Belief Revision Revisited
or:
Thoughts on the Scope of Belief Revision

James Delgrande

Simon Fraser University

Canada

NMR 2006

Windermere, UK, May 31, 2006

Belief Revision

- Current interest in *belief revision* is regarded as going back to the seminal work of Alchourrón, Gärdenfors, and Makinson (e.g. [AGM85]).
- The *AGM Approach* is the dominant paradigm in belief change, and the standard against which approaches are measured.
- Here:
  - ⋆ Examine this picture and
  - ⋆ suggest a reconsideration of what is meant by *belief revision*.
- So: argue for a broader conception of BR.
  - ☞ Note: this isn't a technical question per se, but rather asks what is the scope of the phenomenon of BR.

Overview

- Introduction
- Overview of the AGM approach.
- Three challenges regarding revision:
  - ⋆ Informational Economy
  - ⋆ Inclusion (and Superexpansion)
  - ⋆ Iteration (and inter alia Closure and Success)
- Suggestions regarding the scope of BR.

What is Belief Revision?

Informally, we have an agent, and some new piece of information that is to be incorporated into the agent's set of beliefs.

Example:

Beliefs The person by the workstation is a teaching assistant.

The person by the workstation is a Ph.D. student.

Ph.D. students are graduate students.

Graduate students who are teaching assistants can't hold university fellowships.

Consequence The person by the workstation can't hold a fellowship.

New Information The person by the workstation holds a fellowship.

What is Belief Revision? (ctd)

To keep the KB consistent, it must be *revised*.

- Hence some of the original beliefs must be retracted.
- However one would want to keep some of the initial information.
- *Problem:* Logical considerations alone are not sufficient to answer this question.

Goal: describe belief revision at the *knowledge level*, i.e. on an abstract level, independent of how beliefs are represented and manipulated.

Underlying Assumptions

Categorial Matching The representation of a belief state after change should be of the same format as that prior to change

Consistency Beliefs in belief state should be consistent if possible.

Deductive Closure If the beliefs in a belief state logically entail a sentence ϕ, then ϕ should be included in the state.

- ➥ More generally: Abstract characterization

Principle of Informational Economy The amount of information lost during change should be kept to a minimum

Preference Beliefs considered more important or entrenched should be retained in favour of less important ones.

Domain Propositions used to describe the domain are static.

The AGM Approach

A revision function is a function from a belief state and a sentence to a belief state.

Represent as $K * \alpha$.

- Originally an agent's belief state was given in terms of a deductively closed set of sentences, or a *belief set*. (Also: set of possible worlds, etc.)
- Subsequently recognized for *iterated belief revision* that the formal model of a belief state should include the means necessary for performing belief change.
  - ⋆ So, think of K as comprising an *epistemic state* that includes, in addition to a belief set, "*revision-guiding structures*" (Rott).

    ➥Could write K for the epistemic state and $Bel(K)$ for the underlying contingent beliefs, but for simplicity I'll stick with K.

The AGM Approach (ctd)

Want to incorporate new beliefs in a consistent manner.

Arguably, any rational agent should be bound by the following postulates.

$(K*1)$ $K*\phi$ is a belief set. (Closure)

$(K*2)$ $\phi \in K*\phi$. (Success)

$(K*3)$ $K*\phi \subseteq K+\phi$. (Inclusion)

$(K*4)$ If $\neg\phi \notin K$, then $K+\phi \subseteq K*\phi$. (Preservation)

$(K*5)$ $K*\phi = K_\perp$ iff $\vdash \neg\phi$. (Vacuity)

$(K*6)$ If $\phi \equiv \psi$, then $K*\phi = K*\psi$. (Extensionality)

$(K*7)$ $K*(\phi \wedge \psi) \subseteq (K*\phi)+\psi$. (Superexpansion)

$(K*8)$ If $\neg\psi \notin K*\phi$, then $(K*\phi)+\psi \subseteq K*(\phi \wedge \psi)$. (Subexpansion)

AGM Approach: Ranking Functions

AGM revision functions have been modelled by ranking functions (also called SOS's, OCF's, plausibility orderings, etc.) (Grove, Spohn).

- A *ranking function* is an assignment of non-negative integers to interpretations.
- There must be at least one interpretation assigned value 0.
- The interpretations assigned value 0 characterize the agent's beliefs.
- $K * \alpha$ is characterized by the least set of interpretations in which α is true.

Challenges to the Current View of BR

Present three challenges to the accepted view of revision.

I Informational Economy: Will recapitulate an argument by Hans Rott against *informational economy* and *preference*.

II Contra $K*3$: Argue that there are "reasonable" BR functions that violate $K*3$:

$$K*\phi \subseteq K+\phi.$$

In fact will suggest that there are "reasonable" BR functions violating:

If $\alpha \in K$ then $K*\alpha = K$.

(From [DNP05].)

Challenges to the Current View of BR

III Iterated Revision: Since $K * \alpha$ is a belief set (epistemic state), subsequent revision by β is given by $(K * \alpha) * \beta$.

- AGM doesn't address iteration, but much recent work in BR has.
- The form of iteration is justified by appeal to a *Principle of Recency*

Will argue against this principle, and implicitly $K * 1$ and $K * 2$.

(From [DDL06].)

➥These challenges are not intended to be comprehensive or systematic

☞ Rather these points are intended to suggest a re-examination of the scope of BR.

Challenge I: Informational Economy

Informational Economy When accepting a new piece of information, the agent should minimally change its old beliefs.

> *. . . when we change our beliefs, we want to retain as much as possible of our old beliefs – information in general is not gratuitous and unnecessary losses of information are to be avoided* [Gärdenfors88, P. 49]

There is already a clash ([Levi91]) since the basic postulates characterise *partial meet revision*.

- I.e. revision of K by α relies on the intersection of maximal subsets of K that fail to imply $\neg\alpha$.
  - ⋆ Maximal subsets of K that fail to imply $\neg\alpha$ satisfy informational economy, but
  - ⋆ their intersection does not.

Informational Economy (ctd)

Worse, [Rott 2000] shows the following:

Call Γ a *candidate revision of K by α* if

1. Γ is consistent and logically closed, and
2. $\alpha \in \Gamma$.

Let $\neg\alpha \in K$ and let K_1 and K_2 be two candidate revisions of K by α.

Then K_1 and K_2 cannot be set-theoretically compared in terms of the beliefs on which they differ with K.

☞ I.e. $K_1 \setminus K \not\subset K_2 \setminus K$ and $K_2 \setminus K \not\subset K_1 \setminus K$.

As well, Rott gives a similar argument against *Preference*.

Informational Economy (ctd)

Rott considers various objections to his observations, and gives a reply to each.

- However, we can single out one objection:

  Objection: It is the *epistemic state* to which principles of minimal change should be applied.

  Reply: If an epistemic state is equated with a ranking function on possible worlds (or: OCF, plausibility ordering, SOS, . . .), then existing proposals implementing a notion of minimal change wrt these orderings to date (2002) have poor properties.

- This suggests possible resolutions:

  1. Consider other revision operators wrt ranking functions.
  2. Consider structures other than ranking functions for representing epistemic states.

Challenge II: Gricean Revision

Observe: for $K * \alpha$ there is a coherent reading of revision that "the agent learns *exactly* that α".

Example 1 You own a trucking company that will be sending 3 trucks from the Okanagan Valley to Vancouver over the weekend. There are three separate routes, the Coquihalla Hwy, Hwy 3, and the Trans-Canada Hwy. Each truck will take a different highway. On Monday morning, since the weather has been poor, you expect that no truck was able to make the trip; say $K = Cn(\neg A \wedge \neg B \wedge \neg C)$. On arriving at work you are told that there are goods from the Okanagan waiting to be processed. You conclude that a truck, or trucks, were able to get through, i.e. $K * (A \vee B \vee C) = Cn(A \vee B \vee C)$.

Gricean Revision (ctd)

Example 2 (Hansson) I believe that Cleopatra had a son and a daughter, so $(S \wedge D) \in K$. I am then told that Cleopatra had no children (viz. $K - (S \vee D)$) but then subsequently that she did have a child ($(K - (S \vee D)) + (S \vee D)$). According to the recovery postulate one would now believe that Cleopatra had both a son and a daughter.

Gricean Revision (ctd)

Resolution to Hansson's example:

1. Regard the belief change as $K - (S \vee D) * (S \vee D)$
2. Treat $*$ as *conservative* or *Gricean* change ([DNP05]), denoted $\hat{*}$.
   Idea: For $K' \hat{*} (S \vee D)$, all that one knows about $\{S, D\}$ is $S \vee D$.

- That is, $K \hat{*} \alpha$ is a *conservative extension* of α.
  I.e. for $\beta \in \mathcal{L}(\alpha)$, if $K \hat{*} \alpha \vdash \beta$ then $\alpha \vdash \beta$.
- Thus, $K \hat{*} (p \vee q) \nvdash p$ even if $K \vdash p$.
  ☞So here we have "revision" where knowledge is lost.

Gricean Revision (ctd)

☞We obtain a reasonable belief change operator that non-trivially violates major AGM postulates ($K*3$ and $K*7$).

- In fact, in a slight generalization of the approach, one may lose information in revising by a tautology, viz. asserting that all the agent knows about p is $p \vee \neg p$.

This suggests:

- Revisiting the notion of "belief change" operator, or
- more radically, consider the development of an (object language) in which "operators" may be specified.

Challenge III: Iterated Revision

[DDL06] argues that iterated revision, as generally understood, is a highly restricted version of a more general belief change framework.

Observe: The form of revision, as inherited from the AGM approach, essentially requires the *Principle of Recency* for iterated revision.

- I.e. Since $K * \alpha$ is a belief set, so subsequent iteration by β must be of the form $(K * \alpha) * \beta$.
- Intuition: More recent pieces of information are more reliable.

Argue: Since BR concerns information about static propositions, the order of information is *irrelevant* in the general case.

☞ It *may be* that more recently received items are more reliable, but it *need not* be so.

Iterated Revision:

Revision as Prioritized Merging

Propose that the appropriate framework for revision is "prioritized merging".

- Have a background epistemic state, along with a set of formulas, each with an attached reliability.
- The belief set associated with K can now be regarded as another formula with an attached reliability.

☞But this then is a radically different conception of "revision"!

Taking Stock

☞Have argued for a *broad* conception of belief revision.

- This raises the question:

  What do we mean by belief revision (or: belief change)?

- This is somewhat vague. Better questions:
  - ⋆ How can we delimit the space of belief change operators, and how can we classify different types of belief change operators, in light of the preceding?
  - ⋆ If iterated revision is prioritized merging, what is the difference between BR and merging?

- Key distinction: belief change wrt a single epistemic state vs. multiple epistemic states.

  (If we associate epistemic states with agents, then the distinction is between single and multiple agents.)

Classification of Operators

Before:

Belief Revision	Merging
AGM BR	Belief Set Merging e.g. [KP-P02], [DDL06]
BR with formula SOS's e.g. [NNP96]	OCF merging e.g. [Meyer01]

After:

Single Agent BC	Multi-Agent BC
AGM BR	BR with formula SOS's e.g. [NNP96]
Belief Set Merging e.g. [KP-P02], [DDL06]	OCF merging e.g. [Meyer01]

Classification of Operators

As well, see:

1. *Three scenarios for the revision of epistemic states*, Didier Dubois, NMR 2006.
2. *About time, revision, and update*, Jérôme Lang, NMR 2006.

Epistemic States

If belief change is with respect to an epistemic state, this raises the question

What is an epistemic state?

Possibilities:

1. Ranking function (SOS, OCF).

   Intuition: Since AGM revision functions can be modelled by ranking functions, so a ranking function = an epistemic state and revising an epistemic state (e.g. in repeated revisions) amounts to revising the ordering. (Boutilier, Darwiche/Pearl, Lin/Thielscher, Papini, Spohn, Williams, etc.)

   ☞Experience indicates that this alone is likely too restrictive.

   ☞Problem: expressing minimal change wrt a ranking function.

Epistemic States (ctd)

Two other examples of possible representations for epistemic states:

2. A set of deductively closed theories along with a preference relation on theories (Bochman).

   Intuition: Each closed theory is considered by the agent to be a "serious possibility".

3. An epistemic state consists of a set of possible worlds (giving contingent beliefs) and a *distance* function on pairs of possible worlds (Bennaim, Lehmann, Magidor, Schlechta).

   - A background theory (somehow) induces a distance function.
   - Every set of possible worlds (cf. belief set) induces a ranking function.

   ☞Suggests not that an epistemic state *is* a ranking function, but that an epistemic state *induces* a ranking function.

Epistemic Inputs

- Most work (including work in dynamic logic) accepts that there is a set of distinct belief change operators each reflecting a specific type of input concerning the outside world.
- Another view: Belief change operators can be regarded as directives concerning an agent's doxastic state.

  ☞I.e. as directives that an agent should take on certain beliefs.
- In this latter case, we need a language for expressing (forms of) beliefs.

  ☞Fortunately there are such languages, e.g. (Levesque, Lakemeyer).

Epistemic Inputs (ctd)

For example, our goal could be to define a function $MAKE(ES, \alpha)$ for epistemic state ES and assertion α.

- ☞ Result is an epistemic state wherein α is true.

Examples

- $\mathcal{K} * \alpha$ would be expressed by $MAKE(\mathcal{K},\ B\alpha)$.
- $\mathcal{K} - \alpha$ would be expressed by $MAKE(\mathcal{K},\ \neg B\alpha)$.
- $\mathcal{K} + \alpha$ would be expressed by $MAKE(\mathcal{K},\ B(form(\mathcal{K}) \wedge \alpha))$.
- $\mathcal{K}\ \hat{*}\ \alpha$ would be expressed by $MAKE(\mathcal{K},\ O\langle\mathcal{L}(\alpha)\rangle\alpha)$.

Summary

Theme: Suggest a re-examination of the scope of BR, in particular a broadening of the scope of enquiry.

- General form: epistemic state + input(s) $\Rightarrow$ epistemic state.
- Will necessitate a revisiting of basic assumptions (such as informational economy and the nature of epistemic states).

Of course, intuitions can and should be guided by lessons gained from

- implementations, and
- applications.