Logic and Learning

Day 3: Variations on truth-tracking by Belief Revision

Nina Gierasimczuk

Tsinghua Logic Summer School
June 30th–July 3rd, 2026

Outline

Reminder

Collective truth-tracking by Belief Revision

Bias and Belief Revision

Epistemic spaces and plausibility orders

An *epistemic space* is $\mathbb{S} = (S, \mathcal{O})$, where S a set of possible worlds, and $\mathcal{O} \subseteq \mathcal{P}(S)$ the set of observations. Both S and $\mathcal{O}$ are at most countable.

A *plausibility space* is $\mathbb{B} = (S, \mathcal{O}, \preceq)$, where $\preceq \subseteq S \times S$ is a plausibility preorder (a reflexive and transitive binary relation) on S.

We write $s \prec t$ iff $s \preceq t$ and $t \not\preceq s$; $(s, t) \in \preceq$ (or $s \preceq t$) means that s at least as plausible as t (an arrow from t to s).

Given $\mathbb{S} = (S, \mathcal{O})$, $\mathbb{B}_{\mathbb{S}} = \{(S, \mathcal{O}, \preceq) \mid \preceq \text{ is a preorder on } S\}$.

Knowledge and belief

Let $\mathbb{B} = (S, \mathcal{O}, \preceq)$ be a plausibility space, and let $p \in \mathcal{O}$.

$$\mathbb{B}_{\mathbb{S}} \models Kp \quad \text{iff} \quad S \subseteq p$$

$min(\mathbb{B}) = \{s \in S \mid \text{there is no } t \in S, \text{ s.t. } (t, s) \in \prec\}$.

$$\mathbb{B}_{\mathbb{S}} \models Bp \quad \text{iff} \quad min_{\preceq} S \subseteq p$$

In other words, the set of $\preceq$-minimal possible worlds determines agent's beliefs.

In case S is countably infinite and $\preceq$ is non-well-founded, we have two choices:

1. define a limiting notion of belief: $\mathbb{B}_{\mathbb{S}} \models Bp$ iff $\exists w\ \forall u \preceq w\ \ u \in p$;
2. set belief to be a partial function (if no minimum, belief is undefined).

Observations and data streams

For any $s \in S$, $\mathcal{O}_s = \{O \in \mathcal{O} \mid s \in O\}$ is the set of observations true in s.
We assume that for any $s, t \in S$ if $s \neq t$, then $\mathcal{O}_s \neq \mathcal{O}_t$.

A *stream* of observations ε is an infinite sequence of elements of $\mathcal{O}$.
Let $n \in \mathbb{N}$:

- $\varepsilon[n]$ is an initial segment of ε of length $n+1$;
- ε^n is the n-th element of ε;
- $set(\varepsilon)$ is the set of elements enumerated in ε.

A stream ε is *sound* for s if $set(\varepsilon) \subseteq \mathcal{O}_s$ and it is *complete* for s if $\mathcal{O}_s \subseteq set(\varepsilon)$.

Belief revision methods

A *one-step revision method* is a function R_1, which on the input of a plausibility space and an observation outputs a new, revised, plausibility space.

Definition

- $\mathrm{Cond}_1((S, \mathcal{O}, \preceq), O) = (S', \mathcal{O}^{S'}, \preceq^{S'})$, where $S' = S \cap O$, $\mathcal{O}^{S'} = \{O \cap S' \mid O \in \mathcal{O}\}$ and $\preceq^{S'} = \preceq \cap (S' \times S')$;
- $\mathrm{Lex}_1((S, \mathcal{O}, \preceq), O) = (S, \mathcal{O}, \preceq')$, where $(s, t) \in \preceq'$ if one of the following conditions holds: $(s, t) \in \preceq^{O}$, $(s, t) \in \preceq^{\overline{O}}$ or ($s \in O$ and $t \notin O$);
- $\mathrm{Mini}_1((S, \mathcal{O}, \preceq), O) = (S, \mathcal{O}, \preceq')$, where $\preceq'$ is s.t.: if $s \in min(\preceq)$and $t \notin min(\preceq)$ then $(s, t) \in \preceq'$, and $(s, t) \in \preceq'$ iff $(s, t) \in \preceq$, otherwise.

We iterate revision methods: $R(\mathbb{B}, \sigma * O) := R_1(R(\mathbb{B}, \sigma), O)$, for $R \in \{\mathrm{Cond}, \mathrm{Lex}, \mathrm{Mini}\}$.

Learners

A *learner* is a function $L((S, \mathcal{O}), \sigma) \subseteq S$.

- L learns $s \in S$ if for any ε sound and complete for s, there is a $k \in \mathbb{N}$ such that for all $n \geq k$, $L(\mathbb{S}, \varepsilon[n]) = \{s\}$.
- L learns $\mathbb{S}$ if it learns all $s \in S$.
- $\mathbb{S}$ is *learnable* if there is an L that learns it.

L is *universal* if it can learn any epistemic space that is learnable (by any learning function).

Given a revision method R and a plausibility $\preceq$, a revision-based learner is

$$L_R^{\preceq}((S, \mathcal{O}), \sigma) = min(R((S, \mathcal{O}, \preceq), \sigma)).$$

$\mathbb{S} = (S, \mathcal{O})$ is learnable by R if there is a $\preceq$ on S s.t. $L_R^{\preceq}$ learns $\mathbb{S}$.

Universality results

Definition
L is **universal** if it can learn every epistemic space that is learnable.

	Conditioning	Lexicographic	Minimal
Positive Streams	YES	YES	NO

Occamist plausibility preorder $\preceq_{occ}$ is the so-called specialisation preorder:

$$(s, t) \in \preceq_{occ} \text{ iff } \mathcal{O}_s \subseteq \mathcal{O}_t.$$

Theorem
Minimal revision is not universal.

Theorem
There is no universal belief-revision method under well-foundedness.

A. Baltag, N. Gierasimczuk, S. Smets. Truth tracking by belief revision. Studia Logica 2019.

Outline

Reminder

Collective truth-tracking by Belief Revision

Bias and Belief Revision

Why Collective Truth-Tracking?

- Belief revision models how agents update beliefs in response to new data.
- Classical accounts are based on *rationality postulates* (e.g., consistency after revision).

A complementary perspective:

Does revision lead to truth-learning?

Conditioning and lexicographic revision are truth-tracking in single-agent ideal setup.

In practice, learning is distributed:

- multiple agents,
- partial and heterogeneous information,
- need for aggregation.

Main question:

When do groups of agents collectively track the truth, and which belief-merging methods preserve individual learning power?

Outline of the multi-agent setting

Take finitely many agents, $1..n$, and a plausibility space $(S, \mathcal{O}, \preceq)$, with a designated actual world, $s \in S$.

Each agent preforms their individual inquiry by receiving information from a separate data stream. These individual streams, while individually sound for s, are not necessarily (individually) complete for s. However, taken together the information is complete for s, i.e., the sum of the contents of all of the streams is the set of all and only true observations at s.

A *data profile* is $\Sigma = (\varepsilon_1..\varepsilon_n)$ (inputs to each of the n agents individually).

As before, $\Sigma[k] = (\varepsilon_1[k]..\varepsilon_n[k])$, $\Sigma^k = (\varepsilon_1^k..\varepsilon_n^k)$, $set(\Sigma) = \bigcup_{i=1}^n set(\varepsilon_i)$.

At each stage of their individual inquiry, their plausibility spaces are **merged**, by some plausibility merge function.

Collective Revision

A collective version $\mathbf{R}$ of the single-agent revision method R on E and Σ at step k is defined as follows:

$$\mathbf{R}(E, \Sigma[k]) = \mathbf{R}((\mathbb{B}_1..\mathbb{B}_n), (\varepsilon_1[k]..\varepsilon_n[k])) = (R(\mathbb{B}_1, \varepsilon_1[k])..R(\mathbb{B}_n, \varepsilon_n[k]).$$

Our learners are homogenous (all agents apply the same revision)
and synchronous (all agent apply their updates at the same discrete pace).

Plausibility merge

A plausibility *merge* is a function f that on the input of finitely many plausibility spaces outputs a new (plausibility) space.

$E = (\mathbb{B}_1..\mathbb{B}_n)$ denotes a *plausibility profile* that represents plausibility spaces of our n agents.

- *knowledge-intersection merge* given by $f^{K\cap}(E) = (\bigcap_{i=1}^{n} S_i, \mathcal{O}^{S'}, \preceq^{S'})$.
- *plausibility-intersection merge* given by $f^{P\cap}(E) = (S, \mathcal{O}, \bigcap_{i=1}^{n} \preceq_i)$.
- *majority merge* given by $f^{maj}(E) = (S, \mathcal{O}, maj)$, where $(s,t) \in maj$ if $(s,t) \in \preceq_i$ for majority of plausibility spaces in E.

Collective Learners

A collective learner is $\mathbf{L}(\mathbb{S}, \Sigma[k]) \subseteq S$.
A data profile Σ is sound and complete for $s \in S$ iff $\bigcup_{i=1}^{n} set(\varepsilon_i) = \mathcal{O}_s$.

Collective learner $\mathbf{L}$ learns $s \in S$ if for every Σ sound and complete for s, there is a $k \in \mathbb{N}$, such that for all $n \geq k$, $\mathbf{L}(\mathbb{S}, \Sigma[n]) = \{s\}$.

Merge f is applied to outputs of $\mathbf{R}$ yielding a 'joint plausibility space', which can be inspected for minimal elements that give a (collective) conjecture.

A collective learner based on f, R and $\preceq$ is

$$\mathbf{L}^{\preceq}_{R,f}((S, \mathcal{O}), \Sigma[k]) = min(f(\mathbf{R}((S, \mathcal{O}, \preceq), \Sigma[k]))).$$

Collective revision $\mathbf{R}$ and merge f learn $\mathbb{S}$, if there is a prior $\preceq$, such that $\mathbf{L}^{\preceq}_{R,f}$ learns $\mathbb{S}$.

Conditioning and intersection merge

Conditioning with a sequence of observations can be replaced by conditioning with a single observation: a conjunction (intersection) of the elements of σ.

Lemma

Let $\Sigma = (\sigma_1..\sigma_n)$ be a data profile consisting of n data sequences, and let σ be a data sequence such that $\bigcup_{i=1}^{n} set(\sigma_i) = set(\sigma)$ (i.e., the profile and the sequence enumerate the same set of data). Then we have that

$$f^{K\cap}(\mathbf{Cond}((S,\mathcal{O},\preceq),\Sigma)) = \mathrm{Cond}((S,\mathcal{O},\preceq),\sigma).$$

Proposition

The collective learner $\mathbf{L}$ *based on* $\mathbf{Cond}$ *and* $f^{K\cap}$ *collectively learns* $(S,\mathcal{O},\preceq)$ *if and only if the learner L based on* Cond *learns* $(S,\mathcal{O},\preceq)$.

COLLECTIVE LEARNING WITH LEXICOGRAPHIC REVISION: $f^{P\cap}$?

- $S = \{\underline{s}, t\}$
- $\mathcal{O} = \{p, q\}$, with $p = \{s, t\}$, $q = \{s\}$
- $\preceq = \{(t, s), (s, s), (t, t)\}$

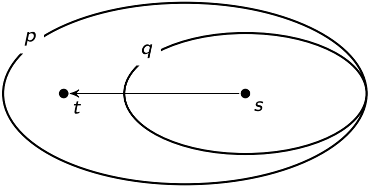

First note that single-agent *Lex* learns $(S, \mathcal{O})$ on $\preceq$ (which is Occam).

Take sound and complete for s data profile $\Sigma = (\varepsilon_1, \varepsilon_2)$ such that:

- $set(\varepsilon_1) = \{p, q\}$;
- $set(\varepsilon_2) = \{p\}$.

From some k on:

$$\text{Lex}(\preceq, \varepsilon_1[k]) = \{(s, t), (s, s), (t, t)\} = \mathbb{B}_{\mathbb{S}1}, \text{ and}$$
$$\text{Lex}(\preceq, \varepsilon_2[k]) = \{(t, s), (s, s), (t, t)\} = \mathbb{B}_{\mathbb{S}2}.$$

$f^{P\cap}(\mathbb{B}_{\mathbb{S}1}, \mathbb{B}_{\mathbb{S}2}) = \{(s, s), (t, t)\}$ does not single out one most plausible world.

So, $\mathbf{L}^{\preceq}_{Lex, P\cap}$ does not learn $(S, \mathcal{O})$.

Collective learning with lexicographic revision: f^{maj}?

- $S = \{\underline{s}, t\}$
- $\mathcal{O} = \{p, q\}$, with $p = \{s, t\}$, $q = \{s\}$
- $\preceq = \{(t, s), (s, s), (t, t)\}$
- $set(\varepsilon_1) = \{p, q\}$
- $set(\varepsilon_2) = \{p\}$
- $set(\varepsilon_3) = \{p\}$

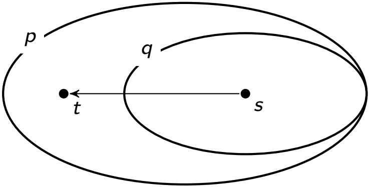

The majority of agents (2 and 3) will forever agree that t is more plausible than s.

$\mathbf{L}^{\preceq}_{Lex,maj}$ will always output $\preceq$.
I.e., it wrongly converges to t on a data profile that is sound and complete for s.

Unlike Cond, Lex is not set driven

Proposition

Lex *is not set-driven.*

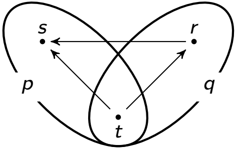

$\mathbb{B}$

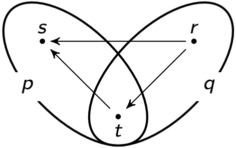

$\mathsf{Lex}(\mathbb{B}, (p))$

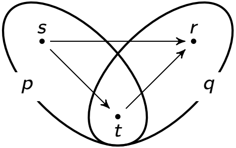

$\mathsf{Lex}(\mathbb{B}, (q))$

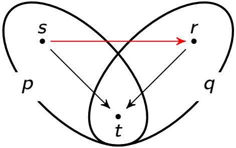

$\mathsf{Lex}(\mathbb{B}, (p, q))$

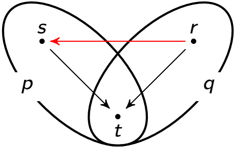

$\mathsf{Lex}(\mathbb{B}, (q, p))$

Descendant merge?

Definition

Two worlds s and t are *comparable* in $\preceq$ if $(s,t) \in \preceq$ or $(t,s) \in \preceq$.
A *descendant merge* is defined in the following way:

$$f^d(\preceq, (\preceq_i)_{i \in \{1,\dots,n\}}) = \bigcup_{i=1}^{n} (\preceq_i \setminus \preceq) \cup \tag{1}$$

$$\{(s,t) \in \preceq \;\mid\; s \text{ and } t \text{ are incomparable in } \bigcup_{i=1}^{n} (\preceq_i \setminus \preceq)\} \tag{2}$$

Unlike previously discussed merging methods, apart from the plausibility profile, f^d takes as an input the initial plausibility.

In a sense, descendant merge is a 'memory-based' merge method.

Descendant merge

Definition

A *descendant* merge is defined in the following way:

$$f^d(\preceq, (\preceq_i)_{i\in\{1..n\}}) = \bigcup_{i\in\{1..n\}} \mathtt{diff}(\preceq, \preceq_i) \cup \tag{3}$$

$$(\mathtt{cons}((\preceq_i)_{i\in\{1..n\}}) \setminus (\preceq \cup \preceq^{-1})) \cup \tag{4}$$

$$\{(s,s) \mid s \in S\} \tag{5}$$

where:

- $\mathtt{diff}(\preceq, \preceq_i) = \{(s,t) \mid (s,t) \in \preceq_i \text{ and } (t,s) \in \preceq\}$;
- $\mathtt{cons}((\preceq_i)_{i\in\{1..n\}}) = \{(s,t) \mid$ there is i s.t. $(s,t) \in \preceq_i$ and there is no j s.t. $(t,s) \in \preceq_j\}$.

Collective learning with Lex and f^d

Example 1

- $S = \{\underline{s}, t\}$
- $\mathcal{O} = \{p, q\}$, with $p = \{s, t\}$, $q = \{s\}$
- $\preceq = \{(t, s), (s, s), (t, t)\}$
- $set(\varepsilon_1) = \{p, q\}$
- $set(\varepsilon_2) = \{p\}$

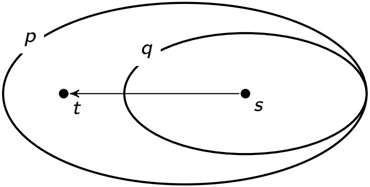

For the step $k \in \mathbb{N}$, such that:

$set(\varepsilon_1[k]) = \{p, q\}$ and $set(\varepsilon_2[k]) = \{p\}$.

We get that:

(1) $\mathtt{diff}(\preceq, \mathsf{Lex}(\preceq, \varepsilon_1[k])) = \{(s, t)\}$ and $\mathtt{diff}(\preceq, \mathsf{Lex}(\preceq, \varepsilon_2[k])) = \emptyset$.
So, $\bigcup_{i \in \{1,2\}} \mathtt{diff}(\preceq, \preceq_i) = \{(s, t)\}$;

(2) $\mathtt{cons}((\preceq_i)_{i \in \{1..n\}}) \setminus (\preceq \cup \preceq^{-1}) = \emptyset$;

(3) after adding all reflexive pairs, we get a preorder, in which s is the most plausible.

Individual streams can only repeat the previously seen information, so the individual outcomes and the collective outcome will not change from that point on.

Collective learning with Lex and f^d

Example 2

- $set(\varepsilon_1) = \{p, q\}$: $\varepsilon_1[1] = (p, q)$
- $set(\varepsilon_2) = \{q, r\}$: $\varepsilon_2[1] = (q, r)$

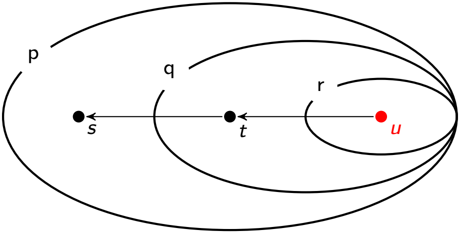

1. $\mathsf{Lex}(\preceq, p) = \preceq$, so $\texttt{diff}(\preceq, \mathsf{Lex}(\preceq, p)) = \emptyset$;
2. $\mathsf{Lex}(\preceq, q) = \{(t, u), (u, s), (t, s), (s, s), (t, t), (u, u)\}$, so $\texttt{diff}(\preceq, \preceq') = \{(u, s), (t, s)\}$.

Together the agents learned: $\emptyset \cup \{(u, s), (t, s)\} = \{(u, s), (t, s)\}$.

They didn't learn anything new about the relationship between u and t, nor about any other pairs (since all states were initially comparable).

Hence, $f^d(\preceq, \mathbf{Lex}(\preceq, \Sigma[0])) = \{(u, s), (t, s), (s, s), (t, t), (u, u)\}$.

- $set(\varepsilon_1) = \{p, q\}$: $\varepsilon_1[1] = (p, q)$
- $set(\varepsilon_2) = \{q, r\}$: $\varepsilon_2[1] = (q, r)$

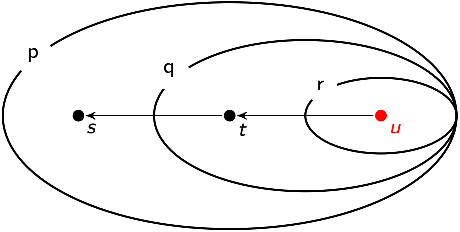

1. $\mathrm{Lex}(\preceq, (p, q)) = \{(u, s), (t, s), (t, u), (s, s), (t, t), (u, u)\}$, so $\mathtt{diff}(\preceq, \mathrm{Lex}(\preceq, (p, q))) = \{(u, s), (t, s)\}$
2. $\mathrm{Lex}(\preceq, (q, r)) = \preceq'' = \{(u, t), (u, s), (t, s), (s, s), (t, t), (u, u)\}$, and so $\mathtt{diff}(\preceq, \preceq'' = \{(u, t), (u, s), (t, s)\}$.

Together they learned $\{(u, t), (u, s), (t, s)\}$.

Hence, $f^d(\preceq, \mathbf{Lex}(\preceq, \Sigma[1])) = \{(u, t), (u, s), (t, s), (s, s), (t, t), (u, u)\}$.
The actual world u is the unique minimal world.

Since the streams can only repeat the previously given information, neither of the sets of minimal elements of the individual preorders will change at any later stage.

Collective learning with Lex and f^d

Example 3

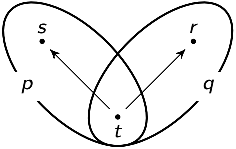

$\mathbb{B}$

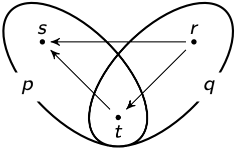

Agent 1: $\mathrm{Lex}(\mathbb{B}, (p))$

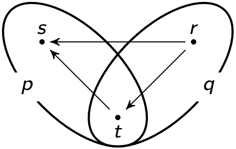

Agent 2: $\mathrm{Lex}(\mathbb{B}, (p))$

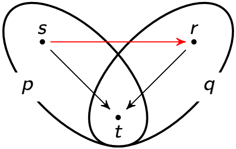

Agent 1: $\mathrm{Lex}(\mathbb{B}, (p, q))$

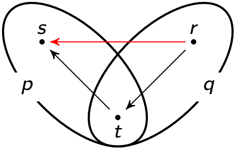

Agent 2: $\mathrm{Lex}(\mathbb{B}, (p, p))$

Example 3, cntd

After the first step, both agents learned that (t, r) and they also learned about a previously incomparable pair: they both have that (s, r). Hence, the f^d at this stage will output $\{(t, r), (s, r), (s, s), (t, t), (r, r)\}$.

In the next step, agent 1 learns (t, s) and (t, r), and (with respect to the initially incomparable worlds) (r, s). Agent 2 only learned (t, r) and with respect to the initially incomparable worlds (s, r). At this stage the union of `diff`-sets for both agents is $\{(t, s), (t, r)\}$. The `cons`-set is however at this stage empty: the learned orderings of worlds incomparable according to the original $\preceq$ are opposite for both agents: (r, s) for agent 1 and (s, r) for agent 2 (red arrows). So, at this stage f^d outputs $\{(t, s), (t, r), (s, s), (t, t), (r, r)\}$. The most plausible world is the actual world t, and the resulting plausibility relation is reflexive and transitive, but not total.

Related work

- Multi-agent belief revision in Dynamic Epistemic Logic [4, 5],

  but we currently abstract away from higher-order epistemic interactions.
- Contrast with belief merging and probabilistic truth-tracking approaches:
  - Integrity-constrained belief merging [6, 7].
  - Truth recovery via Condorcet-style analyses of aggregation operators [8].
  - Truthlikeness-based evaluation when truth is excluded by constraints [9].
  - Worst-case robustness guarantees for belief fusion [10].
- Symbolic truth-tracking work with non-expert information sources [11, 12] relaxes the assumption of sound observation streams; use of conditioning and finite identifiability.

Summary and future work

Summary:

- A general framework for truth-tracking with plausibility revision and merge.
- The role of the interaction between revision and merge.
- Universality of conditioning together with knowledge intersection merge.
- Newly introduced memory-base descendant merge.
- Conjecture about lexicographic revision and the descendant merge.

Future work:

- Universality of Lex and descendant merge.
- Self-correcting streams.
- Relating to distributed computing.
- ...

Literature

A. Baltag, N. Gierasimczuk, and S. Smets. Truth-tracking by belief revision. *Studia Logica*, 107:917–947, 2019.

N. Gierasimczuk. *Knowing One's Limits: Logical Analysis of Inductive Inference*. PhD thesis, Universiteit van Amsterdam, 2010.

A. Baltag, N. Gierasimczuk, and S. Smets. Belief revision as a truth-tracking process. In *Proceedings of TARK XIII*, pages 187–190. ACM, 2011.

A. Baltag and S. Smets. Dynamic belief revision over multi-agent plausibility models. In *Proceedings of LOFT VI*, pages 11–24, 2006.

A. Baltag and S. Smets. A qualitative theory of dynamic interactive belief revision. In H. Arló-Costa, V. F. Hendricks, and J. van Benthem (eds.), *Readings in Formal Epistemology*, pages 813–858. Springer, 2016.

S. Konieczny and R. P. Pérez. Merging with integrity constraints. In *Symbolic and Quantitative Approaches to Reasoning and Uncertainty*, pages 233–244. Springer, 1999.

S. Konieczny and R. P. Pérez. Merging information under constraints: A logical framework. *JLC*, 12:773–808, 2002.

P. Everaere, S. Konieczny, and P. Marquis. The epistemic view of belief merging: Can we track the truth? In *ECAI 2010*, pages 621–626. IOS Press, 2010.

S. D'Alfonso. Belief merging with the aim of truthlikeness. *Synthese*, 193:2013–2034, 2016.

J. Karge and S. Rudolph. The more the worst-case-merrier: A generalized Condorcet jury theorem for belief fusion. In *KR 2022*, pages 205–214, 2022.

J. Singleton and R. Booth. Who's the expert? On multi-source belief change. In *KR 2022*, pages 331–340, 2022.

J. Singleton and R. Booth. Truth-tracking with non-expert information sources. *JAIR*, 81:619–641, 2024.

Outline

Reminder

Collective truth-tracking by Belief Revision

Bias and Belief Revision

Types of bias

We combine revision methods with three kinds of 'cognitive' limitations:

- issues with accepting the input;
- issues with perceiving the input;
- issues with access to belief state.

We investigate (theoretically and practically) their impact on truth-tracking.

The procedure

After the random generation of an epistemic space, one of the states, s, is randomly designated to be the actual world, and a sound and complete data sequence σ for s is generated. Then a plausibility preorder is randomly generated and each of the (biased) revision methods is called to identify s on σ.

The simulation

Each series of tests consisted of 200 runs, the plausibility spaces consisted of 5 possible worlds and 12 observables, and the incoming data sequence was always longer than the number of observables.

Issues with accepting the input (confirmation bias)

Definition

Given an $(S, \mathcal{O})$, the *stubbornness function* is $D : \mathcal{P}(S) \to \mathbb{N}$.

Definition

R_{CB} is defined in the following way:

$$R_{CB}(\mathbb{B}, \lambda) = \mathbb{B},$$

$$R_{CB}(\mathbb{B}, \sigma * p) = \begin{cases} R_1(R_{CB}(\mathbb{B}, \sigma), p) & \text{if} \quad \#p(\sigma) \geq D(\overline{p}), \\ R_{CB}(\mathbb{B}, \sigma) & \text{otherwise.} \end{cases}$$

We obtain $Cond_{CB}$, Lex_{CB}, $Mini_{CB}$.

Results

Proposition
Cond, Lex, Mini are strictly more powerful than $Cond_{CB}$, Lex_{CB}, $Mini_{CB}$.

Corollary
$Cond_{CB}$ and Lex_{CB} are not universal.

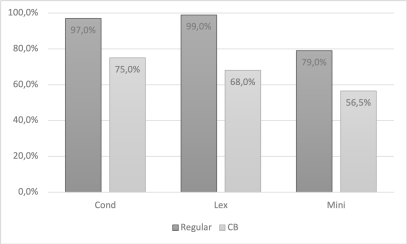

Issues with perceiving the input (framing bias)

Definition

Given $(S, \mathcal{O})$, the *framing function* is $FR : \mathcal{O} \to \mathcal{P}(S)$.

Definition (Framing-bias methods)

We define a framing-biased method in the following way:

$$R_{FR}(\mathbb{B}, \lambda) = \mathbb{B},$$

$$R_{FR}(\mathbb{B}, \sigma * p) = R_1(R_{FR}(\mathbb{B}, \sigma), x), \text{ such that } x \in FR(p).$$

We obtain $Cond_{FR}$, Lex_{FR}, $Mini_{FR}$.

Results

Proposition
$Cond_{FR}$ and Lex_{FR} are not universal.

Proposition
Mini is strictly more powerful than $Mini_{FR}$.

Proposition
Lex_{FR} is universal on fairly framed streams.

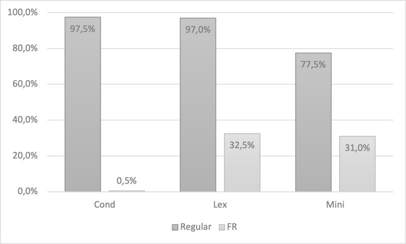

Issues with access to belief state (anchoring bias)

Definition (Anchoring-bias methods)

We define anchoring-biased method in the following way: after receiving p they 'choose' the best world satisfying p. If there are minimal several worlds that are equi-plausible, the method picks one at random.

Additionally, a **resource parameter** (a real between 0 and 100) halves each time a revision takes place, and the process terminates when resource is depleted (< 1).

Results

Proposition

$Cond_{AB}$, Lex_{AB} are not universal.

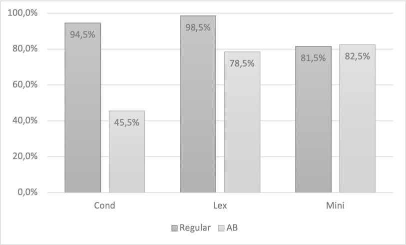

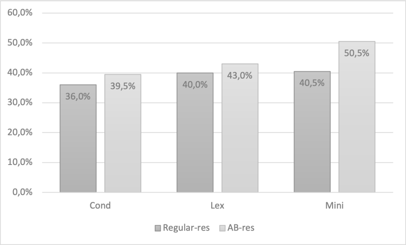

Anchoring ability to select a random world as the candidate for the actual world improves the truth-tracking capability, especially in the case of minimal revision.