

Belief change in branching time: AGM-consistency and iterated revision

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

January 2010. Revised September 2010.

Abstract

We study belief change in the branching-time structures introduced in [8]. First, we identify a property of branching-time frames that is equivalent (when the set of states is finite) to AGM-consistency, which is defined as follows. A frame is AGM-consistent if the partial belief revision function associated with an arbitrary state-instant pair and an arbitrary model based on that frame can be extended to a full belief revision function that satisfies the AGM postulates. Second, we provide a set of modal axioms that characterize the class of AGM-consistent frames within the modal logic introduced in [8]. Third, we introduce a generalization of AGM belief revision functions that allows a clear statement of principles of iterated belief revision and discuss iterated revision both semantically and syntactically.

Keywords: branching time, belief revision, information, iterated belief revision, plausibility ordering.

1 Introduction

The seminal contributions of Hintikka [22] and Alchourrón, Gärdenfors and Makinson (AGM) [1] have given rise to two separate strands in the literature, one dealing

*I am grateful to three anonymous reviewers for helpful and constructive comments. A first draft of this paper was presented at the Workshop on *Reasoning about knowledge and rational action*, University of Liverpool, March 2009 and the Workshop on *Information processing, rational belief change and social interaction*, Dagstuhl, August 2009.

with static belief and the other with belief revision. The analysis of static belief is carried out within the framework of modal logic and Kripke [26] structures, while in the AGM approach beliefs are represented as sets of formulas and belief revision as a function that associates with a set of formulas K (thought of as the initial beliefs) and formula ϕ (thought of as new information) a new belief set K_ϕ^* representing the revised beliefs. In [8] a new framework was proposed aimed at bringing the AGM theory of belief revision within the scope of modal logic by representing the AGM postulates as axioms in a modal language. Since belief revision deals with the interaction of belief and information over time, the proposed framework was based on branching-time temporal logic. Besides the next-time temporal operator, the logic contained a modal operator for belief and a modal operator for information. In this paper we further develop the analysis of [8] by establishing a stronger correspondence between the proposed logic and the AGM theory. In particular, we characterize the conditions under which the “partial” belief revision function induced at a state-time pair (by an arbitrary interpretation of a branching-time belief revision frame) can be extended to a full belief revision function satisfying the AGM postulates. This is done both semantically (Proposition 6) and syntactically (Proposition 9). We also generalize (Sections 5 and 6) the AGM belief revision functions to *iterated* belief revision functions and show that well-known principles of iterated belief revision can be represented as modal axioms. The remainder of this section provides an intuitive review of the branching-time structures introduced in [8] and a more detailed account of the main results.

Static beliefs are usually modeled semantically by means of Kripke frames $\langle \Omega, \mathcal{B} \rangle$, where Ω is a set of states (or possible worlds) and \mathcal{B} is a binary relation on Ω . The interpretation of $\omega \mathcal{B} \omega'$ is that at state ω the agent considers state ω' possible. If S is a set of atomic formulas, one obtains a *model* based on the frame $\langle \Omega, \mathcal{B} \rangle$ by adding a valuation $V : S \rightarrow 2^\Omega$ which specifies, for each atomic formula p , the set of states at which p is true. Rules are then given for determining the truth set of an arbitrary formula ϕ and the agent is said to believe ϕ at a state ω if and only if ϕ is true at every state that she considers possible at ω (that is, if ϕ is true at every state ω' such that $\omega \mathcal{B} \omega'$). The natural way to extend this approach to belief change would be to consider a sequence $\{\langle \Omega, \mathcal{B}_0 \rangle, \langle \Omega, \mathcal{B}_1 \rangle, \dots, \langle \Omega, \mathcal{B}_t \rangle, \dots\}$ of Kripke frames, where \mathcal{B}_t represents the beliefs of the agent at time t . However, such an extension is not sufficient to provide a bridge to the AGM theory of belief revision for two reasons: (1) the stimulus prompting the change in beliefs (which in the AGM theory is an informational input) is not modeled explicitly and (2) some of the AGM postulates require a comparison between the revised belief sets that arise in response to different informational inputs. In order to take care of the latter issue, in [8] branching-time frames were used, where an instant t can have multiple immediate successors. The former issue was dealt with by associating

with each instant t not only a belief relation \mathcal{B}_t but also an information relation \mathcal{I}_t (on the set of states Ω). The interpretation of $\omega\mathcal{I}_t\omega'$ is that at state ω and according to the information received by the agent at time t , state ω' is a possibility. In a model based on the frame, at a state-instant pair (ω, t) the agent believes a formula ϕ if -as usual - ϕ is true at every state ω' such that $\omega\mathcal{B}_t\omega'$; on the other hand, the agent is said to be *informed* that ψ if the truth set of ψ , denoted by $\|\psi\|$, *coincides* with set of states that are reachable from ω by means of the relation \mathcal{I}_t (that is, if $\omega' \in \|\psi\|$ if and only if $\omega\mathcal{I}_t\omega'$). This requirement captures the notion that the agent is informed that, and only that, ψ is the case. This semantic condition was used before in the literature to capture the notion that “all the agent knows is that ψ ” [23, 28]. Thus, in a model based on a branching-time belief revision frame each state-instant pair (ω, t) gives rise to an “initial” belief set K (the set of formulas that the agent believes at (ω, t)) and a collection of potential informational inputs (the formulas of which the agent is informed at the immediate successors of instant t and at state ω) together with the associated new beliefs. Hence each state-instant pair (ω, t) induces a “partial” belief revision function (partial in the sense that not every formula is a potential informational input). We investigate under what conditions such a partial belief revision function can be extended to a “full” AGM belief revision function (full in the sense that the AGM functions consider *every* formula as a potential informational input). We show (Proposition 6) that a necessary and sufficient condition - when the set of states Ω is finite - is that there exist a total pre-order R of Ω that rationalizes belief revision at (ω, t) , in the sense that *both* at instant t *and* at its immediate successors (and at state ω) the states that the agent considers possible (according to the belief relations) are the R -maximal states among the ones that are compatible with the information received. We also provide a property on frames that is equivalent to the existence of such a total pre-order and is directly verifiable on the frame itself. We use the expression “AGM-consistent” to refer to a frame that satisfies any of these three equivalent properties.

In Section 4 we turn to the associated modal logic introduced in [8] and provide a set of axioms that characterizes the class of AGM-consistent branching-time belief revision frames (Proposition 9). Finally, in Sections 5 and 6, we address the issue of iterated belief revision. First, we discuss the semantic and syntactic modal correspondents of some well-known principles of iterated belief revision. Then we introduce a generalization of AGM belief revision functions that can be used to model iterated revision and show that every model based on an AGM-consistent branching-time frame gives rise to such an iterated belief revision function. One advantage of the iterated belief revision functions is that they allow a precise formulation of what an epistemic state is and how an informational input transforms a epistemic state into a new one.

The paper is organized as follows. In the next section we briefly review the AGM theory of belief revision [1], while in Section 3 we review the branching-time belief revision frames introduced in [8] and establish the link between such frames and AGM belief revision functions through the notion of interpretation or model. The main result of this section is Proposition 6 which establishes three equivalent properties of branching-time belief revision frames, one of which is the ability to extend the partial belief revision function induced (at a state-instant pair) by an arbitrary interpretation of the frame to a full AGM belief revision function. In Section 4 we turn to the modal logic introduced in [8], which is the syntactic counterpart of the branching-time belief revision frames, and in Proposition 9 we provide an axiomatic characterization of the class of AGM-consistent frames. Sections 5 and 6 are devoted to a discussion of iterated belief revision and the last section discusses related literature. All the proofs are given in the Appendix.

The original contribution of this paper is contained in Propositions 6 and 9. To the best of our knowledge, the definition of iterated belief revision function put forward in Section 6 is also new.

2 AGM belief revision functions

In this section we briefly review the AGM theory of belief revision [1].¹ Let Φ be the set of formulas of a propositional language based on a countable set S of atomic formulas (or sentence letters).² Given a subset $K \subseteq \Phi$, its PL-deductive closure $[K]^{PL}$ (where ‘PL’ stands for Propositional Logic) is defined as follows: $\psi \in [K]^{PL}$ if and only if there exist $\phi_1, \dots, \phi_n \in K$ (with $n \geq 0$) such that $(\phi_1 \wedge \dots \wedge \phi_n) \rightarrow \psi$ is a tautology (that is, a theorem of Propositional Logic). A set $K \subseteq \Phi$ is *consistent* if $[K]^{PL} \neq \Phi$ (equivalently, if there is no formula ϕ such that both ϕ and $\neg\phi$ belong to $[K]^{PL}$). A set $K \subseteq \Phi$ is *deductively closed* if $K = [K]^{PL}$. A *belief set* is a set $K \subseteq \Phi$ which is deductively closed.

Let K be a consistent belief set representing the agent’s initial beliefs and let $\Psi \subseteq \Phi$ be a set of formulas representing possible items of information. A *belief revision function based on K* is a function $B_K : \Psi \rightarrow 2^\Phi$ (where 2^Φ denotes the set of subsets of Φ) that associates with every formula $\psi \in \Psi$ (thought of as new information) a set $B_K(\psi) \subseteq \Phi$ (thought of as the revised beliefs).³ If $\Psi \neq \Phi$ we

¹For a more detailed account see [16, 20].

²Thus Φ is defined recursively as follows: if $p \in S$ then $p \in \Phi$ and if $\phi, \psi \in \Phi$ then $\neg\phi \in \Phi$ and $(\phi \vee \psi) \in \Phi$. The connectives \wedge and \rightarrow are defined as usual: $\phi \wedge \psi \stackrel{def}{=} \neg(\neg\phi \vee \neg\psi)$ and $\phi \rightarrow \psi \stackrel{def}{=} \neg\phi \vee \psi$.

³In the literature it is common to use the notation $K * \psi$ or K_ψ^* instead of $B_K(\psi)$, but for our purposes the latter notation is clearer.

call B_K a *partial* belief revision function, while if $\Psi = \Phi$ then B_K is called a *full* belief revision function.

Definition 1 Let $B_K : \Psi \rightarrow 2^\Phi$ be a (partial) belief revision function and $B_K^* : \Phi \rightarrow 2^\Phi$ a full belief revision function. We say that B_K^* is an extension of B_K if, for every $\psi \in \Psi$, $B_K^*(\psi) = B_K(\psi)$.

Definition 2 A full belief revision function is called an AGM revision function if it satisfies the following properties, known as the AGM postulates: $\forall \phi, \psi \in \Phi$,

- (AGM1) $B_K(\phi) = [B_K(\phi)]^{PL}$
- (AGM2) $\phi \in B_K(\phi)$
- (AGM3) $B_K(\phi) \subseteq [K \cup \{\phi\}]^{PL}$
- (AGM4) if $\neg\phi \notin K$, then $[K \cup \{\phi\}]^{PL} \subseteq B_K(\phi)$
- (AGM5) $B_K(\phi) = \Phi$ if and only if ϕ is a contradiction
- (AGM6) if $\phi \leftrightarrow \psi$ is a tautology then $B_K(\phi) = B_K(\psi)$
- (AGM7) $B_K(\phi \wedge \psi) \subseteq [B_K(\phi) \cup \{\psi\}]^{PL}$
- (AGM8) if $\neg\psi \notin B_K(\phi)$, then $[B_K(\phi) \cup \{\psi\}]^{PL} \subseteq B_K(\phi \wedge \psi)$.

AGM1 requires the revised belief set to be deductively closed.

AGM2 requires that the information be believed.

AGM3 says that beliefs should be revised minimally, in the sense that no new formula should be added unless it can be deduced from the information received and the initial beliefs.⁴

AGM4 says that if the information received is compatible with the initial beliefs, then any formula that can be deduced from the information and the initial beliefs should be part of the revised beliefs.

AGM5 requires the revised beliefs to be consistent, unless the information ϕ is a contradiction (that is, $\neg\phi$ is a tautology).

AGM6 requires that if ϕ is propositionally equivalent to ψ then the result of revising by ϕ be identical to the result of revising by ψ .

AGM7 and AGM8 are a generalization of AGM3 and AGM4 that

“applies to *iterated* changes of belief. The idea is that if $B_K(\phi)$ is a revision of K [prompted by ϕ] and $B_K(\phi)$ is to be changed by adding further sentences, such a change should be made by using expansions of $B_K(\phi)$ whenever possible. More generally, the minimal change of K to include both ϕ and ψ (that is, $B_K(\phi \wedge \psi)$) ought to be the same as the expansion of $B_K(\phi)$ by ψ , so long as ψ does not contradict the beliefs in $B_K(\phi)$ ” ([20], p. 55; notation changed to match ours).

⁴Note that, for every formula ψ , $\psi \in [K \cup \{\phi\}]^{PL}$ if and only if $(\phi \rightarrow \psi) \in K$ (since, by hypothesis, $K = [K]^{PL}$).

3 Branching-time belief revision frames and models

We now turn to the semantic frames introduced in [8], which provide a way of modeling the evolution of an agent's beliefs over time in response to informational inputs, and establish a connection between these frames and the AGM belief revision functions.

A *next-time branching frame* is a pair $\langle T, \succ \rangle$ where T is a set of instants and \succ is a binary relation on T satisfying the following properties: $\forall t_1, t_2, t_3 \in T$,

1. if $t_1 \succ t_3$ and $t_2 \succ t_3$ then $t_1 = t_2$,
2. if $\langle t_1, \dots, t_n \rangle$ is a sequence in T with $t_i \succ t_{i+1}$, for every $i = 1, \dots, n - 1$, then $t_n \neq t_1$.

The interpretation of $t_1 \succ t_2$ is that t_2 is an *immediate successor* of t_1 or t_1 is the *immediate predecessor* of t_2 : every instant has at most a unique immediate predecessor but can have several immediate successors. We denote the set of immediate successors of $t \in T$ by t^\rightarrow , that is, $t^\rightarrow = \{t' \in T : t \succ t'\}$.

A *branching-time belief-information frame* is a tuple $\langle T, \succ, \Omega, \{\mathcal{I}_t, \mathcal{B}_t\}_{t \in T} \rangle$ where $\langle T, \succ \rangle$ is a next-time branching frame, Ω is a set of states and, for every $t \in T$, \mathcal{I}_t and \mathcal{B}_t are binary relations on Ω , the first representing information and the latter beliefs. The interpretation of $\omega \mathcal{I}_t \omega'$ is that at state ω and time t - according to the information received - it is possible that the true state is ω' . On the other hand, the interpretation of $\omega \mathcal{B}_t \omega'$ is that at state ω and time t , in light of the information received, the agent considers state ω' possible (an alternative expression is “ ω' is a doxastic alternative to ω at time t ”). We shall use the following notation:

$$\mathcal{I}_t(\omega) = \{\omega' \in \Omega : \omega \mathcal{I}_t \omega'\} \text{ and, similarly, } \mathcal{B}_t(\omega) = \{\omega' \in \Omega : \omega \mathcal{B}_t \omega'\}.$$

Thus $\mathcal{I}_t(\omega)$ is the set of states that are reachable from ω according to the relation \mathcal{I}_t and similarly for $\mathcal{B}_t(\omega)$.

Definition 3 A branching-time belief revision frame is a frame $\langle T, \succ, \Omega, \{\mathcal{I}_t, \mathcal{B}_t\}_{t \in T} \rangle$ that satisfies the following properties: $\forall \omega \in \Omega, \forall t, t', t'' \in T$:

1. $\mathcal{B}_t(\omega) \subseteq \mathcal{I}_t(\omega)$
2. $\mathcal{B}_t(\omega) \neq \emptyset$
3. if $t \succ t', t \succ t''$ and $\mathcal{I}_{t'}(\omega) = \mathcal{I}_{t''}(\omega)$ then $\mathcal{B}_{t'}(\omega) = \mathcal{B}_{t''}(\omega)$
4. if $t \succ t'$ and $\mathcal{B}_t(\omega) \cap \mathcal{I}_{t'}(\omega) \neq \emptyset$ then $\mathcal{B}_{t'}(\omega) = \mathcal{B}_t(\omega) \cap \mathcal{I}_{t'}(\omega)$.

Property 1 says that information is believed and Property 2 that beliefs are consistent. The two together imply that $\mathcal{I}_t(\omega) \neq \emptyset$, that is, that information itself is consistent.⁵

Property 3 requires that at any two instants that share the same immediate predecessor, if information is the same then beliefs must be the same. That is, differences in beliefs must be due to differences in information.

Property 4 is called the ‘Qualitative Bayes Rule’ (QBR) in [8], based on the following observation. In a probabilistic setting, let $P_{\omega,t}$ be the probability measure over a set of states Ω representing the agent’s probabilistic beliefs at state ω and instant t , let $F \subseteq \Omega$ be an event representing the information received by the agent at a later instant t' and let $P_{\omega,t'}$ be the posterior probability measure representing the revised beliefs at state ω and instant t' . Bayes’ rule requires that, if $P_{\omega,t}(F) > 0$, then, for every event $E \subseteq \Omega$, $P_{\omega,t'}(E) = \frac{P_{\omega,t}(E \cap F)}{P_{\omega,t}(F)}$. Bayes’ rule thus implies the following (where $\text{supp}(P)$ denotes the support of the probability measure P):

$$\text{if } \text{supp}(P_{\omega,t}) \cap F \neq \emptyset, \text{ then } \text{supp}(P_{\omega,t'}) = \text{supp}(P_{\omega,t}) \cap F.$$

If we set $\mathcal{B}_t(\omega) = \text{supp}(P_{\omega,t})$, $F = \mathcal{I}_{t'}(\omega)$, with $t \mapsto t'$, and $\mathcal{B}_{t'}(\omega) = \text{supp}(P_{\omega,t'})$ then we get Property 4. Thus in a probabilistic setting the proposition “the agent believes that ϕ ” would be interpreted as “the agent assigns probability 1 to the set of states where ϕ is true”.

Figure 1 shows a branching-time belief revision frame. For simplicity, in all the figures we assume that, for every instant t , the information relation \mathcal{I}_t is an equivalence relation (whose equivalence classes are denoted by rectangles) and the belief relation \mathcal{B}_t is transitive and euclidean.⁶ An arrow from ω to ω' means that $\omega' \in \mathcal{B}_t(\omega)$ (or $\omega \mathcal{B}_t \omega'$, that is, ω' is reachable from ω according to the relation \mathcal{B}_t). Note, however, that none of the results below require \mathcal{I}_t to be an equivalence relation (in particular, veridicality of information is *not* assumed), nor do they require \mathcal{B}_t to be transitive and euclidean.

⁵Thus we rule out inconsistent information. As pointed out in [18], it is not clear how one could be informed of a contradiction or, at least, how one could treat a contradiction as information.

⁶ \mathcal{B}_t is transitive if $\omega' \in \mathcal{B}_t(\omega)$ implies that $\mathcal{B}_t(\omega') \subseteq \mathcal{B}_t(\omega)$; it is euclidean if $\omega' \in \mathcal{B}_t(\omega)$ implies that $\mathcal{B}_t(\omega) \subseteq \mathcal{B}_t(\omega')$. Property 2 of Definition 3 is usually referred to as seriality.

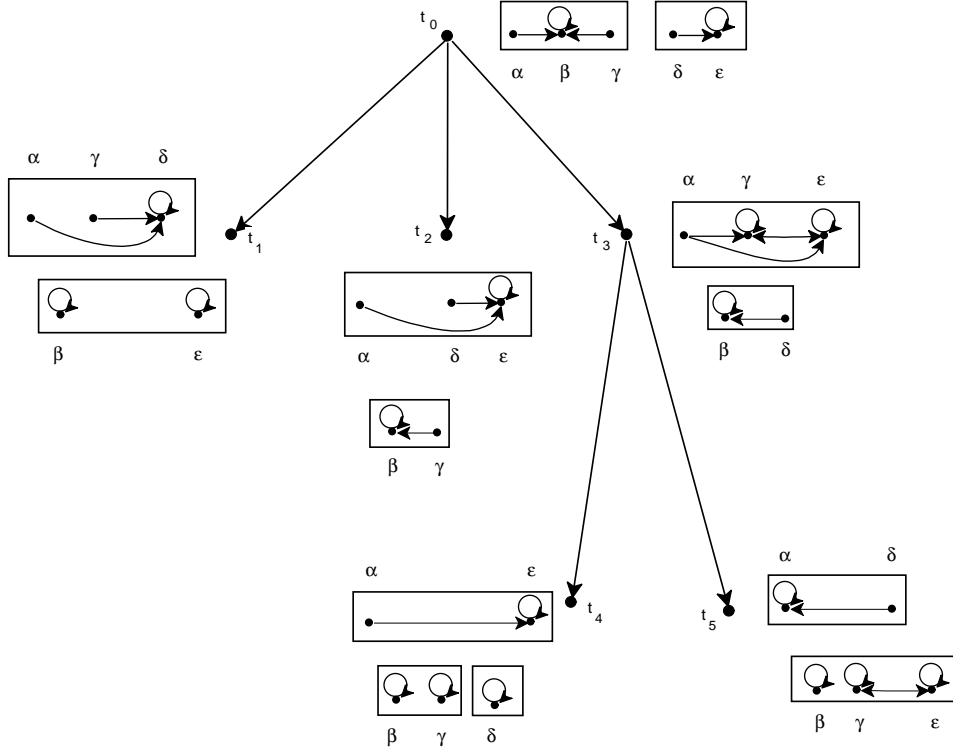


Figure 1

For example, in Figure 1 at state α and instant t_3 the agent is informed that the true state is either α , γ or ε ($\mathcal{I}_{t_3}(\alpha) = \{\alpha, \gamma, \varepsilon\}$) and (incorrectly) believes that it is either γ or ε ($\mathcal{B}_{t_3}(\alpha) = \{\gamma, \varepsilon\}$). At the next instant t_4 (and still at state α) the agent is now informed that the true state is either α or ε ($\mathcal{I}_{t_4}(\alpha) = \{\alpha, \varepsilon\}$) and forms the revised (and still incorrect) belief that the true state is ε . On the other hand, t_5 is an alternative next instant to t_3 and at t_5 (and still at state α) the agent's information is $\mathcal{I}_{t_5}(\alpha) = \{\alpha, \delta\}$ and she forms the revised (and now correct) belief that the true state is α ($\mathcal{B}_{t_5}(\alpha) = \{\alpha\}$). Note that all the properties of Definition 3 are satisfied. In particular the Qualitative Bayes Rule is satisfied everywhere: sometimes vacuously (as is the case at state α and instants t_3 and t_5 where $\mathcal{B}_{t_3}(\alpha) \cap \mathcal{I}_{t_5}(\alpha) = \emptyset$) and sometimes non-trivially (as is the case at state α and instants t_3 and t_4 where $\mathcal{B}_{t_3}(\alpha) \cap \mathcal{I}_{t_4}(\alpha) = \mathcal{B}_{t_4}(\alpha) = \{\varepsilon\}$).

Given a branching-time belief revision frame one obtains a model based on it by adding a valuation that associates with every atomic proposition p the set of states at which p is true. Note that, by defining a valuation this way, we frame the problem as one of belief revision, since the truth value of an atomic proposition

depends only on the state and not on the time.⁷ Let S be a countable set of atomic formulas and Φ the set of propositional formulas built from S . Given a frame $\mathcal{F} = \langle T, \succ, \Omega, \{\mathcal{I}_t, \mathcal{B}_t\}_{t \in T} \rangle$, a *model based on* (or an *interpretation of*) \mathcal{F} is obtained by adding to \mathcal{F} a *valuation* $V : S \rightarrow 2^\Omega$.⁸ Truth of an arbitrary formula $\phi \in \Phi$ at state ω in model \mathcal{M} is denoted by $\omega \models_{\mathcal{M}} \phi$ and is defined recursively as follows:

- (1) for $p \in S$, $\omega \models_{\mathcal{M}} p$ if and only if $\omega \in V(p)$,
- (2) $\omega \models_{\mathcal{M}} \neg\phi$ if and only if $\omega \not\models_{\mathcal{M}} \phi$, and
- (3) $\omega \models_{\mathcal{M}} (\phi \vee \psi)$ if and only if either $\omega \models_{\mathcal{M}} \phi$ or $\omega \models_{\mathcal{M}} \psi$ (or both).

The truth set of formula ϕ in model \mathcal{M} is denoted by $\|\phi\|_{\mathcal{M}}$; thus $\|\phi\|_{\mathcal{M}} = \{\omega \in \Omega : \omega \models_{\mathcal{M}} \phi\}$.

In a given model, the truth of each formula is thus time-independent, that is, the state is sufficient to determine truth (indeed, this is the essence of belief revision as opposed to belief update). Nevertheless, beliefs and information may change over time. For example, while an atomic formula p may be true at state ω (and thus at any state-instant pair (ω, t) , for every instant t), the agent may believe $\neg p$ at state ω and instant t_1 and then - upon being informed that p - switch to believing p at state ω and instant t_2 (with t_2 an immediate successor of t_1). Similarly, what information the agent receives will depend not only on the state but also on time. In other words, if $t \neq t'$ typically the relation \mathcal{B}_t will be different from $\mathcal{B}_{t'}$ (similarly, \mathcal{I}_t will be different from $\mathcal{I}_{t'}$).

Definition 4 *Given a model $\mathcal{M} = \langle T, \succ, \Omega, \{\mathcal{I}_t, \mathcal{B}_t\}_{t \in T}, V \rangle$, a state $\omega \in \Omega$, an instant $t \in T$ and formulas $\phi, \psi \in \Phi$ we say that*

- at (ω, t) the agent is informed that ψ if and only if $\mathcal{I}_t(\omega) = \|\psi\|_{\mathcal{M}}$,
- at (ω, t) the agent believes that ϕ if and only if $\mathcal{B}_t(\omega) \subseteq \|\phi\|_{\mathcal{M}}$.

Note that for information we require *equality* of the two sets (as explained in the Introduction, this captures the idea of ‘being informed precisely that ψ ’ and corresponds to the notion of ‘only knowing’ introduced in [23, 28]), while for belief we use the standard requirement that $\mathcal{B}_t(\omega)$ be a *subset* of the truth set of a formula.

⁷The branching-time structures of Definition 3 can be used to describe either a situation where the objective facts describing the world do not change – so that only the beliefs of the agent change over time – or a situation where both the facts and the doxastic state of the agent change. In the literature the first situation is called *belief revision*, while the latter is called *belief update*: see [25]. In this paper we restrict attention to belief revision.

⁸If instead of belief revision we were interested in belief update, then we would need to define a valuation as a function $V : S \rightarrow 2^{\Omega \times T}$.

Given a model \mathcal{M} and a state-instant pair (ω, t) , according to Definition 4 we can associate with (ω, t) a belief set and a (typically partial) belief revision function as follows. Let

$$K_{\mathcal{M},\omega,t} = \{\phi \in \Phi : \mathcal{B}_t(\omega) \subseteq \|\phi\|_{\mathcal{M}}\}, \quad (1)$$

denote the set of formulas that the agent believes at (ω, t) , that is, his (initial) belief set at (ω, t) . It is straightforward to show that $K_{\mathcal{M},\omega,t}$ is a consistent and deductively closed set. Let

$$\Psi_{\mathcal{M},\omega,t} = \{\psi \in \Phi : \|\psi\|_{\mathcal{M}} = \mathcal{I}_{t'}(\omega) \text{ for some } t' \in t^{\rightarrow}\} \quad (2)$$

be the possible items of information that the agent might receive next time (that is, at some immediate successor of t : recall that $t^{\rightarrow} = \{t' \in T : t \mapsto t'\}$). Finally let $B_{K_{\mathcal{M},\omega,t}} : \Psi_{\mathcal{M},\omega,t} \rightarrow 2^{\Phi}$ be defined as⁹

$$B_{K_{\mathcal{M},\omega,t}}(\psi) = \{\phi \in \Phi : \mathcal{B}_{t'}(\omega) \subseteq \|\phi\|_{\mathcal{M}} \text{ for } t' \in t^{\rightarrow} \text{ with } \mathcal{I}_{t'}(\omega) = \|\psi\|_{\mathcal{M}}\}. \quad (3)$$

That is, if at the immediate successor t' of t the agent is informed that ψ ($\mathcal{I}_{t'}(\omega) = \|\psi\|_{\mathcal{M}}$), then his revised belief set is given by the set of formulas that he believes at (ω, t') : $\{\phi \in \Phi : \mathcal{B}_{t'}(\omega) \subseteq \|\phi\|_{\mathcal{M}}\}$.

For example, consider a model of the frame illustrated in Figure 1 above where, for some atomic formulas p_1, p_2, p_3 and q , $V(p_1) = \{\alpha, \gamma, \delta\} = \mathcal{I}_{t_1}(\alpha)$, $V(p_2) = \{\alpha, \delta, \varepsilon\} = \mathcal{I}_{t_2}(\alpha)$, $V(p_3) = \{\alpha, \gamma, \varepsilon\} = \mathcal{I}_{t_3}(\alpha)$ and $V(q) = \{\gamma\}$. Then the initial beliefs at (α, t_0) are given by the (consistent and deductively closed) set $K_{\alpha,t_0} = \{\phi \in \Phi : \beta \models \phi\}$. The set Ψ_{α,t_0} of potential informational inputs at (α, t_0) is rather small; for example, while $p_1, p_2, p_3 \in \Psi_{\alpha,t_0}$, $(p_1 \vee p_2) \notin \Psi_{\alpha,t_0}$. Thus the associated belief revision function $B_{K_{\alpha,t_0}}$ is a partial belief revision function. As an example we have that $\neg q, p_3 \in B_{K_{\alpha,t_0}}(p_2)$ [because $\mathcal{I}_{t_2}(\alpha) = \|p_2\|$ and $\mathcal{B}_{t_2}(\alpha) = \{\varepsilon\} \subseteq \|\neg q\| \cap \|p_3\|$]; thus, since $\neg q, \neg p_3 \in K_{\alpha,t_0}$ [because $\mathcal{B}_{t_0}(\alpha) = \{\beta\}$ and $\beta \models \neg q$ and $\beta \models \neg p_3$], the agent initially believes both $\neg q$ and $\neg p_3$ and, upon being informed that p_2 (at (α, t_2)) she revises her beliefs by retaining the belief that $\neg q$ but switching from believing that $\neg p_3$ to believing that p_3 . A natural question to ask is whether this partial belief revision function is compatible with the AGM postulates, in the sense that there exists a full belief revision function B_K^* that satisfies the AGM postulates and is an extension of $B_{K_{\alpha,t_0}}$ (see Definition 1). In this case the answer is negative. This can be proved as follows. To simplify the notation we shall drop the subscripts a, t_0 ; thus we write K instead of K_{α,t_0} , B_K

⁹This function is well defined because of Property 3 of Definition 3.

instead of $B_{K_{\alpha, t_0}}$, etc. Suppose that B_K^* is an AGM extension of B_K . Then, since $p_2 \in B_K(p_1)$ and $p_2 \in B_K(p_2)$ (and $B_K(p_1) = B_K^*(p_1)$ and $B_K(p_2) = B_K^*(p_2)$) it follows that¹⁰

$$p_2 \in B_K^*(p_1 \vee p_2). \quad (4)$$

Thus $B_K^*((p_1 \vee p_2) \wedge p_2) = B_K^*(p_1 \vee p_2)$.¹¹ Since $(p_1 \vee p_2) \wedge p_2$ is equivalent to p_2 , by AGM6 $B_K^*((p_1 \vee p_2) \wedge p_2) = B_K^*(p_2)$. Thus (since $B_K(p_2) = B_K^*(p_2)$)

$$B_K^*(p_1 \vee p_2) = B_K(p_2). \quad (5)$$

Since $p_3 \in B_K(p_2)$,

$$[B_K(p_2) \cup \{p_3\}]^{PL} = [B_K(p_2)]^{PL} = B_K(p_2). \quad (6)$$

[It is straightforward to show that, for every $\psi \in \Psi$, $B_K(\psi)$ is deductively closed.] Furthermore, by (5), $p_3 \in B_K^*(p_1 \vee p_2)$. Since $(p_1 \vee p_2)$ is not a contradiction, by AGM5 $B_K^*(p_1 \vee p_2)$ is consistent and thus $\neg p_3 \notin B_K^*(p_1 \vee p_2)$. Hence, by AGM7 and AGM8, $B_K^*((p_1 \vee p_2) \wedge p_3) = [B_K^*((p_1 \vee p_2) \cup \{p_3\})]^{PL}$ and, by (5), the latter is equal to $[B_K(p_2) \cup \{p_3\}]^{PL}$ which, in turn, by (6), is equal to $B_K(p_2)$. Thus

$$B_K^*((p_1 \vee p_2) \wedge p_3) = B_K(p_2). \quad (7)$$

Since $(p_1 \vee p_2) \wedge p_3$ is equivalent to p_3 , by AGM6 $B_K^*((p_1 \vee p_2) \wedge p_3) = B_K^*(p_3)$. Thus, by (7),

$$B_K^*(p_3) = B_K(p_2). \quad (8)$$

Since B_K^* is an extension of B_K , $B_K^*(p_3) = B_K(p_3)$. It follows from this and (8) that $B_K(p_3) = B_K(p_2)$, yielding a contradiction, since $\neg q \in B_K(p_2)$ but $\neg q \notin B_K(p_3)$.

In view of the above example, a natural question to ask is whether there exists a property of branching-time belief revision frames that guarantees that the partial belief revision functions generated by models based on frames that satisfy that

¹⁰This is a consequence of the following result, which is proved in the Appendix (Lemma 13). Let K be a consistent belief set and $B_K : \Phi \rightarrow 2^\Phi$ an AGM belief revision function. Let $\phi, \psi, \chi \in \Phi$ be such that $\chi \in B_K(\phi)$ and $\chi \in B_K(\psi)$. Then $\chi \in B_K(\phi \vee \psi)$.

¹¹Proof: by AGM1, $B_K^*(p_1 \vee p_2) = [B_K^*(p_1 \vee p_2)]^{PL}$. By AGM5, since $(p_1 \vee p_2)$ is not a contradiction, $B_K^*(p_1 \vee p_2) \neq \Phi$. Thus, since $p_2 \in B_K^*(p_1 \vee p_2)$, $\neg p_2 \notin B_K^*(p_1 \vee p_2)$. Hence, by AGM7 and AGM8, $B_K^*((p_1 \vee p_2) \wedge p_2) = [B_K^*(p_1 \vee p_2) \cup \{p_2\}]^{PL} = [B_K^*(p_1 \vee p_2)]^{PL} = B_K^*(p_1 \vee p_2)$.

property are compatible with the AGM postulates.¹² The notion of compatibility with the AGM postulates is made precise in the following definition.

Definition 5 *A branching-time belief revision frame $\mathcal{F} = \langle T, \succrightarrow, \Omega, \{\mathcal{I}_t, \mathcal{B}_t\}_{t \in T} \rangle$ is AGM-consistent at $(\omega, t) \in \Omega \times T$ if, for every model $\mathcal{M} = \langle \mathcal{F}, V \rangle$ based on it, the associated belief revision function $B_{K_{\mathcal{M}, \omega, t}}$ (see (3) above) can be extended (see Definition 1) to a full AGM belief revision function (see Definition 2).*

We showed above that the branching-time belief revision frame illustrated in Figure 1 is *not* AGM consistent at (α, t_0) .

The following proposition, which is proved in the Appendix, builds on results given in [10] and [21].¹³ Note that the Qualitative Bayes Rule (Property 4 of Definition 3) is necessary for the validity of Proposition 6.

A *total pre-order* of Ω is a binary relation $R \subseteq \Omega \times \Omega$ which is complete ($\forall \omega, \omega' \in \Omega$, either $\omega R \omega'$ or $\omega' R \omega$) and transitive ($\forall \omega, \omega', \omega'' \in \Omega$, if $\omega R \omega'$ and $\omega' R \omega''$ then $\omega R \omega''$). We shall interpret $\omega R \omega'$ as “state ω is at least as plausible as state ω' ”. Given a total pre-order R of Ω and a subset $E \subseteq \Omega$, let

$$best_R E \stackrel{def}{=} \{\omega \in E : \omega R \omega', \forall \omega' \in E\}.$$

Thus $best_R E$ is the set of states in E that are most plausible according to R .¹⁴

¹²Why is this a desirable property? By Definition 4, a state-instant pair (ω, t) in a model identifies (1) the agent’s current beliefs, (2) the possible items of information to be received and (3) the agent’s disposition to revise her beliefs in response to those informational inputs. An introspective agent would naturally be worried about the consistency of her disposition to revise her beliefs. In the example just described, by considering counterfactual informational inputs, such as $(p_1 \vee p_2 \vee p_3)$, and her hypothetical response to them, the agent would be able to uncover an inconsistency in her disposition to revise her beliefs in response to the actual informational inputs p_1 , p_2 and p_3 . The notion of AMG-consistency defined below guarantees that no inconsistencies could be detected by contemplating hypothetical information in addition to the actual information.

¹³Both [10] and [21] deal with choice functions $f : \mathcal{E} \rightarrow 2^\Omega$, where \mathcal{E} is a collection of subsets of Ω , satisfying the property that if $E \neq \emptyset$ then $\emptyset \neq f(E) \subseteq E$. Choice functions are used in economics to represent the choices made by an individual when faced with possible menus of alternatives. In [21] a necessary and sufficient condition is given for the rationalizability of a choice function in terms of a preference relation and in [10] choice functions are shown to be interpretable in terms of one-shot belief revision. In the proof given in the Appendix more details are given on how results in [10] and [21] can be extended to branching-time belief revision frames to obtain Proposition 6. In particular, the Qualitative Bayes Rule plays a crucial role.

¹⁴In the literature sometimes the total pre-order is denoted by \succeq and the set $\{\omega \in E : \omega \succeq \omega', \forall \omega' \in E\}$ is referred to as the set of *maximal* elements of E , while some other times the total pre-order is denoted by \preceq and the set $\{\omega \in E : \omega \preceq \omega', \forall \omega' \in E\}$ is referred to as the set of *minimal* elements of E . In order to avoid confusion, we denote the relation by R and refer to the *best* elements of a set.

Proposition 6 Let $\mathcal{F} = \langle T, \succ, \Omega, \{\mathcal{I}_t, \mathcal{B}_t\}_{t \in T} \rangle$ be a branching-time belief revision frame where Ω is finite and let $(\omega, t) \in \Omega \times T$. Then the following conditions are equivalent:

- (a) \mathcal{F} is AGM consistent at (ω, t) .
- (b) There exists a total pre-order $R_{\omega, t}$ of Ω that rationalizes the agent's beliefs at t and at the immediate successors of t (and state ω) in the sense that
 - b1. $\mathcal{B}_t(\omega) = \text{best}_{R_{\omega, t}} \mathcal{I}_t(\omega)$, and
 - b2. for every $t' \in T$ such that $t \succ t'$, $\mathcal{B}_{t'}(\omega) = \text{best}_{R_{\omega, t}} \mathcal{I}_{t'}(\omega)$.
- (c) $\forall u_0, u_1, \dots, u_n \in t^{\rightarrow}$ with $u_n = u_0$ (recall that t^{\rightarrow} is the set of immediate successors of t),

if $\mathcal{I}_{u_{k-1}}(\omega) \cap \mathcal{B}_{u_k}(\omega) \neq \emptyset, \forall k = 1, \dots, n,$
 then $\mathcal{I}_{u_{k-1}}(\omega) \cap \mathcal{B}_{u_k}(\omega) = \mathcal{B}_{u_{k-1}}(\omega) \cap \mathcal{I}_{u_k}(\omega), \forall k = 1, \dots, n.$ (PLS)

A frame that satisfies Property (b) of Proposition 6 is said to be *rationalizable at (ω, t)* and we say that the total pre-order $R_{\omega, t}$ *rationalizes belief revision at (ω, t)* . The branching-time belief revision frame illustrated in Figure 1 is *not* rationalizable at (α, t_0) . In fact, suppose that there is a total pre-order R_{α, t_0} that satisfies (b.1) and (b.2). Let P_{α, t_0} be the corresponding strict order (thus $\omega P_{\alpha, t_0} \omega'$ if and only if $\omega R_{\alpha, t_0} \omega'$ and not $\omega' R_{\alpha, t_0} \omega$). Then, since $\gamma \in \mathcal{I}_{t_1}(\alpha)$ and $\{\delta\} = \mathcal{B}_{t_1}(\alpha) = \text{best}_{R_{\alpha, t_0}} \mathcal{I}_{t_1}(\alpha)$, $\delta P_{\alpha, t_0} \gamma$; similarly, since $\delta \in \mathcal{I}_{t_2}(\alpha)$ and $\{\varepsilon\} = \mathcal{B}_{t_2}(\alpha) = \text{best}_{R_{\alpha, t_0}} \mathcal{I}_{t_2}(\alpha)$, $\varepsilon P_{\alpha, t_0} \delta$. Hence, by transitivity, $\varepsilon P_{\alpha, t_0} \gamma$. However, from $\{\gamma, \varepsilon\} = \mathcal{B}_{t_3}(\alpha) = \text{best}_{R_{\alpha, t_0}} \mathcal{I}_{t_3}(\alpha)$ we get that $\gamma R_{\alpha, t_0} \varepsilon$, yielding a contradiction. Since the frame is not rationalizable at (α, t_0) , it follows from Proposition 6 that it is not AGM-consistent at (α, t_0) , a fact that was proved directly above.

Property PLS of part (c) of Proposition 6 gives a condition on the frame which is necessary and sufficient for the frame to be rationalizable at (ω, t) . To verify that the frame of Figure 1 fails to satisfy Property PLS at (α, t_0) , let $u_0 = u_3 = t_1$, $u_1 = t_3$ and $u_2 = t_2$. Then $\mathcal{I}_{t_1}(\alpha) \cap \mathcal{B}_{t_3}(\alpha) = \{\gamma\} \neq \emptyset$, $\mathcal{I}_{t_3}(\alpha) \cap \mathcal{B}_{t_2}(\alpha) = \{\varepsilon\} \neq \emptyset$ and $\mathcal{I}_{t_2}(\alpha) \cap \mathcal{B}_{t_1}(\alpha) = \{\delta\} \neq \emptyset$, but $\mathcal{B}_{t_1}(\alpha) \cap \mathcal{I}_{t_3}(\alpha) = \emptyset$ and thus $\mathcal{B}_{t_1}(\alpha) \cap \mathcal{I}_{t_3}(\alpha) \neq \mathcal{I}_{t_1}(\alpha) \cap \mathcal{B}_{t_3}(\alpha)$.

Definition 7 A branching-time belief revision frame is AGM-consistent if it is AGM consistent at every state-instant pair (ω, t) .

Thus, by Proposition 6, a frame where Ω is finite is AGM-consistent if and only if it is rationalizable at every state-instant pair (ω, t) . Property *PLS* of part (c) of Proposition 6 provides a way of verifying directly on the frame whether the frame is AGM-consistent.

In an AGM-consistent frame, for every state-instant pair (ω, t) , belief revision can be rationalized by a plausibility ordering of the set of states, in the sense that at t and at the immediate successors of t (and a state ω) the states that the agent considers doxastically possible (that is, according to her beliefs) are the most plausible among the ones that are compatible with the information received. Figure 2 shows an AGM-consistent branching-time belief revision frame. For example, belief revision at (α, t_0) is rationalized by the total pre-order R_{α, t_0} generated by the strict total order $\beta P_{\alpha, t_0} \delta P_{\alpha, t_0} \gamma P_{\alpha, t_0} \alpha$:

$$R_{\alpha, t_0} = \{(\alpha, \alpha), (\beta, \alpha), (\beta, \beta), (\beta, \gamma), (\beta, \delta), (\delta, \alpha), (\delta, \gamma), (\delta, \delta), (\gamma, \gamma), (\gamma, \alpha)\}.$$

Note that, by Proposition 6, AGM-consistency of a frame at a state-instant pair (ω, t) requires the existence of at least one total pre-order that rationalizes belief revision at (ω, t) . Typically, there may be *several* such total pre-orders. For instance, in the frame of Figure 2 at (α, t_0) another possible total pre-order (besides the one mentioned above) is the pre-order R' generated by the strict total order $\beta P'_{\alpha, t_0} \gamma P'_{\alpha, t_0} \alpha P'_{\alpha, t_0} \delta$.

Remark 8 *In an AGM-consistent frame, it is possible that, if t' is an immediate successor of t , any plausibility ordering of Ω that rationalizes belief revision at (ω, t') is necessarily different from any plausibility ordering that rationalizes belief revision at (ω, t) . For example, in the frame of Figure 2 any total pre-order that rationalizes belief revision at (α, t_0) must be such that γ is strictly more plausible than α ,¹⁵ whereas any total pre-order that rationalizes belief revision at (α, t_2) must be such that α is strictly more plausible than γ .¹⁶ Thus the ranking of α and γ is reversed upon moving from (α, t_0) to (α, t_2) .*

Note also that, for a given instant t , if ω and ω' are different states any total pre-order that rationalizes belief revision at (ω, t) may be necessarily different from any total pre-order that rationalizes belief revision at (ω', t) . For example, in Figure 2, any total pre-order that rationalizes belief revision at (β, t_2) must be such that β is strictly more plausible than δ ,¹⁷ whereas any total pre-order that

¹⁵Because $\mathcal{B}_{t_1}(\alpha) = \{\gamma\}$ and $\mathcal{I}_{t_1}(\alpha) = \{\alpha, \gamma\}$.

¹⁶Because $\mathcal{B}_{t_3}(\alpha) = \{\alpha\}$ and $\mathcal{I}_{t_3}(\alpha) = \{\alpha, \gamma\}$. For example, belief revision at (α, t_2) is rationalized by the total pre-order $R_{\alpha, t_2} = \{(\alpha, \alpha), (\alpha, \gamma), (\alpha, \delta), (\beta, \alpha), (\beta, \beta), (\beta, \gamma), (\beta, \delta), (\delta, \delta), (\delta, \gamma), (\gamma, \gamma)\}$, that is, by the strict total order $\beta P_{\alpha, t_2} \alpha P_{\alpha, t_2} \delta P_{\alpha, t_2} \gamma$.

¹⁷Because $\mathcal{B}_{t_3}(\beta) = \{\beta\}$ and $\mathcal{I}_{t_3}(\beta) = \{\beta, \delta\}$.

rationalizes belief revision at (δ, t_2) must be such that δ is strictly more plausible than β .¹⁸

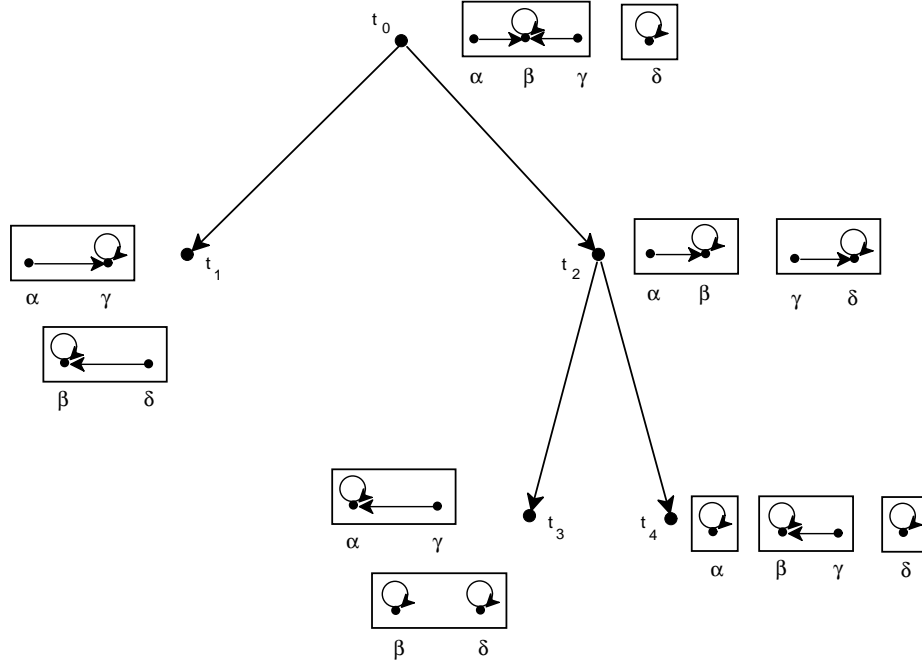


Figure 2

We now turn to a modal-logic characterization of AGM-consistent branching-time belief revision frames.

4 A temporal logic for belief revision

To define the notion of AGM-consistency it was sufficient to consider *propositional* models based on a given branching-time belief revision frame; that is, the language of propositional logic was sufficient. In this section we turn to the more expressive language introduced in [8] - in which belief, information and time are explicitly introduced in the syntax - and provide a syntactic characterization of AGM-consistent belief revision frames within this richer language.

¹⁸Because $\mathcal{B}_{t_3}(\delta) = \{\delta\}$ and $\mathcal{I}_{t_3}(\delta) = \{\beta, \delta\}$. For example, belief revision at (β, t_2) is rationalized by the total pre-order generated by the strict total order $\beta P_{\beta, t_2} \alpha P_{\beta, t_2} \gamma P_{\beta, t_2} \delta$, while belief revision at (δ, t_2) is rationalized the total pre-order generated by the strict total order $\delta P_{\delta, t_2} \beta P_{\delta, t_2} \gamma P_{\delta, t_2} \alpha$.

The language contains the following modal operators: the next-time operator \bigcirc , the belief operator B , the information operator I and the “all state” operator A . The intended interpretation is as follows:

- $\bigcirc\phi$: “at every next instant it will be the case that ϕ ”
- $B\phi$: “the agent believes that ϕ ”
- $I\phi$: “the agent is informed that ϕ ”
- $A\phi$: “it is true at every state that ϕ ”.

The link between the semantics of branching-time belief revision frames and the syntactic language is again given by the addition of a valuation to a frame. As in the previous section, a valuation is a function $V : S \rightarrow 2^\Omega$ which specifies, for every atomic formula $p \in S$ the set of states at which p is true. While for propositional (or Boolean, or non-modal) formulas, states are sufficient to determine truth, the same is not true for modal formulas, since - for instance - there can be a state ω and two different instants t_1 and t_2 such that at (ω, t_1) the agent believes a formula ϕ while at (ω, t_2) she does not, that is, $B\phi$ is true at (ω, t_1) but false at (ω, t_2) . Thus for this more expressive modal language truth of formulas needs to be defined at state-instant pairs (ω, t) .¹⁹ This is done as follows.

Fix a model $\mathcal{M} = \langle T, \succ, \Omega, \{\mathcal{I}_t, \mathcal{B}_t\}_{t \in T}, V \rangle$ as defined in the previous section, where $V : S \rightarrow 2^\Omega$ is a valuation. Given a state ω , an instant t and a formula ϕ , we write $(\omega, t) \models_{\mathcal{M}} \phi$ to denote that ϕ is true at (ω, t) in model \mathcal{M} . Let $\|\phi\|_{\mathcal{M}} \subseteq \Omega \times T$ denote the truth set of ϕ , that is, $\|\phi\|_{\mathcal{M}} = \{(\omega, t) \in \Omega \times T : (\omega, t) \models_{\mathcal{M}} \phi\}$ and let $\|\phi\|_{\mathcal{M}, t} \subseteq \Omega$ denote the set of states at which ϕ is true *at instant* t , that is, $\|\phi\|_{\mathcal{M}, t} = \{\omega \in \Omega : (\omega, t) \models_{\mathcal{M}} \phi\}$. Truth at a state-instant pair (ω, t) is defined recursively as follows.

- if $p \in S$, $(\omega, t) \models p$ if and only if $\omega \in V(p)$.
- $(\omega, t) \models \neg\phi$ if and only if $(\omega, t) \not\models \phi$.
- $(\omega, t) \models \phi \vee \psi$ if and only if either $(\omega, t) \models \phi$ or $(\omega, t) \models \psi$ (or both).
- $(\omega, t) \models \bigcirc\phi$ if and only if $(\omega, t') \models \phi$ for every t' such that $t \succ t'$.
- $(\omega, t) \models B\phi$ if and only if $\mathcal{B}_t(\omega) \subseteq \|\phi\|_{\mathcal{M}, t}$, that is, if $(\omega', t) \models \phi$ for all $\omega' \in \mathcal{B}_t(\omega)$.
- $(\omega, t) \models I\phi$ if and only if ϕ is Boolean and $\mathcal{I}_t(\omega) = \|\phi\|_{\mathcal{M}, t}$, that is, if
 - (1) $(\omega', t) \models \phi$ for all $\omega' \in \mathcal{I}_t(\omega)$, and
 - (2) if $(\omega', t) \models \phi$ then $\omega' \in \mathcal{I}_t(\omega)$.
- $(\omega, t) \models A\phi$ if and only if $\|\phi\|_{\mathcal{M}, t} = \Omega$, that is, if $(\omega', t) \models \phi$ for all $\omega' \in \Omega$.

¹⁹It is shown in [8] (Proposition 5, p. 148) that if ϕ is a Boolean formula then its truth is determined only by the state, that is, for every $\omega \in \Omega$ and $t, t' \in T$, $(\omega, t) \models \phi$ if and only if $(\omega, t') \models \phi$. However, even if ϕ is Boolean, it is possible to have $(\omega, t) \models B\phi$ and $(\omega, t') \not\models B\phi$, that is, the individual might believe ϕ at state ω and instant t but not believe ϕ at the same state but at a different instant t' . Indeed this is the essence of the notion of belief change over time.

Note that, while the other modal operators apply to arbitrary formulas, the information operator is restricted to apply only to Boolean formulas, that is formulas that do not contain modal operators. Boolean formulas represent facts and information is restricted to be about facts.²⁰ We stress again that, while the truth condition for the operator B is the standard one, the truth condition for the operator I is non-standard: instead of simply requiring that $\mathcal{I}_t(\omega) \subseteq \|\phi\|_{\mathcal{M},t}$ we require equality: $\mathcal{I}_t(\omega) = \|\phi\|_{\mathcal{M},t}$ (the reason for this was explained in the Introduction; for further details see [8], where the role of the “all state” operator is also discussed).

A formula ϕ is *valid in a model* if $\|\phi\|_{\mathcal{M}} = \Omega \times T$, that is, if ϕ is true at every state-instant pair (ω, t) . A formula ϕ is *valid in a frame* if it is valid in every model based on it. A property of frames *characterizes* (or is *characterized by*) an axiom if the axiom is valid in every frame that satisfies the property and, conversely, if the frame violates the property then there is a model based on that frame and a state-instant pair at which the axiom is falsified.

Let Δ be an abbreviation for $\neg \bigcirc \neg$ (thus $(\omega, t) \models_{\mathcal{M}} \Delta\phi$ if and only if $(\omega, t') \models_{\mathcal{M}} \phi$ for some t' such that $t \rightsquigarrow t'$),²¹ furthermore, let $\bigwedge_{j=1,\dots,m} \phi_j$ denote the formula $(\phi_1 \wedge \dots \wedge \phi_m)$. In the following proposition (which is proved in the Appendix) all the formulas are restricted to be Boolean, that is, formulas that do not contain any modal operators.

Proposition 9 *The class of AGM-consistent branching-time belief revision frames is characterized by the following axioms (in Axiom 5 we let $\phi_0 = \phi_n$ and $\chi_0 = \chi_n$):*

1. $I\phi \rightarrow B\phi$
2. $B\phi \rightarrow \neg B\neg\phi$
3. $\Delta(I\psi \wedge B\phi) \rightarrow \bigcirc(I\psi \rightarrow B\phi)$
- 4a. $(\neg B\neg\phi \wedge B\psi) \rightarrow \bigcirc(I\phi \rightarrow B\psi)$
- 4b. $\neg B\neg(\phi \wedge \neg\psi) \rightarrow \bigcirc(I\phi \rightarrow \neg B\psi)$
5.
$$\bigwedge_{j=1,\dots,n} \Delta (I\phi_j \wedge \neg B\neg\phi_{j-1} \wedge B\chi_j) \rightarrow$$

$$\bigwedge_{j=1,\dots,n} \bigcirc \left((I\phi_j \rightarrow B(\phi_{j-1} \rightarrow \chi_{j-1})) \wedge (I\phi_{j-1} \rightarrow B(\phi_j \rightarrow \chi_j)) \right)$$

Axiom 1, which corresponds to Property 1 of Definition 3, says that information is believed (if informed that ϕ the agent believes that ϕ) and Axiom 2 says

²⁰A similar (in fact, stronger) restriction is imposed in [27] (p. 175) in the context of dynamic doxastic logic.

²¹In [8] the symbol \diamond was used as a short-hand for $\neg \bigcirc \neg$. However, in temporal logic $\diamond\phi$ is used with the different meaning of ‘eventually ϕ ’. Thus, to avoid confusion, we have switched to the symbol Δ .

that beliefs are consistent (it corresponds to Property 2 of Definition 3). Axiom 3 corresponds to Property 3 of Definition 3, according to which differences in beliefs at immediate successors of an instant must be due to differences in information: if there is a next instant at which the agent is informed that ψ and believes that ϕ then at every next instant it must be the case that, if informed that ψ , she believes that ϕ . Axioms 4a and 4b provide a characterization of Property 4 of Definition 3 (the Qualitative Bayes Rule); the first says that if the agent considers ϕ possible and believes that ψ , then at any next instant at which she is informed that ϕ she must continue to believe that ψ ; the second says that if the agent considers ϕ and $\neg\psi$ possible, then at any next instant at which she is informed that ϕ she cannot believe that ψ . Axiom 5 characterizes Property *PLS* of Proposition 6, which is necessary and sufficient for local rationalizability by a total-preorder.

In the next two sections we show that branching-time belief revision frames, and the associated modal language discussed in this section, can be used to model iterated belief revision.

5 Iterated belief revision in branching-time frames

Branching-time belief revision frames provide a natural setting for studying iterated belief revision, that is, changes in beliefs prompted by a *sequence* of informational inputs. The analysis can be carried out either semantically, within the class of branching-time frames, or syntactically, within the modal language of the previous section; furthermore, the two approaches can be linked via axiomatic characterization results. In this section we will briefly discuss some of the principles of iterated belief revision that have been proposed in the literature,²² while in the next section we provide a generalization of AGM belief revision functions that captures iterated revision and discuss the correspondence between branching-time frames and iterated belief revision functions.

In an AGM-consistent frame a total pre-order associated with a state-instant pair (ω, t) (whose existence is guaranteed by Proposition 6) encodes both the agent's initial beliefs and her disposition to change those beliefs upon receipt of new information. This is what has been called in the literature an *epistemic* or *doxastic state* (see, for example, [13, 29, 34]). AGM-consistency imposes only very weak restrictions on how the epistemic state of the agent can change from (ω, t) to (ω, t') when t' is an immediate successor of t . The following lemma (proved in the Appendix) identifies one such restriction: if $E \subseteq F \subseteq \Omega$ and the agent's beliefs when informed that F do not rule out E , then she will have the same beliefs in the

²²The first analysis of iterated belief revision using the branching-time frames introduced in [8] was carried out in [38].

situation where she is immediately informed that E as in the situation where she is first informed that F and then she is informed that E .²³

Lemma 10 *Let $\mathcal{F} = \langle T, \succ, \Omega, \{\mathcal{I}_t, \mathcal{B}_t\}_{t \in T} \rangle$ be an AGM-consistent frame. Fix an arbitrary state $\omega \in \Omega$ and instants $t, t_1, t_2, t_3 \in T$ such that $t \succ t_1 \succ t_2$ and $t \succ t_3$ (that is, t_1 and t_3 are immediate successors of t and t_2 is an immediate successor of t_1). Then*

$$\begin{aligned} \text{if } \mathcal{I}_{t_3}(\omega) = \mathcal{I}_{t_2}(\omega) \subseteq \mathcal{I}_{t_1}(\omega) \text{ and } \mathcal{B}_{t_1}(\omega) \cap \mathcal{I}_{t_2}(\omega) \neq \emptyset, \\ \text{then } \mathcal{B}_{t_2}(\omega) = \mathcal{B}_{t_3}(\omega). \end{aligned} \quad (REF_{weak})$$

Note that the clause $\mathcal{B}_{t_1}(\omega) \cap \mathcal{I}_{t_2}(\omega) \neq \emptyset$ is crucial: without it the lemma is not true.²⁴ Denote by REF the strengthening of REF_{weak} obtained by dropping the clause $\mathcal{B}_{t_1}(\omega) \cap \mathcal{I}_{t_2}(\omega) \neq \emptyset$ (as before, $t, t_1, t_2, t_3 \in T$ are such that $t \succ t_1 \succ t_2$ and $t \succ t_3$ and $\omega \in \Omega$):

$$\text{if } \mathcal{I}_{t_3}(\omega) = \mathcal{I}_{t_2}(\omega) \subseteq \mathcal{I}_{t_1}(\omega), \text{ then } \mathcal{B}_{t_2}(\omega) = \mathcal{B}_{t_3}(\omega). \quad (REF)$$

Property REF states that “since the subsequent evidence is more specific than the initial evidence (that is, $\mathcal{I}_{t_2}(\omega) \subseteq \mathcal{I}_{t_1}(\omega)$), the later evidence washes away the earlier evidence” ([29], p. 197). Figure 3 shows a locally rationalizable frame that violates Property REF at (α, t) .²⁵ Consider a model based on this frame where, for some atomic formulas p, q and r , $\|p\| = \{\delta\}$, $\|q\| = \{\alpha, \gamma\} = \mathcal{I}_{t_2}(\alpha) = \mathcal{I}_{t_3}(\alpha)$ and $\|r\| = \{\gamma\}$. Then at (α, t) the agent’s disposition to revise her beliefs is such that, if informed that q (which is the case at (α, t_3)) she will believe that r . However, after being informed that $(p \vee q)$ (at (α, t_1) : $\mathcal{I}_{t_1}(\alpha) = \{\alpha, \gamma, \delta\} = \|p \vee q\|$) her disposition changes and, if later she is informed that q (which is the case at (α, t_2)), she will believe that $\neg r$ (despite the fact that information that q is a refinement of the information that $(p \vee q)$).

²³In the following lemma, $E = \mathcal{I}_{t_2}(\omega) = \mathcal{I}_{t_3}(\omega)$ and $F = \mathcal{I}_{t_1}(\omega)$. Note that, although REF_{weak} is a rather weak property and is implied by the AGM postulates, the underlying requirement for iterated belief revision is not uncontroversial: see, for example, [32, 37].

²⁴‘ REF ’ stands for ‘refinement’ (of information). Property REF_{weak} can be derived from the Qualitative Bayes Rule (Property 4 of Definition 3) and the following property, introduced in [8]:

$$\begin{aligned} \text{if } t \succ t_1, t \succ t_3, \mathcal{I}_{t_3}(\omega) \subseteq \mathcal{I}_{t_1}(\omega) \text{ and } \mathcal{B}_{t_1}(\omega) \cap \mathcal{I}_{t_3}(\omega) \neq \emptyset \\ \text{then } \mathcal{B}_{t_3}(\omega) = \mathcal{B}_{t_1}(\omega) \cap \mathcal{I}_{t_3}(\omega). \end{aligned} \quad (CAB)$$

Property CAB is valid in every branching-time belief revision frame which is rationalizable at every state-instant pair and - as shown in [9] - it is characterized by the axioms

$$\Delta(I(\phi \wedge \psi) \wedge B\chi) \rightarrow \bigcirc(I\phi \rightarrow B((\phi \wedge \psi) \rightarrow \chi)) \quad (K7)$$

$$\Delta(I\phi \wedge \neg B\neg(\phi \wedge \psi) \wedge B(\psi \rightarrow \chi)) \rightarrow \bigcirc(I(\phi \wedge \psi) \rightarrow B\chi) \quad (K8).$$

²⁵Belief revision at (α, t) is rationalized by the total pre-order generated by the strict total order $\beta P \delta P \gamma P \alpha$, while belief revision at (α, t_1) is rationalized by any total pre-order that contains the strict component $\delta P \alpha P \gamma$. Note that the ranking of α and γ has been reversed in moving from (α, t) to (α, t_1) .

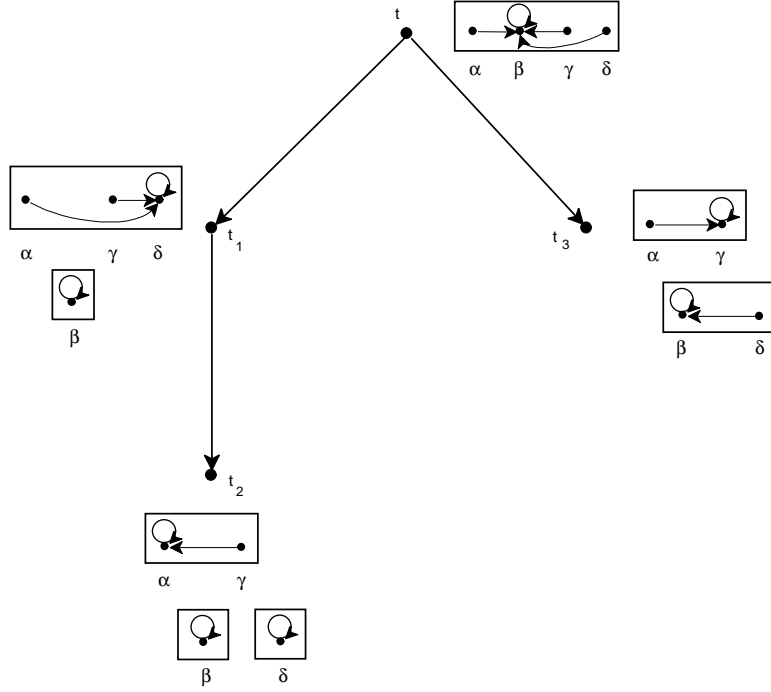


Figure 3

Although not implied by AGM-consistency, Property *REF* captures a principle that is part of most well-known theories of iterated belief revision (see, for example, [11, 12, 13, 24, 29]). It is shown in [38] that Property *REF* is characterized by the following axioms:

$$\begin{aligned}
 A(\psi \rightarrow \phi) \wedge \Delta(I\phi \wedge \Delta(I\psi \wedge B\chi)) &\rightarrow \bigcirc(I\psi \rightarrow B\chi) && \text{Ref}_1 \\
 A(\psi \rightarrow \phi) \wedge \Delta(I\psi \wedge B\chi) &\rightarrow \bigcirc(I\phi \rightarrow \bigcirc(I\psi \rightarrow B\chi)). && \text{Ref}_2
 \end{aligned}$$

A further strengthening of *REF* is given by the following property, which corresponds to the postulate ‘Conjunction’ in [29] (p. 203). It says that if two sequentially received pieces of information are consistent with each other, then they induce the same beliefs as the information consisting of their conjunction. As before let $t, t_1, t_2, t_3 \in T$ be such that $t \succ t_1 \succ t_2$ and $t \succ t_3$ and let $\omega \in \Omega$:

$$\begin{aligned}
 \text{if } \mathcal{I}_{t_2}(\omega) \cap \mathcal{I}_{t_1}(\omega) \neq \emptyset \text{ and } \mathcal{I}_{t_3}(\omega) = \mathcal{I}_{t_2}(\omega) \cap \mathcal{I}_{t_1}(\omega) \\
 \text{then } \mathcal{B}_{t_2}(\omega) = \mathcal{B}_{t_3}(\omega). &&& (\text{REF}_{strong})
 \end{aligned}$$

It is shown in [38] that Property *REF_{strong}* is characterized by the following axioms:

$$\begin{aligned} \neg A \neg(\psi \wedge \phi) \wedge \Delta(I\phi \wedge \Delta(I\psi \wedge B\chi)) &\rightarrow \bigcirc(I(\phi \wedge \psi) \rightarrow B\chi) && Ref_3 \\ \neg A \neg(\psi \wedge \phi) \wedge \Delta(I(\phi \wedge \psi) \wedge B\chi) &\rightarrow \bigcirc(I\phi \rightarrow \bigcirc(I\psi \rightarrow B\chi)). && Ref_4 \end{aligned}$$

The rationale for Property REF_{strong} is that information should be treated cumulatively in the sense that information that E followed by information that F has the same effect on beliefs as information that $E \cap F$ (provided that E and F are compatible, that is, that $E \cap F \neq \emptyset$).

Other principles of iterated belief revision that have been proposed in the literature have corresponding properties in branching-time belief revision frames and can be characterized by modal axioms similar to the ones discussed above: see [38]. Instead of continuing the discussion along these lines, in the next section we go back to the relationship between branching-time frames and AGM belief revision functions and provide a generalization of the latter that can be used to discuss principles of iterated belief revision.

6 Iterated belief revision functions

As in Section 2, let Φ be the set of formulas in the propositional language based on the set S of atomic formulas. Recall that, given a belief set $K \subseteq \Phi$, an AGM belief revision function is a function $B_K : \Phi \rightarrow 2^\Phi$ that associates with every formula $\phi \in \Phi$ (thought of as new information) a revised belief set $B_K(\phi) \subseteq \Phi$, satisfying the AGM postulates (see Definition 2). Several authors (for example [29, 34]) have discussed whether belief revision ought to be thought of as a unary operation (that is, a function taking an informational input $\phi \in \Phi$ and producing a new belief set) or as a binary operation (that is, a function taking a belief set $K \subseteq \Phi$ and an informational input $\phi \in \Phi$ and producing a new belief set). This is an issue that has been raised in the context of iterated belief revision. We propose to model iterated belief revision in terms of a *three*-argument function, that is, a ternary operation. As we shall see, our proposed functions incorporate the belief revision operations suggested in the literature and offer a clear way of stating principles of iterated revision.

Let H be the set of sequences in Φ . If $h = \langle \phi_1, \dots, \phi_n \rangle \in H$ and $\phi \in \Phi$, we denote the sequence $\langle \phi_1, \dots, \phi_n, \phi \rangle \in H$ by $h\phi$. The empty sequence $\langle \rangle$ is denoted by \emptyset and is an element of H . We think of a sequence h as a history of informational inputs received in the past and *up to the moment under consideration*. The first argument of our iterated belief revision functions is a history h . The need to take into account the history of previous informational inputs has been noted in the literature. For instance Rott ([34], p. 398) writes:

“We need to make room for a dependence of the revision function not only on the current belief state, but also on the history of belief changes (previous belief states as well as previous inputs).”

In a similar vein Nayak et al ([29], p. 202) write:

“It is conceivable that at two different times, t_1 and t_2 , an agent has the same set of beliefs but the relative firmness of the beliefs are different. If the agent accepts the same evidence at t_1 and t_2 , the resultant belief sets would be different.”

Presumably, the difference the authors refer to is attributable to the fact that the two different times t_1 and t_2 represent different ways in which the agent arrived at the same set of beliefs, that is, different past histories.

Figure 4 illustrates this possibility by means of an AGM-consistent branching-time belief revision frame.²⁶ Consider a model based on this frame where, for some atomic formulas m, p, q, r and s , $\|m\| = \{\alpha, \beta, \gamma, \delta, \varepsilon\}$, $\|p\| = \{\alpha, \beta, \gamma\}$, $\|q\| = \{\alpha, \beta, \varepsilon\}$, $\|r\| = \{\alpha, \gamma, \varepsilon\}$ and $\|s\| = \{\alpha\}$. Then the agent has the same belief set at (α, t_1) and at (α, t_2) , namely the set $K = \{\phi \in \Phi : \beta \models \phi\}$. However, the same information (at the corresponding next instant), namely that r is the case ($\mathcal{I}_{t_3}(\alpha) = \mathcal{I}_{t_4}(\alpha) = \|r\|$), leads to different beliefs: for instance at (α, t_3) she believes that s while at (α, t_4) believes that $\neg s$ ($\mathcal{B}_{t_3}(\alpha) \subseteq \|s\|$ while $\mathcal{B}_{t_4}(\alpha) \subseteq \|\neg s\|$). This difference in disposition to revise beliefs upon receiving information that r , despite the same “initial” set of beliefs K , can be traced to the different informational history leading to K : the information history at (α, t_1) is given by $\langle m, p \rangle$ while the information history at (α, t_2) is given by $\langle m, q \rangle$.

²⁶It is straightforward to check that the frame of Figure 4 is rationalizable at every state-instant pair. For example, belief revision at (α, t_0) is rationalized by the total pre-order generated by the strict total order $\delta P_{\alpha, t_0} \beta P_{\alpha, t_0} \alpha P_{\alpha, t_0} \gamma P_{\alpha, t_0} \varepsilon$, belief revision at (α, t_1) is rationalized by the total pre-order generated by the strict total order $\beta P_{\alpha, t_1} \alpha P_{\alpha, t_1} \gamma P_{\alpha, t_1} \delta P_{\alpha, t_1} \varepsilon$ and belief revision at (α, t_2) is rationalized by the total pre-order generated by the strict total order $\beta P_{\alpha, t_2} \delta P_{\alpha, t_2} \gamma P_{\alpha, t_2} \alpha P_{\alpha, t_2} \varepsilon$.




Figure 4

The other two arguments in the iterated belief revision functions are a belief set $K \subseteq \Phi$ and an informational input $\phi \in \Phi$. Let \mathbb{K} be the set of deductively closed sets of formulas.

Definition 11 An AGM iterated belief revision function is a function $B : H \times \mathbb{K} \times \Phi \rightarrow 2^\Phi$ that satisfies the AGM postulates: $\forall h \in H, \forall K \in \mathbb{K}, \forall \phi, \psi \in \Phi$

- (AGM1) $B(h, K, \phi) = [B(h, K, \phi)]^{PL}$
- (AGM2) $\phi \in B(h, K, \phi)$
- (AGM3) $B(h, K, \phi) \subseteq [K \cup \{\phi\}]^{PL}$
- (AGM4) if $\neg\phi \notin K$, then $[K \cup \{\phi\}]^{PL} \subseteq B(h, K, \phi)$
- (AGM5) $B(h, K, \phi) = \Phi$ if and only if ϕ is a contradiction
- (AGM6) if $\phi \leftrightarrow \psi$ is a tautology then $B(h, K, \phi) = B(h, K, \psi)$
- (AGM7) $B(h, K, \phi \wedge \psi) \subseteq [B(h, K, \phi) \cup \{\psi\}]^{PL}$
- (AGM8) if $\neg\psi \notin B(h, K, \phi)$, then $[B(h, K, \phi) \cup \{\psi\}]^{PL} \subseteq B(h, K, \phi \wedge \psi)$.

As noted by Nayak et al ([29], p.196) the only restriction that the AGM postulates imply concerning iterated belief revision is the one given in the following lemma, which is the counterpart of Lemma 10.

Lemma 12 Let $B : H \times \mathbb{K} \times \Phi \rightarrow 2^\Phi$ be an AGM iterated belief revision function. Then, for every $h \in H$, $K \in \mathbb{K}$, and $\phi, \psi \in \Phi$

$$\text{if } \neg\psi \notin B(h, K, \phi) \text{ then } B(h\phi, B(h, K, \phi), \psi) = B(h, K, \phi \wedge \psi). \quad (9)$$

The antecedent of (9), namely $\neg\psi \notin B(h, K, \phi)$, says that ψ is compatible with the revised belief set after information that ϕ , when the starting point is given by informational history h and belief set K ; the consequent says that the revised belief set after the further information that ψ , with new starting point given by the updated history $h\phi$ and the revised belief set $B(h, K, \phi)$, coincides with the revised belief set after information that $(\phi \wedge \psi)$, when the starting point is given by informational history h and belief set K . In short: information that ϕ followed by information that ψ produces the same beliefs as the “one step” information that $(\phi \wedge \psi)$, provided that ψ is compatible with the revised beliefs after the first piece of information, namely ϕ .

Property (9) is the counterpart of the semantic property REF_{weak} . The counterpart of the strong version of this property, namely REF_{strong} is obtained by replacing the clause ‘ $\neg\psi \notin B(h, K, \phi)$ ’ with ‘ $(\phi \wedge \psi)$ is a consistent formula’.²⁷

$$\text{if } (\phi \wedge \psi) \text{ is consistent, then } B(h\phi, B(h, K, \phi), \psi) = B(h, K, \phi \wedge \psi). \quad (10)$$

A consequence of (10) is that the order in which two consistent items of information are received is irrelevant:²⁸

$$\begin{aligned} &\text{if } (\phi \wedge \psi) \text{ is consistent,} \\ &\text{then } B(h\phi, B(h, K, \phi), \psi) = B(h\psi, B(h, K, \psi), \phi). \end{aligned} \quad (11)$$

However, (11) is weaker than (10); that is, it is possible for an AGM iterated belief revision function to satisfy (11) but not (10).

Other principles of iterated belief revision that have been proposed in the literature can easily be stated by means of AGM iterated belief revision functions. For instance, Darwiche and Pearl’s postulate DP2 ([13]; see also [29], p. 203) can be stated as follows:

$$\begin{aligned} &\text{if } (\phi \wedge \psi) \text{ is inconsistent while each of } \phi \text{ and } \psi \text{ is consistent,} \\ &\text{then } B(h\phi, B(h, K, \phi), \psi) = B(h, K, \psi). \end{aligned}$$

²⁷The counterpart of the intermediate property REF is: if ψ implies ϕ , then $B(h\phi, B(h, K, \phi), \psi) = B(h, K, \phi \wedge \psi)$.

²⁸Proof. Let $(\phi \wedge \psi)$ be a consistent formula. From (10) we get that $B(h\phi, B(h, K, \phi), \psi) = B(h, K, \phi \wedge \psi)$. Similarly, $B(h\psi, B(h, K, \psi), \phi) = B(h, K, \psi \wedge \phi)$. Since $(\phi \wedge \psi)$ is equivalent to $(\psi \wedge \phi)$, by AGM6 $B(h, K, \phi \wedge \psi) = B(h, K, \psi \wedge \phi)$. Thus $B(h\phi, B(h, K, \phi), \psi) = B(h\psi, B(h, K, \psi), \phi)$.

Rather than restating (within the framework of AGM iterated belief revision functions) the various principles of iterated revision proposed in the literature, we first comment on the philosophical issue of how revision of belief states should be modeled and then turn to the relationship between AGM iterated belief revision functions and branching-time belief revision frames.

Several authors have convincingly argued that a belief state ought to be thought of as comprising *both* the initial set of beliefs *and* the disposition to change those beliefs upon receipt of new information. As Rott ([34], p. 398) puts it,

“an [AGM] revision function does not *revise* a belief state - let alone revise all possible belief states - but *a revision function is a belief state*. Actually, a revision function does not revise anything; in particular, there are no primitive entities in the study of belief revision that could be revised by such a function. Revision functions are themselves the primitive entities of the theory of belief revision.”

Rott goes on to note that, if one accepts this point of view, then one faces the problem of how to represent the *revision of belief states*:

“If unary revision functions are primitive and the appropriate formal representation of doxastic states, how do *they* get revised by propositional inputs?” [*ibidem*]

We argue that the AGM iterated belief revision functions of Definition 11 provide an answer to this question. The function $B : H \times \mathbb{K} \times \Phi \rightarrow 2^\Phi$ can be viewed as a function that transforms a belief state and an informational input into a new belief state, as follows. A *belief state* can be taken to be a triple (h, K, b) where $h \in H$ is a history of previous informational inputs, $K \in \mathbb{K}$ is the current set of beliefs and $b(\cdot) \stackrel{def}{=} B(h, K, \cdot) : \Phi \rightarrow 2^\Phi$ is the one-step revision function obtained from $B : H \times \mathbb{K} \times \Phi \rightarrow 2^\Phi$ by fixing the values of h and K . Upon receipt of information $\phi \in \Phi$, the initial belief state (h, K, b) is transformed into the new belief state (h', K', b') where $h' = h\phi$, $K' = B(h, K, \phi)$ and $b'(\cdot) = B(h', K', \cdot) : \Phi \rightarrow 2^\Phi$.²⁹

We now turn to the relationship between branching-time belief revision frames and AGM iterated belief revision functions. For simplicity we will focus on *rooted*

²⁹Rott’s proposal in [34] is to define iterated belief revision functions as unary operations $* : H \rightarrow 2^\Phi$ taking sequences of input formulas into sets of beliefs. Such functions can be generated by the functions of our Definition 11 as follows: (1) fix a starting point (h, K) , (2) obtain from the sequence of input formulas $\langle \phi_i \rangle_{i=1, \dots, n}$ the sequence $\langle (h_i, K_i) \rangle_{i=1, \dots, n}$ where $h_i = h_{i-1}\phi_i$ and $K_i = B(h_{i-1}, K_{i-1}, \phi_i)$ and then (3) define $*(\langle \phi_i \rangle_{i=1, \dots, n}) = K_n$.

branching-time frames where there is an instant $t_0 \in T$, called the *root*, which has no immediate predecessor and is a predecessor of every other instant (that is, for every $t \in T \setminus \{t_0\}$ there is a sequence $\langle t_0, t_1, \dots, t_n \rangle$ in T such that $t_n = t$ and, for every $i = 1, \dots, n$, $t_{i-1} \succ t_i$).³⁰ Given a branching-time belief revision frame $\mathcal{F} = \langle T, \succ, \Omega, \{\mathcal{I}_t, \mathcal{B}_t\}_{t \in T} \rangle$ and a valuation $V : S \rightarrow 2^\Omega$, let $\mathcal{M} = \langle \mathcal{F}, V \rangle$ be the corresponding model. Then \mathcal{M} gives rise to a partial iterated belief revision function in a natural way, as follows. Associate with every state-instant pair (ω, t) a history $h_{\mathcal{M}, \omega, t}$ and a belief set $K_{\mathcal{M}, \omega, t}$ by letting (as before: see (1)) $K_{\mathcal{M}, \omega, t} = \{\phi \in \Phi : \mathcal{B}_t(\omega) \subseteq \|\phi\|_{\mathcal{M}}\}$ and $h_{\mathcal{M}, \omega, t}$ be the history of past informational inputs up to t , defined as follows. Let $\langle t_0, t_1, \dots, t_n \rangle$ be the path from the root t_0 to t (thus $t_n = t$) and let $\langle \mathcal{I}_{t_0}(\omega), \mathcal{I}_{t_1}(\omega), \dots, \mathcal{I}_{t_n}(\omega) \rangle$ be the corresponding sequence of sets of states reachable from ω by the information relations \mathcal{I}_{t_i} ($i = 0, 1, \dots, n$). For every $i = 0, 1, \dots, n$, let $\Phi_i = \{\phi \in \Phi, \mathcal{I}_{t_i}(\omega) = \|\phi\|_{\mathcal{M}}\}$ and let $h_{\mathcal{M}, \omega, t} = \emptyset$ (recall that \emptyset denotes the empty sequence) if $\Phi_i = \emptyset$ for every $i = 0, 1, \dots, n$, otherwise $h_{\mathcal{M}, \omega, t} = \langle \phi_1, \dots, \phi_m \rangle$ ($m \leq n + 1$) where ϕ_j is an arbitrary selection from $\Phi_j \neq \emptyset$. Finally, if $\phi \in \Phi$ is such such that $\mathcal{I}_{t'}(\omega) = \|\phi\|_{\mathcal{M}}$ for some $t' \in T$ such that $t \succ t'$, let $B(h_{\mathcal{M}, \omega, t}, K_{\mathcal{M}, \omega, t}, \phi) = \{\psi \in \Phi : \mathcal{B}_{t'}(\omega) \subseteq \|\psi\|_{\mathcal{M}}\}$.

As an illustration, consider a model \mathcal{M} based on the frame of Figure 4 where, for some atomic formulas m, p, q and r , $\|m\| = \{\alpha, \beta, \gamma, \delta, \varepsilon\}$, $\|p\| = \{\alpha, \beta, \gamma\}$, $\|q\| = \{\alpha, \beta, \varepsilon\}$ and $\|r\| = \{\alpha, \gamma, \varepsilon\}$. For simplicity we drop the subscript \mathcal{M} . Then

$$\begin{aligned} h_{\alpha, t_0} &= \langle m \rangle & K_{\alpha, t_0} &= \{\phi \in \Phi : \delta \models \phi\} \\ h_{\alpha, t_1} &= \langle m, p \rangle & K_{\alpha, t_1} &= \{\phi \in \Phi : \beta \models \phi\} \\ h_{\alpha, t_2} &= \langle m, q \rangle & K_{\alpha, t_2} &= \{\phi \in \Phi : \beta \models \phi\} \\ h_{\alpha, t_3} &= \langle m, p, r \rangle & K_{\alpha, t_3} &= \{\phi \in \Phi : \alpha \models \phi\} \\ h_{\alpha, t_4} &= \langle m, q, r \rangle & K_{\alpha, t_4} &= \{\phi \in \Phi : \gamma \models \phi\} \end{aligned}$$

and $B(h_{\alpha, t_0}, K_{\alpha, t_0}, p) = K_{\alpha, t_1}$, $B(h_{\alpha, t_0}, K_{\alpha, t_0}, q) = K_{\alpha, t_2}$, $B(h_{\alpha, t_1}, K_{\alpha, t_1}, r) = K_{\alpha, t_3}$ and $B(h_{\alpha, t_2}, K_{\alpha, t_2}, r) = K_{\alpha, t_4}$.

By Proposition 6, the partial iterated belief revision function associated with an arbitrary model based on a frame $\mathcal{F} = \langle T, \succ, \Omega, \{\mathcal{I}_t, \mathcal{B}_t\}_{t \in T} \rangle$ that is rationalizable at every state-instant pair can be extended to a full AGM iterated belief revision function. One can extend the analysis by adding to the AGM postulates appropriate postulates of iterated belief revision and identifying properties of frames that are equivalent to the existence of full AGM iterated belief revision functions

³⁰In a general branching-time frame with no root, instead of identifying a past history with the path from the root to the instant under consideration one would consider a maximal chain of predecessors of that instant.

that (1) satisfy those additional postulates and (2) extend the partial iterated revision functions obtained by interpreting the given frames. We leave this project to future research.

7 Related literature

The branching-time belief revision frames discussed in this paper provide a natural setting for a discussion of iterated belief change both semantically, in terms of property of frames, and syntactically, in terms of modal axioms. The modal logic that we considered is based on three operators: a temporal operator, a belief operator and an information operator. Instead of temporal logic, a number of authors have used dynamic modal logic to model belief revision [14, 15, 19, 27, 33, 35, 36]. This approach is known as *dynamic doxastic logic*. Despite some differences in the proposed logics, the common idea is to think of revision as a dynamic action. Besides the standard belief operator B (representing initial beliefs), these authors introduce, for every Boolean formula ϕ , a revision operator $[\ast\phi]$ with the intended interpretation of $[\ast\phi]\chi$ as “after performing the action of revising by ϕ the individual believes that χ ”. Thus these logics lack an explicit temporal operator and involve an *infinite* number of modal operators (one for each formula ϕ), while our logic uses only three operators.³¹

The branching-time belief revision frames discussed above are a generalization of the Kripke frames used in modeling static beliefs. Indeed, if one considers a sequence of instants $\langle t_1, t_2, \dots, t_n \rangle$ with $t_i \rightsquigarrow t_{i+1}$ (for each $i = 1, \dots, n - 1$), there is an associated sequence $\langle (\Omega, \mathcal{B}_{t_1}), (\Omega, \mathcal{B}_{t_2}), \dots, (\Omega, \mathcal{B}_{t_n}) \rangle$ of Kripke frames, where, for every $i = 1, \dots, n$, the Kripke frame $(\Omega, \mathcal{B}_{t_i})$ represents the agent’s beliefs at instant t_i . The representation of belief change in terms of transformation of a Kripke structure into a new Kripke structure is the key feature of the recent literature on Dynamic Epistemic Logic (DEL) [2, 3, 16, 31]. DEL is a logic based on modal operators that describe operations on Kripke models. These operations, called updates, represent events that involve information being revealed to the agents in a variety of ways, such as through a public or a private announcement. However, time does not play an explicit role in DEL and thus the DEL framework offers very limited flexibility in terms of describing beliefs through time. Another recent approach, where time enters more explicitly, is Epistemic Temporal Logic (ETL) [17, 30]. Epistemic temporal models consist of a set of histories and a binary

³¹In a similar vein, Board [7] proposes a modal logic for belief revision which also uses an infinite number of modal operators: for every formula ϕ , an operator B^ϕ is introduced, representing the hypothetical beliefs of the individual in the case where she learns that ϕ . Thus the interpretation of $B^\phi\psi$ is “upon learning that ϕ , the individual believes that ψ ”.

relation on histories representing the agent's beliefs. Thus there is a connection between epistemic accessibility and the flow of time allowing one to model such properties as memory, perfect recall, etc.³² In ETL models, however, the causes of belief change are given little structure; in particular, information does not play an explicit role. The connections between DEL and ETL have recently been clarified in [4, 5, 6] in the form of representation theorems showing how sequences of models produced by 'product update' in DEL form a special subclass of ETL models.

For further discussion of literature that is somewhat related to the approach proposed in this paper, the reader is referred to [8].

Open issues that are left for future work are (1) the extension of the branching-time belief revision frames (and the associated modal logic) to multi-agent settings, (2) a more comprehensive investigation of principles of iterated belief revision and (3) the integration of time uncertainty into the analysis (see Footnote 32).

A Appendix

First we prove the following lemma (see Footnote 10).

Lemma 13 *Let K be a consistent belief set and $B_K : \Phi \rightarrow 2^\Phi$ an AGM belief revision function. Let $\phi, \psi, \chi \in \Phi$ be such that $\chi \in B_K(\phi)$ and $\chi \in B_K(\psi)$. Then $\chi \in B_K(\phi \vee \psi)$.*

Proof. First we show that

$$(\phi \rightarrow \chi) \in B_K(\phi \vee \psi). \quad (12)$$

If $\neg\phi \in B_K(\phi \vee \psi)$ then, since - by AGM1 - $B_K(\phi \vee \psi)$ is deductively closed and $\neg\phi \rightarrow (\phi \rightarrow \chi)$ is a tautology, $(\phi \rightarrow \chi) \in B_K(\phi \vee \psi)$. If $\neg\phi \notin B_K(\phi \vee \psi)$ then, by AGM7 and AGM8, $B_K((\phi \vee \psi) \wedge \phi) = [B_K(\phi \vee \psi) \cup \{\phi\}]^{PL}$, that is, for every $\xi \in \Phi$,

$$\xi \in B_K((\phi \vee \psi) \wedge \phi) \text{ if and only if } (\phi \rightarrow \xi) \in B_K(\phi \vee \psi). \quad (13)$$

Since $(\phi \vee \psi) \wedge \phi$ is propositionally equivalent to ϕ , by AGM6 $B_K((\phi \vee \psi) \wedge \phi) = B_K(\phi)$. Thus, using (13) and the hypothesis that $\chi \in B_K(\phi)$, we get that $(\phi \rightarrow \chi) \in B_K(\phi \vee \psi)$. A similar proof leads to

³²This can be done in branching-time belief revision frames too, by adding a binary 'time-uncertainty' relation on T . This would allow one to model such phenomena as, for example, forgetting past information or being uncertain about current information.

$$(\psi \rightarrow \chi) \in B_K(\phi \vee \psi). \quad (14)$$

From (12) and (14) and the fact that $B_K(\phi \vee \psi)$ is deductively closed we obtain

$$((\phi \rightarrow \chi) \wedge (\psi \rightarrow \chi)) \in B_K(\phi \vee \psi). \quad (15)$$

Since $((\phi \rightarrow \chi) \wedge (\psi \rightarrow \chi)) \rightarrow ((\phi \vee \psi) \rightarrow \chi)$ is a tautology, it belongs to $B_K(\phi \vee \psi)$. Hence, by (15), $((\phi \vee \psi) \rightarrow \chi) \in B_K(\phi \vee \psi)$. By AGM2, $(\phi \vee \psi) \in B_K(\phi \vee \psi)$. Hence $\chi \in B_K(\phi \vee \psi)$. ■

We now turn to the proof of Proposition 6. First we need some preliminary definitions and results.

Definition 14 A choice structure is a triple $\langle \Omega, \mathcal{E}, f \rangle$ where Ω is a set, $\mathcal{E} \subseteq 2^\Omega$ is a collection of subsets of Ω and $f : \mathcal{E} \rightarrow 2^\Omega$ is a function that satisfies the following properties: $\forall E \in \mathcal{E}$, (1) $f(E) \subseteq E$ and (2) if $E \neq \emptyset$ then $f(E) \neq \emptyset$.

Give a choice structure $\mathcal{C} = \langle \Omega, \mathcal{E}, f \rangle$, a Hansson sequence in \mathcal{C} is a sequence $\langle E_0, \dots, E_n \rangle$ ($n \geq 1$) such that (1) $E_n = E_0$ and, $\forall k = 1, \dots, n$, (2) $E_k \in \mathcal{E}$ and (3) $E_{k-1} \cap f(E_k) \neq \emptyset$.

The following result is due to Hansson ([21], Theorem 7, p. 455).

Proposition 15 Let $\mathcal{C} = \langle \Omega, \mathcal{E}, f \rangle$ be a choice structure. The following are equivalent:

1. there exists a total pre-order $R \subseteq \Omega \times \Omega$ such that, for every $E \in \mathcal{E}$, $f(E) = \text{best}_R E \stackrel{\text{def}}{=} \{\omega \in E : \omega R \omega', \forall \omega' \in E\}$,
2. for every Hansson sequence $\langle E_0, \dots, E_n \rangle$ in \mathcal{C} , $E_{k-1} \cap f(E_k) = f(E_{k-1}) \cap E_k$, $\forall k = 1, \dots, n$.

As we shall see below, by Proposition 15 Property *PLS* of Proposition 6 guarantees the rationalizability of the beliefs at the immediate successors of an instant t (and some state ω). However, our definition of local rationalizability includes the initial beliefs, that is, also the beliefs at (ω, t) . Thus a little more work needs to be done in order to prove the equivalence of (b) and (c) of Proposition 6.

Definition 16 Given two choice structures $\mathcal{C} = \langle \Omega, \mathcal{E}, f \rangle$ and $\mathcal{C}' = \langle \Omega, \mathcal{E}', f' \rangle$, we say that \mathcal{C}' is a QBR-extension of \mathcal{C} by the addition of $O \subseteq \Omega$ (with $O \neq \emptyset$) if (1) $\mathcal{E}' = \mathcal{E} \cup \{O\}$, (2) f' is an extension of f , that is, $\forall E \in \mathcal{E}$, $f'(E) = f(E)$ and (3) $\forall E \in \mathcal{E}$, if $E \cap f'(O) \neq \emptyset$ then $f(E) = E \cap f'(O)$.

Lemma 17 Let $\mathcal{C} = \langle \Omega, \mathcal{E}, f \rangle$ be a choice structure and $\mathcal{C}' = \langle \Omega, \mathcal{E}', f' \rangle$, a QBR-extension of \mathcal{C} by the addition of $O \subseteq \Omega$. Then the following are equivalent:

(A) if $\langle E_0, \dots, E_n \rangle$ is a Hansson sequence in \mathcal{C} then, $\forall k = 1, \dots, n$, $E_{k-1} \cap f(E_k) = f(E_{k-1}) \cap E_k$;

(B) if $\langle E'_0, \dots, E'_n \rangle$ is a Hansson sequence in \mathcal{C}' then, $\forall k = 1, \dots, n$, $E'_{k-1} \cap f'(E'_k) = f'(E'_{k-1}) \cap E'_k$.

Proof. That (B) \Rightarrow (A) is obvious, since the set of Hansson sequences in \mathcal{C}' contains the set of Hansson sequences in \mathcal{C} (they are those where $E'_k \in \mathcal{E}$ for all k). Thus we only need to prove (A) \Rightarrow (B).

Consider first the case where, $\forall E \in \mathcal{E}$, $E \cap f'(O) \neq \emptyset$. Then, by Definition 16, $f(E) = E \cap f'(O)$, $\forall E \in \mathcal{E}$. Define the following relation R' on Ω : for all $x, y \in \Omega$, $xR'y$ if and only if either (1) $x \in f'(O)$ or (2) $x \notin f'(O)$ and $y \notin f'(O)$. Then R' is a total pre-order³³ and, furthermore, for every $E \in \mathcal{E}'$, $f'(E) = \text{best}_{R'} E$.³⁴ Thus, by Proposition 15, (B) holds.

Suppose now that $E \cap f'(O) = \emptyset$ for some $E \in \mathcal{E}$. Let $\mathcal{E}_0 = \{E \in \mathcal{E} : E \cap f'(O) = \emptyset\}$ and let $\Omega_0 = \bigcup_{E \in \mathcal{E}_0} E$. Then $\Omega_0 \cap f'(O) = \emptyset$. By Proposition 15 it follows from

(A) that there is a total pre-order R of Ω such that, for all $E \in \mathcal{E}$, $f(E) = \text{best}_R E$. Fix such a total pre-order R and define the following relation R' on Ω :

$$R' = (R \cap (\Omega_0 \times \Omega_0)) \cup \{(x, y) \in \Omega \times \Omega : x \in f'(O)\} \cup \{(x, y) \in \Omega \times \Omega : y \in \Omega \setminus (\Omega_0 \cup f'(O))\} \quad (16)$$

That is, (i) the elements of $f'(O)$ are the most plausible states, (ii) R' coincides with R on $\Omega_0 \times \Omega_0$ and (iii) the elements of $\Omega \setminus (\Omega_0 \cup f'(O))$ are the least plausible states. We want to show that R' is a total pre-order of Ω and is such that, for every $E \in \mathcal{E}'$, $f'(E) = \text{best}_{R'} E$. If we establish this then, by Proposition 15, (B) holds.

Proof that R' is complete. Fix arbitrary $x, y \in \Omega$. We need to show that either $xR'y$ or $yR'x$. If $x \in f'(O)$ then, by (16), $xR'y$; similarly, if $y \in f'(O)$ then $yR'x$. If $x, y \in \Omega_0$ then it follows from (16) and completeness of R . If $y \in \Omega \setminus (\Omega_0 \cup f'(O))$ then, by (16), $xR'y$; similarly, if $x \in \Omega \setminus (\Omega_0 \cup f'(O))$ then $yR'x$.

³³Proof of completeness. Fix arbitrary $x, y \in \Omega$. We need to show that either $xR'y$ or $yR'x$. If $x \in f'(O)$ then $xR'y$; if $y \in f'(O)$ then $yR'x$; if both $x \notin f'(O)$ and $y \notin f'(O)$ then $xR'y$ and $yR'x$.

Proof of transitivity. Fix arbitrary $x, y, z \in \Omega$ and suppose that $xR'y$ and $yR'z$. We need to show that $xR'z$. If $x \in f'(O)$, then $xR'z$. If $x \notin f'(O)$ then, since $xR'y$, it must be that $y \notin f'(O)$ and thus, since $yR'z$, it must be that also $z \notin f'(O)$. Thus $xR'z$.

³⁴By definition of R' , $\text{best}_{R'} \Omega = f'(O)$. Let $E \in \mathcal{E}$. Then, since $f(E) = E \cap f'(O) = E \cap \text{best}_{R'} \Omega$, $f(E) = \text{best}_{R'} E$ (recall that we are considering the case where, $\forall E \in \mathcal{E}$, $E \cap f'(O) \neq \emptyset$).

Proof that R' is transitive. Fix arbitrary $x, y, z \in \Omega$ and suppose that $xR'y$ and $yR'z$. We need to show that $xR'z$. If $x \in f'(O)$, then, by (16), $xR'z$. Assume that $x \notin f'(O)$. Two cases are possible: (1) $x \in \Omega_0$ and (2) $x \in \Omega \setminus (\Omega_0 \cup f'(O))$. In Case 1, since $xR'y$, it must be that either (1a) $y \in \Omega_0$ or (1b) $y \in \Omega \setminus (\Omega_0 \cup f'(O))$. In Case 1a, since $yR'z$, it must be that either $z \in \Omega_0$, in which case $xR'z$ by (16) and transitivity of R , or $z \in \Omega \setminus (\Omega_0 \cup f'(O))$, in which case $xR'z$ by (16). In Case 1b, since $yR'z$ by (16) it must be that $z \in \Omega \setminus (\Omega_0 \cup f'(O))$ and thus, by (16), $xR'z$. Consider now Case 2, where $x \in \Omega \setminus (\Omega_0 \cup f'(O))$. Then, since $xR'y$, it must be that $y \in \Omega \setminus (\Omega_0 \cup f'(O))$ and thus, since $yR'z$, it must be that also $z \in \Omega \setminus (\Omega_0 \cup f'(O))$. Hence $xR'z$ by (16).

Thus R' is a total pre-order of Ω . It remains to show that, for every $E \in \mathcal{E}'$, $f'(E) = \text{best}_{R'} E$. It is clear from (16) that $f'(O) = \text{best}_{R'} \Omega$ and thus $f'(O) = \text{best}_{R'} O$ (since, by definition of choice structure, $f'(O) \subseteq O \subseteq \Omega$). Thus we only need to show that $f(E) = \text{best}_{R'} E$ for all $E \in \mathcal{E}$. If $E \in \mathcal{E}_0$ (that is, $E \cap f'(O) = \emptyset$) then, since $f(E) = \text{best}_R E$, it follows from (16) that $f(E) = \text{best}_{R'} E$ (since R' and R coincide on $\Omega_0 \times \Omega_0$). Suppose, therefore, that $E \notin \mathcal{E}_0$, that is, $E \cap f'(O) \neq \emptyset$. Then, by Definition 16, $f(E) = E \cap f'(O)$. Hence, since $f'(O) = \text{best}_{R'} \Omega$ and $\text{best}_{R'} \Omega \cap E = \text{best}_{R'} E$ (because $\text{best}_{R'} \Omega \cap E \neq \emptyset$), it follows that $f(E) = \text{best}_{R'} E$. ■

Proof of Proposition 6. Part 1: equivalence of (b) and (c). Fix a branching-time belief revision frame $\langle T, \rightsquigarrow, \Omega, \{\mathcal{I}_t, \mathcal{B}_t\}_{t \in T} \rangle$, an arbitrary state $\hat{\omega}$ and an arbitrary instant \hat{t} . Condition *PLS* states that

$$\begin{aligned} & \forall t_0, t_1, \dots, t_n \in \hat{t}^{\rightarrow} \text{ with } t_n = t_0 \text{ and } n \geq 1, \\ & \text{if } \mathcal{I}_{t_{k-1}}(\hat{\omega}) \cap \mathcal{B}_{t_k}(\hat{\omega}) \neq \emptyset, \forall k = 1, \dots, n, \\ & \text{then } \mathcal{I}_{t_{k-1}}(\hat{\omega}) \cap \mathcal{B}_{t_k}(\hat{\omega}) = \mathcal{B}_{t_{k-1}}(\hat{\omega}) \cap \mathcal{I}_{t_k}(\hat{\omega}), \forall k = 1, \dots, n. \end{aligned} \quad (17)$$

Associate with $(\hat{\omega}, \hat{t})$ the following choice structure $\mathcal{C} = \langle \Omega, \mathcal{E}, f \rangle$: $\mathcal{E} = \{\mathcal{I}_t(\hat{\omega}) : t \in \hat{t}^{\rightarrow}\}$ and, for every $E \in \mathcal{E}$, if $E = \mathcal{I}_t(\hat{\omega})$ for some $t \in \hat{t}^{\rightarrow}$ then $f(E) = \mathcal{B}_t(\hat{\omega})$. Note that the function f is well-defined because of Property 3 of Definition 3. Then (17) can be rewritten as follows (see Definition 14):

$$\begin{aligned} & \text{for every Hansson sequence } \langle E_0, \dots, E_n \rangle \text{ in } \mathcal{C} \\ & E_{j-1} \cap f(E_j) = f(E_{j-1}) \cap E_j, \forall j = 1, \dots, n. \end{aligned} \quad (18)$$

Let $\mathcal{C}' = \langle \Omega, \mathcal{E}', f' \rangle$ be the extension of \mathcal{C} given by $\mathcal{E}' = \mathcal{E} \cup \{\mathcal{I}_{\hat{t}}(\hat{\omega})\}$ and $f'(\mathcal{I}_{\hat{t}}(\hat{\omega})) = \mathcal{B}_{\hat{t}}(\hat{\omega})$ (and, for every $E \in \mathcal{E}$, $f'(E) = f(E)$). Then, by Property 4 of Definition 3, \mathcal{C}' is a QBR extension of \mathcal{C} by the addition of $\mathcal{I}_{\hat{t}}(\hat{\omega})$ (see Definition 16). Thus, by Lemma 17, (18) is equivalent to

$$\begin{aligned} &\text{for every Hansson sequence } \langle E'_0, \dots, E'_n \rangle \text{ in } \mathcal{C}' \\ &E'_{j-1} \cap f'(E'_j) = f'(E'_{j-1}) \cap E'_j, \forall j = 1, \dots, n. \end{aligned} \quad (19)$$

By Proposition 15, (19) is equivalent to the existence of a total pre-order $\hat{R} \subseteq \Omega \times \Omega$ that rationalizes \mathcal{C}' and thus (by construction of \mathcal{C}') \hat{R} that rationalizes belief revision at $(\hat{\omega}, \hat{t})$ (that is, (b) of Proposition 6 is satisfied). ■

Remark 18 *The proofs of Proposition 15 and Lemma 17 do not require Ω to be finite. Thus the equivalence of (b) and (c) of Proposition 6 holds also in the case where Ω is infinite.*

In order to prove the equivalence of (a) and (b) of Proposition 6 we need the following.

Definition 19 *A choice structure $\langle \Omega, \mathcal{E}, f \rangle$ (see Definition 14) is called a U-choice structure ('U' because \mathcal{E} contains the universal set Ω) if (i) $\Omega \in \mathcal{E}$ and (2) $\forall E \in \mathcal{E}, E \neq \emptyset$.*

A U-choice structure $\langle \Omega, \mathcal{E}, f \rangle$ is rationalizable if there exists a total pre-order R of Ω such that, for every $E \in \mathcal{E}$, $f(E) = \text{best}_R E \stackrel{\text{def}}{=} \{\omega \in E : \omega R \omega', \forall \omega' \in E\}$.

A U-choice structure $\langle \Omega, \mathcal{E}, f \rangle$ is AGM-consistent if, for every valuation $V : S \rightarrow 2^\Omega$, the (partial) belief revision function $B_K : \Psi \rightarrow 2^\Phi$ where $K = \{\phi \in \Phi : f(\Omega) \subseteq \|\phi\|\}$, $\Psi = \{\phi \in \Phi : \|\phi\| \in \mathcal{E}\}$ and, for every $\phi \in \Psi$, $B_K(\phi) = \{\chi \in \Phi : f(\|\phi\|\|) \subseteq \|\chi\|\}$, can be extended to a full AGM belief revision function.

The following proposition is proved in [10].

Proposition 20 *A U-choice structure $\langle \Omega, \mathcal{E}, f \rangle$ with Ω finite is AGM-consistent if and only if it is rationalizable.*

We now show that a “local” application of Proposition 20, with some appropriate modifications of the choice structure associated with every state-instant pair (ω, t) , yields a proof of the equivalence between (a) and (b) of Proposition 6.

Proof of Proposition 6. Part 2: equivalence of (a) and (b). Fix a branching-time belief revision frame $\langle T, \succ, \Omega, \{\mathcal{I}_t, \mathcal{B}_t\}_{t \in T} \rangle$, with Ω finite. Fix an arbitrary state $\hat{\omega} \in \Omega$ and an arbitrary instant $\hat{t} \in T$.

Associate with $(\hat{\omega}, \hat{t})$ the following U-choice structure $\mathcal{C} = \langle \Omega, \mathcal{E}, f \rangle$: $\mathcal{E} = \{\Omega\} \cup \{\mathcal{I}_t(\hat{\omega}) : t \in \hat{t}^\rightarrow\}$, $f(\Omega) = \mathcal{B}_{\hat{t}}(\hat{\omega})$ and, for every $E \in \mathcal{E} \setminus \{\Omega\}$, if $E = \mathcal{I}_t(\hat{\omega})$ for some $t \in \hat{t}^\rightarrow$ then $f(E) = \mathcal{B}_t(\hat{\omega})$.³⁵

³⁵As noted above, the function f is well-defined because of Property 3 of Definition 3.

By construction, (a) of Proposition 6 coincides with AGM-consistency of \mathcal{C} (see Definition 19).³⁶

Next we show that (b) of Proposition 6 is equivalent to rationalizability of \mathcal{C} . Suppose that \mathcal{C} is rationalizable and let R be a total pre-order of Ω that rationalizes \mathcal{C} . Then (b.2) of Proposition 6 holds by definition of \mathcal{C} . Furthermore, $\mathcal{B}_{\hat{\omega}}(\hat{\omega}) = f(\Omega) = \text{best}_R \Omega$. Since $\mathcal{B}_{\hat{\omega}}(\hat{\omega}) \subseteq \mathcal{I}_{\hat{\omega}}(\hat{\omega})$, it follows that $\mathcal{B}_{\hat{\omega}}(\hat{\omega}) = \text{best}_R \mathcal{I}_{\hat{\omega}}(\hat{\omega})$ and thus (b.1) holds. Conversely, let R be a total pre-order of Ω that satisfies (b.1) and (b.2). Let $\mathcal{E}_0 = \{E \in \mathcal{E} \setminus \{\Omega\} : E \cap f(\Omega) = \emptyset\}$ and let $\Omega_0 = \bigcup_{E \in \mathcal{E}_0} E$. Then $\Omega_0 \cap f(\Omega) = \emptyset$. Define the following relation R' on Ω :

$$R' = (R \cap (\Omega_0 \times \Omega_0)) \cup \{(x, y) \in \Omega \times \Omega : x \in f(\Omega)\} \cup \{(x, y) \in \Omega \times \Omega : y \in \Omega \setminus (\Omega_0 \cup f(\Omega))\}. \quad (20)$$

Then R' is a total pre-order of Ω (the proof is identical to that given in Lemma 17 for (16), replacing f' with f and O with Ω). We want to show that, for every $E \in \mathcal{E}$, $f(E) = \text{best}_{R'} E$. It is clear from (20) that $f(\Omega) = \text{best}_{R'} \Omega$. Thus we only need to show that $f(E) = \text{best}_{R'} E$ for all $E \in \mathcal{E} \setminus \{\Omega\}$. If $E \in \mathcal{E}_0$ (that is, $E \cap f(\Omega) = \emptyset$) then, since $f(E) = \text{best}_R E$, it follows from (20) that $f(E) = \text{best}_{R'} E$ (since R' and R coincide on $\Omega_0 \times \Omega_0$). Suppose, therefore, that $E \notin \mathcal{E}_0$, that is, $E \cap f(\Omega) \neq \emptyset$. Then, since $f(\Omega) = \text{best}_{R'} \Omega$, $E \cap \text{best}_{R'} \Omega \neq \emptyset$ and thus $E \cap \text{best}_{R'} \Omega = \text{best}_{R'} E$. By Property 4 of Definition 3 (the Qualitative Bayes Rule), $f(E) = E \cap f(\Omega)$.³⁷ Thus $f(E) = \text{best}_{R'} E$.

Since (a) of Proposition 6 is equivalent to AGM-consistency of \mathcal{C} and (b) of Proposition 6 is equivalent to rationalizability of \mathcal{C} , the equivalence of (a) and (b) follows from Proposition 20. ■

Proof of Proposition 9. It is shown in [9] that, for $j = 1, 2$, Axiom j of Proposition 9 characterizes Property j of Definition 3.

Next we show that Axiom 3 of Proposition 9 characterizes Property 3 of Definition 3. Fix an arbitrary frame that satisfies Property 3 of Definition 3, namely if $t \succ t'$, $t \succ t''$ and $\mathcal{I}_{t'}(\omega) = \mathcal{I}_{t''}(\omega)$ then $\mathcal{B}_{t'}(\omega) = \mathcal{B}_{t''}(\omega)$. Fix arbitrary $\hat{\omega} \in \Omega$, $\hat{t} \in T$ and Boolean formulas ϕ and ψ and suppose that $(\hat{\omega}, \hat{t}) \models \Delta(I\psi \wedge B\phi)$. Then there exists a t' such that $\hat{t} \succ t'$ and $(\hat{\omega}, t') \models I\psi \wedge B\phi$, that is, $\mathcal{I}_{t'}(\hat{\omega}) = \|\psi\|_{t'}$ and $\mathcal{B}_{t'}(\hat{\omega}) \subseteq \|\phi\|_{t'}$. We have to show that $(\hat{\omega}, \hat{t}) \models \bigcirc(I\psi \rightarrow B\phi)$. Fix an arbitrary $t \in T$ such that $\hat{t} \succ t$ and suppose that $(\hat{\omega}, t) \models I\psi$. Then $\mathcal{I}_t(\hat{\omega}) = \|\psi\|_t$.

³⁶Given an arbitrary valuation $V : S \rightarrow 2^\Omega$, the initial beliefs and the partial belief revision function associated with $(\hat{\omega}, \hat{t})$ coincide with the initial beliefs and the partial belief revision function associated with \mathcal{C} .

³⁷By definition of \mathcal{C} , $f(\Omega) = \mathcal{B}_{\hat{\omega}}(\hat{\omega})$, $E = \mathcal{I}_t(\hat{\omega})$ for some t such that $\hat{t} \succ t$ and $f(E) = \mathcal{B}_t(\hat{\omega})$. By Property 4 of Definition 3, if $\mathcal{B}_{\hat{\omega}}(\hat{\omega}) \cap \mathcal{I}_t(\hat{\omega}) \neq \emptyset$ then $\mathcal{B}_t(\hat{\omega}) = \mathcal{B}_{\hat{\omega}}(\hat{\omega}) \cap \mathcal{I}_t(\hat{\omega})$.

Since ψ is a Boolean formula, by Proposition 5 in [8], $\|\psi\|_{t'} = \|\psi\|_t$. Hence $\mathcal{I}_{t'}(\hat{\omega}) = \mathcal{I}_t(\hat{\omega})$ and thus, by Property 3 of Definition 3, $\mathcal{B}_{t'}(\hat{\omega}) = \mathcal{B}_t(\hat{\omega})$. Hence $\mathcal{B}_t(\hat{\omega}) \subseteq \|\phi\|_{t'}$. Since ϕ is a Boolean formula, $\|\phi\|_{t'} = \|\phi\|_t$, so that $\mathcal{B}_t(\hat{\omega}) \subseteq \|\phi\|_t$, that is, $(\hat{\omega}, t) \models B\phi$. Hence $(\hat{\omega}, t) \models I\psi \rightarrow B\phi$ and thus, since t was chosen arbitrarily with $\hat{t} \succ t$, $(\hat{\omega}, \hat{t}) \models \bigcirc(I\psi \rightarrow B\phi)$. Conversely, fix a frame that violates Property 3 of Definition 3. Then there exist $\omega \in \Omega$ and $t, t_1, t_2 \in T$ such that $t \succ t_1, t \succ t_2, \mathcal{I}_{t_1}(\omega) = \mathcal{I}_{t_2}(\omega)$ and $\mathcal{B}_{t_1}(\omega) \neq \mathcal{B}_{t_2}(\omega)$. Without loss of generality we can assume that

$$\text{there exists an } \alpha \in \mathcal{B}_{t_2}(\omega) \text{ such that } \alpha \notin \mathcal{B}_{t_1}(\omega) \quad (21)$$

(otherwise renumber the two instants). Construct a model where, for some atomic formulas p and q , $\|p\| = \mathcal{I}_{t_1}(\omega) \times T$ and $\|q\| = \mathcal{B}_{t_1}(\omega) \times T$. Then $(\omega, t_1) \models Ip \wedge Bq$ and thus, since $t \succ t_1$, $(\omega, t) \models \Delta(Ip \wedge Bq)$. Furthermore, since $\mathcal{I}_{t_1}(\omega) = \mathcal{I}_{t_2}(\omega)$, $(\omega, t_2) \models Ip$ and, by (21), $(\omega, t_2) \not\models Bq$, so that $(\omega, t_2) \not\models (Ip \rightarrow Bq)$. Hence, since $t \succ t_2$, $(\omega, t) \not\models \bigcirc(Ip \rightarrow Bq)$ and thus Axiom 3 is falsified at (ω, t) .

It is shown in [9] that Axiom 4a of Proposition 9 (called ND in [9]) is characterized by the following property

$$\text{if } t \succ t' \text{ and } \mathcal{B}_t(\omega) \cap \mathcal{I}_{t'}(\omega) \neq \emptyset \text{ then } \mathcal{B}_{t'}(\omega) \subseteq \mathcal{B}_t(\omega) \quad (22)$$

and Axiom 4b of Proposition 9 (called NA in [9]) is characterized by the following property

$$\text{if } t \succ t' \text{ then } \mathcal{B}_t(\omega) \cap \mathcal{I}_{t'}(\omega) \subseteq \mathcal{B}_{t'}(\omega). \quad (23)$$

Since Property 4 of Definition 3 implies both (22) and (23), it follows that a frame that satisfies Property 4 validates Axioms 4a and 4b. Furthermore, in the presence of Property 1 of Definition 3, the conjunction of (22) and (23) implies Property 4. Thus, in the presence of Property 1, violation of Property 4 implies violation of either (22) or (23) (or both) and thus leads to the possibility of falsifying either Axiom 4a or Axiom 4b (or both).

We conclude the proof of Proposition 9 by showing that Axiom 5 is characterized by Property *PLS* of Proposition 6. Fix a branching-time belief revision frame that satisfies *PLS*, an arbitrary model based on it, arbitrary Boolean formulas ϕ_1, \dots, ϕ_n and χ_1, \dots, χ_n and arbitrary $\hat{\omega} \in \Omega$ and $\hat{t} \in T$ and suppose that (letting $\phi_0 = \phi_n$)

$$(\hat{\omega}, \hat{t}) \models \bigwedge_{j=1, \dots, n} \Delta (I\phi_j \wedge \neg B \neg \phi_{j-1} \wedge B\chi_j) \quad (24)$$

We have to show that, for every $j = 1, \dots, n$ (letting $\phi_0 = \phi_n$ and $\chi_0 = \chi_n$)

$$(\hat{\omega}, \hat{t}) \models \bigcirc \left((I\phi_j \rightarrow B(\phi_{j-1} \rightarrow \chi_{j-1})) \wedge (I\phi_{j-1} \rightarrow B(\phi_j \rightarrow \chi_j)) \right).$$

By (24) there exist $t_1, \dots, t_n \in \hat{t}^{\rightarrow}$ such that

$$\begin{aligned} (\hat{\omega}, t_1) &\models I\phi_1 \wedge \neg B \neg \phi_n \wedge B\chi_1 \quad (\text{recall that } \phi_0 = \phi_n) \text{ and} \\ (\hat{\omega}, t_j) &\models I\phi_j \wedge \neg B \neg \phi_{j-1} \wedge B\chi_j \quad \text{for all } j = 2, \dots, n. \end{aligned} \quad (25)$$

Thus

$$\begin{aligned} \text{(a)} \quad \mathcal{I}_{t_j}(\hat{\omega}) &= \|\phi_j\|_{t_j} \quad \text{for all } j = 1, \dots, n, \\ \text{(b)} \quad \mathcal{B}_{t_j}(\hat{\omega}) \cap \mathcal{I}_{t_{j-1}}(\hat{\omega}) &\neq \emptyset \quad \text{for all } j = 2, \dots, n, \\ \text{(c)} \quad \mathcal{B}_{t_1}(\hat{\omega}) \cap \mathcal{I}_{t_n}(\hat{\omega}) &\neq \emptyset \\ \text{(d)} \quad \mathcal{B}_{t_j}(\hat{\omega}) &\subseteq \|\chi_j\|_{t_j} \quad \text{for all } j = 1, \dots, n. \end{aligned} \quad (26)$$

Fix arbitrary $j \in \{1, \dots, n\}$ and $t \in T$ with $\hat{t} \rightsquigarrow t$. We have to show that if $(\hat{\omega}, t) \models I\phi_j$ then $(\hat{\omega}, t) \models B(\phi_{j-1} \rightarrow \chi_{j-1})$ and if $(\hat{\omega}, t) \models I\phi_{j-1}$ then $(\hat{\omega}, t) \models B(\phi_j \rightarrow \chi_j)$. Suppose first that $(\hat{\omega}, t) \models I\phi_j$, that is, $\mathcal{I}_t(\hat{\omega}) = \|\phi_j\|_t$. Since ϕ_j is a Boolean formula, by Proposition 5 in [8] (p. 148), $\|\phi_j\|_t = \|\phi_j\|_{t_j}$, so that, by (a) of (26), $\mathcal{I}_t(\hat{\omega}) = \mathcal{I}_{t_j}(\hat{\omega})$. It follows from this and Property 3 of Definition 3, that $\mathcal{B}_t(\hat{\omega}) = \mathcal{B}_{t_j}(\hat{\omega})$. Thus, without loss of generality, we can take $t = t_j$. Similarly, if $(\hat{\omega}, t) \models I\phi_{j-1}$ then, without loss of generality, we can take $t = t_{j-1}$. Thus it will be sufficient to show that if $(\hat{\omega}, t_j) \models I\phi_j$ then $(\hat{\omega}, t_j) \models B(\phi_{j-1} \rightarrow \chi_{j-1})$ and if $(\hat{\omega}, t_{j-1}) \models I\phi_{j-1}$ then $(\hat{\omega}, t_{j-1}) \models B(\phi_j \rightarrow \chi_j)$. By (b) and (c) of (26) and property *PLS* we have that (letting $t_0 = t_n$)

$$\mathcal{I}_{t_{j-1}}(\hat{\omega}) \cap \mathcal{B}_{t_j}(\hat{\omega}) = \mathcal{B}_{t_{j-1}}(\hat{\omega}) \cap \mathcal{I}_{t_j}(\hat{\omega}). \quad (27)$$

By (d) of (26), $\mathcal{B}_{t_{j-1}}(\hat{\omega}) \subseteq \|\chi_{j-1}\|_{t_{j-1}}$ and, since χ_{j-1} is a Boolean formula, by Proposition 5 in [8], $\|\chi_{j-1}\|_{t_{j-1}} = \|\chi_{j-1}\|_{t_j}$. Thus

$$\mathcal{B}_{t_{j-1}}(\hat{\omega}) \subseteq \|\chi_{j-1}\|_{t_j}. \quad (28)$$

Hence, by (27) and (28),

$$\mathcal{I}_{t_{j-1}}(\hat{\omega}) \cap \mathcal{B}_{t_j}(\hat{\omega}) \subseteq \|\chi_{j-1}\|_{t_j}. \quad (29)$$

Now (letting $\neg E$ denote the complement E , that is, $\neg E = \Omega \setminus E$),

$$\mathcal{B}_{t_j}(\hat{\omega}) \subseteq \neg \mathcal{I}_{t_{j-1}}(\hat{\omega}) \cup (\mathcal{I}_{t_{j-1}}(\hat{\omega}) \cap \mathcal{B}_{t_j}(\hat{\omega})). \quad (30)$$

By (a) of (26), $\mathcal{I}_{t_{j-1}}(\hat{\omega}) = \|\phi_{j-1}\|_{t_{j-1}}$. Since ϕ_{j-1} is a Boolean formula, $\|\phi_{j-1}\|_{t_{j-1}} = \|\phi_{j-1}\|_{t_j}$. Thus

$$\neg \mathcal{I}_{t_{j-1}}(\hat{\omega}) = \neg \|\phi_{j-1}\|_{t_j} = \|\neg\phi_{j-1}\|_{t_j}. \quad (31)$$

Putting together (30), (31) and (29) we get that $\mathcal{B}_{t_j}(\hat{\omega}) \subseteq \|\neg\phi_{j-1}\|_{t_j} \cup \|\chi_{j-1}\|_{t_j} = \|\phi_{j-1} \rightarrow \chi_{j-1}\|_{t_j}$, that is, $(\hat{\omega}, t_j) \models B(\phi_{j-1} \rightarrow \chi_{j-1})$. The proof that if $(\hat{\omega}, t_{j-1}) \models I\phi_{j-1}$ then $(\hat{\omega}, t_{j-1}) \models B(\phi_j \rightarrow \chi_j)$ is along the same lines.³⁸

Conversely, fix a frame that violates property *PLS*. Then there exist $\hat{\omega} \in \Omega$, $\hat{t} \in T$, $t_1, \dots, t_n \in \hat{t}^{\rightarrow}$, and a $k^* \in \{1, \dots, n\}$ such that (letting $t_0 = t_n$)

$$\begin{aligned} \text{(a)} \quad & \mathcal{I}_{t_{k-1}}(\omega) \cap \mathcal{B}_{t_k}(\omega) \neq \emptyset, \forall k = 1, \dots, n, \\ \text{(b)} \quad & \mathcal{I}_{t_{k^*-1}}(\hat{\omega}) \cap \mathcal{B}_{t_{k^*}}(\hat{\omega}) \neq \mathcal{B}_{t_{k^*-1}}(\hat{\omega}) \cap \mathcal{I}_{t_{k^*}}(\hat{\omega}). \end{aligned} \quad (32)$$

Let $p_1, \dots, p_n, q_1, \dots, q_n$, be atomic formulas and construct a model where, for every $k = 1, \dots, n$, $\|p_k\| = \mathcal{I}_{t_k}(\hat{\omega}) \times T$ and $\|q_k\| = \mathcal{B}_{t_k}(\hat{\omega}) \times T$. Then, by (a) of (32) (letting $p_0 = p_n$)

$$(\hat{\omega}, \hat{t}) \models \bigwedge_{j=1, \dots, n} \Delta (Ip_j \wedge \neg B \neg p_{j-1} \wedge Bq_j). \quad (33)$$

By (b) of (32), either

- (A) there is an $\alpha \in \mathcal{I}_{t_{k^*-1}}(\hat{\omega}) \cap \mathcal{B}_{t_{k^*}}(\hat{\omega})$ such that $\alpha \notin \mathcal{B}_{t_{k^*-1}}(\hat{\omega}) \cap \mathcal{I}_{t_{k^*}}(\hat{\omega})$ or
- (B) there is a $\beta \in \mathcal{B}_{t_{k^*-1}}(\hat{\omega}) \cap \mathcal{I}_{t_{k^*}}(\hat{\omega})$ such that $\beta \notin \mathcal{I}_{t_{k^*-1}}(\hat{\omega}) \cap \mathcal{B}_{t_{k^*}}(\hat{\omega})$.

Consider Case A first. Since $\alpha \in \mathcal{B}_{t_{k^*}}(\hat{\omega})$ and, by Property 1 of Definition 3, $\mathcal{B}_{t_{k^*}}(\hat{\omega}) \subseteq \mathcal{I}_{t_{k^*}}(\hat{\omega})$, it must be that $\alpha \notin \mathcal{B}_{t_{k^*-1}}(\hat{\omega})$, so that $(\alpha, t) \models \neg q_{k^*-1}$, for every $t \in T$. Since $\alpha \in \mathcal{I}_{t_{k^*-1}}(\hat{\omega})$, $(\alpha, t) \models p_{k^*-1}$, for every $t \in T$. Thus $(\alpha, t) \models \neg(p_{k^*-1} \rightarrow q_{k^*-1})$, for every $t \in T$, in particular $(\alpha, t_{k^*}) \models \neg(p_{k^*-1} \rightarrow q_{k^*-1})$. Since $\alpha \in \mathcal{B}_{t_{k^*}}(\hat{\omega})$, it follows that $(\hat{\omega}, t_{k^*}) \models \neg B(p_{k^*-1} \rightarrow q_{k^*-1})$, so that, since $(\hat{\omega}, t_{k^*}) \models Ip_{k^*}$, $(\hat{\omega}, t_{k^*}) \models \neg(Ip_{k^*} \rightarrow B(p_{k^*-1} \rightarrow q_{k^*-1}))$. It follows from this and the fact that $\hat{t} \mapsto t_{k^*}$ that $(\hat{\omega}, \hat{t}) \models \neg \bigcirc (Ip_{k^*} \rightarrow B(p_{k^*-1} \rightarrow q_{k^*-1}))$. This, together with (33) falsifies Axiom 5 of Proposition 9 at $(\hat{\omega}, \hat{t})$.

Now consider Case B. Since $\beta \in \mathcal{B}_{t_{k^*-1}}(\hat{\omega})$ and $\mathcal{B}_{t_{k^*-1}}(\hat{\omega}) \subseteq \mathcal{I}_{t_{k^*-1}}(\hat{\omega})$, it must be that $\beta \notin \mathcal{B}_{t_{k^*}}(\hat{\omega})$, so that $(\beta, t) \models \neg q_{k^*}$, for every $t \in T$. Since $\beta \in \mathcal{I}_{t_{k^*}}(\hat{\omega})$, $(\beta, t) \models p_{k^*}$, for every $t \in T$. Thus $(\beta, t) \models \neg(p_{k^*} \rightarrow q_{k^*})$, for every $t \in T$, in particular $(\beta, t_{k^*-1}) \models \neg(p_{k^*} \rightarrow q_{k^*})$. Since $\beta \in \mathcal{B}_{t_{k^*-1}}(\hat{\omega})$, it follows that $(\hat{\omega}, t_{k^*-1}) \models \neg B(p_{k^*} \rightarrow q_{k^*})$, so that, since $(\hat{\omega}, t_{k^*-1}) \models Ip_{k^*-1}$, $(\hat{\omega}, t_{k^*-1}) \models \neg(Ip_{k^*-1} \rightarrow B(p_{k^*} \rightarrow q_{k^*}))$. It follows from this and the fact that $\hat{t} \mapsto t_{k^*-1}$ that

³⁸By (d) of (26) $\mathcal{B}_{t_j}(\hat{\omega}) \subseteq \|\chi_j\|_{t_j}$ and since χ_j is Boolean, $\|\chi_j\|_{t_j} = \|\chi_j\|_{t_{j-1}}$. Thus, using (27), we get that $\mathcal{B}_{t_{j-1}}(\hat{\omega}) \cap \mathcal{I}_{t_j}(\hat{\omega}) \subseteq \|\chi_j\|_{t_{j-1}}$. Since $\mathcal{B}_{t_{j-1}}(\hat{\omega}) \subseteq \neg \mathcal{I}_{t_j}(\hat{\omega}) \cup (\mathcal{I}_{t_j}(\hat{\omega}) \cap \mathcal{B}_{t_{j-1}}(\hat{\omega}))$ and $\mathcal{I}_{t_j}(\hat{\omega}) = \|\phi_j\|_{t_j} = \|\phi_j\|_{t_{j-1}}$, it follows that $\mathcal{B}_{t_{j-1}}(\hat{\omega}) \subseteq \|\neg\phi_j\|_{t_{j-1}} \cup \|\chi_j\|_{t_{j-1}} = \|\phi_j \rightarrow \chi_j\|_{t_{j-1}}$.

$(\hat{\omega}, \hat{t}) \models \neg \bigcirc (Ip_{k^*-1} \rightarrow B(p_{k^*} \rightarrow q_{k^*}))$. This, together with (33) falsifies Axiom 5 of Proposition 9 at $(\hat{\omega}, \hat{t})$. ■

Proof of Lemma 10. First we prove that a frame which is rationalizable at every state-instant pair satisfies Property *CAB* (see Footnote 24) and then show that Property *CAB*, together with the Qualitative Bayes Rule (Property 4 of Definition 3) implies Property *REF_{weak}*. Fix $\omega \in \Omega$ and $t, t_1, t_3 \in T$ such that $t \succ t_1$, $t \succ t_3$, $\mathcal{I}_{t_3}(\omega) \subseteq \mathcal{I}_{t_1}(\omega)$ and $\mathcal{B}_{t_1}(\omega) \cap \mathcal{I}_{t_3}(\omega) \neq \emptyset$; we want to show that $\mathcal{B}_{t_3}(\omega) = \mathcal{B}_{t_1}(\omega) \cap \mathcal{I}_{t_3}(\omega)$ (this is Property *CAB*). By rationalizability at (ω, t) , there exists a total pre-order R of Ω such that $\mathcal{B}_{t_1}(\omega) = \text{best}_R \mathcal{I}_{t_1}(\omega) \stackrel{\text{def}}{=} \{\omega \in \mathcal{I}_{t_1}(\omega) : \omega R \omega', \forall \omega' \in \mathcal{I}_{t_1}(\omega)\}$ and $\mathcal{B}_{t_3}(\omega) = \text{best}_R \mathcal{I}_{t_3}(\omega) \stackrel{\text{def}}{=} \{\omega \in \mathcal{I}_{t_3}(\omega) : \omega R \omega', \forall \omega' \in \mathcal{I}_{t_3}(\omega)\}$. Since, by hypothesis, $\mathcal{I}_{t_3}(\omega) \subseteq \mathcal{I}_{t_1}(\omega)$ and $\mathcal{B}_{t_1}(\omega) \cap \mathcal{I}_{t_3}(\omega) \neq \emptyset$, $\text{best}_R \mathcal{I}_{t_3}(\omega) = \text{best}_R \mathcal{I}_{t_1}(\omega) \cap \mathcal{I}_{t_3}(\omega)$. Hence $\mathcal{B}_{t_3}(\omega) = \mathcal{B}_{t_1}(\omega) \cap \mathcal{I}_{t_3}(\omega)$.

Next we show that Property *CAB*, together with the Qualitative Bayes Rule (QBR) implies Property *REF_{weak}*. Fix $\omega \in \Omega$ and $t, t_1, t_2, t_3 \in T$ such that $t \succ t_1 \succ t_2$ and $t \succ t_3$ and suppose that $\mathcal{I}_{t_3}(\omega) = \mathcal{I}_{t_2}(\omega) \subseteq \mathcal{I}_{t_1}(\omega)$ and $\mathcal{B}_{t_1}(\omega) \cap \mathcal{I}_{t_2}(\omega) \neq \emptyset$. By QBR, since $\mathcal{B}_{t_1}(\omega) \cap \mathcal{I}_{t_2}(\omega) \neq \emptyset$, $\mathcal{B}_{t_2}(\omega) = \mathcal{B}_{t_1}(\omega) \cap \mathcal{I}_{t_2}(\omega)$. Since $\mathcal{I}_{t_3}(\omega) = \mathcal{I}_{t_2}(\omega)$, $\mathcal{B}_{t_1}(\omega) \cap \mathcal{I}_{t_3}(\omega) \neq \emptyset$ and thus, by Property *CAB*, since $\mathcal{I}_{t_3}(\omega) \subseteq \mathcal{I}_{t_1}(\omega)$, $\mathcal{B}_{t_3}(\omega) = \mathcal{B}_{t_1}(\omega) \cap \mathcal{I}_{t_3}(\omega)$. Hence $\mathcal{B}_{t_2}(\omega) = \mathcal{B}_{t_3}(\omega)$. ■

Proof of Lemma 12. Fix arbitrary $h' \in H$, $K' \in \mathbb{K}$ and $\phi, \psi \in \Phi$. By AGM3 and AGM4, if $\neg\psi \notin K'$ then $B(h', K', \psi) = [K' \cup \{\psi\}]^{PL}$. Thus, letting $h' = h\phi$ and $K' = B(h, K, \phi)$ we get

$$\text{if } \neg\psi \notin B(h, K, \phi) \text{ then } B(h\phi, B(h, K, \phi), \psi) = [B(h, K, \phi) \cup \{\psi\}]^{PL}. \quad (34)$$

By AGM7 and AGM8,

$$\text{if } \neg\psi \notin B(h, K, \phi), \text{ then } [B(h, K, \phi) \cup \{\psi\}]^{PL} = B(h, K, \phi \wedge \psi). \quad (35)$$

Thus, by (34) and (35), if $\neg\psi \notin B(h, K, \phi)$, $B(h\phi, B(h, K, \phi), \psi) = B(h, K, \phi \wedge \psi)$. ■

References

- [1] Alchourrón, Carlos, Peter Gärdenfors and David Makinson, On the logic of theory change: partial meet contraction and revision functions, *The Journal of Symbolic Logic*, 1985, 50: 510-530.

- [2] Baltag, Alexandru and Larry S. Moss, Logics for epistemic programs, *Synthese*, 2004, 139: 165-224.
- [3] van Benthem, Johan, Dynamic logics for belief revision, *Journal of Applied Non-Classical Logics*, 2007, 17: 129–155.
- [4] van Benthem, Johan and Cédric Dégrémont, Multi-agent belief dynamics: bridges between dynamic doxastic and doxastic temporal logics, in Giacomo Bonanno, Wiebe van der Hoek and Benedikt Löwe (editors), *Logic and the foundations of the theory of games and decisions (LOFT8)*, Springer, 2010, 153-175.
- [5] van Benthem Johan, Jelle Gerbrandy, Tomohiro Hoshi, and Eric Pacuit, Merging frameworks for interaction, *Journal of Philosophical Logic*, 2009, 38: 491-526.
- [6] van Benthem, Johan and Eric Pacuit, The tree of knowledge in action: towards a common perspective, in: Guido Governatori, Ian Hodkinson and Yde Venema (editors), *Advances in Modal Logic*, Vol. 6, College Publications, 2006, 87-106.
- [7] Board, Oliver, Dynamic interactive epistemology, *Games and Economic Behavior*, 2004, 49: 49-80.
- [8] Bonanno, Giacomo, Axiomatic characterization of the AGM theory of belief revision in a temporal logic, *Artificial Intelligence*, 2007, 171: 144-160.
- [9] Bonanno, Giacomo, Belief revision in a temporal framework, in Krzysztof R. Apt and Robert van Rooij (editors), *New Perspectives on Games and Interaction*, Texts in Logic and Games Series, Vol. 4, Amsterdam University Press, 2008, 45-79.
- [10] Bonanno, Giacomo, Rational choice and AGM belief revision, *Artificial Intelligence*, 2009, 173: 1194-1203.
- [11] Booth, Richard and Thomas Meyer, Admissible and restrained revision, *Journal of Artificial Intelligence Research*, 2006, 26: 127-151.
- [12] Boutilier, Craig, Iterated revision and minimal change of conditional beliefs, *Journal of Philosophical Logic*, 1996, 25: 263-305.
- [13] Darwiche, Adnan and Judea Pearl, On the logic of iterated belief revision, *Artificial Intelligence*, 1997, 89: 1-29.

- [14] van Ditmarsch, Hans, Belief change and dynamic logic, in: J. Delgrande, J. Lang, H. Rott and J.-M. Tallon (eds.), *Belief change in rational agents*, (Electronic) Proceedings of Dagstuhl Seminar 05321, IBFI, Schloss Dagstuhl, Germany, 2005 (<http://drops.dagstuhl.de/portals/index.php?semnr=05321>).
- [15] van Ditmarsch, Hans, Prolegomena to dynamic logic for belief revision, *Synthese*, 2005, 147: 229-275.
- [16] van Ditmarsch, Hans, Wiebe van der Hoek and Barteld Kooi, *Dynamic epistemic logic*, Springer, 2008.
- [17] Fagin, Ron, Joseph Halpern, Yoram Moses and Moshe Vardi, *Reasoning about knowledge*, MIT Press, 1995.
- [18] Friedman, Nir and Joseph Halpern, Belief revision: a critique, *Journal of Logic, Language, and Information*, 1999, 8: 401–420.
- [19] Fuhrmann, André, On the modal logic of theory change, in André Fuhrmann (editor), *The logic of theory change*, Lecture notes in Artificial Intelligence No. 465, Springer-Verlag, 1991, 259-281.
- [20] Gärdenfors, Peter, *Knowledge in flux: modeling the dynamics of epistemic states*, MIT Press, 1988.
- [21] Hansson, Bengt, Choice structures and preference relations, *Synthese*, 1968, 18: 443-458.
- [22] Hintikka, Jaakko, *Knowledge and belief*, Cornell University Press, 1962.
- [23] Humberstone, Lloyd, The modal logic of ‘all and only’, *Notre Dame Journal of Formal Logic*, 1987, 28: 177-188.
- [24] Jin, Yi and Michael Thielscher, Iterated belief revision, revised, *Artificial Intelligence*, 2007, 171: 1-18.
- [25] Katsuno, Hirofumi and Alberto O. Mendelzon, On the difference between updating a knowledge base and revising it, in Peter Gärdenfors (editor), *Belief revision*, Cambridge University Press, 1992, 183–203.
- [26] Kripke, S., A semantical analysis of modal logic I: normal propositional calculi, *Zeitschrift für Mathematische Logik und Grundlagen der Mathematik*, 1963, 9: 67-96.
- [27] Leitgeb, Hannes and Krister Segerberg, Dynamic doxastic logic: why, how and where to?, *Synthese*, 2007, 155: 167-190.

- [28] Levesque, Hector, All I know: a study in autoepistemic logic, *Artificial Intelligence*, 1990, 5: 263-309.
- [29] Nayak, Abhaya, Maurice Pagnucco and Pavlos Peppas, Dynamic belief revision operators, *Artificial Intelligence*, 2003, 146: 193-228.
- [30] Parikh, Rohit and R. Ramanujam, A knowledge based semantics of messages, *Journal of Logic, Language, and Information*, 2003, 12: 453-467.
- [31] Plaza Jan, Logics of public communications, in M. L. Emrich, M. S. Pfeifer, M. Hadzikadic and Z.W. Ras (editors), *Proceedings of the fourth international symposium on methodologies for intelligent systems*, Oak Ridge National Laboratory, 1989, 201–216. Reprinted in *Synthese*, 2007, 158: 165–179.
- [32] Rabinowicz, Wlodek, Stable revision, or is Preservation worth preserving?, in André Fuhrmann and Hans Rott (editors), *Logic, action and information: essays on logic in philosophy and artificial intelligence*, Berlin, 1995, 101-128.
- [33] de Rijke, Maarten, Meeting some neighbours: a dynamic modal logic meets theories of change and knowledge representation, in: J. van Eijck and A. Visser (editors), *Logic and information flow*, MIT Press, 1994, 170-196.
- [34] Rott, Hans, Coherence and conservatism in the dynamics of belief, *Erkenntnis*, 1999, 50: 387-412.
- [35] Segerberg, Krister, Belief revision from the point of view of doxastic logic, *Bulletin of the IGPL*, 1995, 3: 535-553.
- [36] Segerberg, Krister, Two traditions in the logic of belief: bringing them together, in: H. J. Ohlbach and U. Reyle (editors), *Logic, language and reasoning*, Kluwer Academic Publishers, 1999, 135-147.
- [37] Stalnaker, Robert, Iterated belief revision, *Erkenntnis*, 2009, 70: 189-209.
- [38] Zvesper, Jonathan, How to keep on changing your mind, dynamically, in Johan van Benthem, Shier Ju and Frank Veltman (editors), *A meeting of the minds. Proceedings of the workshop on Logic, Rationality and Interaction*, Texts in Computer Science, Vol. 8, College Publications, 2007, 291-306.