Tutorial on Iterated Belief Change

IJCAI2019
Macao

RICHARD BOOTH
CARDIFF UNIVERSITY

JAKE CHANDLER
LA TROBE UNIVERSITY
LMU

Belief change

- 40+ years of research into the dynamics of belief
- Multidisciplinary venture with independent origins in:
  - database management [Doy79] [FUV83]
  - philosophy [Har76] [Lev77]
- Canonical model introduced in

  Alchourrón, Gärdenfors & Makinson's 'On the Logic of Theory Change" [AGM85]
- Substantial literature (DBPL: 559 hits for "belief revision")

Our topic of interest

- Lack of progress on a shockingly crucial question:

  What kind of beliefs can one arrive at after not one, but *two or more* successive changes in view? (problem of iterated change)
- What we will cover:
  - Introduction to the basic model of AGM
  - Recent developments and key open questions regarding iterations of its two central operations: revision and contraction

Not covered here

- Generalisations and extensions:
  - non AGM-compliant models
  - further operations (e.g. update)
- Connections to other fields:
  - nonmonotonic and conditional logics
  - belief merging and social choice
  - argumentation theory
- Issues of computability, implementation and learning power
- Dynamic doxastic logics, for reasoning about belief change
- Relation to models of probabilistic degrees of belief

Overview

Overview (cont.)

1. SINGLE-STEP BELIEF CHANGE

The problem

- Suppose we are in a state of mind such that we believe:

  The swan caught in the trap is European (A)

  If the swan is European, then it's white all over ($A \to B$)

  and hence, by deductive inference:

  The swan is white all over (B)

The problem (cont.)

- But we then change our views, so that either:

  (1) We come to acquire the belief that $\neg B$

  (The swan *isn't* white all over)

  or

  (2) We come to simply retract the belief that B

- In either case, we can't hold on to *both* A and $A \to B$
- What should our posterior beliefs look like?

More formally

- States of mind are represented by 'doxastic states' $\Psi, \Theta, \ldots$
- Ψ determines a 'belief set' $[\Psi]$: a deductively closed set of sentences drawn from propositional language L
- A revision function $*$ returns the state $\Psi * A$ resulting from adjusting Ψ to accommodate the consistent inclusion of $A \in L$
- A contraction function $\div$ returns the state $\Psi \div A$ resulting from adjusting Ψ to accommodate the retraction of $A \in L$
- What are the contents of $[\Psi * \neg B]$ and $[\Psi \div B]$?

Ordering possibilities

- One intuitive way to tackle the problem:

  Order the maximally specific consistent possibilities (possible worlds) by relative implausibility

  Minimal worlds: worlds taken to be potentially the actual one

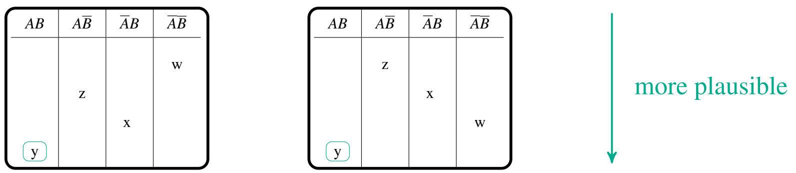

Sample orderings consistent with swan example ($[\Psi] = \mathrm{Cn}(A \wedge B)$)

1-step revision with orderings

- Revision by $\neg B$:

  Posterior minimal worlds = prior minimal $\neg B$-worlds

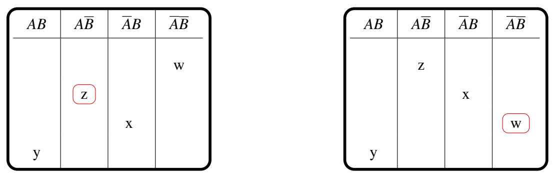

Retain A but lose $A \to B$ ($[\Psi] = \mathrm{Cn}(A \wedge \neg B)$)

Lose A but retain $A \to B$ ($[\Psi] = \mathrm{Cn}(\neg A \wedge \neg B)$)

1-step contraction with orderings

- Contraction by B:

  Posterior minimal worlds = union of prior minimal worlds with prior minimal $\neg B$-worlds

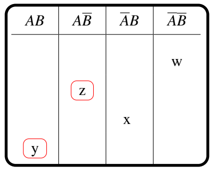

Retain A but lose $A \to B$
($[\Psi] = \mathrm{Cn}(A)$)

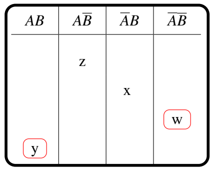

Lose A but retain $A \to B$
($[\Psi] = \mathrm{Cn}(A \leftrightarrow B)$)

More formally

Definition 1

A total preorder (TPO) $\preccurlyeq$ is a relation that is complete ($x \preccurlyeq y$ or $y \preccurlyeq x$) and transitive (if $x \preccurlyeq y$ and $y \preccurlyeq z$, then $x \preccurlyeq z$)

- Ψ is associated with a TPO $\preccurlyeq_\Psi$ over the set of worlds W, representing relative implausibility, such that $B \in [\Psi]$ iff $\min(\preccurlyeq_\Psi, W) \subseteq [\![B]\!]$
- Furthermore:

  $$B \in [\Psi * A] \text{ iff } \min(\preccurlyeq_\Psi, [\![A]\!]) \subseteq [\![B]\!]$$
  $$B \in [\Psi \div A] \text{ iff } \min(\preccurlyeq_\Psi, [\![\neg A]\!]) \cup \min(\preccurlyeq_\Psi, W) \subseteq [\![B]\!]$$

Orderings as mere representations

- One might be skeptical of the claim that we ought to believe all and only what is true in the worlds that one deems the most plausible
- These TPOs can alternatively be viewed as mere representations of constraints on 1-step revision/contraction

  (Analogy: subjective probabilities/utilities as mere representations of constraints on preferences over gambles [Sav54])
- These constraints are principles of minimal change

The AGM postulates for revision

- The AGM postulates for revision (AGM*) of Alchourrón, Gärdenfors and Makinson [AGM85]:

(K1*)	$\mathrm{Cn}([\Psi * A]) \subseteq [\Psi * A]$	
(K2*)	$A \in [\Psi * A]$	(Success)
(K3*)	$[\Psi * A] \subseteq \mathrm{Cn}([\Psi] \cup \{A\})$	
(K4*)	If $\neg A \notin [\Psi]$, then $\mathrm{Cn}([\Psi] \cup \{A\}) \subseteq [\Psi * A]$	
(K5*)	If A is consistent, then so too is $[\Psi * A]$	
(K6*)	If $\mathrm{Cn}(A) = \mathrm{Cn}(B)$, then $[\Psi * A] = [\Psi * B]$	
(K7*)	$[\Psi * (A \wedge B)] \subseteq \mathrm{Cn}([\Psi * A] \cup \{B\})$	
(K8*)	If $\neg B \notin [\Psi * A]$, then $\mathrm{Cn}([\Psi * A] \cup \{B\}) \subseteq [\Psi * (A \wedge B)]$	

The AGM postulates for contraction

- The AGM postulates for contraction (AGM$^{\div}$) [AGM85]:

(K1$^{\div}$) $\mathrm{Cn}([\Psi \div A]) \subseteq [\Psi \div A]$

(K2$^{\div}$) $[\Psi \div A] \subseteq [\Psi]$

(K3$^{\div}$) If $A \notin [\Psi]$, then $[\Psi \div A] = [\Psi]$

(K4$^{\div}$) If $A \notin \mathrm{Cn}(\varnothing)$, then $A \notin [\Psi \div A]$

(K5$^{\div}$) If $A \in [\Psi]$, then $[\Psi] \subseteq \mathrm{Cn}([\Psi \div A] \cup \{A\})$ (Recovery)

(K6$^{\div}$) If $\mathrm{Cn}(A) = \mathrm{Cn}(B)$, then $[\Psi \div A] = [\Psi \div B]$

(K7$^{\div}$) $[\Psi \div A] \cap [\Psi \div B] \subseteq [\Psi \div A \wedge B]$

(K8$^{\div}$) If $A \notin [\Psi \div A \wedge B]$, then $[\Psi \div A \wedge B] \subseteq [\Psi \div A]$

Representation theorems

Theorem 1

The following two claims are equivalent:

(a) For each state Ψ, there exists a unique TPO $\preccurlyeq^{*}_{\Psi}$ over W, such that $B \in [\Psi * A]$ iff $\min(\preccurlyeq^{*}_{\Psi}, [\![A]\!]) \subseteq [\![B]\!]$

(b) $*$ satisfies AGM*[Gro88] [KM91]

Theorem 2

The following two claims are equivalent:

(a) For each state Ψ, there exists a unique TPO $\preccurlyeq^{\div}_{\Psi}$ over W, such that $B \in [\Psi \div A]$ iff $\min(\preccurlyeq^{\div}_{\Psi}, [\![\neg A]\!]) \cup \min(\preccurlyeq^{\div}_{\Psi}, W) \subseteq [\![B]\!]$

(b) $\div$ satisfies AGM$^{\div}$

The Harper & Levi Identities

- Note that nothing so far guarantees that $\preccurlyeq^{\div}_{\Psi} = \preccurlyeq^{*}_{\Psi}$
- This is ensured by the Harper and Levi Identities [Har76] [Lev77]:

(LI) $[\Psi * A] = \mathrm{Cn}([\Psi \div \neg A] \cup \{A\})$

(HI) $[\Psi \div A] = [\Psi] \cap [\Psi * \neg A]$

- Given these:

If the 1-step revision behaviour of Ψ is TPO-representable, then so is its 1-step contraction behaviour (resp. vice versa) and furthermore $\preccurlyeq^{\div}_{\Psi} = \preccurlyeq^{*}_{\Psi}$

The Harper & Levi Identities (cont.)

- Motivation for (LI)
  - Simplest modification of $[\Psi]$ that includes A: $\mathrm{Cn}([\Psi] \cup \{A\})$
  - But $[\Psi]$ and A needn't be consistent: replace $[\Psi]$ by $[\Psi \div \neg A]$
- Informally, (HI) says:
  - Retract anything no longer endorsed, had one come to believe that $\neg A$ (since $\neg A$ considered possible)
  - Retract nothing further and introduce nothing (minimal change)

Conditional beliefs & rational consequence

- 1-step revision behaviour can also be represented by a 'conditional belief set' $[\Psi]_c$
- It is defined via the Ramsey Test:

  (RT) $A \Rightarrow B \in [\Psi]_c$ iff $B \in [\Psi * A]$, with $A, B \in L$
- Call conditional beliefs 'c-beliefs'
- AGM*alternatively amounts to:

  $\mid\!\sim_\Psi = \{\langle A, B\rangle \mid A \Rightarrow B \in [\Psi]_c\}$ is a 'consistency-preserving' rational consequence relation [LM92]

Criticisms of AGM

- We'll assume AGM in what follows, but there have been objections to:
  - $(K4^*)$ [Rab96]
  - $(K5^{\div})$ (Recovery) [Han91] [Han96] [Lev91] (but see [Gla00])
  - $(K7^*)$, $(K7^{\div})$, $(K8^*)$ and $(K8^{\div})$ [Rot04]
- Some have also urged that we consider operators for which $(K2^*)$ and $(K4^{\div})$ do not hold (more later)

2. Iterated revision I: background

The problem

- What about beliefs resulting from sequences of revisions?
- AGM* hardly constrains the result of a sequence of even 2 revisions:

  Only the posterior minimal worlds are determined, but not the rest of the posterior TPO

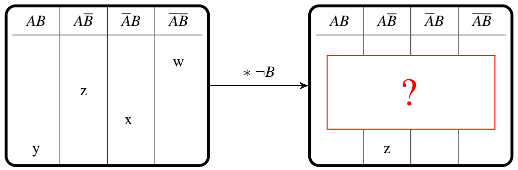

2.1. The postulates of Darwiche & Pearl

The DP postulates: Semantic presentation

- The Darwiche-Pearl (DP) postulates [DP97] constrain sequences of 2 revisions
- In 'semantic' (i.e. TPO) terms:

(C1$^*_\preccurlyeq$) If $x, y \in [\![A]\!]$ then $x \preccurlyeq_{\Psi * A} y$ iff $x \preccurlyeq_\Psi y$

(C2$^*_\preccurlyeq$) If $x, y \in [\![\neg A]\!]$ then $x \preccurlyeq_{\Psi * A} y$ iff $x \preccurlyeq_\Psi y$

(C3$^*_\preccurlyeq$) If $x \in [\![A]\!]$, $y \in [\![\neg A]\!]$ and $x \prec_\Psi y$, then $x \prec_{\Psi * A} y$

(C4$^*_\preccurlyeq$) If $x \in [\![A]\!]$, $y \in [\![\neg A]\!]$ and $x \preccurlyeq_\Psi y$, then $x \preccurlyeq_{\Psi * A} y$

(Moving from $\preccurlyeq_\Psi$ to $\preccurlyeq_{\Psi * A}$ can only involve a 'promotion' of worlds in $[\![A]\!]$ in relation to worlds in $[\![\neg A]\!]$)

The DP postulates: Syntactic presentation

- In 'syntactic' terms, the DP postulates translate into:

(C1*) If $A \in \mathrm{Cn}(B)$ then $[(\Psi * A) * B] = [\Psi * B]$

(C2*) If $\neg A \in \mathrm{Cn}(B)$ then $[(\Psi * A) * B] = [\Psi * B]$

(Adding an intermediate revision by a weaker or contradictory sentence is redundant)

(C3*) If $A \in [\Psi * B]$ then $A \in [(\Psi * A) * B]$

(C4*) If $\neg A \notin [\Psi * B]$ then $\neg A \notin [(\Psi * A) * B]$

(Adding an intermediate revision by A doesn't worsen the status of A or improve the status of $\neg A$)

Presentation in terms of defeaters

- Interpretation with (RT)-style tests for beliefs about defeat:

  B is taken to overrule A in Ψ iff $A \notin [(\Psi * A) * B]$

  B is taken to strictly overrule A in Ψ iff $\neg A \in [(\Psi * A) * B]$

- Similar to Pollock's undercut/rebut relations [Pol87]
- [BC17] show DP to be alternatively presentable as follows:

$$[(\Psi * A) * B] = \begin{cases} [\Psi * A \wedge B], & \text{if } A \in [(\Psi * A) * B] \\ [\Psi * B] \cap [\Psi * A \wedge B], & \text{if } A, \neg A \notin [(\Psi * A) * B] \\ [\Psi * B], & \text{if } \neg A \in [(\Psi * A) * B] \end{cases}$$

Principle (P^*)

- The following strengthening of ($C3^*_{\preccurlyeq}$) and ($C4^*_{\preccurlyeq}$) has also been defended [BM06] [JT07]:

  ($P^*_{\preccurlyeq}$) If $x \in [\![A]\!]$, $y \in [\![\neg A]\!]$ and $x \preccurlyeq_\Psi y$, then $x \prec_{\Psi * A} y$

- Syntactic version:

  (P^*) If $\neg A \notin [\Psi * B]$, then $A \in [(\Psi * A) * B]$

- This characterises the class of 'admissible' operators [BM06]

2.2. CRITICISMS

Introduction

- Two main lines of criticism have been offered against DP:
  - Counterexamples to ($C1^*$) & ($C2^*$)
  - Conflict between each of ($C1^*$)–($C4^*$) and a principle of 'Relevance Sensitivity'

Counterexamples to ($C1^*$)-($C2^*$)

Example 1

Two biased coins are flipped and placed in a box. Based on their bias, I initially believe that both landed heads ($H_1 \wedge H_2$). Two reliable witnesses simultaneously provide independent reports: Carla reports that coin 1 is lying tails up ($\neg H_1$) and Dora that coin 2 is lying tails up ($\neg H_2$). Subsequently, Elmer, who is the most reliable overall, reports that the coin in box 1 is lying heads up (H_1).

Counterexamples to (C1*)-(C2*)

Example 1

Two biased coins are flipped and placed in a box. Based on their bias, I initially believe that both landed heads ($H_1 \wedge H_2$). Two reliable witnesses simultaneously provide independent reports: Carla reports that coin 1 is lying tails up ($\neg H_1$) and Dora that coin 2 is lying tails up ($\neg H_2$). Subsequently, Elmer, who is the most reliable overall, reports that the coin in box 1 is lying heads up (H_1).

- [Sta09] claims: $\neg H_2 \in [(\Psi * \neg H_1 \wedge \neg H_2) * H_1]$
- From $H_2 \in [\Psi]$, $\neg H_1 \notin [\Psi]$, and (K4*): $H_2 \in [\Psi * H_1]$
- But by (C2*): $[\Psi * H_1] = [(\Psi * \neg H_1 \wedge \neg H_2) * H_1]$

Counterexamples to (C1*)-(C2*)

- Rejoinder [CGMW08]: We should consider updates on reports of results ($[(\Psi * Carla(\neg H_1) \wedge Dora(\neg H_2)) * Elmer(H_1)]$)
- See [CGMW08] for another counterexample to (C2*) that they take to be more persuasive.
- [Sta09] and [Mey01] also offer counterexamples to (C1*)

Relevance-sensitivity

- [Par99] argues that revision needs to incorporate a condition of 'relevance-sensitivity'.
- Informally: The aspects of $[\Psi]$ that are 'unrelated' to A ought to be unaffected by revision by A

Definition 2

A sublanguage of L is a language generated by a subset of L's atomic sentences

Definition 3

Where $A \in L$, L_A denotes the unique minimal sublanguage of L that contains a sentence equivalent to A

Relevance-sensitivity (cont.)

- The constraint is formally given by:

(wRS) If $[\Psi] = \mathrm{Cn}(B, C)$, $L_B \cap L_C = \varnothing$ and $A \in L_B$, then $[\Psi * A] \cap \overline{L_B} = [\Psi] \cap \overline{L_B}$

- (wRS) implies the following constraint on $\preccurlyeq_\Psi$ in the special case when $[\![\Psi]\!] = \{x\}$ for some $x \in W$ [PWCF15]:

($\mathrm{wRS}_\preccurlyeq$) If $Diff(x, y) \subset Diff(x, z)$ then $y \prec_\Psi z$

($Diff(x, y)$ = set of atomic propositions on which x, y disagree)

- (wRS) is notably satisfied by the proposal of [Dal88] (see later)

The conflict with DP

Definition 4

A linear order is a TPO such that $x \sim y$ iff $x = y$

- [PFS08] showed:

Theorem 3

For any prior linear order $\preccurlyeq_\Psi$ on a set of 8 worlds, there exists, for each postulate $(\mathrm{C}i^*_\preccurlyeq)$, a sentence A_i such that the posterior TPO $\preccurlyeq_{\Psi * A_i}$ cannot satisfy both $(\mathrm{C}i^*_\preccurlyeq)$ and $(\mathrm{wRS}_\preccurlyeq)$.

- Is this bad news for (wRS) or for DP?

2.3. The nature of states: two strong proposals

Introduction

- So far: states introduced as primitive objects
- Can we say more about them?
- In what follows:
  - 2 proposals that individuate states very coarsely
  - why these proposals are probably not sensible
- Later: some suggestions as to how to incorporate more structure

States as belief sets?

- AGM*+ DP are too weak to entail the following reductionist thesis:

$$\text{If } [\Psi] = [\Theta], \text{ then, for all } A, [\Psi * A] = [\Theta * A]$$

  (Prior beliefs determine posterior beliefs after single revisions)
- This effectively identifies states with belief sets
- It holds for proposals that generate a TPO on the basis of distance between worlds and $[\![[\Psi]]\!]$ [Dal88] [PW18] [Sat88] [Wil86]
- But it seems unduly strong (see also [Han92])

States as belief sets? (cont.)

- Example: Hamming distance [Dal88]

  $d_H(w,w') := |\mathit{Diff}(w,w')|$

  (Recall $\mathit{Diff}(w,w')$ is the set of atomic propositions on which w, w' differ)

  $d_H(w,[\Psi]) := \min_{w' \in [\![\Psi]\!]} d_H(w,w')$
- This imposes the following TPO for the swans example:

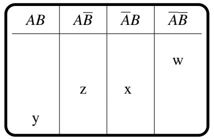

States as total preorders?

- AGM*+ DP also fall short of entailing a weaker claim:

  If $[\Psi]_c = [\Theta]_c$, then $[\Psi * A]_c = [\Theta * A]_c$

  (One's beliefs after single revisions determine one's beliefs after two successive revisions)
- This effectively identifies states with TPOs

Elementary operators

- But recent approaches have taken the last claim for granted
- This is the case of the following 3 popular 'elementary' operators:
  - lexicographic ($*_L$) [NPP03]
  - restrained ($*_R$) [BM06]
  - natural ($*_N$) [Bou96]
- These all *considerably* strengthen DP
- $*_R$ and $*_L$ also satisfy (P*)

Elementary operators: illustration

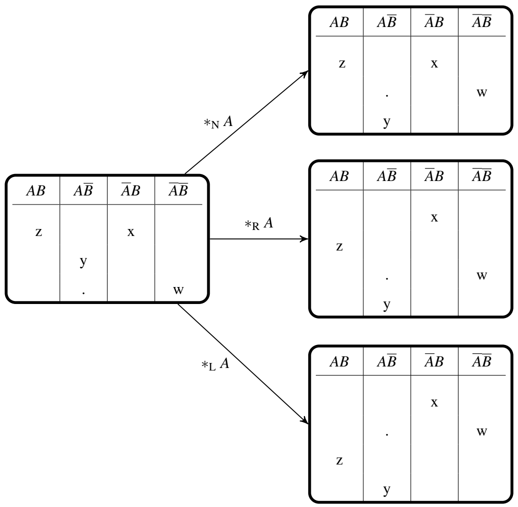

Counterexample

- [BC17] offer the following pair of cases:

Example 2

Bashiir and Ayaan have been invited to a party. Initially unsure as to whether either wanted to attend, I now hear that the venue is located too far out of town for either of them. I also believe that they don't get on and are unlikely to attend the same party.

Example 3

As above, save that I believe that Bashiir and Ayaan have never met and know nothing about each other

Counterexample (cont.)

- Let 'A' = 'Ayaan will attend' and 'B' = 'Bashiir will attend'
- With Ψ being the state after having heard of the party's location:
  - (1) $\neg A \in [\Psi]$ and $\neg B \in [\Psi]$
  - (2) $\neg A \in [\Psi * B]$ and $\neg B \in [\Psi * A]$
  - (3) $A, B \notin [\Psi * A \vee B]$
- Assuming L has atoms $\{A, B\}$, AGM* dictates the following TPO:

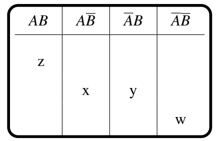

Counterexample (cont.)

- (1)–(3) do *not* however, even given AGM* + DP, determine whether $A \in [(\Psi * A) * B]$
- This seems a good thing:
  - Example 2: potentially, $\neg A \in [(\Psi * A) * B]$, if my belief in their mutual dislike is sufficiently deeply entrenched
  - Example 3: $A \in [(\Psi * A) * B]$
- Upshot: states must be have more structure than TPOs!

Diagnosis

- Core relevant difference:
  - Example 2: B is taken to strictly overrule A
  - Example 3: B is taken to not overrule A
- This distinction does not show up in the 1-step revision dispositions
- Is there a sensible strengthening of AGM* + DP that
  - (i) falls short of reducing states to TPOs and
  - (ii) allows for a neat representation of states and in particular overruling relations?

Two strong principles

- DP constrains the relation between a prior and a posterior conditional belief set
- What about the relation between two posterior conditional belief sets obtained from a common prior by different revisions?
- [BM11] consider general forms of:

($\beta 1{+}^*$) If $A \notin [(\Psi * A) * B]$, then $A \notin [(\Psi * C) * B]$
($\beta 2{+}^*$) If $\neg A \in [(\Psi * A) * B]$, then $\neg A \in [(\Psi * C) * B]$

- These radically strengthen $(\mathrm{C3}^*)$ and $(\mathrm{C4}^*)$ (let $C = \top$)

The operator $\circ$

- [BM11] focus on a more general family of 'non-prioritised revision' operators $\circ$:

  $\circ$ 'promotes' the doxastic status of a sentence while possibly falling short of leading it to be believed (Success fails)

- This operator satisfies the general forms of $(\beta 1{+}^*)$ and $(\beta 2{+}^*)$
- It is representable by associating Ψ with a particular kind of enrichment of $\preccurlyeq_\Psi$

Proper Ordinal Intervals

Definition 5

$\leq_\Psi$ is a proper ordinal interval (POI) assignment to W iff it is a TPO over $W^{\pm} = \{w^i \mid w \in W \text{ and } i \in \{-,+\}\}$ such that:

$$x^+ <_\Psi x^- \text{ (non-zero length)}$$
$$x^+ \leq_\Psi y^+ \text{ iff } x^- \leq_\Psi y^- \text{ (equal length)}$$

- Interpretation: x^+/x^- represent x's position in the light of auspicious/inauspicious inputs
- $\leq_\Psi$ yields TPO $\preccurlyeq_\Psi$ over W by setting: $x \preccurlyeq_\Psi y$ iff $x^+ \leq_\Psi y^+$.

Proper Ordinal Intervals: illustration

- Two different POIs yielding the same TPO:

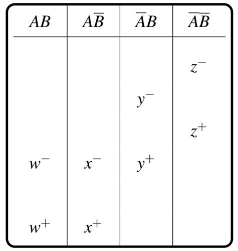

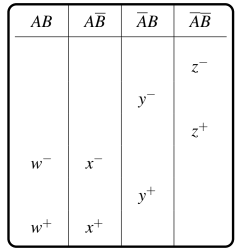

POIs and $\circ$

- Determining $\preccurlyeq_{\Psi\circ A}$ from $\leq_\Psi$:
  (1) For each $A \in L$ define a function r_A by setting, for each world x

$$r_A(x) = \begin{cases} x^+ & \text{if } x \in [\![A]\!] \\ x^- & \text{if } x \in [\![\neg A]\!] \end{cases}$$

  (2) Set $x \preccurlyeq_{\Psi\circ A} y$ iff $r_A(x) \leq_\Psi r_A(y)$

POIs and $\circ$: illustration

- Examples, in which we revise by $\neg A$:

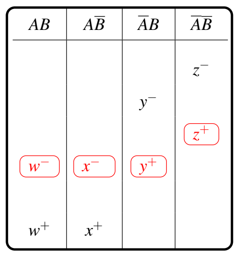

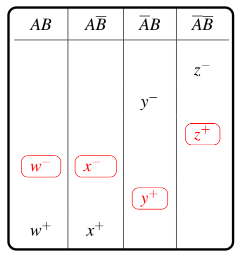

Success fails on the left: $x \notin [\![\neg A]\!]$ but $x \in \min(\preccurlyeq_{\Psi\circ\neg A}, W)$

Special case: improvement operators

- [KP08a] discussed the class of improvement operators
- Characteristic property: $\exists n$, s.t. $A \in [\Psi \circ A^n]$
  (Eventual Success: revising by A enough times renders A believed)

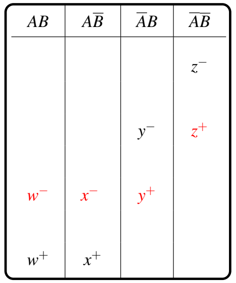

1-improvement revision ($x^- \sim y^+$ whenever $x \prec y$ and there is no z s.t. $x \prec z \prec y$)

Moving to the prioritised case

- Problem 1: The only operators satisfying AGM*, (C1*), (C2*) and ($\beta 1+^*$) are lexicographic revision operators [BC18]
- Problem 2: There is a direct counterexample [BC18]
- Proposal:

$$\preccurlyeq_{\Psi * A} := \preccurlyeq_{\Psi \circ A} *_{\mathrm{N}} A$$

- Rationale:

  Gives us the conditional belief set $[\Psi * A]_c$ that is as similar to $[\Psi \circ A]_c$ as (C1*) and the inclusion of A permit

- Call such operators 'POI revision' operators

Success via naturalisation: illustration

- Example, in which we again revise by $\neg A$:

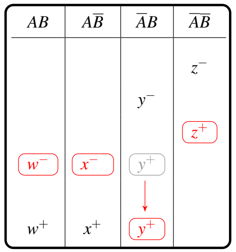

Special cases

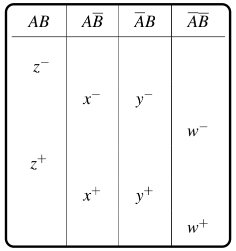

Lexicographic revision ($\forall x, y, x \prec y$ iff $x^+ < y^+ < x^-$)

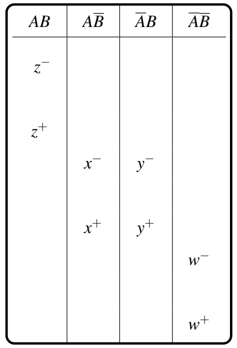

Restrained revision ($\forall x, y,\ x \prec y$ iff $x^- < y^+$)

Two weaker principles

- POI revision only satisfies *weakenings* of $(\beta 1+^*)$ and $(\beta 2+^*)$:

$(\beta 1^*)$ If $A \notin [(\Psi * A) * B]$ and $B \to \neg A \in [\Psi * C]$, then $A \notin [(\Psi * C) * B]$

$(\beta 2^*)$ If $\neg A \in [(\Psi * A) * B]$ and $B \to \neg A \in [\Psi * C]$, then $\neg A \in [(\Psi * C) * B]$

which are in fact *also* satisfied by natural revision

- They remain strengthenings of (C3*) and (C4*)

Representation theorem

Theorem 4

The following two claims are equivalent:

(a) For each state Ψ, there exists a POI assignment $\leq_\Psi$ to W, such that $C \in [(\Psi * A) * B]$ iff $\min(\preccurlyeq_{\Psi * A}, [\![B]\!]) \subseteq [\![C]\!]$, where $\preccurlyeq_{\Psi * A}$ is determined from $\leq_\Psi$ as proposed above

(b) $*$ satisfies (Eq*), (C1*), (C2*), $(\beta 1^*)$, $(\beta 2^*)$ and $(\Omega 1^*)$–$(\Omega 3^*)$

- Where (Eq*) is a conditions of 'irrelevance of syntax' and...

The (Ωi^*)

... (Ωi^*) are conditions on relations of overruling between sentences:

$(\Omega 1^*)$ If $\neg A \notin [\Psi * A \vee B]$ and $A \notin [(\Psi * A) * B]$, then $B \notin [(\Psi * B) * A]$

$(\Omega 2^*)$ If $\neg A \notin [\Psi * A \vee B]$ and $\neg A \in [(\Psi * A) * B]$, then $\neg B \in [(\Psi * B) * A]$

$(\Omega 3^*)$ If $\neg B \in [\Psi * A \vee B]$ and $A \notin [(\Psi * A) * B]$, then $\neg B \in [(\Psi * B) * A]$

The (Ωi^*)

... (Ωi^*) are conditions on relations of overruling between sentences:

$(\Omega 1^*)$ If A is at least as well entrenched as B and B is taken to overrule A, then A is taken to overrule B

$(\Omega 2^*)$ If A is at least as well entrenched as B and B is taken to strictly overrule A, then A is taken to strictly overrule B

$(\Omega 3^*)$ If A is strictly better entrenched than B and B is taken to overrule A, then A is taken to strictly overrule B

The counterexample to states = TPOs

- The difference between examples 2 and 3 translates as:

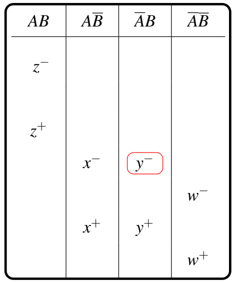

Example 2: $\neg A \in [(\Psi * A) * B]$

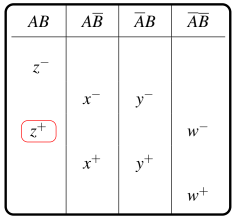

Example 3: $A \in [(\Psi * A) * B]$

Open questions

Open question

What are the semantic conditions under which the overruling relations hold in POIs? What logic of overruling emerges from these?

Open question

[BC18]'s proposal currently only returns a posterior TPO from a prior POI; what about the posterior POI?

Open question

If we weaken $x^+ < x^-$ to $x^+ \leq x^-$, we obtain a class of operators (BOI operators) that also includes $*_{\mathrm{N}}$. Which set of postulates characterises BOI operators?

3.2. Ranking functions

Ranking functions

- A quantitative enrichment of states has also been proposed [Spo14]

Definition 6

A ranking function κ over W, is a function $W \mapsto \mathbb{R}^+ \cup \{\infty\}$, for which $\exists x \in W$ s.t. $\kappa(x) = 0$

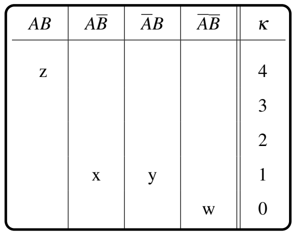

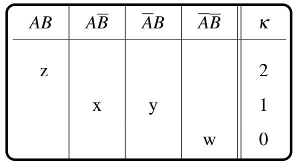

- We'll denote $\{A \in L \mid \forall \kappa(w) = 0, w \models A\}$ by $[\kappa]$

A family of revision operators

- Spohn proposes a parameterised family of revision operators that differ wrt the posterior rank assigned to the negation of the input

Definition 7

Where $r \in \mathbb{R}^+ \cup \{\infty\}$ and $\kappa(A) < \infty$, the ranking function r-revision operator $\circledast^r$ is defined by setting:

$$\kappa \circledast^r A(w) = \begin{cases} \kappa(w) + r - \min(\kappa, [\![\neg A]\!]), & \text{if } w \in [\![\neg A]\!] \\ \kappa(w) - \min(\kappa, [\![A]\!]), & \text{if } w \in [\![A]\!] \end{cases}$$

Ranking theoretic revision: illustration

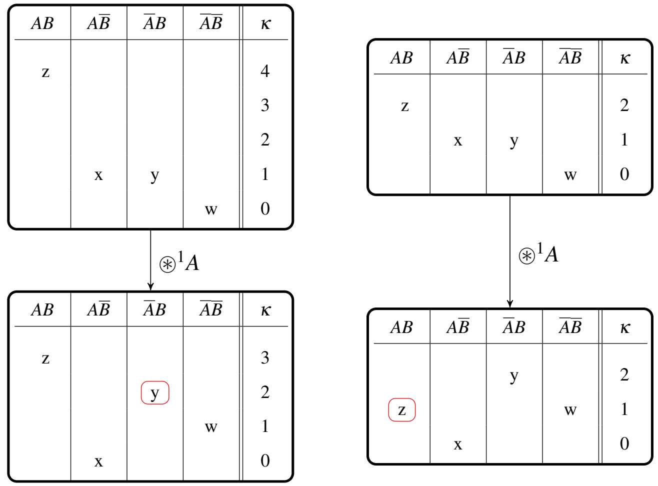

$\neg A \in [(\kappa \circledast^1 A) \circledast^1 B]$

$A \in [(\kappa \circledast^1 A) \circledast^1 B]$

Properties

- The counterparts of $(C1^*_{\preccurlyeq})$ and $(C2^*_{\preccurlyeq})$ are satisfied
- The counterparts of $(C3^*_{\preccurlyeq})$ and $(C4^*_{\preccurlyeq})$ (and indeed, that of $(P^*_{\preccurlyeq})$) are satisfied if, but only if, $A \notin [\kappa]$

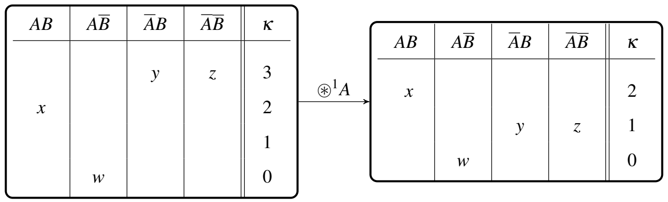

Open question

How does ranking theoretic revision stand in relation to the other properties discussed earlier (e.g. $(\beta 1^*_{\preccurlyeq})$, etc.)?

Axiomatics?

- The association of states with POIs was justified axiomatically;
- This is not so for the association with ranking functions, at least in the case of revision (more on contraction later)

Open question

Can an axiomatic justification of ranking theoretic revision be given?

Further open questions

Open question

What are the semantic conditions under which the overruling relations hold in ranking functions? What logic of overruling emerges from these?

Open question

What is the relation between POIs and ranking functions?

4. ITERATED NON-PRIORITISED REVISION

Introduction

- Already mentioned non-prioritised revision $\circ$ (dropping (K2*))
- For 1-step case: screened [Mak97], credibility- limited [HFCF01], ...
- For iterated case: POI [BM11], improvement [KP08b], credibility-limited [BFKP12], ...
- Here: a very general model, promising to capture various more specific proposals as special cases

Boutilier, Friedman and Halpern (1998)

- [BFH98] propose:

  Rank not worlds, but pairs of worlds and infinite sequences of inputs, which may or may not be true in the relevant world

- A run r is a pair $\langle r_e, r_a \rangle$, with $r_e \in W$ and $r_a \subseteq L^\infty$
- Example:

$$\kappa_\Psi(\langle x, \langle A_1, A_2, \ldots \rangle \rangle) = |\{i \in \mathbb{N} : x \in [\![\neg A_i]\!]\}|$$

(the greater the # of false observations in r_a, the less likely the run)

Boutilier, Friedman and Halpern (1998) (cont.)

- Following a sequence $\vec{S}$ of actual observations made, the new ranking function $\kappa_{\Psi\circ\vec{S}}$ is obtained by *conditioning* (i.e. ∞-revising) κ_Ψ on the set of runs r with $r_a = \vec{S} \cdot \vec{S}'$
- Belief set after $\circ\vec{S}$ determined by set of 'most plausible' worlds:

$$[\![\kappa_{\Psi\circ\vec{S}}]\!] = \{x \in W \mid r_e = x \text{ for some } r \in \kappa^{-1}_{\Psi\circ\vec{S}}(0)\}$$

Example

- Assume 1 atomic sentence A in L with possible values T or F, 3 steps and possible observations $\{A, \neg A\}$

A	$\neg A$	κ
$\langle T, \langle \neg A, \neg A, \neg A\rangle\rangle$	$\langle F, \langle A, A, A\rangle\rangle$	3
$\langle T, \langle \neg A, \neg A, A\rangle\rangle$ $\langle T, \langle \neg A, A, \neg A\rangle\rangle$ $\langle T, \langle A, \neg A, \neg A\rangle\rangle$	$\langle F, \langle A, A, \neg A\rangle\rangle$ $\langle F, \langle A, \neg A, A\rangle\rangle$ $\langle F, \langle \neg A, A, A\rangle\rangle$	2
$\langle T, \langle A, A, \neg A\rangle\rangle$ $\langle T, \langle A, \neg A, A\rangle\rangle$ $\langle T, \langle \neg A, A, A\rangle\rangle$	$\langle F, \langle \neg A, \neg A, A\rangle\rangle$ $\langle F, \langle \neg A, A, \neg A\rangle\rangle$ $\langle F, \langle A, \neg A, \neg A\rangle\rangle$	1
$\langle T, \langle A, A, A\rangle\rangle$	$\langle F, \langle \neg A, \neg A, \neg A\rangle\rangle$	0

Initial state: $[\kappa_\Psi] = Cn(\emptyset)$

Example (cont.)

- Assume 1 atomic sentence A in L with possible values T or F, 3 steps and possible observations $\{A, \neg A\}$

A	$\neg A$	κ
$\langle T, \langle \neg A, \neg A, \neg A\rangle\rangle$	$\langle F, \langle A, A, A\rangle\rangle$	3
$\langle T, \langle \neg A, \neg A, A\rangle\rangle$ $\langle T, \langle \neg A, A, \neg A\rangle\rangle$ $\langle T, \langle A, \neg A, \neg A\rangle\rangle$	$\langle F, \langle A, A, \neg A\rangle\rangle$ $\langle F, \langle A, \neg A, A\rangle\rangle$ $\langle F, \langle \neg A, A, A\rangle\rangle$	2
$\langle T, \langle A, A, \neg A\rangle\rangle$ $\langle T, \langle A, \neg A, A\rangle\rangle$ $\langle T, \langle \neg A, A, A\rangle\rangle$	$\langle F, \langle \neg A, \neg A, A\rangle\rangle$ $\langle F, \langle \neg A, A, \neg A\rangle\rangle$ $\langle F, \langle A, \neg A, \neg A\rangle\rangle$	1
$\langle T, \langle A, A, A\rangle\rangle$	$\langle F, \langle \neg A, \neg A, \neg A\rangle\rangle$	0

After revision by A: $[\kappa_{\Psi\circ A}] = Cn(A)$

Example (cont.)

- Assume 1 atomic sentence A in L with possible values T or F, 3 steps and possible observations $\{A, \neg A\}$

A	$\neg A$	κ
$\langle T, \langle \neg A, \neg A, \neg A\rangle\rangle$	$\langle F, \langle A, A, A\rangle\rangle$	
$\langle T, \langle \neg A, \neg A, A\rangle\rangle$ $\langle T, \langle \neg A, A, \neg A\rangle\rangle$ $\langle T, \langle A, \neg A, \neg A\rangle\rangle$	$\langle F, \langle A, A, \neg A\rangle\rangle$ $\langle F, \langle A, \neg A, A\rangle\rangle$ $\langle F, \langle \neg A, A, A\rangle\rangle$	1
$\langle T, \langle A, A, \neg A\rangle\rangle$ $\langle T, \langle A, \neg A, A\rangle\rangle$ $\langle T, \langle \neg A, A, A\rangle\rangle$	$\langle F, \langle \neg A, \neg A, A\rangle\rangle$ $\langle F, \langle \neg A, A, \neg A\rangle\rangle$ $\langle F, \langle A, \neg A, \neg A\rangle\rangle$	0
$\langle T, \langle A, A, A\rangle\rangle$	$\langle F, \langle \neg A, \neg A, \neg A\rangle\rangle$	

After further revision by $\neg A$: $[\kappa_{(\Psi\circ A)\circ\neg A}] = Cn(\emptyset)$

Example (cont.)

- Assume 1 atomic sentence A in L with possible values T or F, 3 steps and possible observations $\{A, \neg A\}$

A	$\neg A$	κ
$\langle T, \langle \neg A, \neg A, \neg A \rangle\rangle$	$\langle F, \langle A, A, A \rangle\rangle$	
$\langle T, \langle \neg A, \neg A, A \rangle\rangle$ $\langle T, \langle \neg A, A, \neg A \rangle\rangle$ $\langle T, \langle A, \neg A, \neg A \rangle\rangle$	$\langle F, \langle A, A, \neg A \rangle\rangle$ $\langle F, \langle A, \neg A, A \rangle\rangle$ $\langle F, \langle \neg A, A, A \rangle\rangle$	1
$\langle T, \langle A, A, \neg A \rangle\rangle$ $\langle T, \langle A, \neg A, A \rangle\rangle$ $\langle T, \langle \neg A, A, A \rangle\rangle$	$\langle F, \langle \neg A, \neg A, A \rangle\rangle$ $\langle F, \langle \neg A, A, \neg A \rangle\rangle$ $\langle F, \langle A, \neg A, \neg A \rangle\rangle$	0
$\langle T, \langle A, A, A \rangle\rangle$	$\langle F, \langle \neg A, \neg A, \neg A \rangle\rangle$	

After further revision by A: $[\kappa_{((\Psi \circ A) \circ \neg A) \circ A}] = Cn(A)$

Restrictions on the model

- [BFH98] consider additional restrictions on κ_Ψ:
  - Markov: Observations considered independent of observational history, given state of world ($\rightarrow$ compact representation of κ_Ψ)
  - Credibility: Observation A taken to favour every A-world over any $\neg A$-world

    (Entails: $\forall \vec{S}$ s.t. $\exists \vec{S'}$ s.t. $\neg A \notin [\Psi \circ \vec{S} \cdot \vec{S'}]$, there exists n s.t. $A \in [\Psi \circ \vec{S} \cdot A^n]$; cf. improvement operators)
  - Accuracy: False observations considered impossible

    (Entails Success)
- Accuracy + 1 further condition yield full AGM postulates

5. Iterated contraction I: background

5.1. The postulates of Chopra-Ghose-Meyer-Wong

The postulates of Chopra *et al*

- The Chopra *et al* postulates [CGMW08] are very similar to DP:

(C1$^{\div}$) If $\neg A \in \mathrm{Cn}(B)$ then $[(\Psi \div A) * B] = [\Psi * B]$

(C2$^{\div}$) If $A \in \mathrm{Cn}(B)$ then $[(\Psi \div A) * B] = [\Psi * B]$

(Adding an intermediate contraction by a weaker or contradictory sentence is redundant)

(C3$^{\div}$) If $\neg A \in [\Psi * B]$ then $\neg A \in [(\Psi \div A) * B]$

(C4$^{\div}$) If $A \notin [\Psi * B]$ then $A \notin [(\Psi \div A) * B]$

(Adding an intermediate contraction by A doesn't worsen the status of $\neg A$ or improve the status of A)

The postulates of Chopra *et al*, semantic versions

- Their semantic counterparts are given by:

($C1^{\div}_{\preccurlyeq}$) If $x, y \in [\![\neg A]\!]$ then $x \preceq_{\Psi \div A} y$ iff $x \preceq_{\Psi} y$

($C2^{\div}_{\preccurlyeq}$) If $x, y \in [\![A]\!]$ then $x \preceq_{\Psi \div A} y$ iff $x \preceq_{\Psi} y$

($C3^{\div}_{\preccurlyeq}$) If $x \in [\![\neg A]\!]$, $y \in [\![A]\!]$ and $x \prec_{\Psi} y$ then $x \prec_{\Psi \div A} y$

($C4^{\div}_{\preccurlyeq}$) If $x \in [\![\neg A]\!]$, $y \in [\![A]\!]$ and $x \preceq_{\Psi} y$ then $x \preceq_{\Psi \div A} y$

(Moving from $\preccurlyeq_{\Psi}$ to $\preccurlyeq_{\Psi \div A}$ can only involve a 'promotion' of worlds in $[\![\neg A]\!]$ in relation to worlds in $[\![A]\!]$)

A weak analogue of principle (P*)

- Are there analogues to (P*)/ ($P^*_{\preccurlyeq}$)for contraction?
- The following *won't* do, as they conflict with (K2$^{\div}$)

(P$^{\div}$) If $A \notin [\Psi * B]$, then $\neg A \in [(\Psi \div A) * B]$

($P^{\div}_{\preccurlyeq}$) If $x \in [\![\neg A]\!]$, $y \in [\![A]\!]$ and $x \preccurlyeq_{\Psi} y$, then $x \prec_{\Psi \div A} y$

- However, [Spo14] endorses the weaker:

($WP^{\div}_{\preccurlyeq}$) If $\min(\preccurlyeq, W) \subseteq [\![A]\!]$, $x \in [\![\neg A]\!]$, $y \in [\![A]\!]$ and $x \preccurlyeq_{\Psi} y$, then $x \prec_{\Psi \div A} y$

(($P^{\div}_{\preccurlyeq}$) holds when the contraction is a 'genuine' one)

5.2. States as total preorders

Reductive proposals for iterated contraction

- A number of reductive proposals for iterated contraction have been discussed in the literature [NGO07]:
  - lexicographic contraction ($\div_L$)
  - natural contraction ($\div_N$)
  - priority contraction ($\div_P$)
- $\div_N$ and $\div_P$ satisfy $(C1^{\div}_{\preccurlyeq})$–$(C4^{\div}_{\preccurlyeq})$, while $\div_L$ does not, violating $(C3^{\div}_{\preccurlyeq})$ and $(C4^{\div}_{\preccurlyeq})$
- $\div_P$ satisfies $(WP^{\div}_{\preccurlyeq})$, while $\div_L$ and $\div_N$ do not

Reductive proposals for iterated contraction: illustration

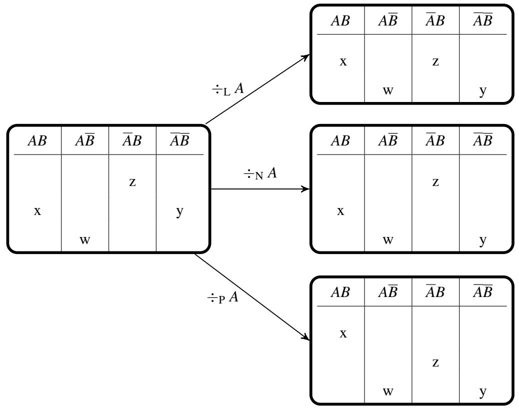

Violation of $(C3^{*}_{\preccurlyeq})/(WP^{\div}_{\preccurlyeq})$

6. ITERATED CONTRACTION II: ENRICHED STATES

Ordinal Intervals

- There is currently no proposal for iterated contraction using POIs

Open question

Can the TQ approach be used to define an interesting contraction operator from POI revision? If so, what are its characteristic properties?

Ranking theoretic contraction

- Spohn has offered a representation theorem for a contraction operator based on the identification of states with ranking functions

Definition 8

The ranking function contraction operator $\ominus$ is defined by:

- If $A \notin [\kappa]$, then $\kappa \ominus A(w) = \kappa(w)$
- If $A \in [\kappa]$, then

$$\kappa \ominus A(w) = \begin{cases} \kappa(w), & \text{if } w \in [\![A]\!] \\ \kappa(w) - \min(\kappa, [\![\neg A]\!]), & \text{if } w \in [\![\neg A]\!] \end{cases}$$

Ranking theoretic contraction: illustration

AB	$A\overline{B}$	$\overline{A}B$	$\overline{A}\overline{B}$	κ
z		y		2
			w	1
	x			0

$\downarrow \ominus A$

AB	$A\overline{B}$	$\overline{A}B$	$\overline{A}\overline{B}$	κ
z				2
		y		1
	x		w	0

Further postulates

- [HS08] propose to supplement $(C1^{\div})$–$(C4^{\div})$ with:

(SV) If $A \notin [\Psi]$, then $[\Psi \div A \cdot \vec{S}] = [\Psi \div \vec{S}]$

Strengthening of $(K4^{\div})$

(PI) If $B \in Cn(A)$ and $B \to A \notin [\Psi \div A]$, then $[\Psi \div A \cdot B \cdot \vec{S}] = [\Psi \div B \to A \cdot B \cdot \vec{S}]$

Rationale not given (what do you make of it?)

(RC) If $A \vee B \in Cn(\varnothing)$, then $[\Psi \div A \cdot B \cdot \vec{S}] = [\Psi \div B \cdot A \cdot \vec{S}]$

Restricted commutativity principle designed to avoid a counterexample to commutativity due to [Han93]

Representation result

Theorem 5

The following two claims are equivalent:

(a) There exists, for each state Ψ, a ranking function κ_Ψ, unique up to a multiplicative constant, such that $[\Psi \div \vec{S}] = [\kappa_\Psi \ominus \vec{S}]$

(b) $\div$ satisfies $AGM^{\div}$, $(C1^{\div})$–$(C4^{\div})$, (SV), (PI) and (RC)

Open question

Which of the new postulates are satisfied/violated by our various operators $\div_L$, $\div_N$, $\div_P$, ...?

7. From iterated revision to contraction and back

Introduction

- (LI) and (HI) relate constraints on 1-step revision & contraction
- What about the iterated case?
- Not much written about this, save a few papers on 2-step change

Extending (HI): triviality result

- One proposal would be to extend (HI) to conditional beliefs:

  (NiHI) $[\Psi \div A]_c = [\Psi]_c \cap [\Psi * \neg A]_c$

- Equivalently:

  $$[(\Psi \div A) * B] = [\Psi * B] \cap [(\Psi * \neg A) * B]$$

- If $B \equiv \top$ then we obtain (HI) as a special case

Theorem 6

Given AGM*, AGM$^{\div}$ and (HI), (NiHI) entails that there do not exist a state Ψ and non-tautologous sentences A and B such that $A \vee B$ is tautologous and $A, B \in [\Psi]$.

Two weak principles

- [BC16] [BCss] outline a family of extensions of (HI) that satisfy the following consequences of (NiHI):

  (LB) $[\Psi \div A]_c \supseteq [\Psi]_c \cap [\Psi * \neg A]_c$

  ('Respect the consensus!')

  (UB) $[\Psi \div A]_c \subseteq [\Psi]_c \cup [\Psi * \neg A]_c$

  ('No making stuff up!')

- Note: (LB) is the half of (NiHI) *not* implicated in Theorem 6

TQ combinators

- They approach the question semantically:

  Finding a suitable TPO aggregation function $\oplus$, such that $\preccurlyeq_{\Psi\div A} = \preccurlyeq_{\Psi} \oplus \preccurlyeq_{\Psi*\neg A}$

- In our domain of interest, (HI), (LB) and (UB) translate directly into

  ($\text{HI}_{\oplus}$) $\min(\preccurlyeq_{1\oplus 2}, W) = \min(\preccurlyeq_1, W) \cup \min(\preccurlyeq_2, W)$

  (SPU) If $x \prec_1 y$ and $x \prec_2 y$ then $x \prec_{1\oplus 2} y$

  (WPU) If $x \preccurlyeq_1 y$ and $x \preccurlyeq_2 y$ then $x \preccurlyeq_{1\oplus 2} y$

  and collectively characterise TeamQueue combinators

The STQ combinator

- Adding a further constraint yields the Synchronous TQ combinator $\oplus_{\text{STQ}}$, which aggregates TPOs in an even-handed manner
- Syntactically, $\oplus_{\text{STQ}}$ gives us:

  (iHIR) $[\Psi \div A]_c = C_R([\Psi]_c \cap [\Psi * \neg A]_c)$

  where C_R is the rational closure function [LM92]
- $C_R(\Gamma)$ has been argued to be the most conservative / least opinionated set of rational conditionals subsuming Γ [Rot91] [HP03] [LM92]

Some results

Theorem 7

If $\div$ is obtained from $*$ by a TQ combinator, then, for $i \in \{1,2,3,4\}$, $(Ci^{*}_{\preccurlyeq})$ entails $(Ci^{\div}_{\preccurlyeq})$

Theorem 8

If $\div$ is obtained from $*$ by the STQ combinator, then $*_R$ and $*_N$ both yield $\div_N$

Theorem 9

If $\div$ is obtained from $*$ by the STQ combinator, then $*_L$ yields the operator $\div_{\text{STQL}} \neq \div_L, \div_P, \div_N$

Open question

Open question

Unless we identify states with TPOs, TQ combination only applies to 2-step change. Can the idea be generalised to finer-grained representations, e.g. ranking functions or POIs?

Extending (LI): triviality result

- Just as (NiHI) led to problems, so does this naive extension of (LI):

(NiLI) $[\Psi * A]_c = \mathrm{Cn}([\Psi \div \neg A]_c \cup \{A\})$

where we extend Cn to subsets Δ of L_c by $\mathrm{Cn}(\Delta) = \Delta \cup \mathrm{Cn}(\Delta \cap L)$.

Theorem 10

$(K2^*)$, $(K2^{\div})$ and (NiLI) entail that there are no consistent belief sets

- Perhaps use a more suitable closure operator, e.g.

(iLIR) $[\Psi * A]_c = \mathrm{C_R}([\Psi \div \neg A]_c \cup \{A\})$

Extending the Levi Identity (cont.)

- [NGOP05] offer the following semantic proposal

$(\mathrm{iLI}{*}_{\preccurlyeq})$ $\preccurlyeq_{\Psi * A} = \preccurlyeq_{(\Psi \div \neg A) * A}$

- Note: this does not reduce iterated revision to iterated contraction!
- How does this fare in conjunction with the syntactic proposal (iLIR)?

Theorem 11

$(\mathrm{iLI}{*}_{\preccurlyeq})$ and (iLIR) entail $* = *_{\mathrm{N}}$

Open question

Open question

For each remaining elementary operator $*$, is there an alternative closure operator C, such that $[\Psi * A]_c = \mathrm{C}([\Psi \div \neg A]_c \cup \{A\})$ iff $\preccurlyeq_{\Psi * A} = \preccurlyeq_{(\Psi \div \neg A) * A}$?

8. ITERATED CHANGE IN CONDITIONAL BELIEFS

Beliefs after a c-revision

- How to extend $*$ to take conditionals as inputs (c-revision)?
- Regarding beliefs resulting from single c-revisions, plausibly:

$$\min(\preccurlyeq_{\Psi * A \Rightarrow B}, W) = \min(\preccurlyeq_{\Psi * A \rightarrow B}, W)$$
$$[\Psi * A \Rightarrow B] = [\Psi * A \rightarrow B]$$

- No extra constraints needed

C-beliefs after a c-revision: material conditionals

- Regarding c-beliefs resulting from single c-revisions, plausibly:

$$\preccurlyeq_{\Psi * A \Rightarrow B} \neq \preccurlyeq_{\Psi * A \rightarrow B}, \text{ i.e. } [\Psi * A \Rightarrow B]_c \neq [\Psi * A \rightarrow B]_c$$

- Indeed, most strengthenings of DP *don't* yield

$$\min(\preccurlyeq_{\Psi * A \rightarrow B}, [\![A]\!]) \subseteq [\![B]\!], \text{ i.e. } A \Rightarrow B \in [\Psi * A \rightarrow B]_c$$

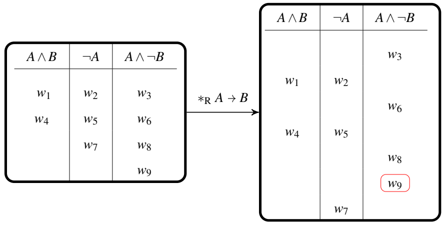

C-beliefs after a c-revision: material conditionals

- But we want 'Success':

(S$_{\preccurlyeq}$) $\min(\preccurlyeq_{\Psi * A \Rightarrow B}, [\![A]\!]) \subseteq [\![B]\!]$

(S) $A \Rightarrow B \in [\Psi * A \Rightarrow B]_c$

The KI postulates

- Beyond (S), the KI postulates [Ker99] provide further constraints
- Semantically:

(KI1$^*_{\preccurlyeq}$) If $x, y \in [\![A \wedge B]\!]$, $x, y \in [\![\neg A]\!]$, or $x, y \in [\![A \wedge \neg B]\!]$, then $x \preccurlyeq_\Psi y$ iff $x \preccurlyeq_{\Psi * A \Rightarrow B} y$

(The relative standing of worlds that 'say the same thing' about $A \Rightarrow B$ (both confirm, both disconfirm or are both irrelevant) is unaffected)

(KI2$^*_{\preccurlyeq}$) If $x \in [\![A \wedge B]\!]$, $y \in [\![A \wedge \neg B]\!]$ and $x \prec_\Psi y$, then $x \prec_{\Psi * A \Rightarrow B} y$

(KI3$^*_{\preccurlyeq}$) If $x \in [\![A \wedge B]\!]$, $y \in [\![A \wedge \neg B]\!]$ and $x \preccurlyeq_\Psi y$, then $x \preccurlyeq_{\Psi * A \Rightarrow B} y$

(Worlds that confirm $A \Rightarrow B$ an only be promoted wrt worlds that disconfirm it)

- The DP postulates are the special cases of these in which $A = \top$

Overview

- A number of further proposals have been suggested. For example:
  (1) Nayak *et al* [NPFP96]
  (2) Boutillier & Goldzsmidt [BG93]
  (3) Kern-Isberner [Ker01]
- The above all commit to particular reductive operator and (2) has further issues

Nayak *et al*

- Most strengthenings of DP do not give us

$$A \Rightarrow B \in [\Psi * A \to B]_c$$

  ... with the exception of $*_L$
- Nayak *et al* simply suggest that $\Psi * A \Rightarrow B = \Psi *_L A \to B$
- So we have (S)
- Furthermore, since KI $\equiv$ DP for non-conditionals and $*_L$ satisfies the latter, the proposal satisfies the former

Boutillier & Goldzsmidt

- Boutillier & Goldzsmidt offer an alternative extension of $*_N$
- It satisfies: If $A \in [\Psi]$, then $B \in [\Psi * A \Rightarrow B]$
- This essentially precludes reasoning by Modus Tollens:

Example 4

I believe that the light in the bathroom next door is on (A), because the light switch in this room is down ($\neg B$). My friend, who lives in the house, tells me that, contrary to what one might expect, when the bathroom light is on, that means that the switch in this room is *up*. So I revise by $A \Rightarrow B$. In doing so, I maintain my belief about the state of the switch ($\neg B$) and conclude that the bathroom light is off ($\neg A$).

Kern-Isberner

- Kern-Isberner's extends a particular family of reductive revision operators ($\neq *_N$, $*_R$ or $*_L$)
- The extension satisfies (S) and KI

Open questions

Open question

Given the disagreement about iterated revision, could one perhaps provide a general way of extending to c-revision any arbitrary iterated revision operator?

Open question

What about c-contraction?

9. Concluding comments

Summary

- Iterated revision hinges on beliefs about overruling
- TPOs are probably not rich enough to capture these
- The two main proposed enrichments of TPOs require further work:
  - POIs: determining posterior POI
  - Ranking functions: axiomatically justifying r-revision
- In both instances, it remains to be seen what logic of overruling emerges from the semantics

Summary (cont.)

- Regarding the relation between iterated revision and contraction:
  - Existing extensions of (HI) only cover 2-step change (assuming states $\neq$ TPOs)
  - Most plausible extension of (LI) doesn't reduce revision to contraction
- Regarding iterated c-revision/c-contraction:
  - very sparse literature
  - common assumption that states = TPOs
- Finally: points of contact with non-mon. logics (closure operators)

Thanks for your time and endurance!

Questions and comments welcome:
`jake.chandler@cantab.net`
`BoothR2@cardiff.ac.uk`

References

[AGM85] Carlos E Alchourrón, Peter Gärdenfors, and David Makinson, *On the logic of theory change: Partial meet contraction and revision functions*, Journal of Symbolic Logic **50** (1985), no. 02, 510–530.

[BC16] Richard Booth and Jake Chandler, *Extending the Harper Identity to iterated belief change*, Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI-16), New York, NY, USA, 9-15 July 2016, 2016, pp. 987–993.

[BC17] Richard Booth and Jake Chandler, *The irreducibility of iterated to single revision*, Journal of Philosophical Logic **46** (2017), no. 4, 405–418.

[BC18] ———, *On strengthening the logic of iterated belief revision: Proper ordinal interval operators*, Proceedings of the 16th International Conference on Principles of Knowledge Representation and Reasoning (KR-08), Tempe, AZ, USA, October 30–November 2, 2018, 2018, pp. 210–219.

References (cont.)

[BCss] ———, *From iterated revision to iterated contraction: Extending the Harper Identity*, Artificial Intelligence (in press).

[BFH98] Craig Boutilier, Nir Friedman, and Joseph Y. Halpern, *Belief revision with unreliable observations*, In Proceedings, Fifteenth National Conference on Artificial Intelligence (AAAI '96, 1998, pp. 127–134.

[BFKP12] Richard Booth, Eduardo Fermé, Sébastien Konieczny, and Ramón Pino Pérez, *Credibility-limited revision operators in propositional logic*, Proceedings of the Thirteenth International Conference on Principles of Knowledge Representation and Reasoning, KR'12, AAAI Press, 2012, pp. 116–125.

[BG93] Craig Boutilier and Moisés Goldszmidt, *Revision by conditional beliefs*, Proceedings of the 11th National Conference on Artificial Intelligence. Washington, DC, USA, July 11-15, 1993., 1993, pp. 649–654.

References (cont.)

[BM06] Richard Booth and Thomas Meyer, *Admissible and restrained revision*, Journal of Artificial Intelligence Research **26** (2006), no. 1, 127–151.

[BM11] ———, *How to revise a total preorder*, Journal of Philosophical Logic **40** (2011), no. 2, 1–46.

[Bou96] Craig Boutilier, *Iterated revision and minimal change of conditional beliefs*, Journal of Philosophical Logic **25** (1996), no. 3, 263–305.

[CGMW08] Samir Chopra, Aditya Ghose, Thomas Meyer, and Ka-Shu Wong, *Iterated belief change and the recovery axiom*, Journal of Philosophical Logic **37** (2008), no. 5, 501–520.

[Dal88] Mukesh Dalal, *Investigations into a theory of knowledge base revision: Preliminary report*, Proceedings of the Seventh AAAI National Conference on Artificial Intelligence, AAAI'88, AAAI Press, 1988, pp. 475–479.

[Doy79] Jon Doyle, *A truth maintenance system*, Artif. Intell. **12** (1979), no. 3, 231–272.

References (cont.)

[DP97] Adnan Darwiche and Judea Pearl, *On the logic of iterated belief revision*, Artificial Intelligence **89** (1997), no. 1, 1–29.

[FUV83] Ronald Fagin, Jeffrey D. Ullman, and Moshe Y. Vardi, *On the semantics of updates in databases*, Proceedings of the Second ACM SIGACT-SIGMOD Symposium on Principles of Database Systems, March 21-23, 1983, Colony Square Hotel, Atlanta, Georgia, USA, 1983, pp. 352–365.

[Gla00] Stephen Murray Glaister, *Recovery recovered*, Journal of Philosophical Logic **29** (2000), no. 2, 171–206.

[Gro88] Adam Grove, *Two modellings for theory change*, Journal of Philosophical Logic **17** (1988), no. 2, 157–170.

[Han91] Sven Ove Hansson, *Belief contraction without recovery*, Studia Logica: An International Journal for Symbolic Logic **50** (1991), no. 2, 251–260.

[Han92] ______, *In defense of the Ramsey Test*, Journal of Philosophy **89** (1992), no. 10, 522.

References (cont.)

[Han93] ______, *Reversing the levi identity*, Journal of Philosophical Logic **22** (1993), no. 6, 637–669.

[Han96] Sven Ove Hansson, *Hidden structures of belief.*, Logic, Action, and Information, 1996, pp. 79–100.

[Har76] William L Harper, *Rational conceptual change*, PSA: Proceedings of the Biennial Meeting of the Philosophy of Science Association, 1976, pp. 462–494.

[HFCF01] Sven Ove Hansson, Eduardo Leopoldo Fermé, John Cantwell, and Marcelo Alejandro Falappa, *Credibility limited revision*, The Journal of Symbolic Logic **66** (2001), no. 4, 1581–1596.

[HP03] Lee C. Hill and Jeff B. Paris, *When maximizing entropy gives the rational closure*, Journal of Logic and Computation **13** (2003), no. 1, 51–68.

References (cont.)

[HS08] Matthias Hild and Wolfgang Spohn, *The measurement of ranks and the laws of iterated contraction*, Artificial Intelligence **172** (2008), no. 10, 1195 – 1218.

[JT07] Yi Jin and Michael Thielscher, *Iterated belief revision, revised*, Artificial Intelligence **171** (2007), no. 1, 1–18.

[Ker99] Gabriele Kern-Isberner, *Postulates for conditional belief revision*, Proceedings of the Sixteenth International Joint Conference on Artificial Intelligence, IJCAI 99, Stockholm, Sweden, July 31 - August 6, 1999. 2 Volumes, 1450 pages, 1999, pp. 186–191.

[Ker01] ______, *Conditionals in nonmonotonic reasoning and belief revision - considering conditionals as agents*, Lecture Notes in Computer Science, vol. 2087, Springer, 2001.

[KM91] Hirofumi Katsuno and Alberto O Mendelzon, *Propositional knowledge base revision and minimal change*, Artificial Intelligence **52** (1991), no. 3, 263–294.

References (cont.)

[KP08a] Sébastien Konieczny and Ramón Pino Pérez, *Improvement operators*, Principles of Knowledge Representation and Reasoning: Proceedings of the Eleventh International Conference, KR 2008, Sydney, Australia, September 16-19, 2008, 2008, pp. 177–187.

[KP08b] ______, *Improvement operators*, Principles of Knowledge Representation and Reasoning: Proceedings of the Eleventh International Conference, KR 2008, Sydney, Australia, September 16-19, 2008, 2008, pp. 177–187.

[Lev77] Isaac Levi, *Subjunctives, dispositions and chances*, Synthèse **34** (1977), no. 4, 423–455.

[Lev91] Isaac Levi, *The fixation of belief and its undoing: Changing beliefs through inquiry*, Cambridge University Press, 1991.

[LM92] Daniel Lehmann and Menachem Magidor, *What does a conditional knowledge base entail?*, Artificial Intelligence **55** (1992), no. 1, 1–60.

[Mak97] D. Makinson, *Screened revision*, Theoria **63** (1997), no. 1-2, 14–23.

[Mey01] Thomas Meyer, *Basic infobase change*, Studia Logica: An International Journal for Symbolic Logic **67** (2001), no. 2, 215–242.

[NGO07] Abhaya C Nayak, Randy Goebel, and Mehmet A Orgun, *Iterated belief contraction from first principles*, Proceedings of the 20th International Joint Conference on Artificial Intelligence (IJCAI-07), Hyderabad, India, January 6–12 2007, 2007, pp. 2568–2573.

[NGOP05] Abhaya C. Nayak, Randy Goebel, Mehmet A. Orgun, and Tam Pham, *Iterated belief change and the Levi Identity*, Belief Change in Rational Agents: Perspectives from Artificial Intelligence, Philosophy, and Economics, 7.-12. August 2005, 2005.

[NPFP96] Abhaya C. Nayak, Maurice Pagnucco, Norman Y. Foo, and Pavlos Peppas, *Learning from conditionals: Judy benjamin's other problems*, 12th European Conference on Artificial Intelligence, Budapest, Hungary, August 11-16, 1996, Proceedings, 1996, pp. 75–79.

[NPP03] Abhaya C Nayak, Maurice Pagnucco, and Pavlos Peppas, *Dynamic belief revision operators*, Artificial Intelligence **146** (2003), no. 2, 193–228.

[Par99] Rohit Parikh, *Logic, language and computation, vol. 2*, Center for the Study of Language and Information, Stanford, CA, USA, 1999, pp. 266–278.

[PFS08] Pavlos Peppas, Anastasios Michael Fotinopoulos, and Stella Seremetaki, *Conflicts between relevance-sensitive and iterated belief revision*, Proceedings of the 2008 Conference on ECAI 2008: 18th European Conference on Artificial Intelligence (Amsterdam, The Netherlands, The Netherlands), IOS Press, 2008, pp. 85–88.

[Pol87] John L. Pollock, *Defeasible reasoning*, Cognitive Science **11** (1987), no. 4, 481–518.

[PW18] Pavlos Peppas and Mary-Anne Williams, *Parametrised difference revision*, Principles of Knowledge Representation and Reasoning: Proceedings of the Sixteenth International Conference, KR 2018, Tempe, Arizona, 30 October - 2 November 2018., 2018, pp. 277–286.

[PWCF15] Pavlos Peppas, Mary-Anne Williams, Samir Chopra, and Norman Y. Foo, *Relevance in belief revision*, Artif. Intell. **229** (2015), 126–138.

[Rab96] Wlodzimierz Rabinowicz, *Stable revision, or is preservation worth preserving?*, Logic, Action, and Information, 1996, pp. 101–128.

[Rot91] Hans Rott, *A nonmonotonic conditional logic for belief revision*, The Logic of Theory Change (André Fuhrmann and Michael Morreau, eds.), Springer, 1991, pp. 135–181.

[Rot04] ______, *A counterexample to six fundamental principles of belief formation*, Synthese **139** (2004), no. 2, 225–240.

[Sat88] Ken Satoh, *Nonmonotonic reasoning by minimal belief revision*, Tech. report, Institute for New Generation Computer Technology, 1988.

[Sav54] L. Savage, *The foundations of statistics*, Wiley, New York, 1954.

[Spo14] Wolfgang Spohn, *The laws of belief - ranking theory and its philosophical applications*, Oxford University Press, 2014.

[Sta09] Robert Stalnaker, *Iterated belief revision*, Erkenntnis (1975-) **70** (2009), no. 2, 189–209.

[Wil86] M.W. Wilkins, *A model-theoretic approach to updating logical databases*, Tech. report, 1986.