Belief revision in a temporal framework

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

June 2, 2006

Extended Abstract for LOFT 2006

1 Introduction

The theory of belief revision deals with (rational) changes in beliefs in response to new information. In the literature a distinction has been drawn between belief revision and belief update (see [6]). The former deals with situations where the objective facts describing the world do not change (so that only the beliefs of the agent change over time), while the letter allows for situations where both the facts and the doxastic state of the agent change over time. We focus on belief revision and propose a temporal framework that allows for iterated revision. We model the notion of "minimal" or "conservative" belief revision by considering logics of increasing strength. We move from one logic to the next by adding one or more axioms and show that the corresponding logic captures more stringent notions of minimal belief revision. The strongest logic that we propose provides a full axiomatization of the well-known AGM theory of belief revision.

2 The basic logic

We consider a propositional language with five modal operators: the next-time operator $\bigcirc$ and it inverse $\bigcirc^{-1}$, the belief operator B, the information operator I and the "all state" operator A. The intended interpretation is as follows:

$\bigcirc\phi$:	"at every next instant it will be the case that ϕ"
$\bigcirc^{-1}\phi$:	"at every previous instant it was the case that ϕ"
$B\phi$:	"the agent believes that ϕ"
$I\phi$:	"the agent is informed that ϕ"
$A\phi$:	"it is true at every state that ϕ".

The formal language is built in the usual way from a countable set S of sentence letters (or atomic propositions), the connectives $\neg$ and $\vee$ (from which the connectives $\wedge$, $\rightarrow$ and $\leftrightarrow$ are defined as usual) and the modal operators. Let $\Diamond\phi \overset{def}{=} \neg \bigcirc \neg\phi$, and $\Diamond^{-1}\phi \overset{def}{=} \neg \bigcirc^{-1} \neg\phi$. Thus the interpretation of $\Diamond\phi$ is "at *some* next instant it will be the case that ϕ " while the interpretation of $\Diamond^{-1}\phi$ is "at some previous instant it was the case that ϕ".

We denote by $\mathbb{L}_0$ the basic logic of belief revision defined by the following axioms and rules of inference.

AXIOMS:

1. All propositional tautologies.
2. Axiom K for $\Box \in \{\bigcirc, \bigcirc^{-1}, B, A\}$:
   $\Box(\phi \rightarrow \psi) \rightarrow (\Box\phi \rightarrow \Box\psi)$.
3. Temporal axioms relating $\bigcirc$ and $\bigcirc^{-1}$:
   $\phi \rightarrow \bigcirc\Diamond^{-1}\phi$ and $\phi \rightarrow \bigcirc^{-1}\Diamond\phi$.
4. Backward Uniqueness axiom:
   $\Diamond^{-1}\phi \rightarrow \bigcirc^{-1}\phi$.
5. S5 axioms for A:
   $A\phi \rightarrow \phi$ and $\neg A\phi \rightarrow A\neg A\phi$.
6. Inclusion axiom for B (note the absence of an analogous axiom for I):
   $A\phi \rightarrow B\phi$.
7. Axioms to capture the non-standard semantics for I:
   $(I\phi \wedge I\psi) \rightarrow A(\phi \leftrightarrow \psi)$ and $A(\phi \leftrightarrow \psi) \rightarrow (I\phi \leftrightarrow I\psi)$.

RULES OF INFERENCE:

1. Modus Ponens: $\frac{\phi, \quad \phi\rightarrow\psi}{\psi}$.
2. Necessitation for $\Box \in \{\bigcirc, \bigcirc^{-1}, A\}$: $\frac{\phi}{\Box\phi}$.

3 The semantics

A *next-time branching frame* is a pair $\langle T, \rightarrowtail \rangle$ where T is a (possibly infinite) set of instants or dates and $\rightarrowtail$ is a binary "precedence" relation on T satisfying the following properties: $\forall t_1, t_2, t_3 \in T$,

(1)	uniqueness	if $t_1 \rightarrowtail t_3$ and $t_2 \rightarrowtail t_3$ then $t_1 = t_2$
(2)	acyclicity	if $\langle t_1, ..., t_n \rangle$ is a sequence with $t_i \rightarrowtail t_{i+1}$ for every $i = 1, ..., n-1$, then $t_n \neq t_1$.

The interpretation of $t_1 \rightarrowtail t_2$ 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.

A *temporal belief revision frame* is a quintuple $\langle T, \rightarrowtail, \Omega, \{\mathcal{B}_t\}_{t\in T}, \{\mathcal{I}_t\}_{t\in T}\rangle$ where $\langle T, \rightarrowtail\rangle$ is a next-time branching frame, Ω is a set of states (or possible worlds) and, for every $t \in T$, $\mathcal{B}_t$ and $\mathcal{I}_t$ are binary relations on Ω.

The interpretation of $\omega\mathcal{B}_t\omega'$ is that at state ω and time t the individual considers state ω' possible (an alternative expression is "ω' is a doxastic alternative to ω at time t"). On the other hand, 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'. We shall use the following notation: $\mathcal{B}_t(\omega) = \{\omega' \in \Omega : \omega\mathcal{B}_t\omega'\}$ and, similarly, $\mathcal{I}_t(\omega) = \{\omega' \in \Omega : \omega\mathcal{I}_t\omega'\}$.

Given a temporal belief revision frame one obtains a *model based on it* by adding a function $V : S \to 2^{\Omega}$ (where 2^{Ω} denotes the set of subsets of Ω) that associates with every atomic proposition q the set of states at which q is true. Note that defining a valuation this way is what frames the problem as one of belief revision, since the truth value of an atomic proposition q depends only on the state and not on the time (belief update would require a valuation to be defined as a function $V : S \to 2^{\Omega\times T}$). Given a model, a state ω, an instant t and a formula ϕ, we write $(\omega, t) \vDash \phi$ to denote that ϕ is true at state ω and time t. Let $\lceil\phi\rceil_t \subseteq \Omega$ denote the set of states at which ϕ is true at time t, that is, $\lceil\phi\rceil_t = \{\omega \in \Omega : (\omega, t) \vDash \phi\}$. Truth at a pair (ω, t) is defined recursively as follows.

if $q \in S$,	$(\omega, t) \vDash q$ if and only if $\omega \in V(q)$.
$(\omega, t) \vDash \neg\phi$	if and only if $(\omega, t) \nvDash \phi$.
$(\omega, t) \vDash \phi \vee \psi$	if and only if either $(\omega, t) \vDash \phi$ or $(\omega, t) \vDash \psi$ (or both).
$(\omega, t) \vDash \bigcirc\phi$	if and only if $(\omega, t') \vDash \phi$ for every t' such that $t \rightarrowtail t'$.
$(\omega, t) \vDash \bigcirc^{-1}\phi$	if and only if $(\omega, t'') \vDash \phi$ for every t'' such that $t'' \rightarrowtail t$.
$(\omega, t) \vDash B\phi$	if and only if $\mathcal{B}_t(\omega) \subseteq \lceil\phi\rceil_t$, that is, if $(\omega', t) \vDash \phi$ for all $\omega' \in \mathcal{B}_t(\omega)$.
$(\omega, t) \vDash I\phi$	if and only if $\mathcal{I}_t(\omega) = \lceil\phi\rceil_t$, that is, if (1) $(\omega', t) \vDash \phi$ for all $\omega' \in \mathcal{I}_t(\omega)$, and (2) if $(\omega', t) \vDash \phi$ then $\omega' \in \mathcal{I}_t(\omega)$.
$(\omega, t) \vDash A\phi$	if and only if $\lceil\phi\rceil_t = \Omega$, that is, if $(\omega', t) \vDash \phi$ for all $\omega' \in \Omega$.

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 $\mathcal{I}_t(\omega) \subseteq \lceil\phi\rceil_t$ we require equality, $\mathcal{I}_t(\omega) = \lceil\phi\rceil_t$. Thus our information operator is similar to the "only knowing" operator discussed in the computer science literature (see [7]). The non-normality of the I operator is what makes it necessary to add the "all state" operator A (see [5]).

We say that *a logic is characterized by a class of frames* if every theorem of the logic is valid in every frame in that class and, conversely, if for every frame that does not belong to that class there is a theorem of the logic which is falsified in a model based on that frame.

Our purpose is to model how the *factual* beliefs of an individual change over time in response to *factual* information. Thus the axioms we introduce are restricted to *Boolean formulas*, which are formulas that do not contain any modal operators. That is, Boolean formulas are defined recursively as follows: (1) every atomic proposition is a Boolean formula, and (2) if ϕ and ψ are Boolean formulas then so are $\neg\phi$ and $(\phi \vee \psi)$. As the following proposition shows, the truth value of a Boolean formula does not change over time: it is only a function of the state.

Proposition 1 *Let ϕ be a Boolean formula. Fix an arbitrary model and suppose that $(\omega, t) \models \phi$. Then, for every $t' \in T$, $(\omega, t') \models \phi$.*

We now turn to axioms of increasing strength that express the notion of minimal change of beliefs.

The first axiom says that if ϕ and ψ are facts (Boolean formulas) and - currently - the agent believes that ϕ and also believes that ψ and his belief that ϕ is non-trivial (in the sense that he considers ϕ possible) then - at every next instant - if he is informed that ϕ it will still be the case that he believes that ψ. That is, if at a next instant he is informed of some fact that he currently believes, then he *cannot drop* any of his current factual beliefs ('W' stands for 'Weak' and 'ND' for 'No Drop'): if ϕ and ψ are Boolean,

$$(B\phi \wedge \neg B\neg\phi \wedge B\psi) \rightarrow \bigcirc(I\phi \rightarrow B\psi). \qquad (WND)$$

The second axiom says that if ϕ and ψ are facts (Boolean formulas) and - currently - the agent believes that ϕ and does not believe that ψ, then - at every next instant - if he is informed that ϕ it will still be the case that he does not believe that ψ. That is, at any next instant at which he is informed of some fact that he believes now he *cannot add* a factual belief that he does not have now ('NA' stands for 'No Add'): if ϕ and ψ are Boolean,

$$(B\phi \wedge \neg B\psi) \rightarrow \bigcirc(I\phi \rightarrow \neg B\psi). \qquad (WNA)$$

Thus, by (WND), no belief can be dropped and, by (WNA), no belief can be added, at any next instant at which the individual is informed of a fact that he currently believes.

Let $\mathbb{L}_1 = \mathbb{L}_0 + WND + WNA$, that is, $\mathbb{L}_1$ is the logic obtained by adding axioms (WND) and (WNA) to $\mathbb{L}_0$.

Proposition 2 *Logic $\mathbb{L}_1$ is characterized by the class of belief revision frames that satisfy the following property:*
$\forall t_1, t_2 \in T$, $\forall \omega \in \Omega$, if $t_1 \rightarrowtail t_2$ and $\mathcal{B}_{t_1}(\omega) \subseteq \mathcal{I}_{t_2}(\omega)$ then $\mathcal{B}_{t_2}(\omega) = \mathcal{B}_{t_1}(\omega)$.

Logic $\mathbb{L}_1$ captures a weak notion of minimal change of beliefs in that it requires the agent not to change his beliefs if he is informed of some fact that he already believes. This requirement is expressed directly in the following axiom (WNC stand for 'Weak No Change'): if ϕ and ψ are Boolean formulas,

$$(I\phi \wedge \Diamond^{-1}(B\phi \wedge \neg B\neg\phi)) \to (B\psi \leftrightarrow \Diamond^{-1}B\psi) \qquad (WNC)$$

WNC says that if the agent is informed of a fact that he non-trivially believed in the immediately preceding past, then he now believes any fact if and only if he believed it then.

Proposition 3 *WNC is a theorem of $\mathbb{L}_1$.*

We now consider a stronger logic than $\mathbb{L}_1$. The following axiom strengthens (WND) by requiring the individual not to drop any of his current factual beliefs at any next instant at which he is informed of some fact that he considers possible (without necessarily believing it: the condition $B\phi$ in the antecedent of (WND) is eliminated): if ϕ and ψ are Boolean,

$$(\neg B\neg\phi \wedge B\psi) \to \bigcirc(I\phi \to B\psi). \qquad (ND)$$

The corresponding strengthening of (WNA) requires that if the individual considers it possible that $(\phi \wedge \neg\psi)$ then at any next instant where he is informed that ϕ he does not believe that ψ: if ϕ and ψ are Boolean,

$$\neg B\neg(\phi \wedge \neg\psi) \to \bigcirc(I\phi \to \neg B\psi). \qquad (NA)$$

The following axiom says that if the agent considers a fact ϕ possible, then he will believe ϕ at any next instant at which he is informed that ϕ ('QA' stands for 'Qualified Acceptance'): if ϕ is Boolean,

$$\neg B\neg\phi \to \bigcirc(I\phi \to B\phi) \qquad (QA)$$

Let $\mathbb{L}_2 = \mathbb{L}_0 + ND + NA + QA$.

Remark 4 *Logic $\mathbb{L}_2$ contains (is a strenghtening of) $\mathbb{L}_1$. In fact, WND is a theorem of logic $\mathbb{L}_0 + ND$ and WNA is a theorem of $\mathbb{L}_0 + NA$*

Proposition 5 *Logic $\mathbb{L}_2$ is characterized by the class of belief revision frames that satisfy the following property:*
$\forall t_1, t_2 \in T$, $\forall \omega \in \Omega$, if $t_1 \rightarrowtail t_2$ and $\mathcal{B}_{t_1}(\omega) \cap \mathcal{I}_{t_2}(\omega) \neq \emptyset$ then
$\mathcal{B}_{t_2}(\omega) = \mathcal{B}_{t_1}(\omega) \cap \mathcal{I}_{t_2}(\omega)$.

We argued in an earlier paper [3] that the above property captures the qualitative content of Bayes' rule. We call the property stated in Proposition 5 *Qualitative Bayes Rule*.

A stronger logic than $\mathbb{L}_2$ is obtained by replacing axiom QA with the following axiom, which requires factual information to be believed ('*Acc*' stands for 'Acceptance'): if ϕ is Boolean,

$$I\phi \to B\phi. \qquad (Acc)$$

We also add the following axiom which says that if the individual receives consistent information then his beliefs are consistent, in the sense that he does not simultaneously believe a formula and its negation ('WC' stands for 'Weak Consistency'): if ϕ is a Boolean formula,

$$(I\phi \wedge \neg A\neg\phi) \rightarrow (B\psi \rightarrow \neg B\neg\psi). \qquad (WC)$$

The next two propositions deal with the relationship between our approach and the AGM theory of belief revision (see [1]).

Let $\mathbb{L}_3 = \mathbb{L}_0 + ND + NA + Acc + WC$. Clearly logic $\mathbb{L}_3$ contains logic $\mathbb{L}_2$.

Definition 6 *An* $\mathbb{L}_3$-frame *is a temporal belief revision frame that satisfies the following properties:*
(1) the Qualitative Bayes Rule,
(2) $\forall\omega \in \Omega$, $\forall t \in T$, $\mathcal{B}_t(\omega) \subseteq \mathcal{I}_t(\omega)$,
(3) $\forall\omega \in \Omega$, $\forall t \in T$, *if* $\mathcal{I}_t(\omega) \neq \varnothing$ *then* $\mathcal{B}_t(\omega) \neq \varnothing$.
An $\mathbb{L}_3$-model *is a model based on an* $\mathbb{L}_3$*-frame.*

Proposition 7 *Logic* $\mathbb{L}_3$ *provides an axiomatization of the first six axioms (the so-called basic axioms) of the AGM theory of belief revision in the sense that both (A) and (B) below hold:*

(A) Fix an arbitrary $\mathbb{L}_3$*-model and let* $t_1, t_2 \in T$ *and* $\alpha \in \Omega$ *be such that* $t_1 \rightarrowtail t_2$ *and* $(\alpha, t_2) \models I\phi$, *with* $\phi \in \Phi^B$ *(where* $\Phi^B \subseteq \Phi$ *denotes the subset of Boolean formulas). Define*
$K = \{\psi \in \Phi^B : (\alpha, t_1) \models B\psi\}$ *and*
$K^*_\phi = \{\psi \in \Phi^B : (\alpha, t_2) \models B\psi\}$.
Then K^*_ϕ *satisfies AGM postulates (K*1)-(K*6).*

(B) Let $K \subseteq \Phi^B$ *be a consistent and deductively closed set and* $\phi \in \Phi^B$. *If* $K^*_\phi \subseteq \Phi^B$ *satisfies AGM postulates (K*1)-(K*6) then there is an* $\mathbb{L}_3$*-model,* $t_1, t_2 \in T$ *and* $\alpha \in \Omega$ *such that*
(B.1) $t_1 \rightarrowtail t_2$
(B.2) $K = \{\psi \in \Phi^B : (\alpha, t_1) \models B\psi\}$
(B.3) $(\alpha, t_2) \models I\phi$
(B.4) $K^*_\phi = \{\psi \in \Phi^B : (\alpha, t_2) \models B\psi\}$
(B.5) if ϕ *is consistent then* $(\beta, t) \models \phi$ *for some* $\beta \in \Omega$ *and* $t \in T$.

Finally, we strengthen $\mathbb{L}_3$ by adding the following two axioms, where, as usual, ϕ, ψ and χ are Boolean formulas:

$$\Diamond(I(\phi \wedge \psi) \wedge B\chi) \rightarrow \bigcirc(I\phi \rightarrow B((\phi \wedge \psi) \rightarrow \chi)) \qquad (K7)$$

$$\Diamond(I\phi \wedge \neg B\neg(\phi \wedge \psi) \wedge B(\psi \rightarrow \chi)) \rightarrow \bigcirc(I(\phi \wedge \psi) \rightarrow B\chi). \qquad (K8)$$

Axiom $(K7)$ says that if there is a next instant where the individual is informed that $\phi \wedge \psi$ and believes that χ, then at every next instant it must be the case that if the individual is informed that ϕ then he must believe that $(\phi \wedge \psi) \rightarrow \chi$.

Axiom $(K8)$ says that if there is a next instant where the individual is informed that ϕ, considers $(\phi \wedge \psi)$ possible and believes that $\psi \rightarrow \chi$, then at every next instant it must be the case that if the individual is informed that $\phi \wedge \psi$ then he believes that χ.

Let $\mathbb{L}_4 = \mathbb{L}_0 + ND + NA + Acc + K7 + K8$. Clearly, logic $\mathbb{L}_4$ contains logic $\mathbb{L}_3$.

Proposition 8 *Logic $\mathbb{L}_4$ provides an axiomatization of the full set of AGM postulates (K*1)-(K*8), in the sense that both (A) and (B) below hold:*

(A) Fix an arbitrary $\mathbb{L}_4$-model and let $t_1, t_2, t_3 \in T$ and $\alpha \in \Omega$ be such that $t_1 \rightarrowtail t_2$, $t_1 \rightarrowtail t_3$, $(\alpha, t_2) \models I\phi$ and $(\alpha, t_3) \models I(\phi \wedge \psi)$ with $\phi, \psi \in \Phi^B$. Define
$K = \{\psi \in \Phi^B : (\alpha, t_1) \models B\psi\}$,
$K^*_\phi = \{\psi \in \Phi^B : (\alpha, t_2) \models B\psi\}$ *and*
$K^*_{\phi\wedge\psi} = \{\chi \in \Phi^B : (\alpha, t_3) \models B\chi\}$.
*Then K^*_ϕ and $K^*_{\phi\wedge\psi}$ satisfy AGM postulates (K*1)-(K*8).*

*(B). Let $K \subseteq \Phi^B$ be a consistent and deductively closed set and $\phi, \psi \in \Phi^B$. If $K^*_\phi, K^*_{\phi\wedge\psi} \subseteq \Phi^B$ satisfy AGM postulates (K*1)-(K*8) then there is an $\mathbb{L}_4$-model, $t_1, t_2, t_3 \in T$ and $\alpha \in \Omega$ such that*
(B.1) $t_1 \rightarrowtail t_2$
(B.2) $K = \{\chi \in \Phi^B : (\alpha, t_1) \models B\chi\}$
(B.3) $(\alpha, t_2) \models I\phi$
(B.4) $K^*_\phi = \{\chi \in \Phi^B : (\alpha, t_2) \models B\chi\}$
(B.5) if ϕ is consistent then $(\beta, t) \models \phi$ for some $\beta \in \Omega$ and $t \in T$
(B.6) $t_1 \rightarrowtail t_3$
(B.7) $(\alpha, t_3) \models I(\phi \wedge \psi)$
(B.8) $K^*_{\phi\wedge\psi} = \{\chi \in \Phi^B : (\alpha, t_3) \models B\chi\}$.

Further extensions of this logic will be explored in future work, in particular extensions obtained by adding axioms that capture the notion of memory.

4 Conclusion

We proposed a temporal logic where information and beliefs are modeled explicitly by means of two modal operators I and B, respectively. This logic can accommodate not only the AGM theory of belief revision but also iterated revision, a topic that has received considerable attention in recent years (see, for example, [8]). There are two advantages to modeling belief revision in a modal framework: (1) one achieves a uniform treatment of static and dynamic beliefs, thus providing a unified theory of both, and (2) the approach allows one to state properties of beliefs in a clear and transparent way by means of syntactic axioms.

Some of the ideas contained in this paper (in particular the modeling of information by means of a non-normal modal operator) were first put forward in [3]. The framework in that paper was different, however, since it was not based on branching-time structures and only two dates were considered with two associated belief operators, B_0 (representing initial beliefs) and B_1 (representing revised beliefs). The main contribution of that paper was a soundness and completeness result for the proposed logic with respect to the class of frames that satisfy the Qualitative Bayes Rule.

For a detailed discussion of literature that is somewhat related to the general approach discussed here, the reader is referred to [3].

A partial list of relevant references is given below. A more complete list is given in the full paper.

References

[1] Alchourron, C., P. Gärdenfors and D. Makinson, On the logic of theory change: partial meet contraction and revision functions, *The Journal of Symbolic Logic*, 1985, 50: 510-530.

[2] Board, O., Dynamic interactive epistemology, *Games and Economic Behavior*, 2004, 49: 49-80.

[3] Bonanno, G. (2005), A simple modal logic for belief revision, *Synthese*, 2005, 147: 193-228 (and *Knowledge, Rationality and Action*, 5-40).

[4] Friedman, N. and J. Halpern, Belief revision: a critique, *Journal of Logic, Language, and Information*, 1999, 8: 401–420.

[5] Goranko, V. and S. Passy, Using the universal modality: gains and questions, *Journal of Logic and Computation*, 1992, 2: 5-30.

[6] Katsuno, H and Mendelzon, A.O. (1991), "Propositional knowledge base revision and minimal change, "*Artificial Intelligence*", 52, 263–294.

[7] Levesque, H. J., All I know: a study in autoepistemic logic, *Artificial Intelligence*, 1990, 5: 263-309.

[8] Nayak, A., M. Pagnucco and P. Peppas, Dynamic belief revision operators, *Artificial Intelligence*, 2003, 146: 193-228.