Making decisions with evidential probability and objective Bayesian calibration inductive logics

Francesco De Pretis
(University of Modena and Reggio Emilia, Italy)
with
William Peden (Lingnan University, Hong Kong)
Mantas Radzvilas(University of Konstanz, Germany)

German Science Foundation project SP 279/21-1 (Project no. 420094936)
Research Grants Council of the Hong Kong SAR, China (Project no. SRFS2122-3H01)

March 2024

Outline

The spread of standard Bayesianism

Bayesian models are being used in an increasingly wide range of research areas:

- increasing accessibility of high-performance computing make Bayesian statistical modelling a viable alternative to frequentist statistical modelling;
- Bayesianism is based on a conceptually and intuitively attractive interpretation of probability;
- standard Bayesianism offers a unified theory of epistemology and decision-making;
- perceived superiority of Bayes factors over null hypothesis testing;
- analytically demonstrable desirable long-term convergence properties of Bayesian models (De Finetti 1980, Zaffora Blando 2022);
- computationally demonstrable desirable short- and medium-term decision-theoretic properties (Radzvilas et al. 2021).

SB and ignorance

A number of philosophical arguments suggest that standard Bayesianism (SB) faces challenges when *representing* ignorance (severe uncertainty), since it conflates two distinct epistemological concepts: *equivocation* and *ignorance*:

- *equivocation* (roughly) – an epistemic state where agent's credences are equally balanced among the relevant events;
- *ignorance* (roughly) – an epistemic state where the agent lacks justified credences about the objective probabilities of the relevant events.

SB and the “paradoxes of ignorance”

In some cases, changes in the extent of the credences’ equivocation and ignorance correspond. However, two “paradoxes of ignorance” – the paradox of ideal evidence and spurious precision – characterise cases where this correspondence breaks down.

Example: A Bayesian-rational agent must predict the outcome of a coin toss. The agent has no information about the coin, and so assigns a uniform prior $(0.5 : \text{heads}, 0.5 : \text{tails})$.

SB and the "paradoxes of ignorance"

In some cases, changes in the extent of the credences' equivocation and ignorance correspond. However, two "paradoxes of ignorance" – the paradox of ideal evidence and spurious precision – characterise cases where this correspondence breaks down.

Example: A Bayesian-rational agent must predict the outcome of a coin toss. The agent has no information about the coin, and so assigns a uniform prior $(0.5 : \text{heads}, 0.5 : \text{tails})$.

- The paradox of ideal evidence: after observing 10 mln. coin tosses with 5 mln. heads and 5 mln. tails, agent's beliefs are represented by the same distribution $(0.5 : \text{heads}, 0.5 : \text{tails})$.

SB and the "paradoxes of ignorance"

In some cases, changes in the extent of the credences' equivocation and ignorance correspond. However, two "paradoxes of ignorance" – the paradox of ideal evidence and spurious precision – characterise cases where this correspondence breaks down.
Example: A Bayesian-rational agent must predict the outcome of a coin toss. The agent has no information about the coin, and so assigns a uniform prior $(0.5 : \text{heads}, 0.5 : \text{tails})$.

- The paradox of ideal evidence: after observing 10 mln. coin tosses with 5 mln. heads and 5 mln. tails, agent's beliefs are represented by the same distribution $(0.5 : \text{heads}, 0.5 : \text{tails})$.
- Spurious precision: Even in situations where the agent has no information about the relevant set of events, agent's credences must be represented with a single probability distribution assigning a precise probability to each event in the set.

Imprecise Frequentism

- In response to these issues, many philosophers have suggested that the probabilities that we use to represent your belief states should be "imprecise" – perhaps measured by some interval value rather than a single real number;

Imprecise Frequentism

- In response to these issues, many philosophers have suggested that the probabilities that we use to represent your belief states should be "imprecise" – perhaps measured by some interval value rather than a single real number;
- imprecise Frequentism (IF) extends SB by representing agent's credence in an event's occurrence via sets of probabilities, known as *IF credence sets* and update them accordingly by a process known as *calibration*.;

Imprecise Frequentism

- In response to these issues, many philosophers have suggested that the probabilities that we use to represent your belief states should be "imprecise" – perhaps measured by some interval value rather than a single real number;
- imprecise Frequentism (IF) extends SB by representing agent's credence in an event's occurrence via sets of probabilities, known as *IF credence sets* and update them accordingly by a process known as *calibration*.;
- the proponents of IB consider a number of decision rules that take into account the whole IF credence set when picking the best action.

IF and decisions under ignorance

- IF offers more representational power for situations of ignorance than SB models;

IF and decisions under ignorance

- IF offers more representational power for situations of ignorance than SB models;
- like SB, IF has some desirable indefinitely long-term convergence properties: under reasonable assumptions, the diverse probabilities in the IF credence set will converge to the actual frequency of the event;

IF and decisions under ignorance

- IF offers more representational power for situations of ignorance than SB models;
- like SB, IF has some desirable indefinitely long-term convergence properties: under reasonable assumptions, the diverse probabilities in the IF credence set will converge to the actual frequency of the event;
- indefinitely long-term convergence properties are of limited use for the decision-theoretic performance assessment of IF models;

IF and decisions under ignorance

- IF offers more representational power for situations of ignorance than SB models;
- like SB, IF has some desirable indefinitely long-term convergence properties: under reasonable assumptions, the diverse probabilities in the IF credence set will converge to the actual frequency of the event;
- indefinitely long-term convergence properties are of limited use for the decision-theoretic performance assessment of IF models;
- belief inertia in Bayesian models (Vallinder 2018, Peden 2022) suggests that higher representational power of IF models may come at a cost of decision-theoretic performance;

IF and decisions under ignorance

- IF offers more representational power for situations of ignorance than SB models;
- like SB, IF has some desirable indefinitely long-term convergence properties: under reasonable assumptions, the diverse probabilities in the IF credence set will converge to the actual frequency of the event;
- indefinitely long-term convergence properties are of limited use for the decision-theoretic performance assessment of IF models;
- belief inertia in Bayesian models (Vallinder 2018, Peden 2022) suggests that higher representational power of IF models may come at a cost of decision-theoretic performance;
- our study extends the agent-based computational methodology for performance assessments of statistical methodologies (Kyburg and Teng 1999, Radzvilas et al 2021) to IF decision models and compares IF and SB performance in a classic decision problem.

Basic decision problem

The decision-maker is observing a sequence of coin tosses. At regular intervals, the system draws a random ticket price $\delta \in [0, 1]$. The decision-maker has to choose an action, the payoff of which depends on the outcome of the next coin toss: bet on ω_h (state where heads obtain) for price δ (action h), bet on ω_t (state where tails obtain) for a price $(1 - \delta)$ (action t), or hold and get a state-independent payoff of 0 (action a).

	ω_h	ω_t
h	$(1-\delta)$	$-\delta$
t	$(\delta-1)$	δ
a	0	0

Figure: Payoff matrix

Basic decision problem

The decision-maker is observing a sequence of coin tosses. At regular intervals, the system draws a random ticket price $\delta \in [0, 1]$. The decision-maker has to choose an action, the payoff of which depends on the outcome of the next coin toss: bet on ω_h (state where heads obtain) for price δ (action h), bet on ω_t (state where tails obtain) for a price $(1 - \delta)$ (action t), or hold and get a state-independent payoff of 0 (action a).

	ω_h	ω_t
h	$(1-\delta)$	$-\delta$
t	$(\delta-1)$	δ
a	0	0

Figure: Payoff matrix

- $\Omega := \{\omega_h, \omega_t\}$ is the set of states, typical element ω_i;
- $C := \{h, t, a\}$ is the set of actions, typical element c;
- $\pi : C \times \Omega \to \mathbb{R}$ is the payoff function that assigns a real number $\pi(c, \omega_i)$ to each $(c, \omega_i) \in C \times \Omega$ according to matrix in Figure 1.

SB credence revision model I

$\mathbb{K} := (\Omega, \Theta, S, m, \kappa, p)$:

- Ω is the set of states;
- $\Theta := \left\{ x \in \mathbb{R} : x \in [0,1] \right\}$ is the set of possible coin biases towards ω_h, typical element θ;
- $m \geq 1$ is the number of observations;
- $S := \{\mathcal{S} \cup \tilde{s}\}$, where $\mathcal{S} := \Omega^m$ and $\tilde{s} := \varnothing$ is the set of observation histories; typical element s;
- $\kappa : \mathcal{S} \to \mathbb{Z}_{\geq 0}$ is the ω_h event counting function, such that, for each $s \in \mathcal{S}$, $\kappa(s) := n\Big(\left\{ (s)_j \in s : j = \omega_h \right\} \Big)$, where $n(\cdot)$ is the cardinality of the set;
- $p : S \to \Delta(\Theta)$ is the credence function with prior $p(\tilde{s})$, such that, for each $s \in \mathcal{S}$ and each $\theta \in \Theta$,

$$p(\theta \mid s) := \frac{p(\theta \mid \tilde{s})\, p(\kappa(s), m \mid \theta)}{p(\kappa(s), m)}.$$

SB credence revision model II

- Beta distribution – a continuous probability distribution on $[0, 1]$;
- $\mathrm{B}\,(\mathrm{a}, \mathrm{b})$ is a beta distribution, its shape is defined by a and b;

SB credence revision model II

- Beta distribution – a continuous probability distribution on $[0, 1]$;
- $\mathrm{B}\left(\mathrm{a}, \mathrm{b}\right)$ is a beta distribution, its shape is defined by a and b;
- $\mathrm{B}\left(1, 1\right)$ is a uniform prior on Θ;
- $\mathrm{B}\left(99, 1\right)$ is a *non-inert* prior biased towards ω_h;
- $\mathrm{B}\left(1, 99\right)$ is a *non-inert* prior biased towards ω_t;

SB credence revision model II

- Beta distribution – a continuous probability distribution on $[0,1]$;
- $\mathrm{B}(\mathrm{a},\mathrm{b})$ is a beta distribution, its shape is defined by a and b;
- $\mathrm{B}(1,1)$ is a uniform prior on Θ;
- $\mathrm{B}(99,1)$ is a *non-inert* prior biased towards ω_h;
- $\mathrm{B}(1,99)$ is a *non-inert* prior biased towards ω_t;
- $p(\theta|\mathrm{a},\mathrm{b}) = \dfrac{\theta^{\mathrm{a}-1}(1-\theta)^{\mathrm{b}-1}}{\mathrm{B}(\mathrm{a},\mathrm{b})}$ is the probability of bias θ given prior $\mathrm{B}(\mathrm{a},\mathrm{b})$;
- $p(\theta \mid \kappa(s), m) := \dfrac{\theta^{\kappa(s)+\mathrm{a}-1}(1-\theta)^{m-\kappa(s)+\mathrm{b}-1}}{\mathrm{B}\big(\mathrm{a}+\kappa(s), \mathrm{b}+m-\kappa(s)\big)}$ probability of bias θ given history s.

SB choice model

$\mathbb{B} := (C, S, \pi, \psi^p)$:

- $C := \{h, t, a\}$ is the set of possible actions;
- S is the set of possible histories;
- $\psi^p : S \to \Delta(\Omega)$ is the aggregate belief function, such that, for any history $s \in S$, $\psi^p(s) \in \Delta(\Omega)$ is such that

$$\psi^p(\omega_h \mid s) := \int_0^1 \theta p(\theta \mid \kappa(s), m)\, d\theta;$$

$$\psi^p(\omega_t \mid s) := 1 - \psi^p(\omega_h \mid s).$$

SB choice model

$\mathbb{B} := (C, S, \pi, \psi^p)$:

- $C := \{h, t, a\}$ is the set of possible actions;
- S is the set of possible histories;
- $\psi^p : S \to \Delta(\Omega)$ is the aggregate belief function, such that, for any history $s \in S$, $\psi^p(s) \in \Delta(\Omega)$ is such that

$$\psi^p(\omega_h \mid s) := \int_0^1 \theta p(\theta \mid \kappa(s), m)\, d\theta;$$

$$\psi^p(\omega_t \mid s) := 1 - \psi^p(\omega_h \mid s).$$

- $E[c \mid \psi^p, s] := \psi^p(\omega_h \mid s)\, \pi(c, \omega_h) + \psi^p(\omega_t \mid s)\, \pi(c, \omega_t)$ is the expected utility of action c given history s;

SB choice model

$\mathbb{B} := (C, S, \pi, \psi^p)$:

- $C := \{h, t, a\}$ is the set of possible actions;
- S is the set of possible histories;
- $\psi^p : S \to \Delta(\Omega)$ is the aggregate belief function, such that, for any history $s \in S$, $\psi^p(s) \in \Delta(\Omega)$ is such that

$$\psi^p(\omega_h \mid s) := \int_0^1 \theta p(\theta \mid \kappa(s), m)\, d\theta;$$

$$\psi^p(\omega_t \mid s) := 1 - \psi^p(\omega_h \mid s).$$

- $E[c \mid \psi^p, s] := \psi^p(\omega_h \mid s)\,\pi(c, \omega_h) + \psi^p(\omega_t \mid s)\,\pi(c, \omega_t)$ is the expected utility of action c given history s;

SB optimal actions

$\mathbf{B}_{P,s} := \left\{c \in C : c \in \arg\max_{c' \in C}\left(E[c' \mid \psi^p, s]\right)\right\}$. *SB* player chooses according to uniform probability distribution (UPD) on $\mathbf{B}_{P,s}$ if $n(\mathbf{B}_{P,s}) > 1$.

IF credence revision model I

Calibration's reasoning is represented using a model $M_F := \{\Omega, \mathrm{A}, \kappa, S, \lambda\}$, where Ω is the set of states, A is the set of considered significance levels with a typical element α, κ is the counting function for ω_h, S is the set of possible histories, and $\varphi : \mathrm{A} \times S \to \mathscr{P}([0,1])$ is the function that assigns, to every significance level-history pair $(\alpha, s) \in \mathrm{A} \times S$, a Clopper-Pearson interval $\varphi(\alpha, s) := (\phi_L, \phi_U) \in \mathscr{P}([0,1])$.

IF credence revision model II

The lower bound $\phi_L \in [0,1]$ and the upper bound $\phi_U \geq \phi_L$ of this interval can be represented in terms of beta distribution quantiles:

$$\phi_L = \mathrm{B}\left(\frac{\alpha}{2}; \kappa(s), m - \kappa(s) + 1\right); \quad (1)$$

$$\phi_U = \mathrm{B}\left(1 - \frac{\alpha}{2}; \kappa(s) + 1, m - \kappa(s)\right). \quad (2)$$

For a significance level α and a history of m coin tosses $s \in S$ with $\kappa(s) \geq 0$ "heads", *Calibration* estimates the actual coin bias to be within the Clopper-Pearson interval $\varphi(\alpha, s)$. Hence, they reject any values $\phi \notin \varphi(\alpha, s)$.

IF credence revision model III

Calibration's expectation-based reasoning about actions can be represented with a model $D_F := \{\Omega, \mathrm{A}, S, C, \pi, \mathcal{E}\}$, where $\mathcal{E} : C \times \mathrm{A} \times S \to \mathcal{P}(\mathbb{R})$ is a function that assigns, to every action-significance level-history combination $(c, \alpha, s) \in C \times \mathrm{A} \times S$, an expected payoff vector $\mathcal{E}(c, \alpha, s) := (E[c \mid \phi_l], \ldots, E[c \mid \phi_u])$, where $\phi_l = \min(\varphi(\alpha, s))$, $\phi_u = \max(\varphi(\alpha, s))$ and each expectation $E[c \mid \phi] := \phi\pi(c, \omega_h) + (1 - \phi)\pi(c, \omega_t)$ for every $\phi \in \varphi(\alpha, s)$.

IF credence revision model IV

For many of the decision rules that define specific calibration players, two important terms will be the minimum and maximum expectations of an action given a confidence interval. Given an action-significance level-history combination $(c, \alpha, s) \in C \times \mathrm{A} \times S$, we define the minimum expectation of an action c as

$$E^{min}_{c|\alpha,s} := \min_{E[c|\phi] \in \mathcal{E}(c,\alpha,s)} \left(E\left[c \mid \phi\right]\right), \tag{3}$$

and the maximum expectation as

$$E^{max}_{c|\alpha,s} := \max_{E[c|\phi] \in \mathcal{E}(c,\alpha,s)} \left(E\left[c \mid \phi\right]\right). \tag{4}$$

Note how $E^{min}_{a|\alpha,s} = E^{max}_{a|\alpha,s} = 0$ for each $(\alpha, s) \in \mathrm{A} \times S$, because a has a guaranteed payoff of zero.

IF choice models: Maximin and Dominance

- $E_{c|P,s} := \left\{ E\left[c|\psi^{p}, s\right] \text{ for each } \psi^{p} \in \Psi^{P} \right\}$ is the set of expected payoffs associated with action c given P and s;
- $E^{min}_{c|P,s} := \min_{E\left[c|\psi^{p'}, s\right] \in E_{c|P,s}} \left(E\left[c|\psi^{p}, s\right] \right)$ is the minimum expected payoff given P;
- $E^{max}_{c|P,s} := \max_{E\left[c|\psi^{p'}, s\right] \in E_{c|P,s}} \left(E\left[c|\psi^{p}, s\right] \right)$ is the maximum expected payoff given P;

IF choice models: Maximin and Dominance

- $E_{c|P,s} := \left\{ E\left[c|\psi^p, s\right] \text{ for each } \psi^p \in \Psi^P \right\}$ is the set of expected payoffs associated with action c given P and s;
- $E^{min}_{c|P,s} := \min_{E\left[c|\psi^{p'}, s\right] \in E_{c|P,s}} \left(E\left[c|\psi^p, s\right] \right)$ is the minimum expected payoff given P;
- $E^{max}_{c|P,s} := \max_{E\left[c|\psi^{p'}, s\right] \in E_{c|P,s}} \left(E\left[c|\psi^p, s\right] \right)$ is the maximum expected payoff given P;

Maximin choice

$\mathbf{M}_{P,s} := \left\{ c \in C : c \in \arg\max_{c' \in C} \left(E^{min}_{c|P,s} \right) \right\}$. UPD on $\mathbf{M}_{P,s}$ if $n\left(\mathbf{M}_{P,s}\right) > 1$.

Dominance choice

$\mathbf{D}_{P,s} := \left\{ c \in C : E^{max}_{c|P,s} > E^{min}_{c'|P,s} \text{ for all } c' \in C \right\}$. UPD on $\mathbf{D}_{P,s}$ if $n\left(\mathbf{D}_{P,s}\right) > 1$.

IF choice models: E-Admissibility, Hurwicz, Regret

E-Admissible choice

$\mathbf{A}_{P,s} := \left\{c \in C : \exists \psi^p \in \Psi^P, c \in \arg\max_{c' \in C} \left(E\left[c' | \psi^p, s\right]\right\}$. UPD on $\mathbf{A}_{P,s}$ if $n\left(\mathbf{A}_{P,s}\right) > 1$.

Hurwicz choice

$\mathbf{H}_{P,s,\alpha} := \left\{c \in C : c \in \arg\max_{c' \in C} \left(\alpha E^{min}_{c'|P,s} + (1-\alpha) E^{max}_{c'|P,s}\right)\right\}$, where $\alpha \in [0,1]$. UPD on $\mathbf{H}_{P,s,\alpha}$ if $n\left(\mathbf{H}_{P,s}\right) > 1$.

Regret choice

$\mathbf{R}_{P,s,\alpha} := \left\{c \in C : c \in \arg\min_{c' \in C} \left(\max_{\psi^p \in \Psi^P} \left(\max_{c' \in C} \left(E\left[c' \mid \psi^p, s\right]\right) - E\left[c \mid \psi^p, s\right]\right)\right)\right\}$. UPD on $\mathbf{R}_{P,s,\alpha}$ if $n\left(\mathbf{R}_{P,s}\right) > 1$.

IF choice models: Opportunity-risk-optimising choice

- $\psi_{P,s}^{min}(\omega_h) := \min_{\psi^{P'} \in \Psi^P} \left(\psi^{P'}(\omega_h \mid s)\right)$ is the lowest belief in ω_h given P and s;
- $\psi_{P,s}^{max}(\omega_h) := \max_{\psi^{P'} \in \Psi^P} \left(\psi^{P'}(\omega_h \mid s)\right)$ is the highest belief in ω_h given P and s;

IF choice models: Opportunity-risk-optimising choice

- $\psi^{min}_{P,s}(\omega_h) := \min_{\psi^{p'} \in \Psi^P} \left(\psi^{p'}(\omega_h \mid s)\right)$ is the lowest belief in ω_h given P and s;
- $\psi^{max}_{P,s}(\omega_h) := \max_{\psi^{p'} \in \Psi^P} \left(\psi^{p'}(\omega_h \mid s)\right)$ is the highest belief in ω_h given P and s;
- $\psi^{avg}_{P,s}(\omega_h) := \frac{\psi^{min}_{P,s}(\omega_h)+\psi^{max}_{P,s}(\omega_h)}{2}$ is the average belief in ω_h given P, s;
- $\psi^{int}_{P,s}(\omega_h) := \frac{\psi^{max}_{P,s}(\omega_h)-\psi^{min}_{P,s}(\omega_h)}{\psi^{max}_{P,\tilde{s}}(\omega_h)-\psi^{min}_{P,\tilde{s}}(\omega_h)}$ is the confidence factor given P and s;

IF choice models: Opportunity-risk-optimising choice

- $\psi^{min}_{P,s}(\omega_h) := \min_{\psi^{p'} \in \Psi^P} \left(\psi^{p'}(\omega_h \mid s)\right)$ is the lowest belief in ω_h given P and s;
- $\psi^{max}_{P,s}(\omega_h) := \max_{\psi^{p'} \in \Psi^P} \left(\psi^{p'}(\omega_h \mid s)\right)$ is the highest belief in ω_h given P and s;
- $\psi^{avg}_{P,s}(\omega_h) := \frac{\psi^{min}_{P,s}(\omega_h) + \psi^{max}_{P,s}(\omega_h)}{2}$ is the average belief in ω_h given P, s;
- $\psi^{int}_{P,s}(\omega_h) := \frac{\psi^{max}_{P,s}(\omega_h) - \psi^{min}_{P,s}(\omega_h)}{\psi^{max}_{P,\tilde{s}}(\omega_h) - \psi^{min}_{P,\tilde{s}}(\omega_h)}$ is the confidence factor given P and s;
- $E\left[c \mid \psi^{avg}_{P,s}(\omega_h)\right] := \pi(c, \omega_h)\,\psi^{avg}_{P,s}(\omega_h) + \pi(c, \omega_t)\left(1 - \psi^{avg}_{P,s}(\omega_h)\right)$ is the average expected payoff of c given P and s;

ORO choice

$$\mathbf{O}_{P,s} := \Big\{ c \in C : c \in \arg\max_{c' \in C} \left(E^{\min}_{c' \mid P,s} \psi^{\text{int}}_{P,s}(\omega_h) + E\left[c' \mid \psi^{\text{avg}}_{P,s}(\omega_h)\right]\left(1 - \psi^{\text{int}}_{P,s}(\omega_h)\right)\right) \Big\}.$$

UPD on $\mathbf{O}_{P,s}$ if $n(\mathbf{O}_{P,s}) > 1$.

IF choice models: Entropy-maximising choice

- $u : \Theta \to \Delta(\Theta)$ is a function assigning a uniform distribution on Θ;
- $\mathcal{E}_s := \left\{ \psi^p \in \Psi^P : \psi^p \in \arg\min_{\psi^{p'} \in \Psi^P} \left| \psi^{p'}(\omega_h \mid s) - \int_0^1 \theta u(\theta)\, d\theta \right| \right\}$ is the set of entropy-maximising belief functions given P and s;

Entropy-maximising choice

$\mathbf{U}_{P,s} := \left\{ c \in C : \exists \psi^p \in \mathcal{E}_s, c \in \arg\max_{c' \in C} \left(\mathrm{E}[c' \mid \psi^p, s] \right) \right\}$. UPD on $\mathbf{U}_{P,s}$ if $n(\mathbf{U}_{P,s}) > 1$.

Games

Game

A game is a sequence of 5 coin tosses. The player bets on the 5th toss;

Big data approach:

- 1000 tests for each player type and each coin bias;
- each test consists of 1000 games, the player makes 1000 bets in each test;
- 1000 randomly generated ticket prices for each test. Each player type faces the same set of ticket prices;
- considered coin biases: 0.1, 0.3, 0.5, 0.7, 0.9;
- for each test with each coin bias, a computer-generated coin toss history, each player type faces the same set of histories;
- each player type bets in 5 million games.

Players

- *Stan* (*SB* player): prior $\mathrm{B}(1,1)$;
- *IF* (Generic *IF* player): *P* with $\underline{p} = \mathrm{B}(1,99)$ and $\overline{p} = \mathrm{B}(99,1)$;
- *IF* player types:
  1. *Dominance*;
  2. *E − Admissibility*;
  3. *Maximin*;
  4. *Pessimist* (Hurwicz with $\alpha = 0.75$);
  5. *Intermediate* (Hurwicz with $\alpha = 0.5$);
  6. *Optimist* (Hurwicz with $\alpha = 0.25$);
  7. *Regret*;
  8. *MaxEnt*;
  9. *ORO*.

Perfect recall within test, no recall between tests: Players retained information about the history of coin tosses from game-to-game within a particular test, but they do not retain any information from one test to another.

Analysis: average frequency per game record

- For each player type and each coin bias, the average frequency per game record over 1000 games and 1000 tests;
- confidence interval for the average frequency per game value at 0.05 significance level computed over 1000 tests, i.e. the mean value plus/minus the related standard error.

Analysis: Wasserstein performance measure I

- $K := (K_1, K_2, K_3, K_4, K_5, K_6, K_7)$ is the set of player choice conditions with a typical element K_i where

Condition	Description
K_1	Directly choosing h.
K_2	Directly choosing t.
K_3	Directly choosing a.
K_4	Randomising between h and t.
K_5	Randomising between h and a.
K_6	Randomising between t and a.
K_7	Randomising between h, t, and a.

Table: Possible player choice conditions. Conditions from K_4 to K_7 never occur for *Stan* and *MaxEnt*, because they never randomise.

Analysis: Wasserstein performance measure I

- $K := (K_1, K_2, K_3, K_4, K_5, K_6, K_7)$ is the set of player choice conditions with a typical element K_i where

Condition	Description
K_1	Directly choosing h.
K_2	Directly choosing t.
K_3	Directly choosing a.
K_4	Randomising between h and t.
K_5	Randomising between h and a.
K_6	Randomising between t and a.
K_7	Randomising between h, t, and a.

Table: Possible player choice conditions. Conditions from K_4 to K_7 never occur for *Stan* and *MaxEnt*, because they never randomise.

- $\Phi :=\{$*Dominance, E-Admissibility, Maximin, Regret, Optimist, Intermediate, Pessimist, ORO, MaxEnt, Stan*$\}$ is the set of player types with a typical element ϕ.

Analysis: Wasserstein performance measure II

- $\sigma : \phi \to \mathscr{P}(K)$ is the choice condition allocation function, such that

$$\sigma(\phi) := \begin{cases} (K_1, K_2, K_3) & \text{if } \phi \in \{Stan, MaxEnt\} \\ K = (K_1, \ldots, K_7) & \text{otherwise.} \end{cases}$$

Analysis: Wasserstein performance measure II

- $\sigma : \phi \to \mathscr{P}(K)$ is the choice condition allocation function, such that

$$\sigma(\phi) := \begin{cases} (K_1, K_2, K_3) & \text{if } \phi \in \{Stan, MaxEnt\} \\ K = (K_1, \ldots, K_7) & \text{otherwise.} \end{cases}$$

- $\Lambda : C \times \Phi \to \mathbb{Y}^{n\left(\sigma(\phi)\right)}$, where $\mathbb{Y} := \{y \in \mathbb{R} : y \in [0,1]\}$, is the scoring function, such that,

$$\Lambda(h, \phi) = \begin{cases} (1,0,0) & \text{if and only if } \phi \in \{MaxEnt,\ Stan\}; \\ \left(1,0,0,\frac{1}{2},0,\frac{1}{2},\frac{1}{3}\right) & \text{otherwise.} \end{cases}$$

$$\Lambda(t, \phi) = \begin{cases} (0,1,0) & \text{if and only if } \phi \in \{MaxEnt,\ Stan\}; \\ \left(0,1,0,\frac{1}{2},\frac{1}{2},0,\frac{1}{3}\right) & \text{otherwise.} \end{cases}$$

$$\Lambda(a, \phi) = \begin{cases} (0,0,1) & \text{if and only if } \phi \in \{MaxEnt,\ Stan\}; \\ \left(0,0,1,0,\frac{1}{2},\frac{1}{2},\frac{1}{3}\right) & \text{otherwise.} \end{cases}$$

Analysis: Wasserstein performance measure III

- $\Gamma := \{250, 1000\}$ is the number of games with a typical element γ;

Analysis: Wasserstein performance measure III

- $\Gamma := \{250, 1000\}$ is the number of games with a typical element γ;
- $W^{\sigma(\phi)}_{c|\theta,\gamma}$ is the data-derived vector representing the mean number of times when each $K_i \in \sigma(\phi)$ was met given c, θ and γ;
- $\mathbf{W}^{\sigma(\phi)} := \left\{W_{h|\theta,\gamma}, W_{t|\theta,\gamma}, W_{a|\theta,\gamma}\right\}$ is the set of performance vectors;

Analysis: Wasserstein performance measure III

- $\Gamma := \{250, 1000\}$ is the number of games with a typical element γ;
- $W^{\sigma(\phi)}_{c|\theta,\gamma}$ is the data-derived vector representing the mean number of times when each $K_i \in \sigma(\phi)$ was met given c, θ and γ;
- $\mathbf{W}^{\sigma(\phi)} := \left\{W_{h|\theta,\gamma}, W_{t|\theta,\gamma}, W_{a|\theta,\gamma}\right\}$ is the set of performance vectors;

$\xi : \Theta \times \Phi \times \Gamma \to [0,1]$ is the Wasserstein performance measure, such that, for any $(\theta, \phi, \gamma) \in \Theta \times \Phi \times \Gamma$,

$$\xi(\theta, \phi, \gamma) := \sqrt{\underbrace{\left[1 - \left(\Lambda(a, \phi) \bullet W^{\sigma(\phi)}_{a|\theta,\gamma}\right)\gamma^{-1}\right]}_{\text{Abstaining factor}}} \cdot$$

$$\sqrt{\left\{\overbrace{1 - \left(2^{-1} - \left|2^{-1} - \theta\right|\right)}^{\substack{\text{Normalizing}\\ \text{and scaling terms}}}\left[\underbrace{\left|\theta - \left(\Lambda(h, \phi) \bullet W^{\sigma(\phi)}_{h|\theta,\gamma}\right)\gamma^{-1}\right|\theta^{-1}}_{\substack{\text{Goodness of fit in predicting}\\ \text{the frequency of heads}}} + \overbrace{\left|1 - \theta - \left(\Lambda(t, \phi) \bullet W^{\sigma(\phi)}_{t|\theta,\gamma}\right)\gamma^{-1}\right|(1 - \theta)^{-1}}^{\substack{\text{Goodness of fit in predicting}\\ \text{the frequency of tails}}}\right]\right\}}$$

Average profit per game I

Overall Leader: *Stan*

IF Leaders: *E-Admissibility, Optimist, Intermediate*.

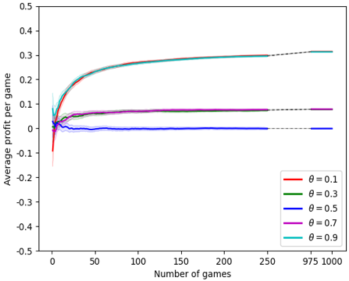

IF Leaders

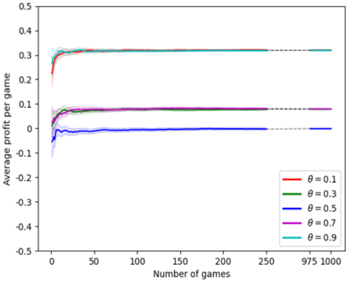

Stan

Average profit per game II

Overall Leader: *Stan*
Worst performer: *Maximin*.

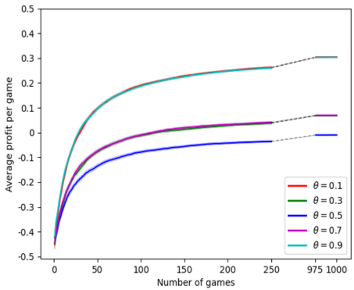

Maximin

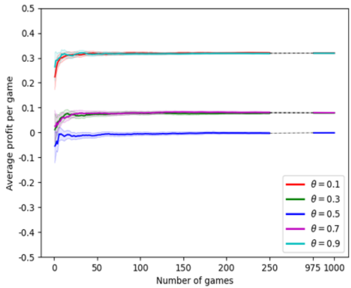

Stan

Wasserstein performance (average frequency)

Coin Bias	Games	Standard Bayesian	IF Leaders	Regret	Dominance	MaxEnt	Pessimist	ORO	Maximin
0.1	250	99.721	95.247	93.436	93.531	90.209	92.483	91.548	87.497
	1000	99.861	98.172	97.483	97.480	96.226	97.123	96.555	95.422
0.3	250	99.955	97.031	94.673	94.665	91.484	92.257	89.766	83.683
	1000	99.950	98.873	97.955	97.949	96.582	97.150	95.486	93.860
0.5	250	99.862	99.814	93.115	93.163	99.712	89.664	86.786	79.488
	1000	99.979	99.977	97.417	97.441	99.972	96.135	94.215	92.322
0.7	250	99.966	97.268	94.916	94.721	91.761	92.479	89.888	83.670
	1000	99.948	98.873	97.967	97.967	96.612	97.163	95.506	93.859
0.9	250	99.764	95.301	93.549	93.508	90.319	92.602	91.671	87.530
	1000	99.939	98.240	97.551	97.554	96.309	97.190	96.630	95.375
Average	**250**	**99.854 ± 0.097**	**96.932 ± 1.636**	**93.938 ± 0.704**	**93.917 ± 0.634**	**92.697 ± 3.490**	**91.897 ± 1.100**	**89.932 ± 1.729**	**84.374 ± 2.926**
Performance	1000	99.935 ± 0.039	98.827 ± 0.635	97.674 ± 0.233	97.678 ± 0.227	97.140 ± 1.395	96.952 ± 0.401	95.678 ± 0.864	94.168 ± 1.128

Table: Performance measurements for 250 games and 1000 games. The players are ordered by their average performances over 250 games (bold).

Conclusion

- The best performing IF players (IF Leaders) are *E-Admissibility, Optimist, Intermediate*.
- No considered IF model can match the performance of *Stan*. At best, IF Leaders can match the performance of *Stan*, but never exceed it. In most cases, *Stan*'s performance is better by a clear margin.
- The divergence in the initial IF credence set P slowed down their convergence to the true coin bias. This sluggishness was a problem even though the chosen set P is far from the most divergent possible in IF epistemology.
- This problem could be mitigated with a less divergent set P. Yet that would make P less adequate from the perspective of modelling ignorance.
- The putative epistemological advantages of IF come at a pragmatic cost, at least in this classic decision problem.

References

- De Finetti, B. (1980). Foresight: Its Logical Laws, Its Subjective Sources. In H. E. Kyburg Jr. & H. Smokler (Eds.), *Studies in subjective probability*. Krieger Publishing Company.
- Ellsberg, D. (1961). Risk, Ambiguity, and the Savage Axioms. *The Quarterly Journal of Economics*, 75(4): 643-669;
- Kyburg, H. E. and Teng, C. M. (1999). Choosing among interpretations of probability. In Laskey, K. B. and Prade, H., editors, *Proceedings of the Fifteenth Conference on Uncertainty in Artificial Intelligence*, UAI'99, page 359-365, San Francisco CA, USA. Morgan Kaufmann Publishers Inc;
- Peden, W. (2022). Evidentialism, Inertia, and Imprecise Probability. *The British Journal for the Philosophy of Science*, pages 1-23
- Radzvilas, M., Peden, W., and De Pretis, F. (2021). A Battle in the Statistics Wars: a simulation-based com- parison of Bayesian, Frequentist and Williamsonian methodologies. *Synthese*, 199(5-6): 13689-13748.
- Vallinder, A. (2018). Imprecise Bayesianism and Global Belief Inertia. *The British Journal for the Philosophy of Science*, 69(4): 1205-1230.
- Zaffora Blando, F. (2022). Bayesian Merging of Opinions and Algorithmic Randomness. *The British Journal for the Philosophy of Science*, pages 1-32.