Logic and Probabilistic Models of Belief Change

Eric Pacuit

Department of Philosophy
University of Maryland, College Park
`pacuit.org`

January 28, 2016

- Computer science: updating databases (Doyle 1979 and Fagin et al. 1983)
- Philosophy (epistemology/philosophy of science):
  - scientific theory change and revisions of probability assignments;
  - belief change (Levi 1977, 1980, Harper 1977) and its rationality.

Carlos **A**lchourrón, Peter **G**ärdenfors, and David **M**akinson.

1985 paper in the Journal of Symbolic Logic.

Starting point of belief revision theory.

C. Alchourrón, P. Gärdenfors and D. Makinson. *On the logic of theory change: Partial meet contraction and revision functions*. Journal of Symbolic Logic, 50, 510 - 530, 1985.

Belief Change

Consider the following beliefs of a rational agent:

- p_1 All Europeans swans are white.
- p_2 The bird caught in the trap is a swan.
- p_3 The bird caught in the trap comes from Sweden.
- p_4 Sweden is part of Europe.

Belief Change

Consider the following beliefs of a rational agent:

- p_1 All Europeans swans are white.
- p_2 The bird caught in the trap is a swan.
- p_3 The bird caught in the trap comes from Sweden.
- p_4 Sweden is part of Europe.

Thus, the agent believes:

- q 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 ($\neg q$).

Belief Change

Consider the following beliefs of a rational agent:

- p_1 All Europeans swans are white.
- p_2 The bird caught in the trap is a swan.
- p_3 The bird caught in the trap comes from Sweden.
- p_4 Sweden is part of Europe.

Thus, the agent believes:

- q The bird caught in the trap is white.

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

Belief Change

Consider the following beliefs of a rational agent:

p_1 All Europeans swans are white.

p_2 The bird caught in the trap is a swan.

p_3 The bird caught in the trap comes from Sweden.

p_4 Sweden is part of Europe.

Thus, the agent believes:

q The bird caught in the trap is white.

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

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

Belief Change

Consider the following beliefs of a rational agent:

- p_1 All Europeans swans are white.
- p_2 The bird caught in the trap is a swan.
- p_3 The bird caught in the trap comes from Sweden.
- p_4 Sweden is part of Europe.

Thus, the agent believes:

- q The bird caught in the trap is white.

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

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

There are several logically consistent ways to incorporate $\neg q$!

Belief Change

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

Belief Change

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

1. The resulting theory squares with the experience;
2. It is simple; and
3. The choices disturb the original theory as little as possible.

Belief Change

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

1. The resulting theory squares with the experience;
2. It is simple; and
3. 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.

The Theory of Belief Revision

C. Alchourrón, P. Gärdenfors and D. Makinson. *On the logic of theory change: Partial meet contraction and revision functions*. Journal of Symbolic Logic, 50, 510 - 530, 1985.

Hans Rott. *Change, Choice and Inference: A Study of Belief Revision and Nonmonotonic Reasoning*. Oxford University Press, 2001.

A.P. Pedersen and H. Arló-Costa. *"Belief Revision."*. In Continuum Companion to Philosophical Logic. Continuum Press, 2011.

Epistemic States

- Belief sets
- (Ellis's belief systems)
- Possible worlds models
- (Doyle's truth maintenance systems)
- Spohn's generalized possible worlds model
- Bayesian models
- Generalized Bayesian models
- (Johnson-Laird mental models)
- ...

Epistemic States

- Belief sets
- (Ellis's belief systems)
- Possible worlds models
- (Doyle's truth maintenance systems)
- Spohn's generalized possible worlds model
- Bayesian models
- Generalized Bayesian models
- (Johnson-Laird mental models)
- ...

J. Halpern. *Reasoning about uncertainty*. The MIT Press, 2003.

AGM

belief := sentence (in some formal language)

beliefs of an agent := a set of such sentences (belief set)

Language of Beliefs in AGM:

propositional logic: propositions p,q,r,...
connectives: negation ($\neg$), conjunction ($\wedge$), disjunction ($\vee$), implication ($\rightarrow$), and equivalence ($\leftrightarrow$).

1. Belief sets should be consistent
2. Belief sets should be closed under logical consequence

Classical Consequence

For any set A of sentences, $Cn(A)$ is the set of logical consequences of A. Cn is a function from sets of sentences to sets of sentences that satisfies the following three conditions:

- $A \subseteq Cn(A)$ (inclusion);
- If $A \subseteq B$, then $Cn(A) \subseteq Cn(B)$ (monotony);
- $Cn(A) = Cn(Cn(A))$ (idempotence)

If p can be derived from A by classical propositional logic, then $p \in Cn(A)$.

K is a belief set just in case $K = Cn(K)$.

K is a belief set just in case $K = Cn(K)$.

- Logical omniscience; explicit vs. implicit beliefs; belief sets are theories.

K is a belief set just in case $K = Cn(K)$.

- Logical omniscience; explicit vs. implicit beliefs; belief sets are theories.
- "A belief set is not what you actually believe, but what you are committed to believe" (Levi 1991).

K is a belief set just in case $K = Cn(K)$.

- Logical omniscience; explicit vs. implicit beliefs; belief sets are theories.
- "A belief set is not what you actually believe, but what you are committed to believe" (Levi 1991).
- Belief bases instead of belief sets. $B_1 = \{p, p \leftrightarrow q\}$, $B_2 = \{p, q\}$.

K is a belief set just in case $K = Cn(K)$.

- Logical omniscience; explicit vs. implicit beliefs; belief sets are theories.
- "A belief set is not what you actually believe, but what you are committed to believe" (Levi 1991).
- Belief bases instead of belief sets. $B_1 = \{p, p \leftrightarrow q\}$, $B_2 = \{p, q\}$. $Cn(B_1) = Cn(B_2)$.

K is a belief set just in case $K = Cn(K)$.

- Logical omniscience; explicit vs. implicit beliefs; belief sets are theories.
- "A belief set is not what you actually believe, but what you are committed to believe" (Levi 1991).
- Belief bases instead of belief sets. $B_1 = \{p, p \leftrightarrow q\}$, $B_2 = \{p, q\}$. $Cn(B_1) = Cn(B_2)$. What happens when we receive the evidence that $\neg p$?

$$K \circ \varphi$$

$K \circ \varphi$

Initial set of beliefs

$$K \circ \varphi$$

Initial set of beliefs

New evidence φ

Belief change operator: $\circ : \mathcal{B} \times \mathcal{L} \to \mathcal{B}$

$$K \circ \varphi$$

Initial set of beliefs

New evidence φ

Minimal Change

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

"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"

(Gardenfors 1988, p. 53).

Keep the *Most Entrenched* Beliefs

"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).

Belief Change

Suppose that K is the current beliefs.

If you give priority to the new information φ:

Belief Change

Suppose that K is the current beliefs.

If you give priority to the new information φ:

1. **Expansion**: $K + \varphi$; φ is added to K giving a new belief set K'.

Belief Change

Suppose that K is the current beliefs.

If you give priority to the new information φ:

1. **Expansion**: $K + \varphi$; φ is added to K giving a new belief set K'.

2. **Contraction**: $K \dot{-} \varphi$; φ is removed from K given a new belief set K'

Belief Change

Suppose that K is the current beliefs.

If you give priority to the new information φ:

1. **Expansion**: $K + \varphi$; φ is added to K giving a new belief set K'.
2. **Contraction**: $K \dot{-} \varphi$; φ is removed from K given a new belief set K'
3. **Revision**: $K * \varphi$; φ is added and other things are removed, so that the resulting new belief set K' is consistent.

Contraction Postulates

$(C1)$ $K \dot{-} \alpha$ is deductively closed

$(C2)$ $K \dot{-} \alpha \subseteq K$

$(C3)$ If $\alpha \not\in K$ or $\vdash \alpha$ then $K \dot{-} \alpha = K$

$(C4)$ If $\not\vdash \alpha$, then $\alpha \not\in K \dot{-} \alpha$

$(C5)$ If $\vdash \alpha \leftrightarrow \beta$, then $K \dot{-} \alpha = K \dot{-} \beta$

$(C6)$ $K \subseteq Cn((K \dot{-} \alpha) \cup \{\alpha\})$

Levi Identity

$$K * \varphi = (K \dot{-} \varphi) + \varphi$$

AGM Postulates

AGM 1: $K * \varphi$ is deductively closed

AGM 2: $\varphi \in K * \varphi$

AGM 3: $K * \varphi \subseteq Cn(K \cup \{\varphi\})$

AGM 4: If $\neg\varphi \notin K$ then $K * \varphi = Cn(K \cup \{\varphi\})$

AGM 5: $K * \varphi$ is inconsistent only if φ is inconsistent

AGM 6: If φ and ψ are logically equivalent then $K * \varphi = K * \psi$

AGM 7: $K * (\varphi \wedge \psi) \subseteq Cn(K * \varphi \cup \{\psi\})$

AGM 8: if $\neg\psi \notin K * \varphi$ then $Cn(K * \varphi \cup \{\psi\}) \subseteq K * (\varphi \wedge \psi)$

Counterexample to AGM 2

$$\varphi \in K * \varphi$$

Counterexample to AGM 2

$$\varphi \in K * \varphi$$

You are walking down a street and see someone holding a sign reading "The World will End Tomorrow", but you don't add this add this to your beliefs.

Counterexample to AGM 2

$$\varphi \in K * \varphi$$

You are walking down a street and see someone holding a sign reading "The World will End Tomorrow", but you don't add this add this to your beliefs.

Two people, Ann and Bob, are reliable sources of information on whether The Netherlands will win the world cup. They are equally reliable.

Counterexample to AGM 2

$$\varphi \in K * \varphi$$

You are walking down a street and see someone holding a sign reading "The World will End Tomorrow", but you don't add this add this to your beliefs.

Two people, Ann and Bob, are reliable sources of information on whether The Netherlands will win the world cup. They are equally reliable. AGM assumes that the most recent evidence that you received takes precedent. Ann says "yes" and a little bit later, Bob says "no".

Counterexample to AGM 2

$$\varphi \in K * \varphi$$

You are walking down a street and see someone holding a sign reading "The World will End Tomorrow", but you don't add this add this to your beliefs.

Two people, Ann and Bob, are reliable sources of information on whether The Netherlands will win the world cup. They are equally reliable. AGM assumes that the most recent evidence that you received takes precedent. Ann says "yes" and a little bit later, Bob says "no". Why should the, possibly arbitrary, order in which you receive the information give more weight to Bob's announcement?

Counterexample to AGM 2

$$\varphi \in K * \varphi$$

You are walking down a street and see someone holding a sign reading "The World will End Tomorrow", but you don't add this add this to your beliefs.

Two people, Ann and Bob, are reliable sources of information on whether The Netherlands will win the world cup. They are equally reliable. AGM assumes that the most recent evidence that you received takes precedent. Ann says "yes" and a little bit later, Bob says "no". Why should the, possibly arbitrary, order in which you receive the information give more weight to Bob's announcement? Is this a counterexample to AGM 2?

Counterexample to AGM 2

$$\varphi \in K * \varphi$$

You are walking down a street and see someone holding a sign reading "The World will End Tomorrow", but you don't add this add this to your beliefs.

Two people, Ann and Bob, are reliable sources of information on whether The Netherlands will win the world cup. They are equally reliable. AGM assumes that the most recent evidence that you received takes precedent. Ann says "yes" and a little bit later, Bob says "no". Why should the, possibly arbitrary, order in which you receive the information give more weight to Bob's announcement? Is this a counterexample to AGM 2? No (Why?)

Rott's Counterexample

AGM 7: $K * (\varphi \wedge \psi) \subseteq Cn(K * \varphi \cup \{\psi\})$

AGM 8: if $\neg\psi \notin K * \varphi$ then $Cn(K * \varphi \cup \{\psi\}) \subseteq K * (\varphi \wedge \psi)$

So, if $\psi \in Cn(\{\varphi\})$, then $K * \varphi = Cn(K * \varphi \cup \{\psi\})$

Rott's Counterexample

There is an appointment to be made in a philosophy department. The position is a metaphysics position, and there are three main candidates: Andrew, Becker and Cortez.

1. Andrew is clearly the best metaphysician, but is weak in logic.
2. Becker is a very good metaphysician, also good in logic.
3. Cortez is a brilliant logician, but weak in metaphysics.

Rott's Counterexample

There is an appointment to be made in a philosophy department. The position is a metaphysics position, and there are three main candidates: Andrew, Becker and Cortez.

1. Andrew is clearly the best metaphysician, but is weak in logic.
2. Becker is a very good metaphysician, also good in logic.
3. Cortez is a brilliant logician, but weak in metaphysics.

Scenario 1: Paul is told by the dean, that the chosen candidate is either Andrew or Becker. Since Andrew is clearly the better metaphysician of the two, Paul concludes that the winning candidate will be Andrew.

Rott's Counterexample

1. Andrew is clearly the best metaphysician, but is weak in logic.
2. Becker is a very good metaphysician, also good in logic.
3. Cortez is a brilliant logician, but weak in metaphysics.

Scenario 2: Paul is told by the dean that the chosen candidate is either Andrew, Becker or Cortez.

Rott's Counterexample

1. Andrew is clearly the best metaphysician, but is weak in logic.
2. Becker is a very good metaphysician, also good in logic.
3. Cortez is a brilliant logician, but weak in metaphysics.

Scenario 2: Paul is told by the dean that the chosen candidate is either Andrew, Becker or Cortez.

" This piece of information sets off a rather subtle line of reasoning. Knowing that Cortez is a splendid logician, but that he can hardly be called a metaphysician, Paul comes to realize that his background assumption that expertise in the field advertised is the decisive criterion for the appointment cannot be upheld. Apparently, competence in logic is regarded as a considerable asset by the selection committee." Paul concludes Becker will be hired.

"...Rott seems to take the point about meta-information to explain why the example conflicts with the theoretical principles,

"...Rott seems to take the point about meta-information to explain why the example conflicts with the theoretical principles, whereas I want to conclude that it shows why the example does not conflict with the theoretical principles, since I take the relevance of the meta-information to show that the conditions for applying the principles in question are not met by the example.

"...Rott seems to take the point about meta-information to explain why the example conflicts with the theoretical principles, whereas I want to conclude that it shows why the example does not conflict with the theoretical principles, since I take the relevance of the meta-information to show that the conditions for applying the principles in question are not met by the example.... I think proper attention to the relation between concrete examples and the abstract models will allow us to reconcile some of the beautiful properties [of the abstract theory of belief revision] with the complexity of concrete reasoning." (Stalnaker, 204)

Counterexamples to Recovery

$$K \subseteq Cn((K \dot{-} \alpha) \cup \{\alpha\})$$

While reading a book about Cleopatra I learned that she had both a son and a daughter. I therefore believe both that Cleopatra had a son (s) and Cleopatra had a daughter (d).

Counterexamples to Recovery

$$K \subseteq Cn((K \dot{-} \alpha) \cup \{\alpha\})$$

While reading a book about Cleopatra I learned that she had both a son and a daughter. I therefore believe both that Cleopatra had a son (s) and Cleopatra had a daughter (d). Later I learn from a well-informed friend that the book in question is just a historical novel. I accordingly contract my belief that Cleopatra had a child ($s \vee d$).

Counterexamples to Recovery

$$K \subseteq Cn((K \dot{-} \alpha) \cup \{\alpha\})$$

While reading a book about Cleopatra I learned that she had both a son and a daughter. I therefore believe both that Cleopatra had a son (s) and Cleopatra had a daughter (d). Later I learn from a well-informed friend that the book in question is just a historical novel. I accordingly contract my belief that Cleopatra had a child ($s \vee d$). However, shortly thereafter I learn from a reliable source that in fact Cleopatra had a child. I thereby reintroduce $s \vee d$ to my collection of beliefs without also returning either s or d.

(Hansson, 1991)

Counterexamples to Recovery (C6)

$$K \subseteq Cn((K \dot{-} \alpha) \cup \{\alpha\})$$

I believed both that George is a criminal (c) and George is a mass murderer (m).

Counterexamples to Recovery (C6)

$$K \subseteq Cn((K \dot{-} \alpha) \cup \{\alpha\})$$

I believed both that George is a criminal (c) and George is a mass murderer (m). Upon receiving certain information I am induced to retract my belief set K by my belief that George is a criminal (c). Of course, I therefore retract my belief set by my belief that George is a mass murderer (m).

Counterexamples to Recovery (C6)

$$K \subseteq Cn((K \dot{-} \alpha) \cup \{\alpha\})$$

I believed both that George is a criminal (c) and George is a mass murderer (m). Upon receiving certain information I am induced to retract my belief set K by my belief that George is a criminal (c). Of course, I therefore retract my belief set by my belief that George is a mass murderer (m). Later I learn that in fact George is a shoplifter (s), so I expand my contracted belief set $K \dot{-} c$ by s to obtain $(K \dot{-} c) + s$.

Counterexamples to Recovery (C6)

$$K \subseteq Cn((K \dot{-} \alpha) \cup \{\alpha\})$$

I believed both that George is a criminal (c) and George is a mass murderer (m). Upon receiving certain information I am induced to retract my belief set K by my belief that George is a criminal (c). Of course, I therefore retract my belief set by my belief that George is a mass murderer (m). Later I learn that in fact George is a shoplifter (s), so I expand my contracted belief set $K \dot{-} c$ by s to obtain $(K \dot{-} c) + s$. As George's being a shoplifter (s) entails his being a criminal (c), $(K \dot{-} c) + c$ is a subset of $(K \dot{-} c) + s$. Yet by Recovery it follows that $K \subseteq (K \dot{-} c) + c$, so m is a member of the expanded belief, so m is a member of the expanded belief set $(K \dot{-} c) + s$.

Counterexamples to Recovery (C6)

$$K \subseteq Cn((K \dot{-} \alpha) \cup \{\alpha\})$$

I believed both that George is a criminal (c) and George is a mass murderer (m). Upon receiving certain information I am induced to retract my belief set K by my belief that George is a criminal (c). Of course, I therefore retract my belief set by my belief that George is a mass murderer (m). Later I learn that in fact George is a shoplifter (s), so I expand my contracted belief set $K \dot{-} c$ by s to obtain $(K \dot{-} c) + s$. As George's being a shoplifter (s) entails his being a criminal (c), $(K \dot{-} c) + c$ is a subset of $(K \dot{-} c) + s$. Yet by Recovery it follows that $K \subseteq (K \dot{-} c) + c$, so m is a member of the expanded belief, so m is a member of the expanded belief set $(K \dot{-} c) + s$. But I do not believe that George is a mass murderer (m). (Hansson, 1996)

Evaluating counterexamples

. . . information about how I learn some of the things I learn, about the sources of my information, or about what I believe about what I believe and don't believe. If the story we tell in an example makes certain information about any of these things relevant, then it needs to be included in a proper model of the story, if it is to play the right role in the evaluation of the abstract principles of the model.

Robert Stalnaker. *Iterated Belief Revision*. Erkenntnis 70, pp. 189 - 209, 2009.

Let K be a belief set and φ a formula.

$K \perp \varphi$ is the **remainder set** of K.

$A \in K \perp \varphi$ iff

1. $A \subseteq K$
2. $\varphi \notin Cn(A)$
3. There is no B such that $A \subset B \subseteq K$ and $\varphi \notin Cn(B)$.

- $K \perp \alpha = \{K\}$ iff $\neg\alpha \notin Cn(K)$
- $K \perp \alpha = \emptyset$ iff $\alpha \in Cn(\emptyset)$
- If $K' \subseteq K$ and $\alpha \notin Cn(K')$ then there is some T such that $K' \subseteq T \in K \perp \alpha$.

A **selection function** γ for K is a function on $K \bot \alpha$ such that:

- If $K \bot \alpha \neq \emptyset$, then $\gamma(K \bot \alpha) \subseteq K \bot \alpha$ and $\gamma(K \bot \alpha) \neq \emptyset$
- If $K \bot \alpha = \emptyset$, then $\gamma(K \bot \alpha) = \{K\}$

Let K be a set of formulas. A function $\dot{-}$ is a **partial meet contraction** for K if there is a selection function γ for K such that for all formula α:

$$K \dot{-} \alpha = \bigcap \gamma(K \perp \alpha)$$

Let K be a set of formulas. A function $\dot{-}$ is a **partial meet contraction** for K if there is a selection function γ for K such that for all formula α:

$$K \dot{-} \alpha = \bigcap \gamma(K \perp \alpha)$$

Then $K * \alpha = Cn(\bigcap \gamma(K \perp \neg\alpha) \cup \{\alpha\})$

Let K be a set of formulas. A function $\dot{-}$ is a **partial meet contraction** for K if there is a selection function γ for K such that for all formula α:

$$K \dot{-} \alpha = \bigcap \gamma(K \perp \alpha)$$

Then $K * \alpha = Cn(\bigcap \gamma(K \perp \neg\alpha) \cup \{\alpha\})$

- γ selects exactly one element of $K \perp \alpha$ (maxichoice contraction)
- γ selects the entire set $K \perp \alpha$ (full meet contraction)

AGM Postulates

AGM 1: $K * \varphi$ is deductively closed

AGM 2: $\varphi \in K * \varphi$

AGM 3: $K * \varphi \subseteq Cn(K \cup \{\varphi\})$

AGM 4: If $\neg\varphi \notin K$ then $K * \varphi = Cn(K \cup \{\varphi\})$

AGM 5: $K * \varphi$ is inconsistent only if φ is inconsistent

AGM 6: If φ and ψ are logically equivalent then $K * \varphi = K * \psi$

AGM 7: $K * (\varphi \wedge \psi) \subseteq Cn(K * \varphi \cup \{\psi\})$

AGM 8: if $\neg\psi \notin K * \varphi$ then $Cn(K * \varphi \cup \{\psi\}) \subseteq K * (\varphi \wedge \psi)$

Theorem (AGM 1985). Let K be a belief set and let $*$ be a function on $\mathcal{L}$. Then

- The function $*$ is a partial meet revision for K if and only if it satisfies the postulates $AGM1$ - $AGM6$
- The function $*$ is a transitively relational partial meet revision for K if and only if it satisfies $AGM1$ - $AGM8$.

Belief Revision: The Semantic View

A. Grove. *Two modelings for theory change*. Journal of Philosophical Logic, 17, pgs. 157 - 170, 1988.

EP. *Dynamic Epistemic Logic II: Logics of information change*. Philosophy Compass, Vol. 8, Iss. 9, pgs. 815 - 833, 2013.

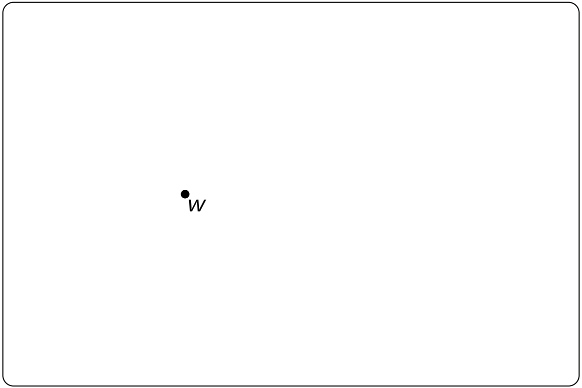

- The set of states, with a distinguished state denoted the "actual world"

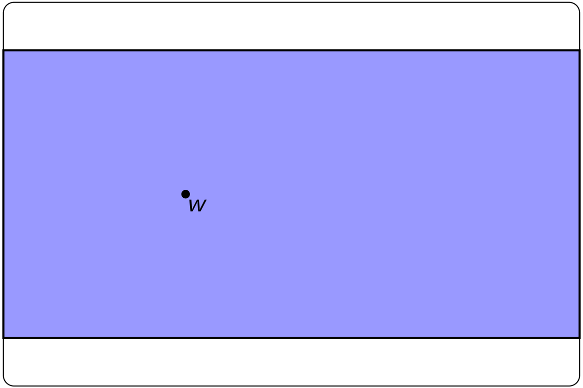

- The set of states, with a distinguished state denoted the "actual world"
- The agent's (hard) information (i.e., the states consistent with what the agent knows)

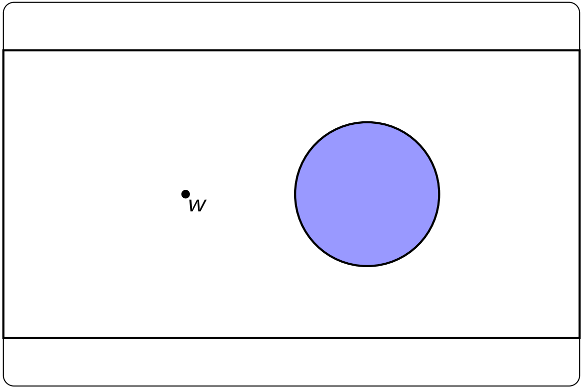

- The agent's (hard) information (i.e., the states consistent with what the agent knows)
- The agent's beliefs (soft information—-the states consistent with what the agent believes)

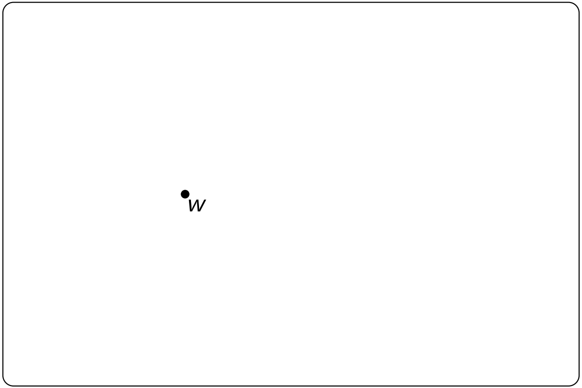

- The states consistent with what the agent knows with a distinguished state (the "actual world")
- Each state is associated with a propositional valuation for the underlying propositional language

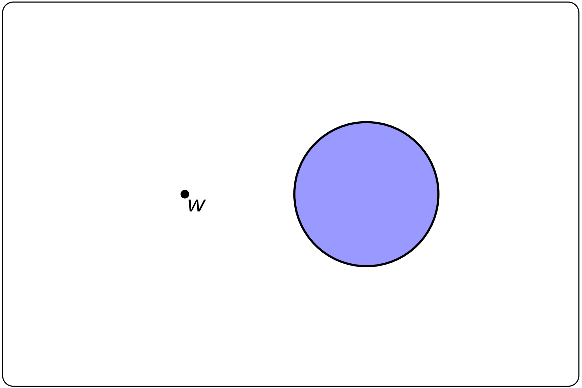

- The agent's beliefs (soft information—-the states consistent with what the agent believes)

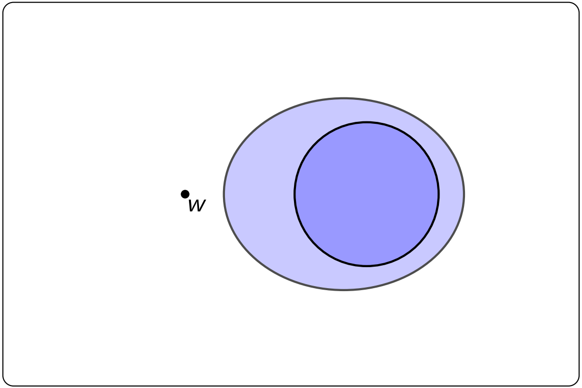

- The agent's beliefs (soft information—the states consistent with what the agent believes)
- The agent's "contingency plan": when the stronger beliefs fail, go with the weaker ones.

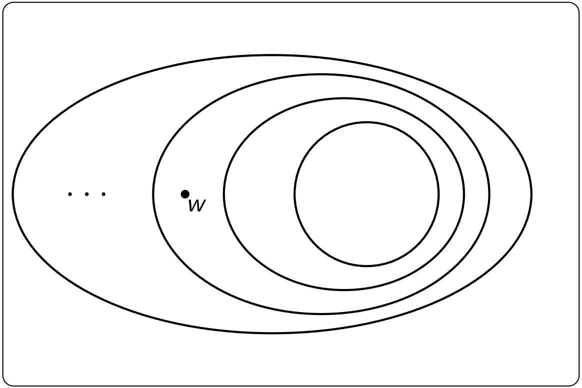

- The agent's beliefs (soft information—the states consistent with what the agent believes)
- The agent's "contingency plan": when the stronger beliefs fail, go with the weaker ones.

Sphere Models

Let W be a set of states, A set $\mathcal{F} \subseteq \wp(W)$ is called a system of spheres provided:

- For each $S, S' \in \mathcal{F}$, either $S \subseteq S'$ or $S' \subseteq S$
- For any $P \subseteq W$ there is a smallest $S \in \mathcal{F}$ (according to the subset relation) such that $P \cap S \neq \emptyset$
- The spheres are non-empty $\bigcap \mathcal{F} \neq \emptyset$ and cover the entire information cell $\bigcup \mathcal{F} = W$

Let $\mathcal{F}$ be a system of spheres on W: for $w, v \in W$, let

$$w \preceq_{\mathcal{F}} v \text{ iff for all } S \in \mathcal{F}, \text{ if } v \in S \text{ then } w \in S$$

Then, $\preceq_{\mathcal{F}}$ is reflexive, transitive, and well-founded.

$w \preceq_{\mathcal{F}} v$ means that no matter what the agent learns in the future, as long as world v is still consistent with his beliefs and w is still epistemically possible, then w is also consistent with his beliefs.

Belief Revision via Plausibility

- $W = \{w_1, w_2, w_3\}$

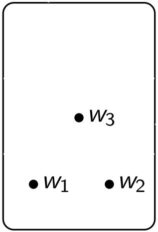

Belief Revision via Plausibility

- $W = \{w_1, w_2, w_3\}$
- $w_1 \preceq w_2$ and $w_2 \preceq w_1$ (w_1 and w_2 are equi-plausbile)
- $w_1 \prec w_3$ ($w_1 \preceq w_3$ and $w_3 \not\preceq w_1$)
- $w_2 \prec w_3$ ($w_2 \preceq w_3$ and $w_3 \not\preceq w_2$)

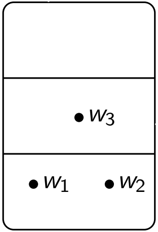

Belief Revision via Plausibility

- $W = \{w_1, w_2, w_3\}$
- $w_1 \preceq w_2$ and $w_2 \preceq w_1$ (w_1 and w_2 are equi-plausbile)
- $w_1 \prec w_3$ ($w_1 \preceq w_3$ and $w_3 \not\preceq w_1$)
- $w_2 \prec w_3$ ($w_2 \preceq w_3$ and $w_3 \not\preceq w_2$)
- $\{w_1, w_2\} \subseteq Min_{\preceq}([w_i])$

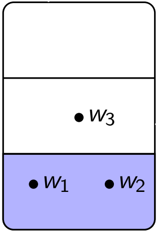

Belief Revision via Plausibility

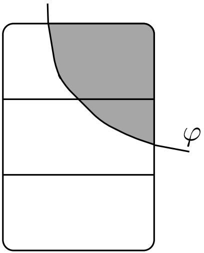

Conditional Belief: $B^{\varphi}\psi$

Belief Revision via Plausibility

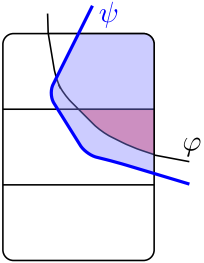

Conditional Belief: $B^\varphi \psi$

$$Min_{\preceq}([\![\varphi]\!]_{\mathcal{M}}) \subseteq [\![\psi]\!]_{\mathcal{M}}$$

Belief Revision via Plausibility

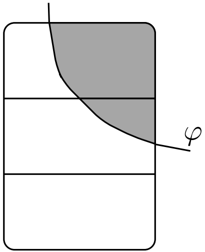

Incorporate the new information φ

Belief Revision via Plausibility

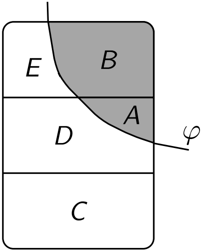

Incorporate the new information φ

Belief Revision via Plausibility

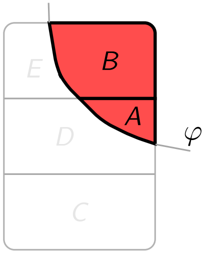

Public Announcement: Information from an infallible source

$(!\varphi)$: $A \prec_i B$

Belief Revision via Plausibility

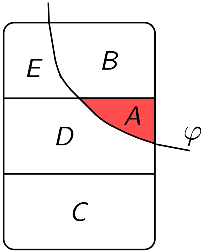

Public Announcement: Information from an infallible source ($!\varphi$): $A \prec_i B$

Conservative Upgrade: Information from a trusted source ($\uparrow\varphi$): $A \prec_i C \prec_i D \prec_i B \cup E$

Belief Revision via Plausibility

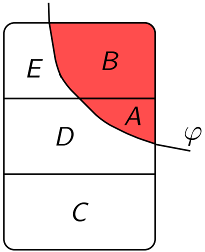

Public Announcement: Information from an infallible source
$(!\varphi)$: $A \prec_i B$

Conservative Upgrade: Information from a trusted source
$(\uparrow\varphi)$: $A \prec_i C \prec_i D \prec_i B \cup E$

Radical Upgrade: Information from a strongly trusted source
$(\Uparrow\varphi)$: $A \prec_i B \prec_i C \prec_i D \prec_i E$

Revision vs. Update

Suppose φ is some incoming information that should be incorporated into the agents beliefs (represented by a theory T).

Revision vs. Update

Suppose φ is some incoming information that should be incorporated into the agents beliefs (represented by a theory T).

A subtle difference:

- If φ describes facts about the current state of affairs
- If φ describes facts that have possible become true only after the original beliefs were formed.

Revision vs. Update

Suppose φ is some incoming information that should be incorporated into the agents beliefs (represented by a theory T).

A subtle difference:

- If φ describes facts about the current state of affairs
- If φ describes facts that have possible become true only after the original beliefs were formed.

Complete vs. incomplete belief sets:
$K = Cn(\{p \vee q\})$ vs. $K = Cn(\{p \vee q, p, q\})$

Revision vs. Update

Suppose φ is some incoming information that should be incorporated into the agents beliefs (represented by a theory T).

A subtle difference:

- If φ describes facts about the current state of affairs
- If φ describes facts that have possible become true only after the original beliefs were formed.

Complete vs. incomplete belief sets:
$K = Cn(\{p \vee q\})$ vs. $K = Cn(\{p \vee q, p, q\})$

Revising by $\neg p$ ($K * \neg p$) vs. Updating by $\neg p$ ($K \diamond \neg p$)

H. Katsuno and A. O. Mendelzon. *Propositional knowledge base revision and minimal change*. Artificial Intelligence, 52, pp. 263 - 294 (1991).

KM Postulates

KM 1: $K \diamond \varphi = Cn(K \diamond \varphi)$

KM 2: $\varphi \in K \diamond \varphi$

KM 3: If $\varphi \in K$ then $K \diamond \varphi = K$

KM 4: $K \diamond \varphi$ is inconsistent iff φ is inconsistent

KM 5: If φ and ψ are logically equivalent then $K \diamond \varphi = K \diamond \psi$

KM 6: $K \diamond (\varphi \wedge \psi) \subseteq Cn(K \diamond \varphi \cup \{\psi\})$

KM 7: If $\psi \in K \diamond \varphi$ and $\varphi \in K \diamond \psi$ then $K \diamond \varphi = K \diamond \psi$

KM 8: If K is complete then $K \diamond (\varphi \wedge \psi) \subseteq K \diamond \varphi \cap K \diamond \psi$

KM 9: $K \diamond \varphi = \bigcap_{M \in Comp(K)} M \diamond \varphi$, where $Comp(K)$ is the class of all complete theories containing K.

Updating and Revising

$$K \diamond \varphi = \bigcap_{M \in Comp(K)} M * \varphi$$

H. Katsuno and A. O. Mendelzon. *On the difference between updating a knowledge base and revising it*. *Belief Revision*, P. Gärdenfors (ed.), pp 182 - 203 (1992).