the literature (see Levesque [17]).*

⁴Note that by making the truth of atomic propositions depend on the state only, rather than on the state and time, we restrict ourselves to situations of belief revision, where the objective state of the world does not change over time: only the epistemic state of the individual changes. The alternative case, where the truth of the atomic propositions is allowed to change over time, is known in the computer science literature as “belief update” (see Katsuno and Mendelzon [14]).

We say that a formula ϕ is *valid in a model* if $\omega \models \phi$ for all $\omega \in \Omega$, that is, if ϕ is true at every state. A formula ϕ is *valid in a frame* if it is valid in every model based on that frame. A logic is *sound* with respect to a class of frames if every theorem of the logic is valid in every frame in that class; it is *complete* with respect to a class of frames if every formula which is valid in every frame in that class is provable in the logic (that is, it is a theorem).

Remark 2 Note that from (Nec_A) and $(Incl)$ one obtains necessitation for B_t as a derived rule of inference: $\frac{\phi}{B_t\phi}$. Furthermore, from necessitation and axiom K one obtains the following derived rule of inference for both A and B_t (usually referred to as rule RK): $\frac{\phi \rightarrow \psi}{A\phi \rightarrow A\psi}$ and $\frac{\phi \rightarrow \psi}{B_t\phi \rightarrow B_t\psi}$. On the other hand, the necessitation rule for I is **not** a rule of inference of logic \mathcal{L} . Indeed necessitation for I_t is not validity preserving⁵; neither is rule RK for I_t .⁶ On the other hand, by Nec_A and I_2 , the following rule for I_t (usually referred to as rule RE): $\frac{\phi \leftrightarrow \psi}{I_t\phi \leftrightarrow I_t\psi}$ is a derived rule of inference of \mathcal{L} .

Note that, despite the non-standard validation rule, axiom K for I_t , namely $I_t\phi \wedge I_t(\phi \rightarrow \psi) \rightarrow I_t\psi$, is trivially valid in every frame.⁷ It follows from the completeness part of Proposition 3 below that axiom K for I_t ($I_t\phi \wedge I_t(\phi \rightarrow \psi) \rightarrow I_t\psi$) is a theorem of logic \mathcal{L} (a syntactic proof is also easily obtained).

The following proposition is an extension of the two-period framework of Bonanno [6]. An outline of the proof is given in the Appendix.

Proposition 3 Logic \mathcal{L} is sound and complete with respect to the class of all frames $\langle \Omega, \{\mathcal{B}_t\}_{t \in \mathbb{N}}, \{\mathcal{I}_t\}_{t \in \mathbb{N}} \rangle$.

We now consider extensions of logic \mathcal{L} . The first extension, denoted by \mathcal{L}_{QBR} , is obtained by adding to \mathcal{L} the following axioms:

$$\begin{aligned} \text{Qualified Acceptance (QA):} & \quad (\neg B_t \neg \phi \wedge I_{t+1}\phi) \rightarrow B_{t+1}\phi \\ \text{Persistence (P):} & \quad (\neg B_t \neg \phi \wedge I_{t+1}\phi) \rightarrow (B_t\psi \rightarrow B_{t+1}\psi) \\ \text{Minimality (M):} & \quad (I_{t+1}\phi \wedge B_{t+1}\psi) \rightarrow B_t(\phi \rightarrow \psi). \end{aligned}$$

One of the axioms of the AGM theory of belief revision is the so-called Success or Acceptance axiom, which requires that information be believed, that is, that it be incorporated in the revised beliefs. Our Qualified Acceptance axiom

⁵If ϕ is a valid formula, then $\|\phi\| = \Omega$. Let $\omega \in \Omega$ be a state where $\mathcal{I}_t(\omega) \neq \Omega$. Then $\omega \not\models I_t\phi$ and therefore $I_t\phi$ is not valid.

⁶That is, from the validity of $\phi \rightarrow \psi$ one cannot infer the validity of $I_t\phi \rightarrow I_t\psi$. To see this, consider the following model: $\Omega = \{\alpha, \beta\}$, $\mathcal{I}_t(\alpha) = \{\alpha\}$, $\mathcal{I}_t(\beta) = \{\beta\}$, $\|p\| = \{\alpha\}$ and $\|q\| = \Omega$. Then $\|p \rightarrow q\| = \Omega$, $\|I_t p\| = \{\alpha\}$, $\|I_t q\| = \emptyset$ and thus $\|I_t p \rightarrow I_t q\| = \{\beta\} \neq \Omega$.

⁷Proof. Fix a frame, an arbitrary model and a state ω . For it to be the case that $\omega \models I_t\phi \wedge I_t(\phi \rightarrow \psi)$ we need $\mathcal{I}_t(\omega) = \|\phi\|$ and $\mathcal{I}_t(\omega) = \|\phi \rightarrow \psi\|$. Now, $\|\phi \rightarrow \psi\| = \|\neg\phi \vee \psi\| = \|\neg\phi\| \cup \|\psi\|$ and therefore we need the equality $\|\phi\| = \|\neg\phi\| \cup \|\psi\|$ to be satisfied. This requires $\|\phi\| = \|\psi\| = \Omega$. Thus if $\mathcal{I}_t(\omega) = \|\phi\| = \|\psi\| = \Omega$, then $\omega \models I_t(\phi \rightarrow \psi) \wedge I_t\phi \wedge I_t\psi$. In every other case, $\omega \not\models I_t\phi \wedge I_t(\phi \rightarrow \psi)$ and therefore the formula $I_t\phi \wedge I_t(\phi \rightarrow \psi) \rightarrow I_t\psi$ is trivially true at ω .

is a weakening of this, in that it requires the individual who has been informed that ϕ to believe ϕ only if, before the receipt of information, he considered ϕ possible (that is, he did not believe $\neg\phi$). The Persistence axiom says that if the individual is informed of something that he previously considered possible, then he continues to believe everything that he believed before. Finally, the Minimality axiom states that beliefs should be revised in a minimal way, in the sense that no new beliefs should be added unless they are implied by the old beliefs and the information received.⁸

By *Qualitative Bayes Rule* (QBR) we mean the following property: $\forall\omega \in \Omega, \forall t \in \mathbb{N}$,

$$\text{if } \mathcal{B}_t(\omega) \cap \mathcal{I}_{t+1}(\omega) \neq \emptyset \text{ then } \mathcal{B}_{t+1}(\omega) = \mathcal{B}_t(\omega) \cap \mathcal{I}_{t+1}(\omega). \quad (\text{QBR})$$

QBR says that if at a state the information received at time $t + 1$ is compatible with the beliefs the individual had at time t , in the sense that there are states that he considered possible at date t and are compatible with the information received at date $t + 1$, then the states that are considered possible according to the revised beliefs at date $t + 1$ are precisely those states. For example, the frame illustrated in Figure 3 of Section 2 satisfies QBR (vacuously after date 1).

In a probabilistic setting, let $P_{t,\omega}$ be the probability measure over Ω representing the individual's beliefs at date t and state ω , let $F \subseteq \Omega$ be an event representing the information received by the individual at date $t + 1$ and let $P_{t+1,\omega}$ be the posterior probability measure representing the revised beliefs at date $t + 1$ and state ω . Bayes' rule says that, if $P_{t,\omega}(F) > 0$, then, for every event $E \subseteq \Omega$, $P_{t+1,\omega}(E) = \frac{P_{t,\omega}(E \cap F)}{P_{t,\omega}(F)}$. Bayes' rule thus implies the following:

$$\text{if } \text{supp}(P_{t,\omega}) \cap F \neq \emptyset, \text{ then } \text{supp}(P_{t+1,\omega}) = \text{supp}(P_{t,\omega}) \cap F$$

where $\text{supp}(P)$ denotes the support of the probability measure P . If we set $\mathcal{B}_t(\omega) = \text{supp}(P_{t,\omega})$, $F = \mathcal{I}_{t+1}(\omega)$ and $\mathcal{B}_{t+1}(\omega) = \text{supp}(P_{t+1,\omega})$ then we get the Qualitative Bayes Rule as stated above. Thus in a probabilistic setting the sentence “the individual believes ϕ ” would be interpreted as “the individual assigns probability 1 to the event $\|\phi\|$ ”.

The following proposition is an extension of the two-period framework of Bonanno [6]. An outline of the proof is given in the Appendix.

⁸The Minimality axiom is trivially satisfied if information is surprising: suppose that at a state, say ω , the individual at time $t + 1$ is informed that ϕ ($\omega \models I_{t+1}\phi$) although he previously believed that ϕ was *not* the case ($\omega \models B_t\neg\phi$). Then, for every formula ψ , the formula $(\phi \rightarrow \psi)$ is trivially true at every state that the individual considered possible at date t and state ω ($\mathcal{B}_t(\omega) \subseteq \|\phi \rightarrow \psi\|$) and therefore he believed it ($\omega \models B_t(\phi \rightarrow \psi)$). Thus the axiom restricts the new beliefs only when the information received is not surprising, that is, only if $\omega \models \neg B_t\neg\phi \wedge I_{t+1}\phi$.

Proposition 4 *Logic \mathcal{L}_{QBR} is sound and complete with respect to the class of frames $\langle \Omega, \{\mathcal{B}_t\}_{t \in \mathbb{N}}, \{\mathcal{I}_t\}_{t \in \mathbb{N}} \rangle$ that satisfy the Qualitative Bayes Rule.*

We now consider stronger logics than \mathcal{L}_{QBR} . Let \mathcal{L}_{AGM} be the logic obtained by adding to \mathcal{L} the following axioms:

Acceptance or Success (A):	$(I_t \phi \wedge \neg A \neg \phi) \rightarrow B_t \phi$
Persistence (P):	$(\neg B_t \neg \phi \wedge I_{t+1} \phi) \rightarrow (B_t \psi \rightarrow B_{t+1} \psi)$
Minimality (M):	$(I_{t+1} \phi \wedge B_{t+1} \psi) \rightarrow B_t(\phi \rightarrow \psi)$
Consistency of beliefs (D_B):	$B_t \phi \rightarrow \neg B_t \neg \phi$
Consistency of information (C_I):	$I_t \phi \rightarrow \neg A \neg \phi$.

The Acceptance axiom says that if the agent is informed of ϕ and ϕ is a consistent formula ($\neg A \neg \phi$), then he believes ϕ . It is clear that Qualified Acceptance can be derived from Acceptance (by axiom *Incl*, $\neg B_t \neg \phi \rightarrow \neg A \neg \phi$) and therefore \mathcal{L}_{AGM} is an extension of \mathcal{L}_{QBR} .

It is well-known that axiom D_B is characterized by seriality of \mathcal{B}_t : $\forall t \in \mathbb{N}$, $\forall \omega \in \Omega$, $\mathcal{B}_t(\omega) \neq \emptyset$.⁹ Given the non-standard validation rule for the I_t operator, the corresponding axiom for I_t , namely $I_t \phi \rightarrow \neg I_t \neg \phi$, is *not* characterized by seriality of \mathcal{I}_t .¹⁰ Seriality of \mathcal{I}_t is characterized by axiom C_I .¹¹

Remark 5 *By axioms C_I and A, the following strengthening of axiom A is a theorem of \mathcal{L}_{AGM} : $I_t \phi \rightarrow B_t \phi$.*

Definition 6 *An AGM frame is a frame $\langle \Omega, \{\mathcal{B}_t\}_{t \in \mathbb{N}}, \{\mathcal{I}_t\}_{t \in \mathbb{N}} \rangle$ that satisfies the following properties: $\forall t \in \mathbb{N}$, $\forall \omega \in \Omega$,*

- (1) $\mathcal{B}_t(\omega) \subseteq \mathcal{I}_t(\omega)$
- (2) if $\mathcal{B}_t(\omega) \cap \mathcal{I}_{t+1}(\omega) \neq \emptyset$ then $\mathcal{B}_{t+1}(\omega) \subseteq \mathcal{B}_t(\omega)$
- (3) $\mathcal{B}_t(\omega) \cap \mathcal{I}_{t+1}(\omega) \subseteq \mathcal{B}_{t+1}(\omega)$
- (4) $\mathcal{B}_t(\omega) \neq \emptyset$.

For example, the frame illustrated in Figure 3 of Section 2 is an AGM frame.

It is easy to show (see Lemma 21 in the Appendix) that the set of AGM frames is a subset of the set of frames that satisfy the Qualitative Bayes Rule. The following proposition is proved in the Appendix.

⁹A formula ϕ is characterized by a property P of frames if (1) ϕ is valid in every frame that satisfies P , and (2) if ϕ is valid in a frame then the frame satisfies P .

¹⁰While it is true that if \mathcal{I}_t is serial ($\forall t \in \mathbb{N}$, $\forall \omega \in \Omega$, $\mathcal{I}_t(\omega) \neq \emptyset$) then the formula $I_t \phi \rightarrow \neg I_t \neg \phi$ is valid, the converse is not true. To see this, fix a frame where, $\forall t \in \mathbb{N}$, $\forall \omega \in \Omega$, $\mathcal{I}_t(\omega) = \emptyset$. Fix an arbitrary model based on this frame and an arbitrary formula ϕ . If $\|\phi\| \neq \emptyset$, then $\omega \not\models I_t \phi$ and therefore $\omega \models I_t \phi \rightarrow \psi$ for every formula ψ (in particular for $\psi = \neg I_t \neg \phi$). On the other hand, if $\|\phi\| = \emptyset$, then $\omega \models I_t \phi$ and, since $\|\neg \phi\| = \Omega$, $\omega \not\models I_t \neg \phi$, that is, $\omega \models \neg I_t \neg \phi$. Thus $I_t \phi \rightarrow \neg I_t \neg \phi$ is valid, despite the fact that \mathcal{I}_t is not serial.

¹¹Proof. Fix a model based on a frame where \mathcal{I}_t is serial, a state α , a date t and a formula ϕ and suppose that $\alpha \models I_t \phi$. Then $\mathcal{I}_t(\alpha) = \|\phi\|$. By seriality of \mathcal{I}_t , there exists a β such that $\beta \in \mathcal{I}_t(\alpha)$. Then $\beta \models \phi$ and, therefore, $\alpha \models \neg A \neg \phi$. Conversely, fix a frame where \mathcal{I}_t is not serial. Then there exists a state α such that $\mathcal{I}_t(\alpha) = \emptyset$. For any model based on this frame and any atomic proposition p , $\|p \wedge \neg p\| = \emptyset$. Thus $\alpha \models I_t(p \wedge \neg p)$ and $\alpha \not\models \neg A \neg(p \wedge \neg p)$, thus axiom C_I is falsified at α with $\phi = (p \wedge \neg p)$.

Proposition 7 *Logic \mathfrak{L}_{AGM} is sound and complete with respect to the class of AGM frames.*

The next proposition establishes the relationship between logic \mathfrak{L}_{AGM} and the AGM theory of belief revision. First we recall the AGM postulates, which were developed within the framework of belief sets. Let Φ be the set of formulas in a propositional language. Given a subset $S \subseteq \Phi$, its PL-deductive closure $[S]^{PL}$ (where ‘PL’ stands for ‘Propositional Logic’) is defined as follows: $\psi \in [S]^{PL}$ if and only if there exist $\phi_1, \dots, \phi_n \in S$ such that $(\phi_1 \wedge \dots \wedge \phi_n) \rightarrow \psi$ is a truth-functional tautology (that is, a theorem of Propositional Logic). A *belief set* is a set $K \subseteq \Phi$ such that $K = [K]^{PL}$. A belief set K is consistent if $K \neq \Phi$ (equivalently, if there is no formula ϕ such that both ϕ and $\neg\phi$ belong to K). Given a consistent belief set K (thought of as the initial beliefs of the individual) and a formula ϕ (thought of as a new piece of information), the *revision of K by ϕ* , denoted by K_ϕ^* , is a subset of Φ that satisfies the following conditions, known as the AGM postulates:

- (K*1) K_ϕ^* is a belief set
- (K*2) $\phi \in K_\phi^*$
- (K*3) $K_\phi^* \subseteq [K \cup \{\phi\}]^{PL}$
- (K*4) if $\neg\phi \notin K$, then $[K \cup \{\phi\}]^{PL} \subseteq K_\phi^*$
- (K*5) $K_\phi^* = \Phi$ if and only if ϕ is a contradiction
- (K*6) if $\phi \leftrightarrow \psi$ is a tautology then $K_\phi^* = K_\psi^*$
- (K*7) $K_{\phi \wedge \psi}^* \subseteq [K_\phi^* \cup \{\psi\}]^{PL}$
- (K*8) if $\neg\psi \notin K_\phi^*$, then $[K_\phi^* \cup \{\psi\}]^{PL} \subseteq K_{\phi \wedge \psi}^*$

The set of postulates (K*1) through (K*6) is called the *basic set* of postulates for belief revision (Gärdenfors, [10] p. 55). The next proposition, proved in the Appendix, shows that every model of logic \mathfrak{L}_{AGM} satisfies the basic set of AGM postulates.

Proposition 8 *Fix an arbitrary model that validates the axioms of logic \mathfrak{L}_{AGM} (by Proposition 7 any model based on an AGM frame is such a model). Fix an arbitrary state ω , an arbitrary date t and let $K = \{\psi : \omega \models B_t\psi\}$. Suppose that there is a formula ϕ such that $\omega \models I_{t+1}\phi$ and define $K_\phi^* = \{\psi : \omega \models B_{t+1}\psi\}$. Then K_ϕ^* satisfies postulates (K*1) to (K*6).*

AGM postulates (K*7) and (K*8), require that the revision of K that includes both information ϕ and information ψ (that is, $K_{\phi \wedge \psi}^*$) ought to be the same as the expansion of K_ϕ^* by ψ , so long as ψ does not contradict the beliefs in K_ϕ^* . In our framework we are able to model, at every date and state, only the information that is actually received by the individual and cannot capture the counterfactual of how the individual would have modified his beliefs if he had received a different piece of information. Thus we cannot compare the revised beliefs that the individual holds after first receiving information ϕ and

subsequently information ψ with the beliefs he would have had if he had been simultaneously informed of both ϕ and ψ . However, by strengthening the logic further we can indirectly capture the full set of AGM postulates.

In the literature it has been shown that there is an equivalence between the AGM postulates and the notion of a plausibility ordering of the set of states (see Grove [12] and Board [4]). A *plausibility ordering* of Ω is a binary relation \succsim on Ω that is complete ($\forall \omega, \omega' \in \Omega$, either $\omega \succsim \omega'$ or $\omega' \succsim \omega$) and transitive (if $\omega \succsim \omega'$ and $\omega' \succsim \omega''$ then $\omega \succsim \omega''$). An initial belief set K can be thought of as a pair (α, E) in a model, with $\alpha \in E \subseteq \Omega$, by taking K to be the set of formulas that are true at every state $\omega \in E$. If ϕ represents new information, then the revised belief set K_ϕ^* is defined as the set of formulas that are true at the states that (1) satisfy ϕ and (2) are closest to α according to the plausibility relation \succsim . The next proposition shows that when information becomes more refined over time the belief revision history of the agent can indeed be rationalized by a plausibility ordering over the set of states.

Given a plausibility ordering \succsim of Ω and a subset $X \subseteq \Omega$, we denote by $\min(\succsim, X)$ the set $\{\omega \in X : \omega \succsim \omega' \text{ for all } \omega' \in X\}$.

Definition 9 A frame $\langle \Omega, \{\mathcal{B}_t\}_{t \in \mathbb{N}}, \{\mathcal{I}_t\}_{t \in \mathbb{N}} \rangle$ satisfies the property of information refinement over time if, for every $t \in \mathbb{N}$ and $\omega \in \Omega$, $\mathcal{I}_{t+1}(\omega) \subseteq \mathcal{I}_t(\omega)$.

For example, the frame illustrated in Figure 3 of Section 2 is such a frame.

Definition 10 Given a frame $\langle \Omega, \{\mathcal{B}_t\}_{t \in \mathbb{N}}, \{\mathcal{I}_t\}_{t \in \mathbb{N}} \rangle$ and a state $\omega \in \Omega$, we call the sequence $\{\mathcal{B}_t(\omega), \mathcal{I}_t(\omega)\}_{t \in \mathbb{N}}$ a belief revision history. The belief revision history $\{\mathcal{B}_t(\omega), \mathcal{I}_t(\omega)\}_{t \in \mathbb{N}}$ is consistent, successful and refined if it satisfies the following properties: $\forall t \in \mathbb{N}$, (1) $\mathcal{B}_t(\omega) \neq \emptyset$, (2) $\mathcal{B}_t(\omega) \subseteq \mathcal{I}_t(\omega)$ and (3) $\mathcal{I}_{t+1}(\omega) \subseteq \mathcal{I}_t(\omega)$.

Definition 11 A plausibility relation \succsim on Ω rationalizes the belief revision history $\{\mathcal{B}_t(\omega), \mathcal{I}_t(\omega)\}_{t \in \mathbb{N}}$ if, for every $t \in \mathbb{N}$, $\mathcal{B}_t(\omega) = \min(\succsim, \mathcal{I}_t(\omega))$.

The following proposition is proved in the Appendix.

Proposition 12 Let $\mathcal{F} = \langle \Omega, \{\mathcal{B}_t\}_{t \in \mathbb{N}}, \{\mathcal{I}_t\}_{t \in \mathbb{N}} \rangle$ be a frame such that, for every $\omega \in \Omega$, the corresponding belief revision history $\{\mathcal{B}_t(\omega), \mathcal{I}_t(\omega)\}_{t \in \mathbb{N}}$ is consistent, successful and refined (see Definition 10). Then, for every $\omega \in \Omega$, there exists a plausibility relation \succsim on Ω that rationalizes the belief history $\{\mathcal{B}_t(\omega), \mathcal{I}_t(\omega)\}_{t \in \mathbb{N}}$ if and only if \mathcal{F} is an AGM frame.

For example, the frame illustrated in Figure 3 of Section 2 satisfies the hypotheses of Proposition 12. Indeed every belief revision history $\{\mathcal{B}_t(\omega), \mathcal{I}_t(\omega)\}_{t \in \{0,1,2,3\}}$ with $\omega \in \{B, V, M, F\}$ is rationalized by the following plausibility ordering: $\succsim = \{(B, B), (B, V), (B, M), (B, F), (V, B), (V, V), (V, M), (V, F), (M, M), (M, F), (F, F)\}$.

The class of frames considered in Proposition 12 can be axiomatized as follows. Let \mathfrak{L}_{RP} (where ‘RP’ stands for ‘Rationalizable by a Plausibility relation’) be the logic obtained from \mathfrak{L}_{AGM} by adding the following axiom:

$$R. \quad (I_t\phi \wedge I_{t+1}\psi \wedge \neg A\neg\phi \wedge \neg A\neg\psi) \rightarrow A(\psi \rightarrow \phi)$$

The following proposition is proved in the Appendix.

Proposition 13 *Logic \mathfrak{L}_{RP} is sound and complete with respect to the class of AGM frames that satisfy information refinement over time (see Definition 9).*

By Proposition 12, in this class of frames every belief history can be rationalized by a plausibility ordering.

The hypothesis of information refinement ($\mathcal{I}_{t+1}(\omega) \subseteq \mathcal{I}_t(\omega)$) is crucial for Proposition 12, as the example illustrated in Figure 4 shows. The frame of Figure 4 is not an AGM frame, since it violates the Qualitative Bayes Rule: $\mathcal{B}_1(\alpha) \cap \mathcal{I}_2(\alpha) = \{\beta\} \neq \emptyset$ and yet $\mathcal{B}_2(\alpha) = \{\beta, \gamma\}$. However the belief revision history $\{\mathcal{B}_t(\alpha), \mathcal{I}_t(\alpha)\}_{t \in \{0,1,2\}}$ is rationalized by the following plausibility relation $\succsim = \{(\alpha, \alpha), (\beta, \alpha), (\beta, \beta), (\beta, \gamma), (\gamma, \alpha), (\gamma, \beta), (\gamma, \gamma)\}$.




Figure 4

A non-AGM frame which is rationalizable

4 Conclusion

The notions of static belief and of belief revision have been studied extensively in the literature. However, there is a surprising lack of uniformity in the two

approaches. In the philosophy and logic literature, starting with Hintikka's [13] seminal contribution, the notion of *static* belief has been studied mainly within the context of modal logic. The study of *belief revision*, on the other hand, has mainly followed the AGM approach where beliefs are modeled as sets of formulas in a given syntactic language and the issue is how a belief set ought to be modified when new information, represented by a formula ϕ , becomes available. A yet different approach can be found in the economics and game theory literature, where it is standard to represent beliefs by means of a probability measure over a set of states and belief revision is modeled using Bayes' rule. With a few exceptions, the tools of modal logic have not been explicitly employed in the analysis of the interaction of belief and information over time. In this paper we have proposed a unifying framework for static beliefs and belief revision by bringing belief revision under the umbrella of modal logic. For a detailed discussion of the relationship between our approach and the existing literature the reader is referred to Bonanno [6].

A Appendix

Outline of the proof of Proposition 3. Soundness is easily proved. The completeness proof is first carried out with respect to the class of augmented frames, which are defined as follows.

Definition 14 *An augmented frame is a collection $\langle \Omega, \{\mathcal{B}_t\}_{t \in \mathbb{N}}, \{\mathcal{I}_t\}_{t \in \mathbb{N}}, \mathcal{A} \rangle$ obtained by adding an equivalence relation \mathcal{A} to a regular frame $\langle \Omega, \{\mathcal{B}_t\}_{t \in \mathbb{N}}, \{\mathcal{I}_t\}_{t \in \mathbb{N}} \rangle$ with the additional requirement that $\mathcal{B}_t \subseteq \mathcal{A}$ for every $t \in \mathbb{N}$.*

Let \mathbb{M} be the set of maximally consistent sets (MCS) of formulas of \mathcal{L} . Define the following binary relations $\mathcal{A}, \mathcal{B}_t \subseteq \mathbb{M} \times \mathbb{M}$: $\alpha \mathcal{A} \beta$ if and only if $\{\phi : A\phi \in \alpha\} \subseteq \beta$ and $\alpha \mathcal{B}_t \beta$ if and only if $\{\phi : B_t \phi \in \alpha\} \subseteq \beta$. The relation \mathcal{A} is an equivalence relation because of axioms T_A and $\bar{5}_A$ and, for every $t \in \mathbb{N}$, \mathcal{B}_t is a subrelation of \mathcal{A} because of axiom *Incl*. Furthermore, the following lemma is a consequence of axioms I_1 and I_2 .

Lemma 15 *Let $\alpha, \beta \in \mathbb{M}$ be such that $\alpha \mathcal{A} \beta$ and let ϕ be a formula such that $I\phi \in \alpha$ and $\phi \in \beta$. Then, for every formula ψ , if $I\psi \in \alpha$ then $\psi \in \beta$, that is, $\{\psi : I\psi \in \alpha\} \subseteq \beta$.*

The definition of the relation \mathcal{I}_t is more complicated, because of the non-standard validation rule for the operator I_t . Let ω_0 be an arbitrary object such that $\omega_0 \notin \mathbb{M}$, that is, ω_0 can be anything but a MCS. Define the relation \mathcal{I}_t on $\mathbb{M} \cup \{\omega_0\}$ as follows: $\alpha \mathcal{I}_t \beta$ if and only if (1) $\alpha \in \mathbb{M}$ and (2) *either* $\beta \in \mathbb{M}$ and, for some formula ϕ , $I_t \phi \in \alpha$ and $\phi \in \beta$, *or* for all ϕ , $I_t \phi \notin \alpha$, and $\beta = \omega_0$.

The structure $\langle \mathbb{M} \cup \{\omega_0\}, \{\mathcal{B}_t\}_{t \in \mathbb{N}}, \{\mathcal{I}_t\}_{t \in \mathbb{N}}, \mathcal{A} \rangle$ so defined is an augmented frame. For every $\alpha \in \mathbb{M}$, let $\mathcal{A}(\alpha) = \{\omega \in \mathbb{M} : \alpha \mathcal{A} \omega\}$. Consider the canonical model based on this frame defined by $\|p\| = \{\omega \in \mathbb{M} : p \in \omega\}$, for every atomic proposition p . For every formula ϕ define $\|\phi\|$ according to the semantic rules

given in Section 3, with the following modified truth conditions for the operators I_t and A : $\alpha \models I_t\phi$ if and only if $\mathcal{I}_t(\alpha) = \|\phi\| \cap \mathcal{A}(\alpha)$ and $\alpha \models A\phi$ if and only if $\mathcal{A}(\alpha) \subseteq \|\phi\|$.

The crucial step in the completeness proof is the following Truth Lemma (for a proof see Goranko and Passy [11] and Bonanno [6]).

Lemma 16 (*Truth Lemma*). *For every $\omega \in \mathbb{M}$ and for every formula ϕ , $\omega \models \phi$ if and only if $\phi \in \omega$.*

Note, in particular, that $\alpha \models I_t\phi$ can hold in two different cases: either (1) $I_t\phi \in \alpha$ and there exists a $\beta \in \mathbb{M}$ such that $\phi \in \beta$ (in which case, by Lemmas 15 and 16, $\mathcal{I}_t(\alpha) = \|\phi\| \cap \mathcal{A}(\alpha) \subseteq \mathbb{M}$) or (2) $I_t\phi \in \alpha$ and for every $\omega \in \mathbb{M}$, $\phi \notin \omega$ (thus, by Lemma 16, $\|\phi\| = \emptyset$), which implies that $\neg\phi$ is a theorem of \mathcal{L} . Note that, by definition of \mathcal{I}_t , in the latter case, $\mathcal{I}_t(\alpha) = \emptyset$. (A third possibility is $I_t\phi \notin \alpha$ for every formula ϕ , in which case, by definition of \mathcal{I}_t , $\mathcal{I}_t(\alpha) = \{\omega_0\}$ with $\omega_0 \notin \mathbb{M}$.)

With the aid of Lemma 16 it can be shown that logic \mathcal{L} is complete with respect to the class of augmented frames $\langle \Omega, \{\mathcal{B}_t\}_{t \in \mathbb{N}}, \{\mathcal{I}_t\}_{t \in \mathbb{N}}, \mathcal{A} \rangle$. To complete the proof of Proposition 3, namely that logic \mathcal{L} is sound and complete with respect to the class of frames $\langle \Omega, \{\mathcal{B}_t\}_{t \in \mathbb{N}}, \{\mathcal{I}_t\}_{t \in \mathbb{N}}, \rangle$, we only need to invoke the result (Chellas, 1984, Theorem 3.12, p. 97) that completeness with respect to the class of augmented frames (where \mathcal{A} is an equivalence relation) implies completeness with respect to the generated sub-frames (where \mathcal{A} is the universal relation). The latter are precisely what we called frames. In a frame where the relation \mathcal{A} is the universal relation the semantic rule $\alpha \models I_t\phi$ if and only if $\mathcal{I}_t(\alpha) = \|\phi\| \cap \mathcal{A}(\alpha)$ becomes $\alpha \models I_t\phi$ if and only if $\mathcal{I}_t(\alpha) = \|\phi\|$ and the semantic rule $\alpha \models A\phi$ if and only if $\mathcal{A}(\alpha) \subseteq \|\phi\|$ becomes $\alpha \models A\phi$ if and only if $\|\phi\| = \Omega$, since $\mathcal{A}(\alpha) = \Omega$. ■

Proof of Proposition 4. (A) Validity. Fix a frame $\langle \Omega, \{\mathcal{B}_t\}_{t \in \mathbb{N}}, \{\mathcal{I}_t\}_{t \in \mathbb{N}} \rangle$ that satisfies QBR, that is, $\forall \omega \in \Omega, \forall t \in \mathbb{N}$, if $\mathcal{B}_t(\omega) \cap \mathcal{I}_{t+1}(\omega) \neq \emptyset$ then $\mathcal{B}_{t+1}(\omega) = \mathcal{B}_t(\omega) \cap \mathcal{I}_{t+1}(\omega)$. By Proposition 3 it is enough to show that the three axioms QA , P and M are valid in it. Fix an arbitrary model based on this frame and arbitrary state ω , date t and formulas ϕ and ψ . First we show that $\omega \models (\neg B_t \neg \phi \wedge I_{t+1} \phi) \rightarrow B_{t+1} \phi$. Suppose that $\omega \models \neg B_t \neg \phi \wedge I_{t+1} \phi$. Then $\mathcal{I}_{t+1}(\omega) = \|\phi\|$ and $\mathcal{B}_t(\omega) \cap \mathcal{I}_{t+1}(\omega) \neq \emptyset$. By QBR, $\mathcal{B}_{t+1}(\omega) \subseteq \mathcal{I}_{t+1}(\omega)$ and therefore $\omega \models B_{t+1} \phi$. Thus axiom QA is valid. Next we show that $\omega \models (\neg B_t \neg \phi \wedge I_{t+1} \phi) \rightarrow (B_t \psi \rightarrow B_{t+1} \psi)$. Suppose that $\omega \models \neg B_t \neg \phi \wedge I_{t+1} \phi \wedge B_t \psi$. Then $\mathcal{I}_{t+1}(\omega) = \|\phi\|$, $\mathcal{B}_t(\omega) \cap \mathcal{I}_{t+1}(\omega) \neq \emptyset$ and $\mathcal{B}_t(\omega) \subseteq \|\psi\|$. By QBR $\mathcal{B}_{t+1}(\omega) \subseteq \mathcal{B}_t(\omega)$ and therefore $\omega \models B_{t+1} \psi$. Thus axiom P is valid. Finally we show that $\omega \models (I_{t+1} \phi \wedge B_{t+1} \psi) \rightarrow B_t(\phi \rightarrow \psi)$. Suppose that $\omega \models I_{t+1} \phi \wedge B_{t+1} \psi$. Then $\mathcal{I}_{t+1}(\omega) = \|\phi\|$ and $\mathcal{B}_{t+1}(\omega) \subseteq \|\psi\|$. Fix an arbitrary $\omega' \in \mathcal{B}_t(\omega)$. If $\omega' \models \neg \phi$, then $\omega' \models \phi \rightarrow \psi$; if $\omega' \models \phi$, then $\omega' \in \mathcal{B}_t(\omega) \cap \mathcal{I}_{t+1}(\omega)$ and by QBR $\mathcal{B}_t(\omega) \cap \mathcal{I}_{t+1}(\omega) \subseteq \mathcal{B}_{t+1}(\omega)$, so that $\omega' \models \psi$ and therefore $\omega' \models \phi \rightarrow \psi$. Hence $\mathcal{B}_t(\omega) \subseteq \|\phi \rightarrow \psi\|$, that is, $\omega \models B_t(\phi \rightarrow \psi)$. Thus axiom M is valid.

(B) Completeness. Let \mathbb{M}_{QBR} be the set of maximally consistent sets (MCS) of formulas of \mathcal{L}_{QBR} . By Proposition 3 we only need to show that the frame

associated with the canonical model satisfies QBR. First we show that

$$\forall t \in \mathbb{N}, \forall \omega \in \mathbb{M}_{QBR}, \text{ if } \mathcal{B}_t(\omega) \cap \mathcal{I}_{t+1}(\omega) \neq \emptyset \text{ then } \mathcal{B}_{t+1}(\omega) \subseteq \mathcal{I}_{t+1}(\omega).$$

Fix an arbitrary $\alpha \in \mathbb{M}_{QBR}$ and suppose that $\mathcal{B}_t(\alpha) \cap \mathcal{I}_{t+1}(\alpha) \neq \emptyset$. Let $\beta \in \mathcal{B}_t(\alpha) \cap \mathcal{I}_{t+1}(\alpha)$. Since $\mathcal{B}_t(\alpha) \subseteq \mathbb{M}_{QBR}$, $\beta \in \mathbb{M}_{QBR}$ and therefore, by definition of \mathcal{I}_t , there exists a formula ϕ such that $I_{t+1}\phi \in \alpha$ and $\phi \in \beta$. Since $\beta \in \mathcal{B}_t(\alpha)$, $\neg B_t \neg \phi \in \alpha$ (see Chellas [7] Theorem 5.6, p. 172). Thus $(I_{t+1}\phi \wedge \neg B_t \neg \phi) \in \alpha$. Since Qualified Acceptance is a theorem, $(I_{t+1}\phi \wedge \neg B_t \neg \phi) \rightarrow B_{t+1}\phi \in \alpha$. Thus $B_{t+1}\phi \in \alpha$. We want to show that $\mathcal{B}_{t+1}(\alpha) \subseteq \mathcal{I}_{t+1}(\alpha)$. Fix an arbitrary $\gamma \in \mathcal{B}_{t+1}(\alpha)$. By definition of \mathcal{B}_{t+1} , $\{\psi : B_{t+1}\psi \in \alpha\} \subseteq \gamma$. In particular, since $B_{t+1}\phi \in \alpha$, $\phi \in \gamma$. By definition of \mathcal{I}_{t+1} , since $I_{t+1}\phi \in \alpha$ and $\phi \in \gamma$, $\gamma \in \mathcal{I}_{t+1}(\alpha)$.

Next we show that

$$\forall t \in \mathbb{N}, \forall \omega \in \mathbb{M}_{QBR}, \text{ if } \mathcal{B}_t(\omega) \cap \mathcal{I}_{t+1}(\omega) \neq \emptyset \text{ then } \mathcal{B}_{t+1}(\omega) \subseteq \mathcal{B}_t(\omega).$$

Fix an arbitrary $\alpha \in \mathbb{M}_{QBR}$ and suppose that $\mathcal{B}_t(\alpha) \cap \mathcal{I}_{t+1}(\alpha) \neq \emptyset$. Let $\beta \in \mathcal{B}_t(\alpha) \cap \mathcal{I}_{t+1}(\alpha)$. As shown above, there exists a ϕ such that $I_{t+1}\phi \in \alpha$, $\phi \in \beta$ and $\neg B_t \neg \phi \in \alpha$. By Persistence, for every formula ψ , $(I_{t+1}\phi \wedge \neg B_t \neg \phi) \rightarrow (B_t\psi \rightarrow B_{t+1}\psi) \in \alpha$. Thus

$$(B_t\psi \rightarrow B_{t+1}\psi) \in \alpha. \quad (1)$$

Fix an arbitrary $\gamma \in \mathcal{B}_{t+1}(\alpha)$. Then, by definition of \mathcal{B}_{t+1} , $\{\psi : B_{t+1}\psi \in \alpha\} \subseteq \gamma$. We want to show that $\gamma \in \mathcal{B}_t(\alpha)$, that is, that $\{\psi : B_t\psi \in \alpha\} \subseteq \gamma$. Let ψ be such that $B_t\psi \in \alpha$. By (1) $B_{t+1}\psi \in \alpha$ and therefore $\psi \in \gamma$.

Finally we show that

$$\forall t \in \mathbb{N}, \forall \omega \in \mathbb{M}_{QBR}, \mathcal{B}_t(\omega) \cap \mathcal{I}_{t+1}(\omega) \subseteq \mathcal{B}_{t+1}(\omega).$$

Fix arbitrary $\alpha, \beta \in \mathbb{M}_{QBR}$ such that $\beta \in \mathcal{B}_t(\alpha) \cap \mathcal{I}_{t+1}(\alpha)$. Then there exists a ϕ such that $I_{t+1}\phi \in \alpha$ and $\phi \in \beta$. Fix an arbitrary $\gamma \in \mathcal{B}_t(\alpha) \cap \mathcal{I}_{t+1}(\alpha)$. We want to show that $\gamma \in \mathcal{B}_{t+1}(\alpha)$, that is, that $\{\psi : B_{t+1}\psi \in \alpha\} \subseteq \gamma$. Let ψ be an arbitrary formula such that $B_{t+1}\psi \in \alpha$. Then $(I_{t+1}\phi \wedge B_{t+1}\psi) \in \alpha$. By Minimality, $(I_{t+1}\phi \wedge B_{t+1}\psi) \rightarrow B_t(\phi \rightarrow \psi) \in \alpha$. Thus $B_t(\phi \rightarrow \psi) \in \alpha$. Since $\gamma \in \mathcal{B}_t(\alpha)$, $(\phi \rightarrow \psi) \in \gamma$. Since $I_{t+1}\phi \in \alpha$ and $\gamma \in \mathcal{I}_t(\alpha)$, $\phi \in \gamma$. Hence $\psi \in \gamma$. ■

In order to prove Proposition 7 we first note that Definition 6 is equivalent to the following definition (indeed properties (1) and (4) of Definition 6 imply properties (1), (4) and (5) of the following definition, and *vice versa*).

Definition 17 *An AGM frame is a frame $\langle \Omega, \{\mathcal{B}_t\}_{t \in \mathbb{N}}, \{\mathcal{I}_t\}_{t \in \mathbb{N}} \rangle$ that satisfies the following properties: $\forall t \in \mathbb{N}, \forall \omega \in \Omega$,*

- (1) *if $\mathcal{I}_t(\omega) \neq \emptyset$ then $\mathcal{B}_t(\omega) \subseteq \mathcal{I}_t(\omega)$*
- (2) *if $\mathcal{B}_t(\omega) \cap \mathcal{I}_{t+1}(\omega) \neq \emptyset$ then $\mathcal{B}_{t+1}(\omega) \subseteq \mathcal{B}_t(\omega)$*

- (3) $\mathcal{B}_t(\omega) \cap \mathcal{I}_{t+1}(\omega) \subseteq \mathcal{B}_{t+1}(\omega)$
- (4) $\mathcal{B}_t(\omega) \neq \emptyset$
- (5) $\mathcal{I}_t(\omega) \neq \emptyset$.

Proof of Proposition 7. By Proposition 4, it is sufficient to show that axioms A , D_B and C_I are valid in every AGM frame and that the canonical model satisfies properties (1), (4) and (5) of Definition 6. The proof of validity is straightforward and is omitted. The proof that, since D_B is an axiom, the canonical model satisfies property (4) of Definition 6 is well-known (see Blackburn *et al* [3] p. 203). Let \mathbb{M}_{AGM} be the set of maximally consistent sets of logic \mathfrak{L}_{AGM} . In order to prove that property (1) holds we first start with the augmented canonical model $\langle \mathbb{M}_{AGM} \cup \{\omega_0\}, \{\mathcal{B}_t\}_{t \in \mathbb{N}}, \{\mathcal{I}_t\}_{t \in \mathbb{N}}, \mathcal{A} \rangle$ (where \mathcal{A} is an equivalence relation) and show that it satisfies the following property: $\forall \omega \in \mathbb{M}_{AGM}, \forall t \in \mathbb{N}$, if $\mathcal{I}_t(\omega) \cap \mathcal{A}(\omega) \neq \emptyset$ then $\mathcal{B}_t(\omega) \subseteq \mathcal{I}_t(\omega)$. Fix arbitrary $\alpha \in \mathbb{M}_{AGM}$ and $t \in \mathbb{N}$ and suppose that $\mathcal{I}_t(\alpha) \cap \mathcal{A}(\alpha) \neq \emptyset$. Let $\beta \in \mathcal{I}_t(\alpha) \cap \mathcal{A}(\alpha)$. Since $\beta \in \mathcal{A}(\alpha)$, $\beta \in \mathbb{M}_{AGM}$. Therefore, since $\beta \in \mathcal{I}_t(\alpha)$, by definition of \mathcal{I}_t there exists a ϕ such that $I_t\phi \in \alpha$ and $\phi \in \beta$. Thus $I_t\phi \wedge \neg A\neg\phi \in \alpha$. Since $(I_t\phi \wedge \neg A\neg\phi) \rightarrow B_t\phi$ is an axiom of \mathfrak{L}_{AGM} , $(I_t\phi \wedge \neg A\neg\phi) \rightarrow B_t\phi \in \alpha$. Thus $B_t\phi \in \alpha$. Fix an arbitrary $\gamma \in \mathcal{B}_t(\alpha)$. Then, by definition of \mathcal{B}_t , $\gamma \models \phi$ and, since \mathcal{B}_t is a subrelation of \mathcal{A} , $\gamma \in \mathcal{A}(\alpha)$. Thus $\gamma \in \|\phi\| \cap \mathcal{A}(\alpha)$ and therefore $\gamma \in \mathcal{I}_t(\alpha)$, since, by Lemmas 15 and 16, $\mathcal{I}_t(\alpha) = \|\phi\| \cap \mathcal{A}(\alpha)$. Taking the subframe generated by α , the above property becomes property (1) of Definition 17, since $\mathcal{I}_t(\omega) \cap \mathcal{A}(\omega) \neq \emptyset$ is equivalent to $\mathcal{I}_t(\omega) \neq \emptyset$ because \mathcal{A} is now the universal relation. Finally we show that property (5) of Definition 17 is satisfied in the canonical model. Fix arbitrary $\alpha \in \mathbb{M}_{AGM}$ and $t \in \mathbb{N}$. Suppose, by contradiction, that $\mathcal{I}_t(\alpha) = \emptyset$. Let p be an atomic proposition. Then, by lemma 16, $I_t(p \wedge \neg p) \in \alpha$, since $\|p \wedge \neg p\| = \emptyset$. Since $I_t\phi \rightarrow \neg A\neg\phi$ is an axiom of \mathfrak{L}_{AGM} , $I_t(p \wedge \neg p) \rightarrow \neg A\neg(p \wedge \neg p) \in \alpha$. Thus $\neg A\neg(p \wedge \neg p) \in \alpha$ implying that there exists a $\beta \in \mathbb{M}_{AGM}$ such that $(p \wedge \neg p) \in \beta$, contradicting the definition of maximally consistent set. ■

Proof of Proposition 8. (K*1). We need to show that K_ϕ^* is a belief set, that is, $K_\phi^* = [K_\phi^*]^{PL}$. Clearly, $K_\phi^* \subseteq [K_\phi^*]^{PL}$, since $\psi \rightarrow \psi$ is a tautology. Thus we only need to show that $[K_\phi^*]^{PL} \subseteq K_\phi^*$. Let $\psi \in [K_\phi^*]^{PL}$, i.e. there exist $\phi_1, \dots, \phi_n \in K_\phi^*$ such that $(\phi_1 \wedge \dots \wedge \phi_n) \rightarrow \psi$ is a tautology. Thus, by Necessitation of B_{t+1} (cf. Remark 2), $B_{t+1}((\phi_1 \wedge \dots \wedge \phi_n) \rightarrow \psi)$ is a theorem of \mathfrak{L}_{AGM} and therefore it is valid in the given model. Hence $\omega \models B_{t+1}((\phi_1 \wedge \dots \wedge \phi_n) \rightarrow \psi)$. By definition of K_ϕ^* , since $\phi_1, \dots, \phi_n \in K_\phi^*$, $\omega \models B_{t+1}(\phi_1 \wedge \dots \wedge \phi_n)$. By axiom K_B , $\omega \models B_{t+1}((\phi_1 \wedge \dots \wedge \phi_n) \rightarrow \psi) \wedge B_{t+1}(\phi_1 \wedge \dots \wedge \phi_n) \rightarrow B_{t+1}\psi$. Thus $\omega \models B_{t+1}\psi$, that is, $\psi \in K_\phi^*$.

(K*2). Since $I_{t+1}\phi \rightarrow B_{t+1}\phi$ is a theorem of \mathfrak{L}_{AGM} (see Remark 5), it is valid in the given model and, therefore, $\omega \models I_{t+1}\phi \rightarrow B_{t+1}\phi$. Thus, since $\omega \models I_{t+1}\phi$, $\omega \models B_{t+1}\phi$, that is, $\phi \in K_\phi^*$.

(K*3). Let $\psi \in K_\phi^*$, i.e. $\omega \models B_{t+1}\psi$. By Minimality, $\omega \models (I_{t+1}\phi \wedge B_{t+1}\psi) \rightarrow B_t(\phi \rightarrow \psi)$. By hypothesis, $\omega \models I_{t+1}\phi$. Thus $\omega \models B_t(\phi \rightarrow \psi)$, that is, $(\phi \rightarrow \psi) \in K$. Hence $\{\phi, (\phi \rightarrow \psi)\} \in K \cup \{\phi\}$ and, since $(\phi \wedge (\phi \rightarrow \psi)) \rightarrow \psi$ is a tautology, $\psi \in [K \cup \{\phi\}]^{PL}$.

(K*4). Suppose $\neg\phi \notin K$, that is, $\omega \models \neg B_t\neg\phi$. By Persistence, for every formula ψ , $\omega \models (I_{t+1}\phi \wedge \neg B_t\neg\phi) \rightarrow (B_t\psi \rightarrow B_{t+1}\psi)$. Thus, since by hypothesis $\omega \models I_{t+1}\phi$,

$$\omega \models (B_t\psi \rightarrow B_{t+1}\psi) \text{ for every formula } \psi. \quad (2)$$

Let $\chi \in [K \cup \{\phi\}]^{PL}$, that is, there exist $\phi_1, \dots, \phi_n \in K \cup \{\phi\}$ such that $(\phi_1 \wedge \dots \wedge \phi_n) \rightarrow \chi$ is a tautology. We want to show that $\chi \in K_\phi^*$, i.e. $\omega \models B_{t+1}\chi$. Since $(\phi_1 \wedge \dots \wedge \phi_n) \rightarrow \chi$ is a tautology, $\omega \models B_t((\phi_1 \wedge \dots \wedge \phi_n) \rightarrow \chi)$. If $\phi_1, \dots, \phi_n \in K$, then $\omega \models B_t(\phi_1 \wedge \dots \wedge \phi_n)$ and therefore $\omega \models B_t\chi$. Thus, by (2), $\omega \models B_{t+1}\chi$. If $\phi_1, \dots, \phi_n \notin K$, then, by renumbering the formulas if necessary, we can assume that $\phi_1 = \phi$ and $\phi_2, \dots, \phi_n \in K$. In this case we have $\omega \models B_t(\phi_2 \wedge \dots \wedge \phi_n)$ and $\omega \models B_t((\phi_2 \wedge \dots \wedge \phi_n) \rightarrow (\phi \rightarrow \chi))$ since $(\phi \wedge \phi_2 \wedge \dots \wedge \phi_n) \rightarrow \chi$ is a tautology and it is equivalent to $(\phi_2 \wedge \dots \wedge \phi_n) \rightarrow (\phi \rightarrow \chi)$. Thus $\omega \models B_t(\phi \rightarrow \chi)$. Hence, by (2) (with $\psi = (\phi \rightarrow \chi)$), $\omega \models B_{t+1}(\phi \rightarrow \chi)$. By Remark 5 $\omega \models I_{t+1}\phi \rightarrow B_{t+1}\phi$. From this and the hypothesis that $\omega \models I_{t+1}\phi$ it follows that $\omega \models B_{t+1}\phi$. By axiom K_B , $\omega \models B_{t+1}(\phi \rightarrow \chi) \wedge B_{t+1}\phi \rightarrow B_{t+1}\chi$. Thus $\omega \models B_{t+1}\chi$.

(K*5). We have to show that $K_\phi^* \neq \Phi$, unless ϕ is a contradiction (that is, $\neg\phi$ is a tautology). The possibility of contradictory information is ruled out by axiom C_I .¹² Thus we only need to show that $K_\phi^* \neq \Phi$. Let p be an atomic proposition. By Consistency of beliefs, $B_{t+1}p \rightarrow \neg B_{t+1}\neg p$; thus if $p \in K_\phi^*$ then $\neg p \notin K_\phi^*$ and therefore $K_\phi^* \neq \Phi$.

(K*6). We have to show that if $\phi \leftrightarrow \psi$ is a tautology then $K_\phi^* = K_\psi^*$. If $\phi \leftrightarrow \psi$ is a tautology, then $\|\phi \leftrightarrow \psi\| = \Omega$, that is, $\|\phi\| = \|\psi\|$. Thus, by axiom I_2 , $\omega \models I\phi$ if and only if $\omega \models I\psi$. Hence, by definition of K_ϕ^* , $K_\phi^* = K_\psi^*$.

In order to prove Proposition 12 we need some preliminary results.

Lemma 18 *Let \succsim be a complete and transitive binary relation on Ω and $X \subseteq Y \subseteq \Omega$. If $\min(\succsim, Y) \cap X \neq \emptyset$ then $\min(\succsim, X) = \min(\succsim, Y) \cap X$.*

Proof. First we show that $\min(\succsim, Y) \cap X \subseteq \min(\succsim, X)$. If $\min(\succsim, Y) \cap X = \emptyset$ there is nothing to prove. Therefore let $\beta \in \min(\succsim, Y) \cap X$. Then $\beta \in X$ and $\beta \succsim \gamma$ for all $\gamma \in Y$. Since $X \subseteq Y$, it follows that $\beta \in \min(\succsim, X)$. Next we show that if $\min(\succsim, Y) \cap X \neq \emptyset$ then $\min(\succsim, X) \subseteq \min(\succsim, Y) \cap X$. Let $\beta \in \min(\succsim, X)$. Fix an arbitrary $\gamma \in \min(\succsim, X)$. Then $\gamma \in X$ and $\gamma \succsim \beta$. Suppose that $\gamma \notin \min(\succsim, Y)$. Then there exists a $\delta \in Y$ such that $\delta \prec \gamma$ (that is, $\delta \succsim \gamma$ and $\gamma \not\succsim \delta$). By transitivity (since $\gamma \succsim \beta$), $\delta \prec \beta$, contradicting the fact that $\beta \in \min(\succsim, Y)$. ■

¹²Our logic \mathfrak{L}_{AGM} excludes the possibility of contradictory information. Indeed, as pointed out by Friedman and Halpern [8], it is not clear how one could be informed of a contradiction.

Lemma 19 Let $\langle \Omega, \{\mathcal{B}_t\}_{t \in \mathbb{N}}, \{\mathcal{I}_t\}_{t \in \mathbb{N}} \rangle$ be a QBR frame that satisfies information refinement over time (that is, $\forall t \in \mathbb{N}, \forall \omega \in \Omega, \mathcal{I}_{t+1}(\omega) \subseteq \mathcal{I}_t(\omega)$). Fix $\alpha \in \Omega$ and let $\{\mathcal{B}_t(\alpha), \mathcal{I}_t(\alpha)\}_{t \in \mathbb{N}}$ be the corresponding belief revision history. Let $T_0, T_1 \in \mathbb{N}$ be such that $T_0 < T_1$ and, $\forall t \in \mathbb{N}$ with $T_0 < t \leq T_1$, $\mathcal{B}_{T_0}(\alpha) \cap \mathcal{I}_t(\alpha) \neq \emptyset$. Then, $\forall t \in \mathbb{N}$ with $T_0 < t \leq T_1$, $\mathcal{B}_t(\alpha) = \mathcal{B}_{T_0}(\alpha) \cap \mathcal{I}_t(\alpha)$.

Proof. We prove it by induction. The statement is clearly true for $t = T_0 + 1$, since by hypothesis $\mathcal{B}_{T_0}(\alpha) \cap \mathcal{I}_{T_0+1}(\alpha) \neq \emptyset$ and thus, by QBR, $\mathcal{B}_{T_0+1}(\alpha) = \mathcal{B}_{T_0}(\alpha) \cap \mathcal{I}_{T_0+1}(\alpha)$. If $T_1 = T_0 + 1$ there is nothing else to prove. Suppose therefore that $T_1 > T_0 + 1$ and proceed with the induction step: suppose that the statement is true for every $t \in \mathbb{N}$ with $T_0 < t \leq T$ (with $T < T_1$). We want to show that it is true for $t = T + 1$. By the induction hypothesis, $\mathcal{B}_T(\alpha) = \mathcal{B}_{T_0}(\alpha) \cap \mathcal{I}_T(\alpha)$. Thus $\mathcal{B}_T(\alpha) \cap \mathcal{I}_{T+1}(\alpha) = \mathcal{B}_{T_0}(\alpha) \cap \mathcal{I}_T(\alpha) \cap \mathcal{I}_{T+1}(\alpha)$. Since $\mathcal{I}_{T+1}(\alpha) \subseteq \mathcal{I}_T(\alpha)$, $\mathcal{I}_{T+1}(\alpha) \cap \mathcal{I}_T(\alpha) = \mathcal{I}_{T+1}(\alpha)$. Thus

$$\mathcal{B}_T(\alpha) \cap \mathcal{I}_{T+1}(\alpha) = \mathcal{B}_{T_0}(\alpha) \cap \mathcal{I}_{T+1}(\alpha). \quad (3)$$

By hypothesis, $\mathcal{B}_{T_0}(\alpha) \cap \mathcal{I}_{T+1}(\alpha) \neq \emptyset$. Hence, by (3),

$$\mathcal{B}_T(\alpha) \cap \mathcal{I}_{T+1}(\alpha) \neq \emptyset. \quad (4)$$

It follows from QBR that

$$\mathcal{B}_{T+1}(\alpha) = \mathcal{B}_T(\alpha) \cap \mathcal{I}_{T+1}(\alpha). \quad (5)$$

From (3) and (5) we get that $\mathcal{B}_{T+1}(\alpha) = \mathcal{B}_{T_0}(\alpha) \cap \mathcal{I}_{T+1}(\alpha)$. ■

Lemma 20 Given a belief revision history $\mathcal{H} = \{\mathcal{B}_t(\alpha), \mathcal{I}_t(\alpha)\}_{t \in \mathbb{N}}$ which is consistent, successful and refined (see Definition 10) there exists a binary relation \succsim that rationalizes \mathcal{H} if and only if \mathcal{H} is qualitatively Bayesian, that is, if and only if, $\forall t \in \mathbb{N}$, if $\mathcal{B}_t(\alpha) \cap \mathcal{I}_{t+1}(\alpha) \neq \emptyset$ then $\mathcal{B}_{t+1}(\alpha) = \mathcal{B}_t(\alpha) \cap \mathcal{I}_{t+1}(\alpha)$.

Proof. First we prove that if \succsim rationalizes $\{\mathcal{B}_t(\alpha), \mathcal{I}_t(\alpha)\}_{t \in \mathbb{N}}$ then $\{\mathcal{B}_t(\alpha), \mathcal{I}_t(\alpha)\}_{t \in \mathbb{N}}$ is qualitatively Bayesian. Fix an arbitrary $t \in \mathbb{N}$ such that $\mathcal{B}_t(\alpha) \cap \mathcal{I}_{t+1}(\alpha) \neq \emptyset$. By hypothesis, $\mathcal{B}_t(\alpha) = \min(\succsim, \mathcal{I}_t(\alpha))$ and $\mathcal{B}_{t+1}(\alpha) = \min(\succsim, \mathcal{I}_{t+1}(\alpha))$. Since $\mathcal{I}_{t+1}(\alpha) \subseteq \mathcal{I}_t(\alpha)$ it follows from Lemma 18 (with $X = \mathcal{I}_{t+1}(\alpha)$ and $Y = \mathcal{I}_t(\alpha)$) that $\mathcal{B}_{t+1}(\alpha) = \mathcal{B}_t(\alpha) \cap \mathcal{I}_{t+1}(\alpha)$.

Next we prove that if $\{\mathcal{B}_t(\alpha), \mathcal{I}_t(\alpha)\}_{t \in \mathbb{N}}$ is qualitatively Bayesian then there exists a complete and transitive binary relation \succsim on Ω that rationalizes it. Define the function $rank : \Omega \rightarrow \mathbb{N}$ as follows:

$$\begin{aligned} rank(\omega) &= 0 && \text{if } \omega \in \mathcal{B}_0(\alpha) \\ &= \infty && \text{if } \omega \in \Omega \setminus \bigcup_{t \in \mathbb{N}} \mathcal{B}_t(\alpha) \\ &= t && \text{if } \omega \in \mathcal{B}_t(\alpha) \text{ and } \mathcal{B}_{t-1}(\alpha) \cap \mathcal{I}_t(\alpha) = \emptyset \end{aligned}$$

First we show that this function's domain is indeed Ω . Fix arbitrary $t' \in \mathbb{N} \setminus \{0\}$ and $\omega \in \mathcal{B}_{t'}(\alpha)$. If $\mathcal{B}_{t'-1}(\alpha) \cap \mathcal{I}_{t'}(\alpha) = \emptyset$ then $rank(\omega) = t'$; if $\mathcal{B}_{t'-1}(\alpha) \cap \mathcal{I}_{t'}(\alpha) \neq \emptyset$ let $T = \{t \in \mathbb{N} : t < t' \text{ and } \mathcal{B}_{t-1}(\alpha) \cap \mathcal{I}_t(\alpha) = \emptyset\}$. If $T = \emptyset$, then, by Lemma 19, $\omega \in \mathcal{B}_0(\alpha)$ and therefore $rank(\omega) = 0$. If $T \neq \emptyset$, then, by Lemma 19,

$rank(\omega) = \hat{t}$ where $\hat{t} = \max T$. Now define the binary relation \lesssim on Ω as follows: $\omega \lesssim \omega'$ if and only if $rank(\omega) \leq rank(\omega')$. Clearly \lesssim is complete and transitive. Now we show that \lesssim rationalizes $\{\mathcal{B}_t(\alpha), \mathcal{I}_t(\alpha)\}_{t \in \mathbb{N}}$. Fix an arbitrary t_0 . We want to show that $\mathcal{B}_{t_0}(\alpha) = \min(\lesssim, \mathcal{I}_{t_0}(\alpha))$, that is, for every $\omega \in \mathcal{B}_{t_0}(\alpha)$ and for every $\omega' \in \mathcal{I}_{t_0}(\alpha)$, $rank(\omega) \leq rank(\omega')$. If $t_0 = 0$ it follows from the fact that (1) $\mathcal{B}_0(\alpha) \subseteq \mathcal{I}_0(\alpha)$ and (2) by construction $rank(\omega) \geq 0$ for every $\omega \in \Omega$ and $rank(\omega) = 0$ if and only if $\omega \in \mathcal{B}_0(\alpha)$. Suppose therefore that $t_0 > 0$. We need to consider two cases.

CASE 1: for every $t \leq t_0$, $\mathcal{B}_{t-1}(\alpha) \cap \mathcal{I}_t(\alpha) \neq \emptyset$. Then, by Lemma 19, $\mathcal{B}_{t_0}(\alpha) = \mathcal{B}_0(\alpha) \cap \mathcal{I}_{t_0}(\alpha)$ and therefore $rank(\omega) = 0$ for every $\omega \in \mathcal{B}_{t_0}(\alpha)$.

CASE 2: for some $t \leq t_0$ $\mathcal{B}_{t-1}(\alpha) \cap \mathcal{I}_t(\alpha) = \emptyset$. Let $T = \{t \in \mathbb{N} : t \leq t_0 \text{ and } \mathcal{B}_{t-1}(\alpha) \cap \mathcal{I}_t(\alpha) = \emptyset\}$ and let $t' = \min T$. Then, by Lemma 19, for every $\omega \in \mathcal{B}_t(\alpha)$ with $t < t'$, $rank(\omega) = 0$, and, by construction, for every $\omega \in \Omega \setminus \mathcal{B}_0(\alpha)$, $rank(\omega) \geq t'$ and, for every $\omega \in \mathcal{B}_{t'}(\alpha)$, $rank(\omega) = t'$. Thus $\mathcal{B}_{t'}(\alpha) = \min(\lesssim, \mathcal{I}_{t'}(\alpha))$. If $t_0 = t'$ the proof is complete. Suppose therefore that $t_0 > t'$. If for every t with $t_0 < t \leq t'$, $\mathcal{B}_{t'}(\alpha) \cap \mathcal{I}_t(\alpha) \neq \emptyset$ then, by Lemma 19, $\mathcal{B}_{t_0}(\alpha) = \mathcal{B}_{t'}(\alpha) \cap \mathcal{I}_{t_0}(\alpha)$ and thus, for every $\omega \in \mathcal{B}_{t_0}(\alpha)$, $rank(\omega) = t'$ and hence $\mathcal{B}_{t_0}(\alpha) = \min(\lesssim, \mathcal{I}_{t_0}(\alpha))$. Otherwise repeat the argument: define $T = \{t \in \mathbb{N} : t \leq t_0, t > t' \text{ and } \mathcal{B}_{t-1}(\alpha) \cap \mathcal{I}_t(\alpha) = \emptyset\}$ and let $t'' = \min T'$ and show that $\mathcal{B}_{t''}(\alpha) = \min(\lesssim, \mathcal{I}_{t''}(\alpha))$ and $rank(\omega) \geq t''$ for every $\omega \in \mathcal{I}_t(\alpha)$ with $t \geq t''$. By a finite repetition of this argument and Lemma 19 it follows that $\mathcal{B}_{t_0}(\alpha) = \min(\lesssim, \mathcal{I}_{t_0}(\alpha))$. ■

Lemma 21 *Let $\mathcal{F} = \langle \Omega, \{\mathcal{B}_t\}_{t \in \mathbb{N}}, \{\mathcal{I}_t\}_{t \in \mathbb{N}} \rangle$ be a frame such that, for every $\omega \in \Omega$, the corresponding belief revision history $\{\mathcal{B}_t(\omega), \mathcal{I}_t(\omega)\}_{t \in \mathbb{N}}$ is consistent ($\mathcal{B}_t(\omega) \neq \emptyset$) and successful ($\mathcal{B}_t(\omega) \subseteq \mathcal{I}_t(\omega)$). Then \mathcal{F} is an AGM frame if and only if it satisfies the Qualitative Bayes Rule.*

Proof. Let \mathcal{F} be an AGM frame. Fix arbitrary $t \in \mathbb{N}$ and $\omega \in \Omega$ and suppose that $\mathcal{B}_t(\omega) \cap \mathcal{I}_{t+1}(\omega) \neq \emptyset$. Then, by property (2) of Definition 17, $\mathcal{B}_{t+1}(\omega) \subseteq \mathcal{B}_t(\omega)$. By property (1) (since $\mathcal{I}_{t+1}(\omega) \neq \emptyset$), $\mathcal{B}_{t+1}(\omega) \subseteq \mathcal{I}_{t+1}(\omega)$. Thus $\mathcal{B}_{t+1}(\omega) \subseteq \mathcal{B}_t(\omega) \cap \mathcal{I}_{t+1}(\omega)$. By property (3), $\mathcal{B}_t(\omega) \cap \mathcal{I}_{t+1}(\omega) \subseteq \mathcal{B}_{t+1}(\omega)$. Thus $\mathcal{B}_{t+1}(\omega) = \mathcal{B}_t(\omega) \cap \mathcal{I}_{t+1}(\omega)$. Conversely, let $\mathcal{F} = \langle \Omega, \{\mathcal{B}_t\}_{t \in \mathbb{N}}, \{\mathcal{I}_t\}_{t \in \mathbb{N}} \rangle$ be a frame such that, for every $\omega \in \Omega$, the corresponding belief revision history $\{\mathcal{B}_t(\omega), \mathcal{I}_t(\omega)\}_{t \in \mathbb{N}}$ is consistent (so that \mathcal{F} satisfies property (4) of Definition 17), successful (so that \mathcal{F} satisfies properties (1) and (5) of Definition 17) and satisfies the Qualitative Bayes Rule, that is, for every $t \in \mathbb{N}$, $\mathcal{B}_{t+1}(\omega) = \mathcal{B}_t(\omega) \cap \mathcal{I}_{t+1}(\omega)$. Then \mathcal{F} satisfies properties (2) and (3) of Definition 17 and therefore is an AGM frame. ■

Proof of Proposition 12. Let $\mathcal{F} = \langle \Omega, \{\mathcal{B}_t\}_{t \in \mathbb{N}}, \{\mathcal{I}_t\}_{t \in \mathbb{N}} \rangle$ be a frame such that, for every $\omega \in \Omega$, the corresponding belief revision history $\{\mathcal{B}_t(\omega), \mathcal{I}_t(\omega)\}_{t \in \mathbb{N}}$ is consistent, successful and refined. By Lemma 21, \mathcal{F} is an AGM frame if and only if it satisfies the Qualitative Bayes Rule and by Lemma 20 this is necessary and sufficient for every belief history to be rationalizable by a plausibility ordering on Ω . ■

Proof of Proposition 13. By Proposition 7 it is enough to show that axiom R is valid in every AGM frame that satisfies information refinement and that the canonical model satisfies information refinement.

(A) Validity. Fix a frame that satisfies the information refinement property: $\forall t \in \mathbb{N}, \forall \omega \in \Omega, \mathcal{I}_{t+1}(\omega) \subseteq \mathcal{I}_t(\omega)$. Fix an arbitrary model based on it, arbitrary $t \in \mathbb{N}, \alpha \in \Omega$ and formula ϕ and suppose that $\alpha \models I_t \phi \wedge I_{t+1} \psi \wedge \neg A \neg \phi \wedge \neg A \neg \psi$. Then $\mathcal{I}_t(\alpha) = \|\phi\|$ and $\mathcal{I}_{t+1}(\alpha) = \|\psi\|$. Thus, by the property, $\|\psi\| \subseteq \|\phi\|$ and therefore $\|\psi \rightarrow \phi\| = \Omega$. Hence $\omega \models A(\psi \rightarrow \phi)$.

(B) Completeness. Let \mathbb{M}_{RP} be the set of MCS of logic \mathfrak{L}_{RP} and let $\langle \mathbb{M}_{RP} \cup \{\omega_0\}, \{\mathcal{B}_t\}_{t \in \mathbb{N}}, \{\mathcal{I}_t\}_{t \in \mathbb{N}}, \mathcal{A} \rangle$ be the augmented canonical model. First we show that this augmented model satisfies the following property:

$$WIR_1 \quad \forall t \in \mathbb{N}, \forall \omega \in \mathbb{M}_{RP}, \text{ if } \mathcal{I}_t(\omega) \cap \mathcal{A}(\omega) \neq \emptyset \text{ and } \mathcal{I}_{t+1}(\omega) \cap \mathcal{A}(\omega) \neq \emptyset \\ \text{ then } \mathcal{I}_{t+1}(\omega) \cap \mathcal{A}(\omega) \subseteq \mathcal{I}_t(\omega) \cap \mathcal{A}(\omega).$$

Fix arbitrary $t \in \mathbb{N}$ and $\alpha \in \mathbb{M}_{RP}$ and suppose that $\mathcal{I}_t(\alpha) \cap \mathcal{A}(\alpha) \neq \emptyset$ and $\mathcal{I}_{t+1}(\alpha) \cap \mathcal{A}(\alpha) \neq \emptyset$. Let $\beta \in \mathcal{I}_t(\alpha) \cap \mathcal{A}(\alpha)$ and $\gamma \in \mathcal{I}_{t+1}(\alpha) \cap \mathcal{A}(\alpha)$. Then, by definition of \mathcal{I}_t , there exists a formula ϕ such that $I_t \phi \in \alpha$ and $\phi \in \beta$. Thus, since $\beta \in \mathcal{A}(\alpha)$, $I_t \phi \wedge \neg A \neg \phi \in \alpha$. Similarly, there exists a ψ such that $I_{t+1} \psi \in \alpha$ and $\psi \in \gamma$, so that $I_{t+1} \psi \wedge \neg A \neg \psi \in \alpha$. Since R is an axiom, $(I_t \phi \wedge I_{t+1} \psi \wedge \neg A \neg \phi \wedge \neg A \neg \psi) \rightarrow A(\psi \rightarrow \phi) \in \alpha$. Hence $A(\psi \rightarrow \phi) \in \alpha$. Fix an arbitrary $\delta \in \mathcal{I}_{t+1}(\alpha) \cap \mathcal{A}(\alpha)$ (this set is non-empty, because γ belongs to it). Since $\delta \in \mathcal{A}(\alpha)$ and $A(\psi \rightarrow \phi) \in \alpha$, $(\psi \rightarrow \phi) \in \delta$. We want to show that $\delta \in \mathcal{I}_t(\alpha)$. Since $I_{t+1} \psi \in \alpha$, $\mathcal{I}_{t+1}(\alpha) = \|\psi\| \cap \mathcal{A}(\alpha)$, so that $\psi \in \delta$. From this and the fact that $(\psi \rightarrow \phi) \in \delta$, it follows that $\phi \in \delta$ and therefore, since $\mathcal{I}_t(\alpha) = \|\phi\| \cap \mathcal{A}(\alpha)$, $\delta \in \mathcal{I}_t(\alpha)$. Taking the sub-frame generated by α , the above property WIR_1 becomes (letting $\Omega = \mathcal{A}(\alpha)$)

$$WIR_2 \quad \forall t \in \mathbb{N}, \forall \omega \in \Omega \text{ if } \mathcal{I}_t(\omega) \neq \emptyset \text{ and } \mathcal{I}_{t+1}(\omega) \neq \emptyset, \text{ then } \mathcal{I}_{t+1}(\omega) \subseteq \mathcal{I}_t(\omega)$$

since $\mathcal{I}_t(\omega) \cap \mathcal{A}(\omega) \neq \emptyset$ is equivalent to $\mathcal{I}_t(\omega) \neq \emptyset$ because \mathcal{A} is now the universal relation. Together with the fact that in the canonical model for the logic \mathfrak{L}_{AGM} (which is contained in the logic \mathfrak{L}_{RP}) the relation \mathcal{I}_t is serial ($\forall t \in \mathbb{N}, \forall \omega \in \Omega, \mathcal{I}_t(\omega) \neq \emptyset$) property WIR_2 becomes the information refinement property: $\forall t \in \mathbb{N}, \forall \omega \in \Omega, \mathcal{I}_{t+1}(\omega) \subseteq \mathcal{I}_t(\omega)$. ■

References

- [1] Alchourron, C., P. Gärdenfors and D. Makinson, On the logic of theory change: partial meet contraction and revision functions, *The Journal of Symbolic Logic*, 1985, 50: 510-530.
- [2] Battigalli, P., Strategic independence and perfect Bayesian equilibria, *Journal of Economic Theory*, 1996, 70: 201-234.
- [3] Blackburn, P., M. de Rijke and Y. Venema, *Modal logic*, Cambridge University Press, 2001.
- [4] Board, O., Dynamic interactive epistemology, *Games and Economic Behavior*, 2004, 49: 49-80.

- [5] Bonanno, G., Rational belief equilibria, *Economic Notes*, 1993, 22: 430-463.
- [6] Bonanno, G., A simple modal logic for belief revision, *Knowledge, Rationality and Action*, 2005, 147(2): 193-228.
- [7] Chellas, B., *Modal logic: an introduction*, Cambridge University Press, 1984.
- [8] Friedman, N. and J. Halpern, Belief revision: a critique, *Journal of Logic, Language, and Information*, 1999, 8: 401-420.
- [9] Fudenberg, D. and J. Tirole, Perfect Bayesian equilibrium and sequential equilibrium, *Journal of Economic Theory*, 1991, 53: 236-260.
- [10] Gärdenfors, P., *Knowledge in flux: modeling the dynamics of epistemic states*, MIT Press, 1988.
- [11] Goranko, V. and S. Passy, Using the universal modality: gains and questions, *Journal of Logic and Computation*, 1992, 2: 5-30.
- [12] Grove, A., Two modellings for theory change, *Journal of Philosophical logic*, 1988, 17: 157-170.
- [13] Hintikka, J., *Knowledge and belief*, Cornell University Press, 1962.
- [14] Katsuno, H. and Mendelzon, A. O., On the difference between updating a knowledge and revising it, in: P. Gärdenfors, ed., *Belief Revision*, Volume 29 in Cambridge Tracts in Theoretical Computer Science, 1992, Cambridge University Press, pp. 83-203.
- [15] Kreps, D. and R. Wilson, Sequential equilibria, *Econometrica*, 50: 863-894.
- [16] Kripke, S., A semantical analysis of modal logic I: normal propositional calculi, *Zeitschrift für Mathematische Logik und Grundlagen der Mathematik*, 1963, 9: 67-96.
- [17] Levesque, H. J., All I know: a study in autoepistemic logic, *Artificial Intelligence*, 1990, 5: 263-309.
- [18] Nayak, A., M. Pagnucco and P. Peppas, Dynamic belief revision operators, *Artificial Intelligence*, 2003, 146: 193-228