Bridging Bayesian probability and AGM revision via stability principles

Chris Mierzewski
(with Alexandru Baltag)

Institute for Logic, Language and Computation

Overview

Compare the behaviour of AGM revision and Bayesian conditioning.

Overview

Compare the behaviour of AGM revision and Bayesian conditioning.

Plan

- Stability
- The tracking problem
- No-Go Theorem (Lin&Kelly)
- Threshold-raising
- Recovering AGM from Bayes through Maximum Entropy

Tom & Johan

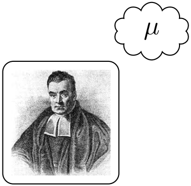

Tom & Johan

$$\mu \xrightarrow[\text{Bayesian conditioning}]{|_X} \mu_X$$

Tom & Johan

$$\mu \xrightarrow[\text{Bayesian conditioning}]{|_X} \mu_X$$

$$K$$

Tom & Johan

$$\mu \xrightarrow[\text{Bayesian conditioning}]{|_X} \mu_X$$

$$K \xrightarrow[\text{revision operator}]{*_X} K^*X$$

Tom & Johan

$$
\begin{tikzcd}[column sep=huge, row sep=huge]
\mu \arrow[r, "|_X", "\text{Bayesian conditioning}"'] \arrow[d, "\text{acceptance}"] & \mu_X \arrow[d, "\text{acceptance}"'] \\
K \arrow[r, "*_X", "\text{revision operator}"'] & K^* X
\end{tikzcd}
$$

Harmony and acceptance

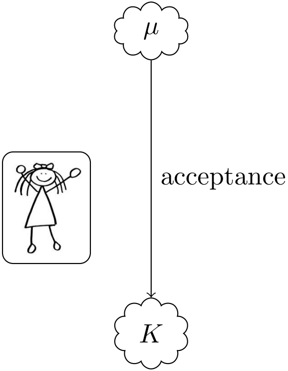

Harmony and acceptance

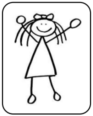

$$\begin{tikzcd}
\mu \arrow[rr, "|_X", "\text{Bayesian conditioning}"'] \arrow[d, "\text{acceptance}"] & & \mu_X \\
K \arrow[rr, "*_X", "\text{revision operator}"'] & & K^* X
\end{tikzcd}$$

Harmony and acceptance

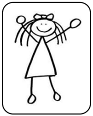

$$
\begin{tikzcd}
\mu \arrow[rr, "|_X", "\text{Bayesian conditioning}"'] \arrow[d, "\text{acceptance}"] & & \mu_X \arrow[d, "\text{acceptance}"'] \\
K \arrow[rr, "*_X", "\text{revision operator}"'] & & K^*X
\end{tikzcd}
$$

Setup

- Probability spaces $(\Omega, \mathfrak{A}, \mu)$. Propositions $X, Y, ... \in \mathfrak{A}$.
- $\Delta_{\mathfrak{A}}$ is the set of all probability measures on $\mathfrak{A}$.

Setup

- Probability spaces $(\Omega, \mathfrak{A}, \mu)$. Propositions $X, Y, \ldots \in \mathfrak{A}$.
- $\Delta_{\mathfrak{A}}$ is the set of all probability measures on $\mathfrak{A}$.
- Acceptance rule: map $\alpha : \Delta_{\mathfrak{A}} \to \mathfrak{A}$. The agent accepts $X \in \mathfrak{A}$ if and only if $\alpha(\mu) \subseteq X$: i.e. $\alpha(\mu)$ is the strongest accepted proposition.

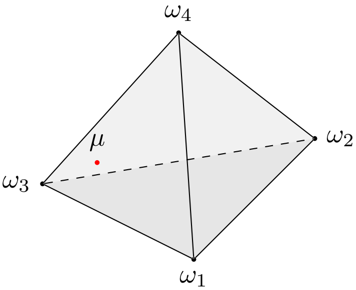

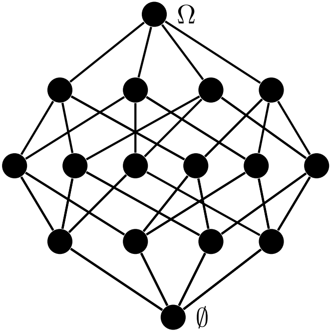

Setup

- Probability spaces $(\Omega, \mathfrak{A}, \mu)$. Propositions $X, Y, ... \in \mathfrak{A}$.
- $\Delta_{\mathfrak{A}}$ is the set of all probability measures on $\mathfrak{A}$.
- Acceptance rule: map $\alpha : \Delta_{\mathfrak{A}} \to \mathfrak{A}$. The agent accepts $X \in \mathfrak{A}$ if and only if $\alpha(\mu) \subseteq X$: i.e. $\alpha(\mu)$ is the strongest accepted proposition.

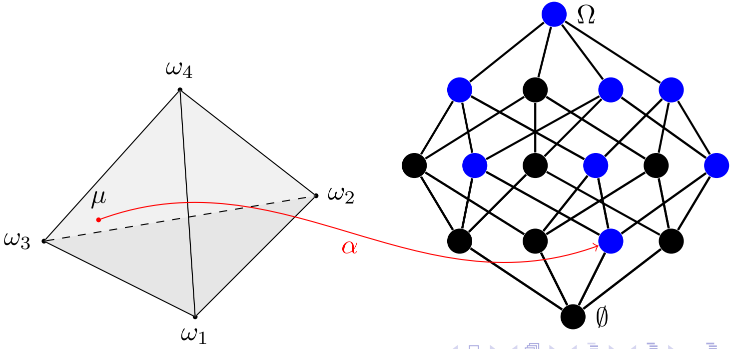

Setup 2

- Qualitative revision operators $* : \mathfrak{A} \times \mathfrak{A} \to \mathfrak{A}$: first variable represents the current strongest accepted proposition, and the second the new revision input.

For any belief state $K \in \mathfrak{A}$ and propositions X, Y, the revision $*$ is AGM-compliant if

- $K^*X \subseteq X$
- $K \cap X \subseteq K^*X$ (Inclusion)
- If $K \cap X \neq \emptyset$, then $K^*X \subseteq K \cap X$ (Preservation)
- If $K^*X = \emptyset$ then $K = \emptyset$ or $X = \emptyset$
- $(K^*X) \cap Y \subseteq K^*(X \cap Y)$
- If $(K^*X) \cap Y \neq \emptyset$, then $K^*(X \cap Y) \subseteq (K^*X) \cap Y$

Tracking

Tracking

A qualitative revision policy A maps each $\mu \in \Delta_{\mathfrak{A}}$ to a proposition $\alpha(\mu)$ and a revision operator $*$ applicable to that proposition. It *tracks* Bayesian conditioning if:

$$\forall \mu \in \Delta_{\mathfrak{A}}, \forall X \in \mathfrak{A} \text{ with } \mu(X) > 0, \; \alpha(\mu)^* X = \alpha(\mu_X).$$

$$\begin{tikzcd}
\mu \arrow[rr, "|_X", "\text{Bayesian conditioning}"'] \arrow[d, "\alpha", "\text{acceptance rule}"'] & & \mu_X \arrow[d, "\alpha"] \\
\alpha(\mu) \arrow[rr, "*_X", "\text{revision operator}"'] & & \alpha(\mu)^* X = \alpha(\mu_X)
\end{tikzcd}$$

Old problems with Locke

Lockean rule λ_t with threshold t:

$$(\rightarrow) \quad \text{If } \lambda_t(\mu) \subseteq X \text{ then } \mu(X) \geq t$$

$$(\leftarrow) \quad \text{If } \mu(X) \geq t \text{ then } \lambda_t(\mu) \subseteq X$$

Old problems with Locke

Lockean rule λ_t with threshold t:

$$(\rightarrow)\quad \text{If } \lambda_t(\mu) \subseteq X \text{ then } \mu(X) \geq t$$
$$(\leftarrow)\quad \text{If } \mu(X) \geq t \text{ then } \lambda_t(\mu) \subseteq X$$

- Acceptance must be reasonable: it must avoid Lottery-style paradoxes, but it should not require measure 1.
- Leitgeb: keep the $(\rightarrow)$-direction of the Lockean thesis, but restrict the $(\leftarrow)$-direction.

Stability (Leitgeb)

Stability

Let $(\Omega, \mathfrak{A}, \mu)$ a probability space, and $t \in (0.5, 1]$.

A set $X \in \mathfrak{A}$ is (μ, t)-*stable* if and only if $\forall Y \in \mathfrak{A}$ such that $X \cap Y \neq \emptyset$ and $\mu(Y) > 0$, we have $\mu_Y(X) \geq t$.

Stability (Leitgeb)

Stability

Let $(\Omega, \mathfrak{A}, \mu)$ a probability space, and $t \in (0.5, 1]$.
A set $X \in \mathfrak{A}$ is (μ, t)-*stable* if and only if $\forall Y \in \mathfrak{A}$ such that $X \cap Y \neq \emptyset$ and $\mu(Y) > 0$, we have $\mu_Y(X) \geq t$.

- Interpretation: a proposition X is (μ, t)-stable if *only* learning a proposition *inconsistent* with X can bring the probability of X below the threshold: robustness under new information.
- Note that the probability of stable propositions is always above the threshold.

Stability (Leitgeb)

Two proposed norms for acceptance:

the Stability Principle (SP) : *"Given a threshold t and $\mu \in \Delta_{\mathfrak{A}}$, the strongest accepted proposition must be a (μ, t)-stable set in $\mathfrak{A}$."*

The strongest accepted proposition must be robust.

Stability (Leitgeb)

Two proposed norms for acceptance:

the Stability Principle (SP) : *"Given a threshold t and $\mu \in \Delta_{\mathfrak{A}}$, the strongest accepted proposition must be a (μ, t)-stable set in $\mathfrak{A}$."*

The strongest accepted proposition must be robust.

Relativised Lockean Principle (RLP): "*Accept as many propositions X with $\mu(X) \geq t$ as is possible without violating* (SP)."

We want to believe as many propositions above the threshold as we can, to remain close to Lockean intuitions.

Stability (Leitgeb)

Two proposed norms for acceptance:

the Stability Principle (SP) : *"Given a threshold t and $\mu \in \Delta_{\mathfrak{A}}$, the strongest accepted proposition must be a (μ, t)-stable set in $\mathfrak{A}$."*

The strongest accepted proposition must be robust.

Relativised Lockean Principle (RLP): "*Accept as many propositions X with $\mu(X) \geq t$ as is possible without violating* (SP)."

We want to believe as many propositions above the threshold as we can, to remain close to Lockean intuitions.

$\rightarrow$ Together this suggests: "*pick the logically strongest ($\subseteq$-least) stable set*". Does one always exist?

Systems of Spheres

Well-ordering (Leitgeb)

Let $\mu \in \Delta_{\mathfrak{A}}$ a σ-additive measure, $t \in (0.5, 1]$. Then the set $\mathfrak{S}^{\tau}_{<1}(\mu) := \{X \in \mathfrak{A} \,|\, \mu(X) < 1 \text{ and } X \text{ is } (\mu, t)\text{-stable}\}$ is well-ordered by set inclusion, and has order type at most ω.

Systems of Spheres

Well-ordering (Leitgeb)

Let $\mu \in \Delta_{\mathfrak{A}}$ a σ-additive measure, $t \in (0.5, 1]$. Then the set $\mathfrak{S}^{\tau}_{<1}(\mu) := \{X \in \mathfrak{A} \,|\, \mu(X) < 1 \text{ and } X \text{ is } (\mu, t)\text{-stable}\}$ is well-ordered by set inclusion, and has order type at most ω.

We do need σ-additivity here to exclude infinitely descending chains: (suppose we have a chain $X_0 \supset ...X_n \supset X_{n+1} \supset ...$ with X_0 a (μ, t)-stable set and $\mu(X_0) < 1$. Let $A_i := X_i \setminus X_{i+1}$. Then $\mu(X_0) = \lim_{n\to\infty}(\sum_{i=0}^{n} \mu(A_i))$. The limit of partial sums converges, so $\lim_{n\to\infty}(\mu(A_n)) = 0$. Take the sequence $\langle \mu(X_0 \,|\, X_0^c \cup A_i) : i \in \mathbb{N}\rangle$. Each term is equal to $\frac{\mu(A_i)}{\mu(X_0^c)+\mu(A_i)}$ with $\mu(A_i) \to 0$. So there is some N with $\mu(X_0 \,|\, X_0^c \cup A_N) < t$. So X_0 is not (μ, t)-stable after all; contradiction.)

Systems of Spheres

- Problem with stable sets of measure one: a least set of measure 1 may not exist.

Systems of Spheres

- Problem with stable sets of measure one: a least set of measure 1 may not exist.
- Leitgeb: fix it by postulation. Restrict attention to spaces having such a set (Least Certain Set property).

Systems of Spheres

- Problem with stable sets of measure one: a least set of measure 1 may not exist.
- Leitgeb: fix it by postulation. Restrict attention to spaces having such a set (Least Certain Set property).

Spheres

Let $(\Omega, \mathfrak{A}, \mu)$ a space satisfying *(LCS)*, $t \in (0.5, 1]$. Let S_∞ the least measure-1 set in $\mathfrak{A}$. Then the set $\mathfrak{S}^\tau(\mu) := \mathfrak{S}^\tau_{<1}(\mu) \cup \{S_\infty\}$ is well-ordered by set-inclusion.

- Fairly severe restriction, but things work fine for regular spaces, countable full powerset algebras...

The τ-rule

The τ-rule

For any probability measure μ on $\mathfrak{A}$ which satisfies (LCS), and any $t \in (0.5, 1]$, let $\mathfrak{S}^{\tau}(\mu)$ the system of spheres generated by μ. Then we define the map $\tau_t : \Delta_{\mathfrak{A}} \to \mathfrak{A}$ as

$$\tau_t(\mu) := \min_{\subseteq} \mathfrak{S}^{\tau}(\mu)$$

The τ-rule

The τ-rule

For any probability measure μ on $\mathfrak{A}$ which satisfies (LCS), and any $t \in (0.5, 1]$, let $\mathfrak{S}^{\tau}(\mu)$ the system of spheres generated by μ. Then we define the map $\tau_t : \Delta_{\mathfrak{A}} \to \mathfrak{A}$ as

$$\tau_t(\mu) := \min_{\subseteq} \mathfrak{S}^{\tau}(\mu)$$

- No Lottery paradox.

The τ-rule

The τ-rule

For any probability measure μ on $\mathfrak{A}$ which satisfies (LCS), and any $t \in (0.5, 1]$, let $\mathfrak{S}^\tau(\mu)$ the system of spheres generated by μ. Then we define the map $\tau_t : \Delta_{\mathfrak{A}} \to \mathfrak{A}$ as

$$\tau_t(\mu) := \min_{\subseteq} \mathfrak{S}^\tau(\mu)$$

- No Lottery paradox.
- $\mathfrak{S}^\tau(\mu)$ is a system of spheres: it generates a ranking (total preorder) on Ω. Via Grove's Theorem, each system of spheres generates and AGM revision operator.

The τ-rule

The τ-rule

For any probability measure μ on $\mathfrak{A}$ which satisfies (LCS), and any $t \in (0.5, 1]$, let $\mathfrak{S}^{\tau}(\mu)$ the system of spheres generated by μ. Then we define the map $\tau_t : \Delta_{\mathfrak{A}} \to \mathfrak{A}$ as

$$\tau_t(\mu) := \min_{\subseteq} \mathfrak{S}^{\tau}(\mu)$$

- No Lottery paradox.
- $\mathfrak{S}^{\tau}(\mu)$ is a system of spheres: it generates a ranking (total preorder) on Ω. Via Grove's Theorem, each system of spheres generates and AGM revision operator.

→ Leitgeb's τ-rule (1) is plausible as an acceptance principle, (2) offers a nice connection with AGM.

Acceptance zones

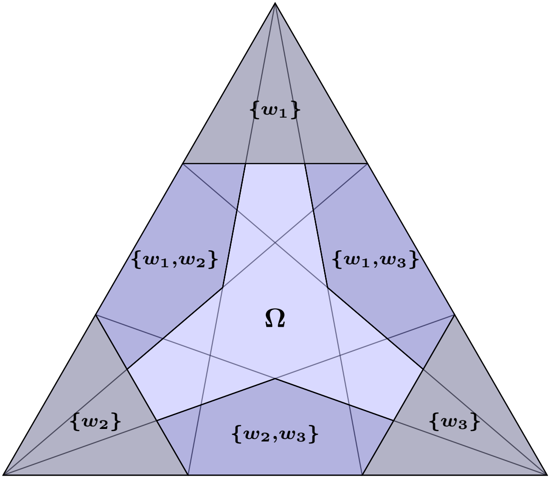

Figure: Acceptance zones for Leitgeb's τ-rule with $t = 2/3$ and $|\Omega| = 3$.

Tracking

Observation

Let $\mu \in \Delta_{\mathfrak{A}}$, $t \in (0.5, 1]$, and $*$ the AGM revision generated by τ_t. Then $\forall X \in \mathfrak{A}$ with $\mu(X) > 0$, the set $\tau(\mu)^* X$ is (μ_X, t)-stable.

- So the AGM-revised set $\tau(\mu)^* X$ is always stable after conditioning. Is it always the *least* stable set?

Tracking

Observation

Let $\mu \in \Delta_{\mathfrak{A}}$, $t \in (0.5, 1]$, and $*$ the AGM revision generated by τ_t. Then $\forall X \in \mathfrak{A}$ with $\mu(X) > 0$, the set $\tau(\mu)^* X$ is (μ_X, t)-stable.

- So the AGM-revised set $\tau(\mu)^* X$ is always stable after conditioning. Is it always the *least* stable set? **No.**

Tracking

Observation

Let $\mu \in \Delta_{\mathfrak{A}}$, $t \in (0.5, 1]$, and $*$ the AGM revision generated by τ_t. Then $\forall X \in \mathfrak{A}$ with $\mu(X) > 0$, the set $\tau(\mu)^* X$ is (μ_X, t)-stable.

- So the AGM-revised set $\tau(\mu)^* X$ is always stable after conditioning. Is it always the *least* stable set? **No.**

- $\Omega := \{\omega_1, ..., \omega_4\}$ and $\mathfrak{A}$ the full power set algebra over Ω.
- Set $t = 0.7$.
- Take $\mu = (0.5, 0.12, 0.05, 0.33)$ and $X := \{\omega_1, \omega_2, \omega_3\}$. Then

$\rightarrow$ $\tau(\mu) = \{\omega_1, \omega_2, \omega_4\}$ and so $\tau(\mu)^* X = \{\omega_1, \omega_2\} = \tau(\mu) \cap X$.

$\rightarrow$ conditioning on X gives $\mu_X \approx (0.746, 0.179, 0.075, 0)$, and we get $\tau(\mu_X) = \{\omega_1\}$. So $\tau(\mu_X) \subset \tau(\mu)^* X$: conditioning raises the probability of ω_1 just enough to make it (μ_X, t)-stable.

No-Go Theorem (Lin&Kelly)

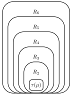

Figure: System of spheres centered on X.

In this case, we have $\tau(\mu_X) \subset \tau(\mu)^* X$; further, the revision $\tau(\mu) \to \tau(\mu_X)$ is not AGM, as it fails Inclusion.

- This example also shows how *tracking fails* for τ.

No-Go Theorem (Lin&Kelly)

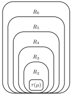

Figure: System of spheres centered on X.

In this case, we have $\tau(\mu_X) \subset \tau(\mu)^* X$; further, the revision $\tau(\mu) \to \tau(\mu_X)$ is not AGM, as it fails Inclusion.

- This example also shows how *tracking fails* for τ. Special case of:

No-Go Theorem (Lin&Kelly)

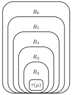

Figure: System of spheres centered on X.

In this case, we have $\tau(\mu_X) \subset \tau(\mu)^* X$; further, the revision $\tau(\mu) \to \tau(\mu_X)$ is not AGM, as it fails Inclusion.

- This example also shows how *tracking fails* for τ. Special case of:

The No-Go Theorem (Lin&Kelly)

Let $|\Omega| > 2$, $\mathfrak{A}$ a field of sets over Ω, and let $\alpha : \Delta_{\mathfrak{A}} \to \mathfrak{A}$ be any *sensible* acceptance rule. Then no AGM revision policy based on α tracks Bayesian conditioning.

Tracking failure

- How tracking fails for the τ-rule: in general, conditioning + acceptance results in a logically **stronger** belief set than acceptance + AGM revision. No commutativity whenever it is strictly stronger.

$$\begin{tikzcd}
\mu \arrow[r, "|_X", "\text{Bayes}"'] \arrow[d, "\tau", "\text{acceptance}"'] & \mu_X \\
\tau(\mu) \arrow[r, "*_X", "\text{AGM revision}"'] & \tau(\mu)^* X
\end{tikzcd}$$

Tracking failure

- How tracking fails for the τ-rule: in general, conditioning + acceptance results in a logically **stronger** belief set than acceptance + AGM revision. No commutativity whenever it is strictly stronger.

$$
\begin{tikzcd}
\mu \arrow[rr, "|_X", "\text{Bayes}"'] \arrow[d, "\tau", "\text{acceptance}"'] & & \mu_X \arrow[d, "\tau"] \\
\tau(\mu) \arrow[r, "*_X", "\text{AGM revision}"'] & \tau(\mu)^* X \supset & \tau(\mu_X)
\end{tikzcd}
$$

No go.

- Close, but not quite.

Does not comMute!!

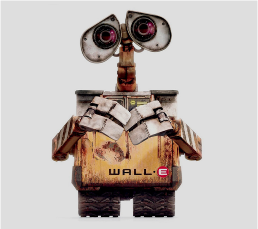

Picture

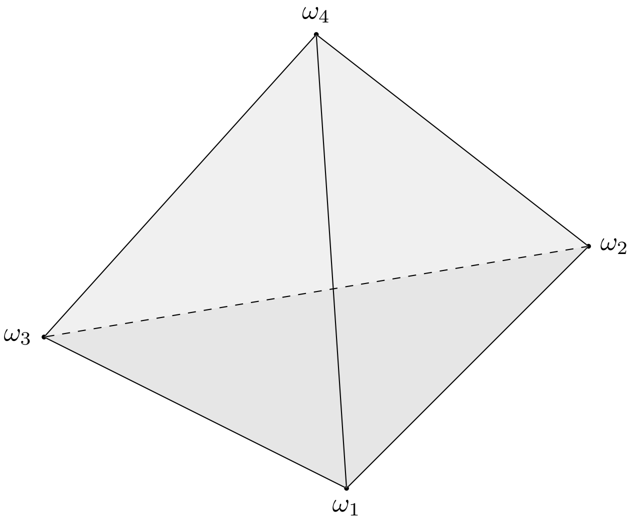

Picture

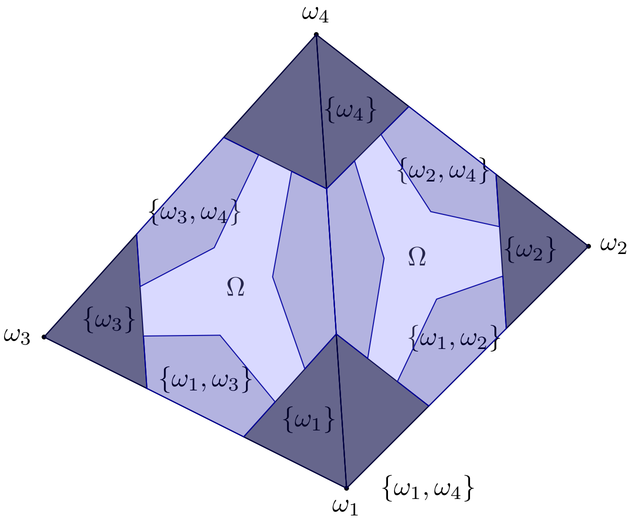
ω_4
$\{\omega_4\}$
$\{\omega_2, \omega_4\}$
$\{\omega_3, \omega_4\}$
Ω
Ω
$\{\omega_2\}$
ω_2
ω_3
$\{\omega_3\}$
$\{\omega_1, \omega_2\}$
$\{\omega_1, \omega_3\}$
$\{\omega_1\}$
$\{\omega_1, \omega_4\}$
ω_1

Picture

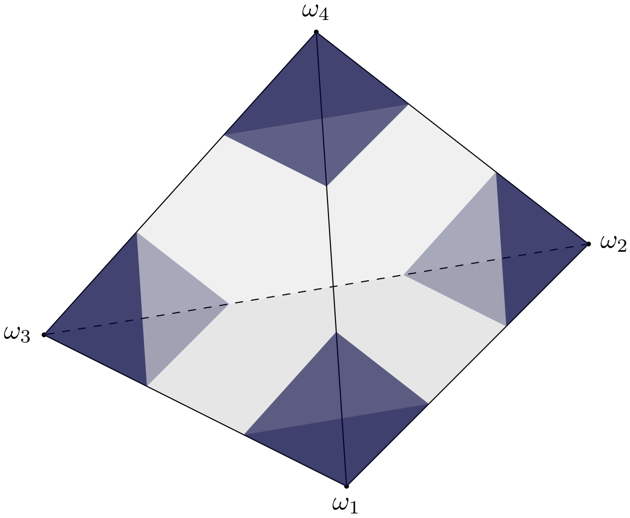

Picture

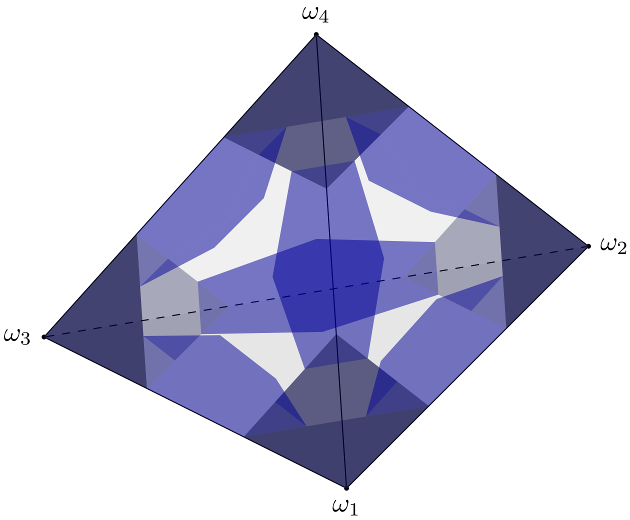

Picture

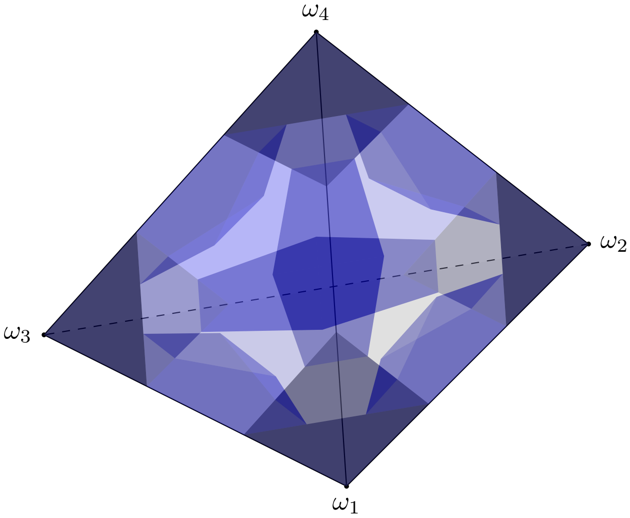

Urns example

You are given an urn. You know that it is either of the type **A** – containing 30% black marbles and 70% white marbles, or **B** – containing 70% black and 30% white. Suppose you draw (with replacement) 10 marbles form the urn. How many black marbles would you have to draw to be convinced your urn is of type **A** ?

- 0,1 or 2 black marbles, for a threshold of 0.5. But drawing 3 marbles yields disagreement between conditioning and revision: on the Bayesian side, you then believe your urn is of type **A**. On the AGM side, you are undecided.

All this assuming a 50-50 prior for urns **A, B** *and using a binomial distribution to compute conditional probabilities.*

Threshold-raising

- When tracking fails for the τ-rule, selecting $\tau(\mu)^* X$ as strongest accepted proposition goes against the Lockean principle (RLP), since $\tau(\mu_X)$ is logically stronger.

Threshold-raising

- When tracking fails for the τ-rule, selecting $\tau(\mu)^*X$ as strongest accepted proposition goes against the Lockean principle (RLP), since $\tau(\mu_X)$ is logically stronger.
- Can we 'force' agreement by changing the threshold?
- Idea: When tracking fails for a threshold t, raise the threshold to a new value $q > t$, so that the AGM-revised set becomes the least stable set for q.

Threshold-raising

- When tracking fails for the τ-rule, selecting $\tau(\mu)^*X$ as strongest accepted proposition goes against the Lockean principle (RLP), since $\tau(\mu_X)$ is logically stronger.
- Can we 'force' agreement by changing the threshold?
- Idea: When tracking fails for a threshold t, raise the threshold to a new value $q > t$, so that the AGM-revised set becomes the least stable set for q. That is: $\tau_q(\mu_X) = \tau_t(\mu)^*X$, and so the revision $\tau_t(\mu) \mapsto \tau_q(\mu_X)$ is AGM.

Threshold-raising

We want a q such that $\tau_q(\mu_X) = \tau_t(\mu)^* X$.

$$
\begin{tikzcd}
\mu \arrow[r, "|_X", "\text{Bayes}"'] \arrow[d, "\tau_t", "\text{acceptance}"'] & \mu_X \arrow[d, "\tau_t"] \\
\tau_t(\mu) \arrow[r, "*_X", "\text{AGM revision}"'] & \tau_t(\mu)^* X \supseteq \tau_t(\mu_X)
\end{tikzcd}
$$

Threshold-raising

We want a q such that $\tau_q(\mu_X) = \tau_t(\mu)^* X$.

$$
\begin{tikzcd}
\mu \arrow[r, "|_X", "\text{Bayes}"'] \arrow[d, "\tau_t", "\text{acceptance}"'] & \mu_X \arrow[d, "\tau_{\color{red}q}"'] \arrow[dr, "\tau_t"] & \\
\tau_t(\mu) \arrow[r, "*_X", "\text{AGM revision}"'] & \tau_t(\mu)^* X \supset & \tau_t(\mu_X)
\end{tikzcd}
$$

Threshold-raising

We want a q such that $\tau_q(\mu_X) = \tau_t(\mu)^* X$.

$$
\begin{tikzcd}
\mu \arrow[r, "|_X", "\text{Bayes}"'] \arrow[d, "\tau_t", "\text{acceptance}"'] & \mu_X \arrow[d, "\tau_{\color{red}q}"'] \arrow[dr, "\tau_t"] & \\
\tau_t(\mu) \arrow[r, "*_X", "\text{AGM revision}"'] & \tau_t(\mu)^* X \supset & \tau_t(\mu_X)
\end{tikzcd}
$$

- Works only if $\exists q \in (0.5, 1]$, $\tau_t(\mu)^* X$ is the least (μ_X, q)-stable set, while $\tau_t(\mu_X)$ is not stable for q.

Threshold-raising

We want a q such that $\tau_q(\mu_X) = \tau_t(\mu)^* X$.

$$\begin{tikzcd}[column sep=huge, row sep=huge]
\mu \arrow[r, "|_X", "\text{Bayes}"'] \arrow[d, "\tau_t", "\text{acceptance}"'] & \mu_X \arrow[d, "\tau_{\color{red}q}"'] \arrow[d, bend left, "\tau_t"] \\
\tau_t(\mu) \arrow[r, "*_X", "\text{AGM revision}"'] & \tau_t(\mu)^* X \supset \tau_t(\mu_X)
\end{tikzcd}$$

- Works only if $\exists q \in (0.5, 1]$, $\tau_t(\mu)^* X$ is the least (μ_X, q)-stable set, while $\tau_t(\mu_X)$ is not stable for q.
- When do we have such a threshold q?

Stability comes in degrees

Degree of stability

The degree of stability of $X \in \mathfrak{A}$ with respect to a measure $\mu \in \Delta_{\mathfrak{A}}$, denoted $\mathcal{S}(\mu, X)$ is defined as:

$$\mathcal{S}(\mu, X) := \sup\{q \in [0,1] \mid X \text{ is } (\mu, q)\text{-stable}\}$$

When $\mu(X) > 0$, we have $\mathcal{S}(\mu, X) = \inf\{\mu_Y(X) \mid \mu(Y) > 0,\ X \cap Y \neq \emptyset\}$, and

$$X \text{ is } (\mu, t)\text{-stable if and only if } \mathcal{S}(\mu, X) \geq t$$

Stability comes in degrees

Degree of stability

The degree of stability of $X \in \mathfrak{A}$ with respect to a measure $\mu \in \Delta_{\mathfrak{A}}$, denoted $\mathcal{S}(\mu, X)$ is defined as:

$$\mathcal{S}(\mu, X) := \sup\{q \in [0,1] \mid X \text{ is } (\mu, q)\text{-stable}\}$$

When $\mu(X) > 0$, we have $\mathcal{S}(\mu, X) = \inf\{\mu_Y(X) \mid \mu(Y) > 0,\ X \cap Y \neq \emptyset\}$, and

$$X \text{ is } (\mu, t)\text{-stable if and only if } \mathcal{S}(\mu, X) \geq t$$

One can raise the threshold to “correct” the revision process *only if* $\mathcal{S}(\mu_X, \tau(\mu_X)) < \mathcal{S}(\mu_X, \tau(\mu)^* X)$

Incorrigible

We have $t = 0.7$, a distribution $\mu = (0.5, 0.12, 0.05, 0.33)$, and $X = \{\omega_1, \omega_2, \omega_3\}$. Here $\tau_t(\mu) = \{\omega_1, \omega_2, \omega_4\}$. Then $\mu_X \approx (0.746, 0.179, 0.075, 0)$, and tracking fails since $\tau_t(\mu_X) = \{\omega_1\}$ and $\tau_t(\mu)^* X = \{\omega_1, \omega_2\}$. We have the following degrees of stability with respect to μ_X:

Incorrigible

We have $t = 0.7$, a distribution $\mu = (0.5, 0.12, 0.05, 0.33)$, and $X = \{\omega_1, \omega_2, \omega_3\}$. Here $\tau_t(\mu) = \{\omega_1, \omega_2, \omega_4\}$. Then $\mu_X \approx (0.746, 0.179, 0.075, 0)$, and tracking fails since $\tau_t(\mu_X) = \{\omega_1\}$ and $\tau_t(\mu)^* X = \{\omega_1, \omega_2\}$. We have the following degrees of stability with respect to μ_X:

$$\mathcal{S}(\tau_t(\mu_X)) = \frac{\mu_X(\omega_1)}{\mu_X(\omega_1) + \mu_X(\Omega \setminus \{\omega_1\})} = \frac{0.746}{1} = 0.746$$

Incorrigible

We have $t = 0.7$, a distribution $\mu = (0.5, 0.12, 0.05, 0.33)$, and $X = \{\omega_1, \omega_2, \omega_3\}$. Here $\tau_t(\mu) = \{\omega_1, \omega_2, \omega_4\}$. Then $\mu_X \approx (0.746, 0.179, 0.075, 0)$, and tracking fails since $\tau_t(\mu_X) = \{\omega_1\}$ and $\tau_t(\mu)^* X = \{\omega_1, \omega_2\}$. We have the following degrees of stability with respect to μ_X:

$$\mathcal{S}(\tau_t(\mu_X)) = \frac{\mu_X(\omega_1)}{\mu_X(\omega_1) + \mu_X(\Omega \setminus \{\omega_1\})} = \frac{0.746}{1} = 0.746$$

And

$$\mathcal{S}(\tau_t(\mu)^* X) = \frac{\mu_X(\omega_2)}{\mu_X(\omega_2) + \mu_X(\Omega \setminus \{\omega_1, \omega_2\})} = \frac{0.179}{0.179 + 0.075} \approx 0.705$$

Incorrigible

We have $t = 0.7$, a distribution $\mu = (0.5, 0.12, 0.05, 0.33)$, and $X = \{\omega_1, \omega_2, \omega_3\}$. Here $\tau_t(\mu) = \{\omega_1, \omega_2, \omega_4\}$. Then $\mu_X \approx (0.746, 0.179, 0.075, 0)$, and tracking fails since $\tau_t(\mu_X) = \{\omega_1\}$ and $\tau_t(\mu)^* X = \{\omega_1, \omega_2\}$. We have the following degrees of stability with respect to μ_X:

$$\mathcal{S}(\tau_t(\mu_X)) = \frac{\mu_X(\omega_1)}{\mu_X(\omega_1) + \mu_X(\Omega \setminus \{\omega_1\})} = \frac{0.746}{1} = 0.746$$

And

$$\mathcal{S}(\tau_t(\mu)^* X) = \frac{\mu_X(\omega_2)}{\mu_X(\omega_2) + \mu_X(\Omega \setminus \{\omega_1, \omega_2\})} = \frac{0.179}{0.179 + 0.075} \approx 0.705$$

So any threshold that makes $\tau_t(\mu)^* X$ stable also makes $\tau_t(\mu_X)$ stable. We cannot force agreement by threshold-raising.

Thresholds

- Lockean vs. stability thresholds
- The threshold-raising method does not always work, and for all sufficiently large probability spaces there exist 'non-correctible' counterexamples[†].

[†]Not negligible. E.g. in a probability simplex, the measures vulnerable to such non-correctible cases form a neighbourhood of positive Lebesgue measure.

Thresholds

- Lockean vs. stability thresholds
- The threshold-raising method does not always work, and for all sufficiently large probability spaces there exist 'non-correctible' counterexamples†.
- Non-correctible cases show an incompatibility between the Lockean and Stability principles and the τ-generated AGM revision, even if one allows thresholds to vary.

†Not negligible. E.g. in a probability simplex, the measures vulnerable to such non-correctible cases form a neighbourhood of positive Lebesgue measure.

Thresholds

- Lockean vs. stability thresholds
- The threshold-raising method does not always work, and for all sufficiently large probability spaces there exist 'non-correctible' counterexamples†.
- Non-correctible cases show an incompatibility between the Lockean and Stability principles and the τ-generated AGM revision, even if one allows thresholds to vary.
- The cautiousness of AGM revision does not mix well with the fine-grained nature of probability measures.

†Not negligible. E.g. in a probability simplex, the measures vulnerable to such non-correctible cases form a neighbourhood of positive Lebesgue measure.

AGM from Bayes via Max Entropy

- What if the probabilistic representation of the agent's credal state is not fully specified?

AGM from Bayes via Max Entropy

- What if the probabilistic representation of the agent's credal state is not fully specified? (e.g., information loss, incomplete description)

AGM from Bayes via Max Entropy

- What if the probabilistic representation of the agent's credal state is not fully specified? (e.g., information loss, incomplete description)
- Suppose the agent only has a qualitative description of her credal state, but is strictly committed to Bayesian conditioning as an update method.

AGM from Bayes via Max Entropy

- What if the probabilistic representation of the agent's credal state is not fully specified? (e.g., information loss, incomplete description)
- Suppose the agent only has a qualitative description of her credal state, but is strictly committed to Bayesian conditioning as an update method.
- Would like to: use τ to find a measure representing the credal state

AGM from Bayes via Max Entropy

- What if the probabilistic representation of the agent's credal state is not fully specified? (e.g., information loss, incomplete description)
- Suppose the agent only has a qualitative description of her credal state, but is strictly committed to Bayesian conditioning as an update method.
- Would like to: use τ to find a measure representing the credal state; use Bayes

AGM from Bayes via Max Entropy

- What if the probabilistic representation of the agent's credal state is not fully specified? (e.g., information loss, incomplete description)
- Suppose the agent only has a qualitative description of her credal state, but is strictly committed to Bayesian conditioning as an update method.
- Would like to: use τ to find a measure representing the credal state; use Bayes; then τ again

AGM from Bayes via Max Entropy

- What if the probabilistic representation of the agent's credal state is not fully specified? (e.g., information loss, incomplete description)
- Suppose the agent only has a qualitative description of her credal state, but is strictly committed to Bayesian conditioning as an update method.
- Would like to: use τ to find a measure representing the credal state; use Bayes; then τ again. This generates a qualitative revision.

Max Entropy Principle

Entropy: $\mathcal{H}(\mu) = \sum_{\omega \in \Omega} -\mu(\omega) \log \mu(\omega)$

Max Entropy Principle

$$\text{Entropy: } \mathcal{H}(\mu) = \sum_{\omega \in \Omega} -\mu(\omega) \log \mu(\omega)$$

- Entropy as a measure of uncertainty.

Max Entropy Principle

$$\text{Entropy: } \mathcal{H}(\mu) = \sum_{\omega \in \Omega} -\mu(\omega) \log \mu(\omega)$$

- Entropy as a measure of uncertainty.

Maximum Entropy Principle (MEP):

If all that is known to the agent is that a probability distribution lies within some zone $\mathcal{N} \subseteq \Delta_{\mathfrak{A}}$, the agent selects a distribution with maximal entropy among those in $\mathcal{N}$, if such exist.

Max Entropy Principle

Entropy: $\mathcal{H}(\mu) = \sum_{\omega \in \Omega} -\mu(\omega) \log \mu(\omega)$

- Entropy as a measure of uncertainty.

Maximum Entropy Principle (MEP):
If all that is known to the agent is that a probability distribution lies within some zone $\mathcal{N} \subseteq \Delta_{\mathfrak{A}}$, the agent selects a distribution with maximal entropy among those in $\mathcal{N}$, if such exist.

- Max Entropy distribution thought to be least biased representation of the agent's credal state, given the constraints.

Max Entropy Principle

Entropy: $\mathcal{H}(\mu) = \sum_{\omega \in \Omega} -\mu(\omega) \log \mu(\omega)$

- Entropy as a measure of uncertainty.

Maximum Entropy Principle (MEP):
If all that is known to the agent is that a probability distribution lies within some zone $\mathcal{N} \subseteq \Delta_{\mathfrak{A}}$, the agent selects a distribution with maximal entropy among those in $\mathcal{N}$, if such exist.

- Max Entropy distribution thought to be least biased representation of the agent's credal state, given the constraints.

Start with a qualitative representation. Use τ and MEP to find a measure representing the credal state; use Bayes; then τ again.

Max Entropy Principle

Entropy: $\mathcal{H}(\mu) = \sum_{\omega \in \Omega} -\mu(\omega) \log \mu(\omega)$

- Entropy as a measure of uncertainty.

Maximum Entropy Principle (MEP):
If all that is known to the agent is that a probability distribution lies within some zone $\mathcal{N} \subseteq \Delta_{\mathfrak{A}}$, the agent selects a distribution with maximal entropy among those in $\mathcal{N}$, if such exist.

- Max Entropy distribution thought to be least biased representation of the agent's credal state, given the constraints.

Start with a qualitative representation. Use τ and MEP to find a measure representing the credal state; use Bayes; then τ again. The resulting revision is always AGM.

AGM from Bayes via Max Entropy

For finite probability spaces:

Let $K \neq \emptyset$ a proposition in $\mathfrak{A}$, and τ the Leitgeb rule for some fixed $t \in [0.5, 1)$. Then there is a unique maximal entropy distribution $\mu \in \Delta_{\mathfrak{A}}$ such that $\tau(\mu) = K$. Moreover, for any positive probability $X \in \mathfrak{A}$, we have $\tau(\mu_X) = K^*X$, where $*$ is the AGM revision operator generated by $\mathfrak{S}^\tau(\mu)$.

AGM from Bayes via Max Entropy

For finite probability spaces:

Let $K \neq \emptyset$ a proposition in $\mathfrak{A}$, and τ the Leitgeb rule for some fixed $t \in [0.5, 1)$. Then there is a unique maximal entropy distribution $\mu \in \Delta_{\mathfrak{A}}$ such that $\tau(\mu) = K$. Moreover, for any positive probability $X \in \mathfrak{A}$, we have $\tau(\mu_X) = K^*X$, where $*$ is the AGM revision operator generated by $\mathfrak{S}^\tau(\mu)$.

K

AGM from Bayes via Max Entropy

For finite probability spaces:

Let $K \neq \emptyset$ a proposition in $\mathfrak{A}$, and τ the Leitgeb rule for some fixed $t \in [0.5, 1)$. Then there is a unique maximal entropy distribution $\mu \in \Delta_{\mathfrak{A}}$ such that $\tau(\mu) = K$. Moreover, for any positive probability $X \in \mathfrak{A}$, we have $\tau(\mu_X) = K^* X$, where $*$ is the AGM revision operator generated by $\mathfrak{S}^\tau(\mu)$.

$$\begin{tikzcd} \mu \\ K \arrow[u, "\text{Max Entropy in } \tau^{-1}(K)"] \end{tikzcd}$$

AGM from Bayes via Max Entropy

For finite probability spaces:

Let $K \neq \emptyset$ a proposition in $\mathfrak{A}$, and τ the Leitgeb rule for some fixed $t \in [0.5, 1)$. Then there is a unique maximal entropy distribution $\mu \in \Delta_{\mathfrak{A}}$ such that $\tau(\mu) = K$. Moreover, for any positive probability $X \in \mathfrak{A}$, we have $\tau(\mu_X) = K^* X$, where $*$ is the AGM revision operator generated by $\mathfrak{S}^\tau(\mu)$.

$$\begin{tikzcd}
\mu \arrow[r, "|_X", "\text{Bayesian conditioning}"'] & \mu_X \\
K \arrow[u, "\text{Max Entropy in } \tau^{-1}(K)"] &
\end{tikzcd}$$

AGM from Bayes via Max Entropy

For finite probability spaces:

> Let $K \neq \emptyset$ a proposition in $\mathfrak{A}$, and τ the Leitgeb rule for some fixed $t \in [0.5, 1)$. Then there is a unique maximal entropy distribution $\mu \in \Delta_{\mathfrak{A}}$ such that $\tau(\mu) = K$. Moreover, for any positive probability $X \in \mathfrak{A}$, we have $\tau(\mu_X) = K^*X$, where $*$ is the AGM revision operator generated by $\mathfrak{S}^\tau(\mu)$.

$$
\begin{tikzcd}
\mu \arrow[rr, "|_X"', "\text{Bayesian conditioning}"'] & & \mu_X \arrow[d, "\tau\ \text{acceptance}"] \\
K \arrow[u, "\text{Max Entropy in } \tau^{-1}(K)"] & & \tau(\mu_X)
\end{tikzcd}
$$

AGM from Bayes via Max Entropy

For finite probability spaces:

Let $K \neq \emptyset$ a proposition in $\mathfrak{A}$, and τ the Leitgeb rule for some fixed $t \in [0.5, 1)$. Then there is a unique maximal entropy distribution $\mu \in \Delta_{\mathfrak{A}}$ such that $\tau(\mu) = K$. Moreover, for any positive probability $X \in \mathfrak{A}$, we have $\tau(\mu_X) = K^* X$, where $*$ is the AGM revision operator generated by $\mathfrak{S}^\tau(\mu)$.

$$
\begin{tikzcd}
\mu \arrow[rr, "|_X", "\text{Bayesian conditioning}"'] & & \mu_X \arrow[d, "\tau \text{ acceptance}"] \\
K \arrow[u, "\text{Max Entropy in } \tau^{-1}(K)"] \arrow[rr, dashed, "*_X", "\text{AGM}"'] & & K^* X = \tau(\mu_X)
\end{tikzcd}
$$

AGM from Bayes via Max Entropy

How it works - a sketch:

- Restrict attention to rank-uniform measures in $\tau^{-1}(K)$.

Any $\mu \in \Delta_{\mathfrak{A}}$ is entropy-dominated by some rank-equivalent, rank-uniform probability measure.

AGM from Bayes via Max Entropy

How it works - a sketch:

- Restrict attention to rank-uniform measures in $\tau^{-1}(K)$.

Any $\mu \in \Delta_{\mathfrak{A}}$ is entropy-dominated by some rank-equivalent, rank-uniform probability measure.

- The desired maximal entropy measure in $\tau^{-1}(K)$ is the rank-uniform measure μ with two ranks which assigns the least possible measure to K.
- Finally, the resulting revision $K \mapsto \tau(\mu_X)$ is always AGM because:

If μ is rank-uniform, then for any $X \in \mathfrak{A}$, the revision $\tau(\mu) \mapsto \tau(\mu_X)$ is the AGM revision generated by $\mathfrak{S}^t(\mu)$.

AGM from Bayes via Max Entropy

$$
\begin{tikzcd}[column sep=huge, row sep=huge]
\mu \arrow[rrr, "|_X"] \arrow[dd, bend right, "\tau"'] & & & \mu_X \arrow[d] \arrow[dd, bend left, "\tau"] \\
\mathfrak{S} \arrow[u, "\text{Max Entropy}"'] \arrow[rrr, "\text{Restriction to } X"'] \arrow[d, "\min"'] & & & \mathfrak{S} \restriction X \arrow[d, "\min"] \\
\tau(\mu) = \min \mathfrak{S} \arrow[rrr, dashed, "\text{AGM}"] & & & \tau(\mu_X)
\end{tikzcd}
$$

Figure: Recovering AGM revision from a plausibility ordering.

$\hookrightarrow$ Reduces to a convex optimisation problem, with linear inequality constraints given by the stability requirement.

AGM from Bayes via Max Entropy

- We have seen AGM revision is too coarse-grained to fully track Bayesian conditioning: it cannot deal with retaining too much information about the probability measure.

AGM from Bayes via Max Entropy

- We have seen AGM revision is too coarse-grained to fully track Bayesian conditioning: it cannot deal with retaining too much information about the probability measure.
- But with an incomplete probabilistic description, AGM can emerge from the τ-rule + two probabilistic principles. Slogan: under incomplete information, "AGM = τ-rule + Maximum Entropy + Bayesian Conditioning".

AGM from Bayes via Max Entropy

- We have seen AGM revision is too coarse-grained to fully track Bayesian conditioning: it cannot deal with retaining too much information about the probability measure.
- But with an incomplete probabilistic description, AGM can emerge from the τ-rule + two probabilistic principles. Slogan: under incomplete information, "AGM = τ-rule + Maximum Entropy + Bayesian Conditioning".
- One can prove a similar result if more qualitative information is retained: e.g plausibility orderings.

AGM from Bayes via Max Entropy

- We have seen AGM revision is too coarse-grained to fully track Bayesian conditioning: it cannot deal with retaining too much information about the probability measure.
- But with an incomplete probabilistic description, AGM can emerge from the τ-rule + two probabilistic principles. Slogan: under incomplete information, "AGM = τ-rule + Maximum Entropy + Bayesian Conditioning".
- One can prove a similar result if more qualitative information is retained: e.g plausibility orderings.
- $\rightsquigarrow$ *How much* information must be lost for AGM to emerge from conditioning in this way? Many geometric/information-theoretic questions.

Further questions, applications

- Logics
- Games
- Qualitative probability

Conclusion

- Stability principles offer a nice acceptance rule which avoids the Lottery paradox and is closely related to AGM revision.
- Perfect tracking is impossible for AGM; one can approximate it, but it comes at a cost.
- AGM revision could be seen as a special case of Bayesian reasoning under the constraint of incomplete information.

References

Alchourrón, C., Gärdenfors, P. and Makinson, D., 1985, "On the logic of theory change: partial meet contraction and revision functions", Journal of Symbolic Logic, 50(2): pp.510–530.

Halpern, J., 2003, *Reasoning About Uncertainty*, MIT Press, Cambridge.

Jaynes, E. T, 1957, "Information Theory and Statistical Mechanics", Physical Review Series II 106 (4): pp. 620-663.

Kelly, K. and Lin, H., 2012, "Propositional Reasoning that Tracks Probabilistic Reasoning", Journal of Philosophical Logic, 41(6): pp. 957–981.

References 2

- Kelly, K. and Lin, H., 2012, “A geo-logical solution to the lottery-paradox, with applications to conditional logic”, Synthese, Vol. 186 (2), pp. 531–575.
- Leitgeb, H., 2013, “Reducing Belief Simpliciter to Degrees of Belief”, Annals of Pure and Applied Logic, 164, pp. 1338–1389.
- Leitgeb, H, 2012, “A Joint Theory of Belief and Probability: Comparing the Theories”, talk given at the MCMP Round Table on Acceptance, February 3rd, 2012. Slides at URL=<http://www.mcmp.philosophie.uni-muenchen.de/events/archive/2012/add_act/r_t_accept/round_tabl

References 3

- Paris J.B., 1994, *The uncertain reasoner's companion: a mathematical perspective*, Cambridge Tracts in Theoretical Computer Science 39, Cambridge University Press.
- Roman, S., 1997, *Coding and Information Theory*, Graduate Texts in Mathematics, Springer.