

Bonanno, Giacomo

Working Paper

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

Working Paper, No. 10-1

Provided in Cooperation with:

University of California Davis, Department of Economics

Suggested Citation: Bonanno, Giacomo (2010) : Belief change in branching time: AGM-consistency and iterated revision, Working Paper, No. 10-1, University of California, Department of Economics, Davis, CA

This Version is available at:

<https://hdl.handle.net/10419/58382>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

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

Abstract

We study belief change in the branching-time structures introduced in [4]. First, we identify a property of branching-time frames that is equivalent 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 [4]. 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

In [4] belief change over time was modeled by means of branching-time structures; a corresponding modal logic with operators for next-time, information and belief was proposed and some aspects of the relationship between this logic and the AGM theory of belief revision ([1]) were discussed. In this paper we establish a stronger correspondence between the semantics of branching-time frames and AGM belief revision and address the issue of iterated belief revision, both syntactically and semantically.

The addition of a valuation to a branching-time frame gives rise - for every state-instant pair (ω, t) - to an “initial” belief set K (the agent’s beliefs at (ω, t)) and a partial belief revision function based on K (constructed from the agent’s beliefs at the immediate successors of instant t and at state ω). We investigate under what conditions such a partial belief revision function can be extended to a full AGM revision function. We find 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 are the R -maximal states among the ones that are compatible with the information received. We then provide a set of axioms that characterizes this class of branching-time belief revision frames within the modal logic introduced in [4]. Finally, 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 a rationalizable 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 a doxastic state is and how an informational input transforms a doxastic state into a new one.

2 Branching-time belief revision frames

The semantic frames discussed in this section provide a way of modeling the evolution of an agent’s beliefs over time in response to informational inputs.

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. If $t \in T$ we denote the set of immediate successors of t by t^\succ , that is, $t^\succ = \{t' \in T : t \succ t'\}$.

A *branching-time belief-information frame* is a tuple $\langle T, \succ, \Omega, \{\mathcal{B}_t, \mathcal{I}_t\}_{t \in T} \rangle$ 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 1 A branching-time belief revision frame is a frame $\langle T, \succ, \Omega, \{\mathcal{B}_t, \mathcal{I}_t\}_{t \in T} \rangle$ 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.¹

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 [4], 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 $t \succ t'$, and $\mathcal{B}_{t'}(\omega) = \text{supp}(P_{\omega,t'})$ then we get Property 4. Thus in a probabilistic setting the proposition “at instant t

¹Thus we rule out inconsistent information. As pointed out by Friedman and Halpern [11], it is not clear how one could be informed of a contradiction or, at least, how one could treat a contradiction as information.

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 euclidean². 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.




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)

² \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 1 of Definition 1 is usually referred to as seriality.

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 1 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 $\mathcal{B}_{t_3}(\alpha) \cap \mathcal{I}_{t_4}(\alpha) = \mathcal{B}_{t_4}(\alpha) = \{\varepsilon\}$).

Next we relate branching-time belief revision frames to the AGM theory of belief revision ([1]), which is reviewed in the following section.³

3 AGM belief revision functions

Let Φ be the set of formulas of a propositional language based on a countable set S of atomic formulas.⁴ 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).⁵ If $\Psi \neq \Phi$ we call B_K a *partial* belief revision function, while if $\Psi = \Phi$ then B_K is called a *full* belief revision function.

Definition 2 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 $\psi \in \Psi$, $B_K^*(\psi) = B_K(\psi)$.

³For a more detailed account see [12] or [10].

⁴Thus 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 $\phi \rightarrow \psi \stackrel{def}{=} \neg\phi \vee \psi$.

⁵In 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.

Definition 3 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.⁶

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 $B_K(\phi)$ ” ([12], p. 55; notation changed to match ours).

4 Branching-time models and AGM belief revision

We now return to the semantic structures of Definition 1 and interpret them by adding a valuation that associates with every atomic proposition $p \in S$ the set of

⁶Note that, for every formula ψ , $\psi \in [K \cup \{\phi\}]^{PL}$ if and only if $(\phi \rightarrow \psi) \in K$ (since, by hypothesis, $K = [K]^{PL}$).

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 (see Footnote 4). Given a branching-time belief revision frame $\mathcal{F} = \langle T, \succ, \Omega, \{\mathcal{B}_t, \mathcal{I}_t\}_{t \in T} \rangle$, a *model based on* (or an *interpretation of*) \mathcal{F} is obtained by adding to \mathcal{F} a *valuation* $V : S \rightarrow 2^\Omega$ (where 2^Ω denotes the set of subsets of Ω).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 $\|\phi\|_{\mathcal{M}} = \{\omega \in \Omega : \omega \models_{\mathcal{M}} \phi\}$.

Definition 4 Given a model $\mathcal{M} = \langle T, \succ, \Omega, \{\mathcal{B}_t, \mathcal{I}_t\}_{t \in T}, V \rangle$, 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 (this corresponds to the notion of ‘all the agent knows’: see [4] for a discussion and references), while for belief we impose the standard requirement that $\mathcal{B}_t(\omega)$ be a *subset* of the truth set of a formula.

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

⁷In principle, the branching-time structures of Definition 1 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 [15]). We restrict attention to belief revision.

⁸If instead of belief revision we were interested in belief update (see Footnote 7), then we would need to define a valuation as a function $V : S \rightarrow 2^{\Omega \times T}$.

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

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 function. As an example we have that $\neg q, p_3 \in B_{K_{\alpha,t_0}}(p_2)$; thus, since $\neg q, \neg p_3 \in K_{\alpha,t_0}$ 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 maintaining 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 $B_{K_{\alpha,t_0}}^*$ that satisfies the AGM postulates and is an extension of $B_{K_{\alpha,t_0}}$ (see Definition 2). 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 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 $B_K(p_2) = B_K^*(p_2)$)

⁹This function is well defined because of Property 3 of Definition 1.

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,

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

Since $p_3 \in B_K(p_2)$,¹²

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

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 , $B_K^*(p_3) = B_K(p_3)$. Thus it follows from (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 property are compatible with the AGM postulates. The notion of compatibility with the AGM postulates is made precise in the following definition.

Definition 5 *A branching-time belief revision frame $\mathcal{F} = \langle T, \succ, \Omega, \{\mathcal{B}_t, \mathcal{I}_t\}_{t \in T} \rangle$ 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 2) to a full AGM belief revision function (see Definition 3).*

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, extends results given in [6] and [14]. Note that the Qualitative Bayes Rule (Property 4 of Definition 1) is crucial for the validity of Proposition 6.

by AGM7 and AGM8, $B_K^*((p_1 \vee p_2) \wedge p_2) = [B_K^*(p_1 \vee p_2) \cup \{p_2\}]^{PL} = [B_K^*(p_1 \vee p_2)]^{PL} = B_K^*(p_1 \vee p_2)$.

¹²It is straightforward to show that, for every $\phi \in \Psi$, $B_K(\phi)$ is deductively closed.

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 $E \subseteq \Omega$, let¹³

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

Thus $\text{best}_R E$ is the set of states in E that are most plausible according to R .

Proposition 6 *Let $\mathcal{F} = \langle T, \succ, \Omega, \{\mathcal{B}_t, \mathcal{I}_t\}_{t \in T} \rangle$ be a branching-time belief revision frame (see Definition 1) 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),

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

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 $\mathcal{B}_{t_1}(\alpha) = \text{best}_{R_{\alpha, t_0}} \mathcal{I}_{t_1}(\alpha) = \{\delta\}$, $\delta P_{\alpha, t_0} \gamma$; similarly, since $\delta \in \mathcal{I}_{t_2}(\alpha)$ and $\mathcal{B}_{t_2}(\alpha) = \text{best}_{R_{\alpha, t_0}} \mathcal{I}_{t_2}(\alpha) = \{\varepsilon\}$, $\varepsilon P_{\alpha, t_0} \delta$. Hence, by transitivity, $\varepsilon P_{\alpha, t_0} \gamma$. However, from $\mathcal{B}_{t_3}(\alpha) = \text{best}_{R_{\alpha, t_0}} \mathcal{I}_{t_3}(\alpha) = \{\gamma, \varepsilon\}$ we get that $\gamma R_{\alpha, t_0} \varepsilon$, yielding

¹³In 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 \leq and the set $\{\omega \in E : \omega \leq \omega', \forall \omega' \in E\}$ 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.

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 necessary and sufficient condition for a branching-time belief revision frame to be rationalizable at (ω, t) . To verify that the frame of Figure 1 is not rationalizable at (α, t_0) using this property, 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 frame is locally rationalizable if it is rationalizable at every state-instant pair (ω, t) ; it is AGM-consistent if it is AGM consistent at every (ω, t) .*

Thus, by Proposition 6, a frame where Ω is finite is locally rationalizable if and only if it is AGM-consistent.

In a locally rationalizable 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 a locally rationalizable (and thus 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 $\beta P_{\alpha, t_0} \delta P_{\alpha, t_0} \gamma P_{\alpha, t_0} \alpha$:

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

Remark 8 *In a locally rationalizable frame, it is possible that, if t' is an immediate successor of t , the plausibility ordering of Ω at (ω, t') is different from the plausibility ordering 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 ,¹⁴ whereas any total pre-order that rationalizes belief revision at (α, t_2) must be such that α is strictly more plausible than γ .¹⁵ 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 the total pre-order that rationalizes belief revision at (ω, t) may be different from the total pre-order that rationalizes belief revision at (ω', t) . For example, in Figure 2,

¹⁴Because $\mathcal{B}_{t_1}(\alpha) = \{\gamma\}$ and $\mathcal{I}_{t_1}(\alpha) = \{\alpha, \gamma\}$.

¹⁵Because $\mathcal{B}_{t_3}(\alpha) = \{\alpha\}$ 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 $\beta P_{\alpha, t_2} \alpha P_{\alpha, t_2} \delta P_{\alpha, t_2} \gamma$.

any total pre-order that rationalizes belief revision at (β, t_2) must be such that β is strictly more plausible than δ ,¹⁶ whereas any total pre-order that rationalizes belief revision at (δ, t_2) must be such that δ is strictly more plausible than β .¹⁷




Figure 2

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

5 A temporal logic for belief revision

We briefly review the modal language introduced in [4], which 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:

¹⁶Because $\mathcal{B}_{t_3}(\beta) = \{\beta\}$ and $\mathcal{I}_{t_3}(\beta) = \{\beta, \delta\}$.

¹⁷Because $\mathcal{B}_{t_3}(\delta) = \{\delta\}$ 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 $\delta P_{\delta, t_2} \beta P_{\delta, t_2} \gamma P_{\delta, t_2} \alpha$.

- $\bigcirc\phi$: “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 ”.

Fix a model $\mathcal{M} = \langle T, \succ, \Omega, \{\mathcal{B}_t, \mathcal{I}_t\}_{t \in T}, V \rangle$, 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 (ω, t) is defined as usual for $p \in S$ (where S is the set of atomic formulas), $\neg\phi$ and $(\phi \vee \psi)$. For the modal formulas we have

- $(\omega, t) \models_{\mathcal{M}} \bigcirc\phi$ if and only if $(\omega, t') \models_{\mathcal{M}} \phi$ for every t' such that $t \succ t'$
- $(\omega, t) \models_{\mathcal{M}} B\phi$ if and only if $\mathcal{B}_t(\omega) \subseteq \|\phi\|_{\mathcal{M}, t}$
- $(\omega, t) \models_{\mathcal{M}} I\phi$ if and only if $\mathcal{I}_t(\omega) = \|\phi\|_{\mathcal{M}, t}$
- $(\omega, t) \models_{\mathcal{M}} A\phi$ if and only if $\|\phi\|_{\mathcal{M}, t} = \Omega$.

Note 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: $\mathcal{I}_t(\omega) = \|\phi\|_{\mathcal{M}, t}$ (for an explanation see [4], where the role of the “all state” operator is also discussed). Note also that, while the other modal operators apply to arbitrary formulas, the information operator is restricted to apply only to pure Boolean formulas, that is formulas that do not contain any modal operators.¹⁸ Pure Boolean formulas represent facts and information is restricted to be about facts.

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 \diamond be an abbreviation for $\neg \bigcirc \neg$ (thus $(\omega, t) \models_{\mathcal{M}} \diamond\phi$ if and only if $(\omega, t') \models_{\mathcal{M}} \phi$ for some t' such that $t \succ t'$); 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 pure Boolean, that is, formulas that do not contain any modal operators.

¹⁸A similar (in fact, stronger) restriction is imposed in [16] in the context of dynamic doxastic logic (p. 175).

Proposition 9 *The class of locally rationalizable 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. $\Diamond(I\psi \wedge B\phi) \rightarrow \bigcirc(I\psi \rightarrow B\phi)$
- 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. $\bigwedge_{j=1,\dots,n} \Diamond(I\phi_j \wedge \neg B\neg\phi_{j-1} \wedge B\chi_j) \rightarrow$
 $\bigwedge_{j=1,\dots,n} \bigcirc\left(\left(I\phi_j \rightarrow B(\phi_{j-1} \rightarrow \chi_{j-1})\right) \wedge \left(I\phi_{j-1} \rightarrow B(\phi_j \rightarrow \chi_j)\right)\right)$

Axiom 1 says that information is believed and Axiom 2 that beliefs are consistent. Axiom 3 corresponds to Property 3 of Definition 1. Axioms 4a and 4b provide a characterization of Property 4 of Definition 1 (the Qualitative Bayes Rule). Axiom 5 characterizes Property *PLS* of Proposition 6.

6 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,¹⁹ while in the next section we provide a generalization of AGM belief revision functions (see Definition 3) that captures iterated revision and discuss the correspondence between branching-time frames and iterated belief revision functions.

In a locally rationalizable frame the total pre-order associated with a state-instant pair (ω, t) 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, [9], [17], [19]). AGM-consistency (which, by Proposition 6, is equivalent to local rationalizability)

¹⁹The first analysis of iterated belief revision using the branching-time frames introduced in [4] was carried out by Zvesper [21].

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 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 .²⁰

Lemma 10 *Let $\mathcal{F} = \langle T, \succ, \Omega, \{\mathcal{B}_t, \mathcal{I}_t\}_{t \in T} \rangle$ be a locally rationalizable 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*²¹

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

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. 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, 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 $\omega \in \Omega$:

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

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” ([17], p. 197). Figure 3 shows a locally rationalizable frame that violates Property REF at (α, t) .²² 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

²⁰In 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, [20] and [18].

²¹ REF stands for ‘refinement’ (of information). Property REF_{weak} can be derived from the Qualitative Bayes Rule (Property 4 of Definition 1) and the following property, introduced in [4]:

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

Property CAB is valid in every locally rationalizable frame and, as shown in [5], it is characterized by the axioms

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

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

²²Belief 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) .

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 $(p \vee q)$).




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, [7], [8], [9], [13], [17]). It is shown in [21] that Property *REF* is characterized by the following axioms:

$$\begin{aligned}
 A(\psi \rightarrow \phi) \wedge \diamond(I\phi \wedge \diamond(I\psi \wedge B\chi)) &\rightarrow \bigcirc(I\psi \rightarrow B\chi) && Ref_1 \\
 A(\psi \rightarrow \phi) \wedge \diamond(I\psi \wedge B\chi) &\rightarrow \bigcirc(I\phi \rightarrow \bigcirc(I\psi \rightarrow B\chi)) && Ref_2
 \end{aligned}$$

A further strengthening of *REF* is given by the following property, which corresponds to the postulate ‘Conjunction’ in Nayak et al ([17], 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 $\omega \in \Omega$:

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

It is shown in [21] that Property REF_{strong} is characterized by the following axioms:

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

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 [21]. 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.

7 Iterated belief revision functions

As in Section 3, let Φ be the set of formulas in a 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 3). Several authors 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).²³ 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

²³See, for example, [17] and [19].

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 ([19], 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 ([17], 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 frame.²⁴ 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_1) is given by $\langle m, q \rangle$.

²⁴It is straightforward to check that the frame of Figure 4 is locally rationalizable. 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 $\beta P_{\alpha, t_2} \gamma P_{\alpha, t_2} \alpha P_{\alpha, t_2} \delta P_{\alpha, t_2} \varepsilon$.




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 ([17], 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 update 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 .

(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’.²⁵

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

A consequence of (10) is that the order in which two consistent items of information are received is irrelevant:²⁶

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

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 ([9]; see also [17], p. 203) can be stated as follows:

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

²⁵The counterpart of the intermediate property REF is: if ψ implies ϕ , then $B(h\phi, B(h, K, \phi), \psi) = B(h, K, \phi \wedge \psi)$.

²⁶Proof. 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 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 ([19], 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. A *belief state* can be taken to be a triple $(h, K, b : \Phi \rightarrow 2^\Phi)$ where h is a history of previous informational inputs, 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 $b'(\cdot) = B(h\phi, B(h, K, \phi), \cdot) : \Phi \rightarrow 2^\Phi$.²⁷

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* 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,

²⁷Rott’s proposal in [19] 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 $*(\langle \phi_i \rangle_{i=1, \dots, n}) = K_n$.

for every $i = 1, \dots, n$, $t_{i-1} \rightsquigarrow t_i$.²⁸ Given a frame $\mathcal{F} = \langle T, \rightsquigarrow, \Omega, \{\mathcal{B}_t, \mathcal{I}_t\}_{t \in T} \rangle$ and a valuation $V : S \rightarrow 2^\Omega$, let \mathcal{M} be the corresponding model. Then \mathcal{M} gives rise to a partial iterated belief revision function in a natural way. 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 and let $\langle \mathcal{I}_{t_0}(\omega), \mathcal{I}_{t_1}(\omega), \dots, \mathcal{I}_{t_n}(\omega) \rangle$. 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 $\Phi_{j-1} \neq \emptyset$. Finally, if $\phi \in \Phi$ is such such that $\mathcal{I}_{t'}(\omega) = \|\phi\|_{\mathcal{M}}$ for some $t' \in T$ such that $t \rightsquigarrow t'$, 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

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

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 $\mathcal{F} = \langle T, \rightsquigarrow, \Omega, \{\mathcal{B}_t, \mathcal{I}_t\}_{t \in T} \rangle$ that is locally rationalizable 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 that (i) satisfy those additional postulates and (ii) extend the partial iterated revision functions obtained by interpreting the given frames. We leave this project to future research.

²⁸In 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.

8 Concluding remarks

The branching-time frames discussed in this paper provide a natural setting for a discussion of belief change both semantically, in terms of property of frames, and syntactically, in terms of modal axioms.²⁹ Furthermore, a correspondence between interpretations of branching-time frames and AGM belief revision functions can also be established, thereby providing a link to the vast literature on belief revision and iterated belief revision.

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

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

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 $\chi \in B_K(\phi \vee \psi)$. ■

²⁹For a different, but related, approach see [2] and [3]

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 ([14], 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-order³⁰ and, furthermore, for every $E \in \mathcal{E}'$, $f'(E) = \text{best}_{R'} E$.³¹ 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 $yR'x$.

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).

³⁰Proof 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 $xR'z$.

³¹By 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, $\forall E \in \mathcal{E}, E \cap f'(O) \neq \emptyset$).

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 $\langle T, \succ, \Omega, \{\mathcal{B}_t, \mathcal{I}_t\}_{t \in T} \rangle$, 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 1. 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 $\mathcal{E}' = \mathcal{E} \cup \{\mathcal{I}_t(\hat{\omega})\}$ and $f'(\mathcal{I}_t(\hat{\omega})) = \mathcal{B}_t(\hat{\omega})$. Then, by Property 4 of Definition 1, \mathcal{C}' is a QBR extension of \mathcal{C} by the addition of $\mathcal{I}_t(\hat{\omega})$ (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 $\phi \in \Psi$, $B_K(\phi) = \{\psi \in \Phi : f(\|\phi\|) \subseteq \|\psi\|\}$, can be extended to a full AGM belief revision function.

The following proposition is proved in [6].

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.

A “pointwise” 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 $\langle T, \succ, \Omega, \{\mathcal{B}_t, \mathcal{I}_t\}_{t \in T} \rangle$, 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 $f(E) = \mathcal{B}_t(\hat{\omega})$.³²

By construction, (a) of Proposition 6 is equivalent to AGM-consistency of \mathcal{C} (see Definition 19).³³

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, $\mathcal{B}_{\hat{t}}(\hat{\omega}) = f(\Omega) = \text{best}_R \Omega$. Since $\mathcal{B}_{\hat{t}}(\hat{\omega}) \subseteq \mathcal{I}_{\hat{t}}(\hat{\omega})$, it follows that $\mathcal{B}_{\hat{t}}(\hat{\omega}) = \text{best}_R \mathcal{I}_{\hat{t}}(\hat{\omega})$ and thus (b.1) holds. Conversely, let R be a total pre-order of Ω that satisfies (b.1)

³²As noted above, the function f is well-defined because of Property 3 of Definition 1.

³³Given 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} .

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 1 (the Qualitative Bayes Rule), $f(E) = E \cap f(\Omega)$.³⁴ 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 [5] that, for $j = 1, 2$, Axiom j of Proposition 9 characterizes Property j of Definition 1.

Next we show that Axiom 3 of Proposition 9 characterizes Property 3 of Definition 1. Fix an arbitrary frame that satisfies Property 3 of Definition 1, 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 pure Boolean formulas ϕ and ψ and suppose that $(\hat{\omega}, \hat{t}) \models \diamond(I\psi \wedge B\phi)$. 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 $(\hat{\omega}, \hat{t}) \models \circ(I\psi \rightarrow B\phi)$. Fix an arbitrary $t \in T$ such that $\hat{t} \succ t$ and suppose that $(\hat{\omega}, t) \models I\psi$. Then $\mathcal{I}_t(\hat{\omega}) = \|\psi\|_t$. Since ψ is a pure Boolean formula, by Proposition 5 in [4], $\|\psi\|_{t'} = \|\psi\|_t$. Hence $\mathcal{I}_{t'}(\hat{\omega}) = \mathcal{I}_t(\hat{\omega})$ and thus, by Property 3 of Definition 1, $\mathcal{B}_{t'}(\hat{\omega}) = \mathcal{B}_t(\hat{\omega})$. 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 $\hat{t} \succ t$, $(\hat{\omega}, \hat{t}) \models \circ(I\psi \rightarrow B\phi)$. Conversely, fix a frame that violates Property 3 of Definition 1. 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

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

³⁴By definition of \mathcal{C} , $f(\Omega) = \mathcal{B}_{\hat{t}}(\hat{\omega})$, $E = \mathcal{I}_t(\hat{\omega})$ for some t such that $\hat{t} \succ t$ and $f(E) = \mathcal{B}_t(\hat{\omega})$. By Property 4 of Definition 1, if $\mathcal{B}_{\hat{t}}(\hat{\omega}) \cap \mathcal{I}_t(\hat{\omega}) \neq \emptyset$ then $\mathcal{B}_t(\hat{\omega}) = \mathcal{B}_{\hat{t}}(\hat{\omega}) \cap \mathcal{I}_t(\hat{\omega})$.

(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 $t \succ t_1$, $(\omega, t) \models \diamond(Ip \wedge Bq)$. 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 [5] that Axiom 4a of Proposition 9 (called ND in [5]) 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 [5]) 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 1 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 1, 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 pure 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 $\phi_0 = \phi_n$)

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

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

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

Fix arbitrary $j \in \{1, \dots, n\}$ and $t \in T$ with $\hat{t} \succ t$. 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, $\mathcal{I}_t(\hat{\omega}) = \|\phi_j\|_{t_j}$. Since ϕ_j is a pure Boolean formula, by Proposition 5 in [4], $\|\phi_j\|_t = \|\phi_j\|_{t_j}$, 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 1, 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 pure Boolean formula, by Proposition 5 in [4], $\|\chi_{j-1}\|_{t_{j-1}} = \|\chi_{j-1}\|_{t_j}$. 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.³⁵

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 $t_0 = t_n$)

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

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 $p_0 = p_n$)

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

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

³⁵By (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 $\mathcal{B}_{t_{j-1}}(\hat{\omega}) \subseteq \|\neg\phi_j\|_{t_{j-1}} \cup \|\chi_j\|_{t_{j-1}} = \|\phi_j \rightarrow \chi_j\|_{t_{j-1}}$.

Proof of Lemma 10. First we prove that every locally rationalizable frame satisfies Property *CAB* (see Footnote 21) and then show that Property *CAB*, together with the Qualitative Bayes Rule (Property 4 of Definition 1) 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 local rationalizability, 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] van Benthem, Johan, Dynamic logics for belief change, *Journal of applied non-classical logics*, 2007, 17: 129–155.

- [3] van Benthem, Johan and Cédric Dégrémont, Multi-agent belief dynamics: bridges between dynamic doxastic and doxastic temporal logics, in Bonanno Giacomo, Wiebe van der Hoek and Benedikt Löwe (editors), *Logic and the foundations of the theory of games and decisions (LOFT8)*, Springer, 2010, 153-175.
- [4] Bonanno, Giacomo, Axiomatic characterization of the AGM theory of belief revision in a temporal logic, *Artificial Intelligence*, 2007, 171: 144-160.
- [5] 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.
- [6] Bonanno, Giacomo, Rational choice and AGM belief revision, *Artificial Intelligence*, 2009, 173: 1194-1203.
- [7] Booth, Richard and Thomas Meyer, Admissible and restrained revision, *Journal of Artificial Intelligence Research*, 2006, 26:127-151.
- [8] Boutilier, Craig, Iterated revision and minimal change of conditional beliefs, *Journal of Philosophical Logic*, 1996, 25: 263-305.
- [9] Darwiche, Adnan and Judea Pearl, On the logic of iterated belief revision, *Artificial Intelligence*, 1997, 89: 1-29.
- [10] van Ditmarsch, Hans, Wiebe van der Hoek and Barteld Kooi, *Dynamic epistemic logic*, Springer, 2008.
- [11] Friedman, Nir and Joseph Halpern, Belief revision: a critique, *Journal of Logic, Language, and Information*, 1999, 8: 401-420.
- [12] Gärdenfors, Peter, *Knowledge in flux: modeling the dynamics of epistemic states*, MIT Press, 1988.
- [13] Jin, Yi and Michael Thielscher, Iterated belief revision, revised, *Artificial Intelligence*, 2007, 171: 1-18.
- [14] Hansson, Bengt, Choice structures and preference relations, *Synthese*, 1968, 18: 443-458.
- [15] 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.

- [16] Leitgeb, Hannes and Krister Segerberg, Dynamic doxastic logic: why, how and where to?, *Synthese*, 2007, 155: 167-190.
- [17] Nayak, Abhaya, Maurice Pagnucco and Pavlos Peppas, Dynamic belief revision operators, *Artificial Intelligence*, 2003, 146: 193-228.
- [18] 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.
- [19] Rott, Hans, Coherence and conservatism in the dynamics of belief, *Erkenntnis*, 1999, 50: 387-412.
- [20] Stalnaker, Robert, Iterated belief revision, *Erkenntnis*, 2009, 70: 189-209.
- [21] 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.