# Conditional Logic and Belief Revision

Ginger Schultheis (vks@mit.edu) and David Boylan (dboylan@mit.edu) January 2017

## *History*

The formal study of belief revision grew out out of two research traditions:

- ▷ Computer science: computer scientists develop procedures by which databases can be updated.

Jon Doyle (1979) on truth maintenance systems, and Fagin et al (1983) on database priorities.

- ▷ Philosophy
  - ○ Defeasible reasoning (e.g., inductive inference): what is it for one belief to be a *prima facia* reason for another? what is it for belief to be a *defeater* for another?
  - ○ Scientific theory change: a *scientific corpus* is a body of information that's accepted as true. A scientific corpus can change if new data is brought forward; or if a *new theory* is introduced that explains the data just as well. Are there general rules governing these changes?

## *AGM Theory*

### Overview

- ▷ Beliefs are represented by a set of sentences: the **belief set**. This set is logically closed—every sentence that follows logically from this set is already in the set.
- ▷ There are three types of belief change:
  - ○ **Contraction**: a belief set $K$ is superseded by a set $K \div p$, a subset of $K$ that does not contain $p$.
  - ○ **Expansion**: a belief set $K$ is superseded by a set $K + p$, the smallest logically closed set containing $K$ and $p$.
  - ○ **Revision**: a sentence $p$ is added to $K$, and other sentences are removed when needed to make the new set $K^* + p$ consistent.
- ▷ Language and notation:
  - ○ Sentences of the language are represented by lowercase letters ($p$, $q$, $r$ ... ); the language also contains the truth-functional connectives: $\neg, \wedge, \vee, \rightarrow$ , $\leftrightarrow$; sets of sentences are represented by uppercase letters ($A$, $B$, $C$ ... )
  - ○ For any set of sentences $A$, define Cn(A) as the set of logical consequences of $A$. Cn is a *consequence operation* because it satisfies three conditions:
    - ▷ **Inclusion**: $A \subseteq Cn(A)$
    - ▷ **Monotony**: If $A \subseteq B$, then $Cn(A) \subseteq Cn(B)$
    - ▷ **Iteration**: $Cn(A) = Cn(Cn(A))$

- Cn(∅) is the set of tautologies.

## Expansion

- The expansion of $K$ by $p$ is denoted $K+p$ and is defined: $K+p = Cn(K \cup \{p\})$

## Contraction

Expansion is easy. Contraction is harder. When removing a sentence $p$ from $K$, other sentences will often have to be removed to make sure $K \div p$ is closed under consequence. But there will often be many ways to to remove sentences from $K$, each of which leaves with a different consistent belief set.

An example: K contains $p, q, p \wedge q \rightarrow r$, and their logical consequences. We want to contract $K$ by removing $r$. At least one of $p$, $q$, or $p \wedge q \rightarrow r$ must also be deleted. How do we make this choice?

There are two things to do when giving a theory of contraction. The first is to formulate a set of *postulates* that any contraction operation should satisfy. The second is to construct an actual contraction operation (or a class of operations) that satisfies those postulates.

- AGM Contraction Postulates
  - **Closure**: $K \div p$ = Cn($K \div p$)
  - **Success**: If p ∉ Cn(∅), then p ∉ Cn($K \div p$)
  - **Inclusion**: $K \div p \subseteq K$
  - **Vacuity**: If $p \notin Cn(K)$, then $K \div p = K$
  - **Extensionality**: If $p \leftrightarrow q \in Cn(\emptyset)$, then $K \div p = K \div q$.
  - **Recovery**: $K \subseteq (K \div p) + p$

The contracted set is logically closed.

If $p$ is not a tautology, then the contracted set shouldn't imply $p$.

The contracted set is a subset of the original.

Logically equivalent sentences should be treated alike in contraction.

Any beliefs you lose when you remove $p$ will be regained if you later expand $K$ with $p$

- Contraction operation: Partial meet contraction
  - For any set $A$ and any sentence $p$, $A \bot p$ is the **remainder set** of $A$: the set of maximal subsets of $A$ that do not imply $p$. More precisely, a set $B$ is a member of $A \bot p$ just in case:
    - $B \subseteq A$
    - $p \notin B$
    - There is no set $C$ such that $B \subset C \subseteq A$ and $p \notin C$
  - The members of $A \bot p$ are called $p$-**remainders** of A.
  - Should we identify $K \div p$ with *one of the members* of $K \bot p$? No. Why not?
  - **Partial meet contraction** tells us to use a **selection function** $\gamma$ that selects the 'best' elements of $K \bot p$. Then identify $K \div p$ with the *intersection* of these. Formally:
    - $\div$ is a partial meet contraction operation just in case: $K \div p = \bigcap \gamma(K \bot p)$

Partial meet contraction has been shown to satisfy the six postulates mentioned above.

- Recovery
  - Recovery is controversial. Here's a worrisome example:

> I believe that George is mass murderer ($m$) and hence that George is a criminal ($c$). Then I'm told that George is not a criminal, so I drop that belief as well as my belief that he is a mass murderer. I am then told that George is a shoplifter ($s$).
>
> My belief set now is the expansion of $K \div c$ by $s$. Since $s$ implies $c$, $(K \div c) + c \subseteq (K \div c) + s$. By recovery, $(K \div c) + c$ includes $m$. It follows that $(K \div c) + s$ includes $m$.
>
> Since I previously believed George was a mass murderer, I can no longer believe him a shoplifter without believing him a mass murderer.

<u>Revision</u>

▷ Revision Operation

- The revision operator * has two tasks: (i) add the new belief $p$ to the belief set $K$, and (ii) ensure that the resulting belief set $K * p$ is consistent (as long as $p$ is consistent).
- We accomplish (i) by expanding $K$ with $p$; we accomplish (ii) by contracting $K$ by $\neg p$ before adding $p$. This two-step process is reflected in the following definition of the revision operator:
  - ▷ **Levi Identity**: $K * p = (K \div \neg p) + p$
- If $\div$ is the partial meet contraction operator, then * is the **partial meet revision** operator. This is the standard revision operator in the AGM model.

▷ Revision Postulates: An operator * is a partial meet revision operator iff it satisfies the following six postulates:

- **Closure**: $K * p = Cn(K * p)$
- **Success**: $p \in K * p$
- **Inclusion**: $K * p \subseteq K + p$
- **Vacuity**: If $\neg p \notin K$, then $K * p = K + p$
- **Consistency**: $K * p$ is consistent if $p$ is.
- **Extensionality**: If $p \leftrightarrow q \in Cn(\emptyset)$, then $K * p = K * q$

The revised set is logically closed.

If $p$ is consistent, $p$ should be in the revised set.

The revised set should be a subset of the original set expanded by $p$.

If $p$ is consistent, nothing is contracted.

Extensionally equivalent sentences are treated alike in revision.

▷ Informational Economy
The main principle underling AGM revision is known as the **principle of informational economy** (also called the **principle of conservativity**).

## *Plausibility Models*

Modelling a belief state with sets of sentences is not so intuitive nor the easiest to work with. But we can use the semantics for conditionals to give possible worlds theories of belief revision.

As we will see, these theories turn out to be equivalent to AGM.

<u>Ordering Models</u>

- As with conditionals, we use an ordering on worlds. This time it represents how plausible worlds are taken to be.

Spheres are one way of representing the plausiblity ordering.

- Our models are ordered triples, $\langle W, \preceq, V \rangle$.
  - W is a set of worlds.
  - V is an evaluation function: it assigns truth values to sentences at worlds.
  - $\preceq$ is a relation with the following properties:
    1. For any two worlds $w, w'$ either $w \preceq w'$ or $w' \preceq w$.
    2. $\preceq$ is **transitive**.
    3. $\preceq$ is **reflexive**.
    4. For any sentence $\phi$, if $[\![\phi]\!]$ is consistent with $W$, then there is some $w'$ s.t. ($\preceq$ is **well-founded**.)
       (If $[\![\phi]\!]$ is consistent with $W$, i.e. is non-empty, then there is a set of **best** $\phi$-worlds.)

Here we don't take $\preceq$ to be a three place function: it simply compares the plausibility of two worlds. But in the full theory, it will be: we want what's plausible to vary from world to world.

- The idea for the belief operator is that the agent believes $\alpha$ conditional on $\phi$ just in case all the most plausible $\phi$-worlds are $\alpha$-worlds.
  (Sound familiar?)
- More formally, let $Min(\phi) = \{w : w \in [\![\phi]\!]$ and $\neg\exists w' : w' \prec w$ and $w' \in [\![\phi]\!]\}$. We then say
  $\mathcal{M}, w \vDash B^{\phi}\alpha$ iff $Min(\phi) \subseteq [\![\alpha]\!]$
- We let unconditional belief in $\phi$ be a special case of belief, the case of believing $\phi$ conditional on a tautology.
  More formally: $\mathcal{M}, w \vDash B\phi$ iff $\mathcal{M}, w \vDash B^{\top}\phi$ iff $Min(W) \subseteq [\![\phi]\!]$

## *Correspondences*

By now you might have the feeling that this all looks a lot like AGM. You would be right: the theories are equivalent.

- First, we need some notation. We will be moving back and forth between sets of setences (closed under consequence) and sets of worlds.
  Where $S$ is a set of sentences, let's say that $[\![S]\!]$ is the set of worlds which make true every sentence in $S$.
  Where $W$ is a set of worlds, let's say that $t(W)$ is the set of sentences true at every world in $W$.
- That is we can prove the following theorems:
  - **Theorem 1 (Grove)**: Let $\preceq$ be an ordering centered on $[\![K]\!]$. If we define $K * p$ to be t(f(p)), then the AGM revision postulates are satisfied.
  - **Theorem 2 (Grove)**: Let $*$ be some revision function satisfying the AGM revision postulates. Then, given a set $K$, we can construct an ordering centered on $[\![K]\!]$ such that $K * p = t(f(p))$.

## *Strengthening and Other Interesting Features*

- Like conditional logic, this system is not strongly monotonic; that is, we do not have as an axiom:

- **Strong monotonicity** $B_i^{\phi\alpha} \supset B_i^{\phi\wedge\psi}\alpha$

  Can you show why not?

▷ However, we do have some weaker monotonicity principles. In particular we have:

- **Cautious Monotonicity**: $B^{\phi}\alpha \wedge B^{\phi}\beta \supset B^{\phi\wedge\beta}\alpha$

  Suppose that $B^{\phi}\alpha \wedge B^{\phi}\beta$. Then in all the closest $\phi$ worlds $\alpha$ and $\beta$ are true. Call this set S. S is the set of closest $\phi \wedge \beta$ worlds. We already said that all these worlds are $\alpha$ worlds, so $B^{\phi\wedge\beta}\alpha$ is true.

- **Rational Monotonicity**: $(B^{\phi}\alpha \wedge \neg B^{\phi}\neg\beta) \supset B^{\phi\wedge\beta}\alpha$.

  Yesterday we showed that if $\leq$ is almost connected, then RM holds. But this result is will show the same thing for the belief revision version of RM.

  All we need to do is show that in this case $\leq$ is indeed almost connected, but this follows straightforwardly from the first property we imposed on $\leq$.

▷ A related interesting question is whether $B(\mathit{If}\ \phi\ \mathit{then}\ \psi) \vDash B^{\phi}\psi$.

Various *triviality results* about this have been proved. Such results show that AGM plus the Ramsey test entail that there cannot be three propositions which are pairwise consistent but jointly inconsistent.

This is called a triviality result because models not allowing for three propositions in this way are in a sense trivial.

▷ Note in this connection:

$B(\phi \supset \psi) \nvDash B^{\phi}\psi$

Counterexample: Suppose all the closest worlds are $\neg\phi$ worlds and that the closest $\phi$ worlds are $\neg\psi$ worlds.

▷ This is not that shocking: the material conditional is almost surely not the indicative conditional.