Logic and Artificial Intelligence

Lecture 14

Arthur Paul Pedersen

Department of Philosophy
Carnegie Mellon University

www.andrew.cmu.edu/∼ppederse
apaulpedersen@cmu.edu

October 17, 2011

The Basic Problem of Belief Revision

Consider the following beliefs of a rational agent:

(α) All Europeans swans are white.
(β) The bird caught in the trap is a swan.
(γ) The bird caught in the trap comes from Sweden.
(δ) Sweden is part of Europe.

Thus, the agent believes:

(ϵ) The bird caught in the trap is white.

Now suppose the rational agent—for example, You—learn that the bird caught in the trap is black ($\neg\epsilon$).

The Basic Problem of Belief Revision

Consider the following beliefs of a rational agent:

(α) All Europeans swans are white.
(β) The bird caught in the trap is a swan.
(γ) The bird caught in the trap comes from Sweden.
(δ) Sweden is part of Europe.

(ϵ) The bird caught in the trap is white.

Question: How should the agent incorporate $\neg\epsilon$ into his belief state to obtain a *consistent* belief state?

Problem: Logical considerations alone are insufficient to answer this question!

Why?

The Basic Problem of Belief Revision

Consider the following beliefs of a rational agent:

(α) All Europeans swans are white.
(β) The bird caught in the trap is a swan.
(γ) The bird caught in the trap comes from Sweden.
(δ) Sweden is part of Europe.

(ϵ) The bird caught in the trap is white.

Question: How should the agent incorporate $\neg\epsilon$ into his belief state to obtain a *consistent* belief state?

Problem: Logical considerations alone are insufficient to answer this question!

There are several logically distinct ways to incorporate $\neg\epsilon$!

The Basic Problem of Belief Revision

Consider the following beliefs of a rational agent:

(α) All Europeans swans are white.
(β) The bird caught in the trap is a swan.
(γ) The bird caught in the trap comes from Sweden.
(δ) Sweden is part of Europe.

(ϵ) The bird caught in the trap is white.

Question: How should the agent incorporate $\neg\epsilon$ into his belief state to obtain a *consistent* belief state?

Problem: Logical considerations alone are insufficient to answer this question!

What extralogical factors serve to determine what beliefs to give up and what beliefs to retain?

Quine and Minimal Mutilation

Quine's "Two Dogmas of Empiricism" (1951) undermined two presumptions of classical as well as modern empiricism:

(i) There is a sharp distinction between analytic judgments and synthetic judgments; and
(ii) Every meaningful sentence can be reduced to a construction upon observation reports.

Quine then discusses *belief revision*, describing how we should, and for the most part do, accommodate scientific lore in the face of belief (or theory) contravening experiences.

Belief revision is a *matter of choice*, and the choices are to be made in such a way that:

(a) The resulting theory squares with the experience;

(b) It is simple; and

(c) The choices disturb the original theory as little as possible.

Quine then discusses *belief revision*, describing how we should, and for the most part do, accommodate scientific lore in the face of belief (or theory) contravening experiences.
pause

Belief revision is a *matter of choice*, and the choices are to be made in such a way that:

(a) The resulting theory squares with the experience;
(b) It is simple; and
(c) The choices disturb the original theory as little as possible.

Research has relied on the following related *guiding ideas*:

(1) When accepting a new piece of information, an agent should aim at a minimal change of his old beliefs.
(2) If there are different ways to effect a belief change, the agent should give up those beliefs which are least entrenched.

References to Minimal Mutilation 1

When accepting a new piece of information, an agent should aim at a minimal change of his old beliefs.

> *The concept of contraction leads us to the concept of minimal change of belief, or briefly, revision (Makinson 1985, p. 352).*

> *The criterion of informational economy demands that as few beliefs as possible be given up so that the change is in some sense a minimal change of K to accommodate for A (Gärdenfors 1988, p. 53).*

> *The amount of information lost in a belief change should be kept minimal (Gärdenfors and Rott 1995, p. 38).*

> *At the center of the AGM theory [of theory change] are a number of approaches to giving formal substance to the maxim [of minimal mutilation: keep incisions into theories as small as possible!] (Fuhrmann 1997, p. 17).*

> *The hallmark of the AGM postulates is the principle of minimal belief change, that is, the need to preserve as much of earlier beliefs as possible and to add only those beliefs which are absolutely compelled by the revision specified (Darwiche and Pearl 1997, p. 2).*

References to Minimal Mutilation 2

If there are different ways to effect a belief change, the agent should give up those beliefs which are least entrenched.

> *When a belief set K is contracted (or revised), the sentences in K that are given up are those with the lowest epistemic entrenchment (Gärdenfors 1988, p. 87).*

> *The guiding idea for the construction is that when a knowledge system K is revised or contracted, the sentences in K which are given up are those having the lowest degrees of epistemic entrenchment (Gärdenfors and Makinson 1988, p. 88).*

> *In so far as some beliefs are considered more important or entrenched than others, one should retract the least important ones (Gärdenfors and Rott 1995, p. 38).*

References to Minimal Mutilation 2

If there are different ways to effect a belief change, the agent should give up those beliefs which are least entrenched. [1]

> *...when it comes to choosing between candidates for removal, the least entrenched ones ought to be given up (Fuhrmann 1997, p. 24).*

> *A hallmark of the AGM theory is its commitment to the principle of informational economy: beliefs are only given up when there are no less entrenched candidates.... If one of two beliefs must be retracted in order to accommodate some new fact, the less entrenched belief will be relinquished, while the more entrenched persists (Boutilier 1996, pp. 264-265).*

[1] The foregoing has been compiled in (Rott 2000)

This Week

We will focus on the AGM paradigm of belief change.

1. Three Epistemic Attitudes and Three Epistemic Changes
2. Belief Contraction
3. Belief Revision
4. Propositional Models of Belief Change
5. Belief Change and Rational Choice?

Reference.

"Belief Revision." A.P. Pedersen & H. Arló-Costa. In L. Horsten and R. Pettigrew, editors, *Continuum Companion to Philosophical Logic*. Continuum Press, 2011.

Three Epistemic Attitudes

Belief recognizes *belief, disbelief, and suspense.*

(α) Life exists on Earth.

(β) Life exists on the moon.

(γ) Life exists on Saturn's moon.

Three Epistemic Attitudes

Belief observes logical norms.

B_{Cn} Closure under logical consequence.

$B_{\wedge}$ Closure under conjunction.

B_{Con} Consistency.

Three Epistemic Changes

In the AGM framework, an agent's belief state is represented by a set of sentences K satisfying B_{Cn} and $B_{\bigwedge}$, called a *belief set*.

(i) In *expansion*, a sentence ϕ is added to a belief set K to obtain an expanded belief set $K + \phi$.

(ii) In *revision*, a sentence ϕ is added to a belief set K to obtain a revised belief set $K * \phi$ in a way that preserves logical consistency.

(iii) In contraction, a sentence ϕ is removed from K to obtain a contracted belief set $K \dot{-} \phi$ that does not include ϕ.

Revision can be reduced to contraction via the so-called *Levi identity*, according to which the revision of a belief set K with a sentence ϕ is identical to the contraction $K \dot{-} \neg\phi$ expanded by ϕ:

$$K * \phi \quad = \quad (K \dot{-} \neg\phi) + \phi.$$

Preliminaries

We presuppose a propositional language $\mathcal{L}$ with connectives $\neg, \wedge, \vee, \rightarrow, \leftrightarrow$.

For($\mathcal{L}$) denotes the set of formulae of $\mathcal{L}$;

$a, b, c, \ldots, p, q, r, \ldots$ denote propositional variables of $\mathcal{L}$;

$\alpha, \beta, \delta, \ldots, \phi, \psi, \chi, \ldots$ denote arbitrary formulae of $\mathcal{L}$;

$\Gamma, \Delta, \Theta, \ldots, \Lambda, \Sigma, \Psi, \ldots$ denote arbitrary sets of formula.

Consequence Operators

We assume that $\mathcal{L}$ is governed by a Tarskian consequence operation $\mathrm{Cn} : \mathscr{P}(\mathrm{For}(\mathcal{L})) \rightarrow \mathscr{P}(\mathrm{For}(\mathcal{L}))$ such that:

(i) $\Gamma \subseteq \mathrm{Cn}(\Gamma)$;

(ii) If $\Gamma \subseteq \Delta$, then $\mathrm{Cn}(\Gamma) \subseteq \mathrm{Cn}(\Delta)$;

(iii) $\mathrm{Cn}(\mathrm{Cn}(\Gamma)) \subseteq \mathrm{Cn}(\Gamma)$.

In addition, the operator Cn is assumed to satisfy the following conditions:

(iv) $\mathrm{Cn}_0(\Gamma) \subseteq \mathrm{Cn}(\Gamma)$, where Cn_0 is the classical consequence operation;

(v) If $\phi \in \mathrm{Cn}(\Gamma)$, then there is some finite $\Gamma_0 \subseteq \Gamma$ such that $\phi \in \mathrm{Cn}(\Gamma_0)$;

(vi) If $\phi \in \mathrm{Cn}(\Gamma \cup \{\psi\})$, then $\psi \rightarrow \phi \in \mathrm{Cn}(\Gamma)$.

Conditions (i)-(vi) are respectively called *Inclusion, Monotony, Idempotence, Supraclassicality, Compactness*, and *Deduction*.

Consequence Operators

Γ is called *logically closed* with respect to Cn if $\mathrm{Cn}(\Gamma) = \Gamma$.

While in logical parlance logically closed sets are called *theories*, the belief revision literature has adopted its own terminology, calling theories *belief sets*.

A belief set is usually denoted by K.

The usual epistemological interpretation of theories is as *commitment sets*, representing the doxastic commitments of a rational agent (Levi 1991).

$\mathrm{Cn}(\emptyset)$ consists of the agent's *underlying, firm, full beliefs.*

Elementary Properties of Cn

Proposition

Let $\Gamma, \Delta \subseteq \text{For}(\mathcal{L})$, *and let* $\phi, \psi \in \text{For}(\mathcal{L})$. *Then* Cn *satisfies the following properties:*

(i) $\text{Cn}(\text{Cn}(\Gamma)) = \text{Cn}(\Gamma)$.

(ii) *If* ϕ *and* ψ *are truth-functionally equivalent, then*

$$\phi \in \text{Cn}(\Gamma) \text{ if and only if } \psi \in \text{Cn}(\Gamma).$$

(iii) $\text{Cn}(\Gamma \cup \Delta) = \text{Cn}(\Gamma \cup \text{Cn}(\Delta))$.

(iv) *If* $\Gamma \subseteq \Delta \subseteq \text{Cn}(\Gamma)$, *then* $\text{Cn}(\Gamma) = \text{Cn}(\Delta)$.

(v) $\text{Cn}(\Gamma \cup \{\phi\}) \cap \text{Cn}(\Gamma \cup \{\psi\}) = \text{Cn}(\Gamma \cup \{\phi \vee \psi\})$.

(vi) $\text{Cn}(\Gamma \cup \{\phi\}) \cap \text{Cn}(\Gamma \cup \{\neg\phi\}) = \text{Cn}(\Gamma)$.

(vii) $\text{Cn}(\Gamma \cup \{\phi, \psi\}) = \text{Cn}(\Gamma \cup \{\phi \wedge \psi\})$.

(viii) *For every nonempty* $I \subseteq \mathscr{P}(\text{For}(\mathcal{L}))$,

$$\text{Cn}\left(\bigcap_{\Gamma \in I} \text{Cn}(\Gamma)\right) = \bigcap_{\Gamma \in I} \text{Cn}(\Gamma).$$

Remainder Sets

Definition

Let K be a collection of formulae and α be a formula. The α-*remainder set* of K, $K\perp\alpha$, is the collection of subsets Γ of For($\mathcal{L}$) such that:

(i) $\Gamma \subseteq K$;

(ii) $\alpha \notin \text{Cn}(\Gamma)$;

(iii) There is no set Δ such that $\Gamma \subset \Delta \subseteq K$ and $\alpha \notin \text{Cn}(\Delta)$.

A member of $K\perp\alpha$ is called an α-*remainder* of K. We let $K\perp\mathcal{L} := \{K\perp\alpha : \alpha \in \text{For}(\mathcal{L})\}$.

Example. Let $\mathcal{L} = \{p, q\}$, $\text{Cn} = \text{Cn}_0$, and $K = \text{Cn}(\{p, q\})$ Identify:

$K\perp(p \wedge q)$

$K\perp p$

$K\perp(q \rightarrow p)$

$$K\bot(p \wedge q) = \{\mathrm{Cn}(\{p \leftrightarrow q\}), \mathrm{Cn}(\{p\}), \mathrm{Cn}(\{q\})\}$$
$$K\bot p = \{\mathrm{Cn}(\{p \leftrightarrow q\}), \mathrm{Cn}(\{q\})\}$$
$$K\bot(q \rightarrow p) = \{\mathrm{Cn}(\{q\})\}.$$

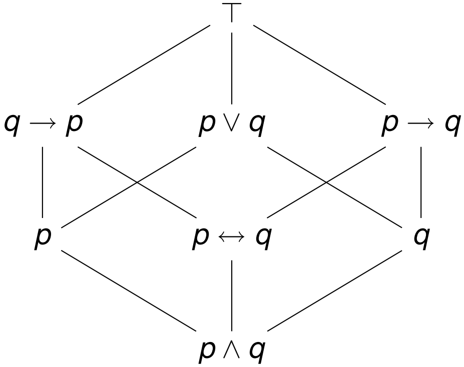

Properties of Remainder Sets

Definition

Let K be a belief set and α be a formula. The α-*remainder set* of K, $K \perp \alpha$, is the collection of subsets Γ of For($\mathcal{L}$) such that:

(i) $\Gamma \subseteq K$;

(ii) $\alpha \notin \text{Cn}(\Gamma)$;

(iii) There is no set Δ such that $\Gamma \subset \Delta \subseteq K$ and $\alpha \notin \text{Cn}(\Delta)$.

A member of $K \perp \alpha$ is called an α-*remainder* of K. We let $K \perp \mathcal{L} := \{K \perp \alpha : \alpha \in \text{For}(\mathcal{L})\}$.

Some important properties:

(a) $K \perp \alpha = \{K\}$ if and only if $\alpha \notin \text{Cn}(K)$;

(b) $K \perp \alpha = \emptyset$ if and only if $\alpha \in \text{Cn}(\emptyset)$.

The *Upper Bound Property*:

(c) If $\Gamma \subseteq K$ and $\alpha \notin \text{Cn}(\Gamma)$, then there is some Δ such that $\Gamma \subseteq \Delta \in K \perp \alpha$.

Selection Functions

Definition

Let K be a belief set. A *selection function* for K is a function γ on $K{\perp}\mathcal{L}$ such that for all formulae α:

(i) If $K{\perp}\alpha \neq \emptyset$, then:

  (a) $\gamma(K{\perp}\alpha) \subseteq K{\perp}\alpha$, and

  (b) $\gamma(K{\perp}\alpha) \neq \emptyset$;

(ii) If $K{\perp}\alpha = \emptyset$, then $\gamma(K{\perp}\alpha) = \{K\}$.

Example.

$$K{\perp}(p \wedge q) = \{\mathrm{Cn}(\{p \leftrightarrow q\}), \mathrm{Cn}(\{p\}), \mathrm{Cn}(\{q\})\}$$
$$K{\perp}p = \{\mathrm{Cn}(\{p \leftrightarrow q\}), \mathrm{Cn}(\{q\})\}$$

Consider:

$$\gamma(K{\perp}(p \wedge q)) = \{\mathrm{Cn}(\{p\})\}$$
$$\gamma(K{\perp}p) = \{\mathrm{Cn}(\{p \leftrightarrow q\}), \mathrm{Cn}(\{q\})\}$$

Partial Meet Contraction

Definition

Let K be a set of formulae. A function $\dot{-}$ on For($\mathcal{L}$) is a *partial meet contraction* for K if there is a selection function γ for K such that for all formulae α:

$$K \dot{-} \alpha \quad = \quad \bigcap \gamma(K \bot \alpha).$$

Example.

$$K \bot (p \wedge q) = \{\mathrm{Cn}(\{p \leftrightarrow q\}), \mathrm{Cn}(\{p\}), \mathrm{Cn}(\{q\})\}$$
$$K \bot p = \{\mathrm{Cn}(\{p \leftrightarrow q\}), \mathrm{Cn}(\{q\})\}$$

Consider:

$$\gamma(K \bot (p \wedge q)) = \{\mathrm{Cn}(\{p\})\}$$
$$\gamma(K \bot p) = \{\mathrm{Cn}(\{p \leftrightarrow q\}), \mathrm{Cn}(\{q\})\}$$

So:

$$K \dot{-} (p \wedge q) = \bigcap \gamma(K \bot p \wedge q) = \mathrm{Cn}(\{p\})$$
$$K \dot{-} p = \gamma(K \bot p) = \mathrm{Cn}(\{q\})$$

Two limiting cases of partial meet contraction are of special interest:

(i) γ selects exactly one element of $K \perp \alpha$ (*maxichoice contraction*);
(ii) γ selects the entire set $K \perp \alpha$ (*full meet contraction*).

So we have:

(i) $K \dot{-} \alpha = \{\Gamma\}$ for some $\Gamma \in K \perp \alpha$;
(ii) $K \dot{-} \alpha = \bigcap K \perp \alpha$.