# Tight Confidence Sequences From Universal Portfolio

## Based on Orabona and Jun (2024), IEEE Transactions on Information Theory

Xiaoqi (Shirley) Liu

E-Values Reading Group
Department of Statistics, University of Oxford

10 March 2026

# Problem setup

# Problem setup

We observe bounded data $X_1, X_2, \ldots$ with

$$0 \leq X_t \leq 1, \qquad \mathbb{E}[X_t \mid X_{1:t-1}] = \mu \quad \text{a.s.}$$

# Problem setup

We observe bounded data $X_1, X_2, \ldots$ with

$$0 \leq X_t \leq 1, \qquad \mathbb{E}[X_t \mid X_{1:t-1}] = \mu \quad \text{a.s.}$$

**Goal:** a confidence sequence (CS), i.e., a sequence of confidence intervals $I_t = [\ell_t(X_{1:t}), u_t(X_{1:t})]$ such that

$$\mathbb{P}\big(\mu \in I_t,\ \forall t \geq 1\big) \geq 1 - \delta.$$

# Problem setup

We observe bounded data $X_1, X_2, \ldots$ with

$$0 \leq X_t \leq 1, \qquad \mathbb{E}[X_t \mid X_{1:t-1}] = \mu \quad \text{a.s.}$$

**Goal:** a confidence sequence (CS), i.e., a sequence of confidence intervals $I_t = [\ell_t(X_{1:t}), u_t(X_{1:t})]$ such that

$$\mathbb{P}\big(\mu \in I_t,\ \forall t \geq 1\big) \geq 1 - \delta.$$

**Main contribution of the paper:** tight CSs based on regret bounds for betting/universal portfolio algorithms.

# Outline

# Outline

# Testing by betting

# Testing by betting

- Start with $\text{Wealth}_0 = 1$

# Testing by betting

- Start with $\text{Wealth}_0 = 1$
- Bet on continuous coin outcomes: $c_t \in [-1, 1]$ (possibly adversarial)

# Testing by betting

- Start with $\text{Wealth}_0 = 1$
- Bet on continuous coin outcomes: $c_t \in [-1, 1]$ (possibly adversarial)
- Bet a signed fraction $\beta_t \in [-1, 1]$ of current wealth on coin being positive ($\beta_t > 0$) or negative ($\beta_t < 0$). Update wealth

$$\text{Wealth}_t = \text{Wealth}_{t-1}(1 + \beta_t c_t).$$

# Testing by betting

- Start with $\text{Wealth}_0 = 1$
- Bet on continuous coin outcomes: $c_t \in [-1, 1]$ (possibly adversarial)
- Bet a signed fraction $\beta_t \in [-1, 1]$ of current wealth on coin being positive ($\beta_t > 0$) or negative ($\beta_t < 0$). Update wealth

$$\text{Wealth}_t = \text{Wealth}_{t-1}(1 + \beta_t c_t).$$

Under $H_0$ coin is fair, i.e., $\mathbb{E}[c_t \mid c_{1:t-1}] = 0$, and $(\text{Wealth}_i)_{i=1}^t$ is a nonnegative supermartingale, i.e.,

$$\mathbb{E}[\text{Wealth}_t \mid \text{Wealth}_{1:t-1}] = \text{Wealth}_{t-1}.$$

# Testing by betting

- Start with $\text{Wealth}_0 = 1$
- Bet on continuous coin outcomes: $c_t \in [-1, 1]$ (possibly adversarial)
- Bet a signed fraction $\beta_t \in [-1, 1]$ of current wealth on coin being positive ($\beta_t > 0$) or negative ($\beta_t < 0$). Update wealth

$$\text{Wealth}_t = \text{Wealth}_{t-1}(1 + \beta_t c_t).$$

Under $H_0$ coin is fair, i.e., $\mathbb{E}[c_t \mid c_{1:t-1}] = 0$, and $(\text{Wealth}_i)_{i=1}^t$ is a nonnegative supermartingale, i.e.,

$$\mathbb{E}[\text{Wealth}_t \mid \text{Wealth}_{1:t-1}] = \text{Wealth}_{t-1}.$$

Thus by Ville's inequality, we obtain the time-uniform error guarantee:

$$\mathbb{P}\left(\sup_{t \geq 0} \text{Wealth}_t \geq \tfrac{1}{\delta}\right) \leq \delta.$$

# Testing by betting

- Start with $\text{Wealth}_0 = 1$
- Bet on continuous coin outcomes: $c_t \in [-1, 1]$ (possibly adversarial)
- Bet a signed fraction $\beta_t \in [-1, 1]$ of current wealth on coin being positive ($\beta_t > 0$) or negative ($\beta_t < 0$). Update wealth
$$\text{Wealth}_t = \text{Wealth}_{t-1}(1 + \beta_t c_t).$$

Under $H_0$ coin is fair, i.e., $\mathbb{E}[c_t \mid c_{1:t-1}] = 0$, and $(\text{Wealth}_i)_{i=1}^t$ is a nonnegative supermartingale, i.e.,

$$\mathbb{E}[\text{Wealth}_t \mid \text{Wealth}_{1:t-1}] = \text{Wealth}_{t-1}.$$

Thus by Ville's inequality, we obtain the time-uniform error guarantee:

$$\mathbb{P}\left(\sup_{t \geq 0} \text{Wealth}_t \geq \tfrac{1}{\delta}\right) \leq \delta.$$

Testing by betting can be generalised through e-values & e-processes (Ramdas and Wang, 2025)

# From testing by betting to mean inference

# From testing by betting to mean inference

Consider random variables $X_1, X_2, \dots \in [0, 1]$ such that $\mathbb{E}[X_t \mid X_{1:t-1}] = \mu$ a.s..

# From testing by betting to mean inference

Consider random variables $X_1, X_2, \dots \in [0,1]$ such that $\mathbb{E}[X_t \mid X_{1:t-1}] = \mu$ a.s..

For each candidate mean $m \in [0,1]$, define the continuous coin

$$c_t(m) := X_t - m \in [-m, 1-m] \subseteq [-1, 1].$$

# From testing by betting to mean inference

Consider random variables $X_1, X_2, \cdots \in [0,1]$ such that $\mathbb{E}[X_t \mid X_{1:t-1}] = \mu$ a.s..

For each candidate mean $m \in [0,1]$, define the continuous coin

$$c_t(m) := X_t - m \in [-m, 1-m] \subseteq [-1, 1].$$

Under $H_0$, $\mu = m$, i.e., $\mathbb{E}[c_t(m) \mid X_{1:t-1}] = 0$.

# From testing by betting to mean inference

Consider random variables $X_1, X_2, \dots \in [0,1]$ such that $\mathbb{E}[X_t \mid X_{1:t-1}] = \mu$ a.s..

For each candidate mean $m \in [0,1]$, define the continuous coin

$$c_t(m) := X_t - m \in [-m, 1-m] \subseteq [-1,1].$$

Under $H_0$, $\mu = m$, i.e., $\mathbb{E}[c_t(m) \mid X_{1:t-1}] = 0$.

Build a wealth process $W_t(m)$ for each $m$, where by Ville,

$$\mathbb{P}\left(\exists t \geq 0 : \text{Wealth}_t(\mu) \geq \frac{1}{\delta}\right) \leq \delta.$$

# From testing by betting to mean inference

Consider random variables $X_1, X_2, \dots \in [0,1]$ such that $\mathbb{E}[X_t \mid X_{1:t-1}] = \mu$ a.s..

For each candidate mean $m \in [0,1]$, define the continuous coin

$$c_t(m) := X_t - m \in [-m, 1-m] \subseteq [-1, 1].$$

Under $H_0$, $\mu = m$, i.e., $\mathbb{E}[c_t(m) \mid X_{1:t-1}] = 0$.

Build a wealth process $W_t(m)$ for each $m$, where by Ville,

$$\mathbb{P}\left(\exists t \geq 0 : \mathrm{Wealth}_t(\mu) \geq \frac{1}{\delta}\right) \leq \delta.$$

Define the CS as the set of $m \in [0,1]$ that haven't been rejected yet:

$$I_t := \left\{ m \in [0,1] : \sup_{0 \leq s \leq t} \mathrm{Wealth}_s(m) < \frac{1}{\delta} \right\}.$$

# From testing by betting to mean inference

Consider random variables $X_1, X_2, \dots \in [0,1]$ such that $\mathbb{E}[X_t \mid X_{1:t-1}] = \mu$ a.s..

For each candidate mean $m \in [0,1]$, define the continuous coin

$$c_t(m) := X_t - m \in [-m, 1-m] \subseteq [-1,1].$$

Under $H_0$, $\mu = m$, i.e., $\mathbb{E}[c_t(m) \mid X_{1:t-1}] = 0$.

Build a wealth process $W_t(m)$ for each $m$, where by Ville,

$$\mathbb{P}\left(\exists t \geq 0 : \text{Wealth}_t(\mu) \geq \tfrac{1}{\delta}\right) \leq \delta.$$

Define the CS as the set of $m \in [0,1]$ that haven't been rejected yet:

$$I_t := \left\{ m \in [0,1] : \sup_{0 \leq s \leq t} \text{Wealth}_s(m) < \tfrac{1}{\delta} \right\}.$$

$\Rightarrow$ **Which betting algorithm should we use?**

# Outline

# Warm-up: the Krichevsky–Trofimov (KT) betting algorithm

# Warm-up: the Krichevsky–Trofimov (KT) betting algorithm

A classic choice is the KT bettor that bets

$$\beta_t = \frac{\sum_{i=1}^{t-1} c_i}{t}.$$

# Warm-up: the Krichevsky–Trofimov (KT) betting algorithm

A classic choice is the KT bettor that bets

$$\beta_t = \frac{\sum_{i=1}^{t-1} c_i}{t}.$$

Worst-case guarantee (Orabona and Pal, 2016): for *any* sequence of $c_1, c_2, \dots \in [-1, 1]$, we have

$$\ln \mathrm{Wealth}_t \ \geq \frac{(\sum_{i=1}^{t} c_i)^2}{2t} \ - \ \frac{1}{2} \ln(4t).$$

# Warm-up: the Krichevsky–Trofimov (KT) betting algorithm

A classic choice is the KT bettor that bets

$$\beta_t = \frac{\sum_{i=1}^{t-1} c_i}{t}.$$

Worst-case guarantee (Orabona and Pal, 2016): for *any* sequence of $c_1, c_2, \dots \in [-1, 1]$, we have

$$\ln \mathrm{Wealth}_t \ \geq \ \frac{(\sum_{i=1}^{t} c_i)^2}{2t} \ - \ \tfrac{1}{2} \ln(4t).$$

Plug in $c_t = X_t - \mu$ and combine with Ville:

$$\mathbb{P}\left( \exists t \geq 0 : \ \frac{(\sum_{i=1}^{t} (X_i - \mu))^2}{2t} \geq \ln(\tfrac{1}{\delta}) + \tfrac{1}{2} \ln(4t) \right) \leq \delta.$$

# Warm-up: the Krichevsky–Trofimov (KT) betting algorithm

A classic choice is the KT bettor that bets

$$\beta_t = \frac{\sum_{i=1}^{t-1} c_i}{t}.$$

Worst-case guarantee (Orabona and Pal, 2016): for *any* sequence of $c_1, c_2, \cdots \in [-1, 1]$, we have

$$\ln \text{Wealth}_t \ \geq \frac{(\sum_{i=1}^{t} c_i)^2}{2t} \ - \ \tfrac{1}{2} \ln(4t).$$

Plug in $c_t = X_t - \mu$ and combine with Ville:

$$\mathbb{P}\left(\exists t \geq 0 : \ \frac{(\sum_{i=1}^{t}(X_i - \mu))^2}{2t} \geq \ln(\tfrac{1}{\delta}) + \tfrac{1}{2}\ln(4t)\right) \leq \delta.$$

Inverting the inequality gives

$$\mathbb{P}\left(\exists t \geq 0 : \left|\mu - \frac{1}{t}\sum_{i=1}^{t} X_i\right| \geq \sqrt{\frac{2\ln\frac{2\sqrt{t}}{\delta}}{t}}\right) \leq \delta.$$

# Why KT is not the end of the story

# Why KT is not the end of the story

Issues with KT when we have $c_t(m) = X_t - m \in [-m, 1-m] \subseteq [-1, 1]$:

# Why KT is not the end of the story

Issues with KT when we have $c_t(m) = X_t - m \in [-m, 1-m] \subseteq [-1, 1]$:

- KT gives a time-uniform Hoeffding-type bound. It only uses boundedness, ignoring variance.

# Why KT is not the end of the story

Issues with KT when we have $c_t(m) = X_t - m \in [-m, 1-m] \subseteq [-1, 1]$:

- KT gives a time-uniform Hoeffding-type bound. It only uses boundedness, ignoring variance.
- KT bets lie in the small interval $\beta_t = \frac{1}{t}\sum_{i=1}^{t-1} c_i \in [-1, 1]$, whereas the admissible betting range is $\beta_t \in \left[-\frac{1}{1-m}, \frac{1}{m}\right]$ due to $1 + \beta_t c_t(m) \geq 0$ for all $c_t(m) \in [-m, 1-m]$.

# Why KT is not the end of the story

Issues with KT when we have $c_t(m) = X_t - m \in [-m, 1-m] \subseteq [-1, 1]$:

- KT gives a time-uniform Hoeffding-type bound. It only uses boundedness, ignoring variance.
- KT bets lie in the small interval $\beta_t = \frac{1}{t}\sum_{i=1}^{t-1} c_i \in [-1, 1]$, whereas the admissible betting range is $\beta_t \in \left[-\frac{1}{1-m}, \frac{1}{m}\right]$ due to $1 + \beta_t c_t(m) \geq 0$ for all $c_t(m) \in [-m, 1-m]$.

  $\Rightarrow$ KT under-utilises the allowable range near the boundaries $m = 1$ or $1$.

# Why KT is not the end of the story

Issues with KT when we have $c_t(m) = X_t - m \in [-m, 1-m] \subseteq [-1, 1]$:

- KT gives a time-uniform Hoeffding-type bound. It only uses boundedness, ignoring variance.
- KT bets lie in the small interval $\beta_t = \frac{1}{t}\sum_{i=1}^{t-1} c_i \in [-1, 1]$, whereas the admissible betting range is $\beta_t \in \left[-\frac{1}{1-m}, \frac{1}{m}\right]$ due to $1 + \beta_t c_t(m) \geq 0$ for all $c_t(m) \in [-m, 1-m]$.

  $\Rightarrow$ KT under-utilises the allowable range near the boundaries $m = 1$ or $1$.

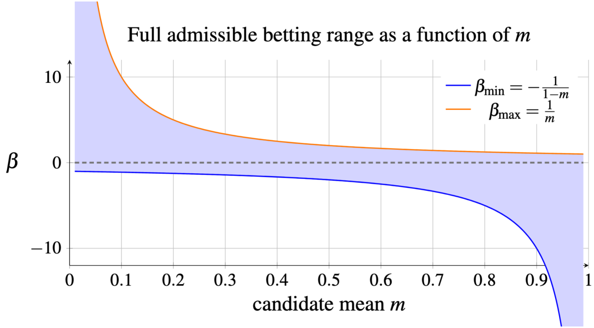

# Why KT is not the end of the story

Issues with KT when we have $c_t(m) = X_t - m \in [-m, 1-m] \subseteq [-1, 1]$:

- KT gives a time-uniform Hoeffding-type bound. It only uses boundedness, ignoring variance.
- KT bets lie in the small interval $\beta_t = \frac{1}{t}\sum_{i=1}^{t-1} c_i \in [-1, 1]$, whereas the admissible betting range is $\beta_t \in \left[-\frac{1}{1-m}, \frac{1}{m}\right]$ due to $1 + \beta_t c_t(m) \geq 0$ for all $c_t(m) \in [-m, 1-m]$.

  $\Rightarrow$ KT under-utilises the allowable range near the boundaries $m = 1$ or $1$.

$\Rightarrow$ **Can we do better than KT?**

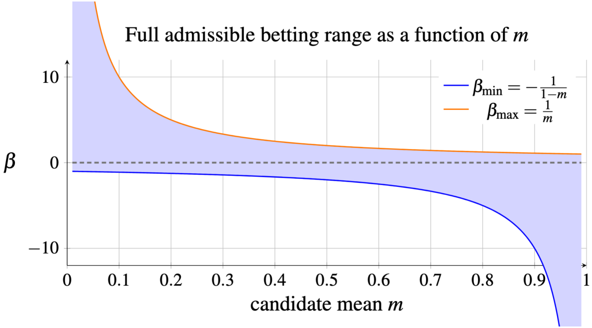

# Outline

# Outline

# A reformulation using online learning tools

# A reformulation using online learning tools

To build CS, for each $t \geq 0$, collect $m \in [0, 1]$ such that

# A reformulation using online learning tools

To build CS, for each $t \geq 0$, collect $m \in [0, 1]$ such that

$$\ln \mathrm{Wealth}_t(\beta, m) < \ln \frac{1}{\delta} .$$

# A reformulation using online learning tools

To build CS, for each $t \geq 0$, collect $m \in [0,1]$ such that

$$\ln \mathrm{Wealth}_t(\beta_t^*, m) - \underbrace{\Big( \ln \mathrm{Wealth}_t(\beta_t^*, m) - \ln \mathrm{Wealth}_t(\beta, m) \Big)}_{=:\mathrm{Regret}_t(m) \text{ (associated with log-wealth)}} < \ln \tfrac{1}{\delta} \,,$$

where $\beta_t^*$ is best fixed bet in hindsight:

$$\beta_t^\star(m) \in \operatorname*{arg\,max}_{\beta \in \left[-\frac{1}{1-m}, \frac{1}{m}\right]} \ln \mathrm{Wealth}_t(\beta, m), \quad \ln \mathrm{Wealth}_t(\beta, m) = \sum_{i=1}^{t} \ln\big(1 + \beta(X_i - m)\big).$$

# A reformulation using online learning tools

To build CS, for each $t \geq 0$, collect $m \in [0,1]$ such that

$$\ln \text{Wealth}_t(\beta_t^*, m) - \underbrace{\Big(\ln \text{Wealth}_t(\beta_t^*, m) - \ln \text{Wealth}_t(\beta, m)\Big)}_{=:\text{Regret}_t(m) \text{ (associated with log-wealth)}} < \ln \tfrac{1}{\delta} \,,$$

where $\beta_t^*$ is best fixed bet in hindsight:

$$\beta_t^\star(m) \in \operatorname*{arg\,max}_{\beta \in \left[-\frac{1}{1-m}, \frac{1}{m}\right]} \ln \text{Wealth}_t(\beta, m), \quad \ln \text{Wealth}_t(\beta, m) = \sum_{i=1}^{t} \ln\big(1 + \beta(X_i - m)\big).$$

The paper proves inequalities of the schematic form

$$\mathbb{P}\Big(\ln \text{Wealth}_t(\beta_t^*, \mu) \;\leq\; \text{Regret}_t(\mu) + \ln \tfrac{1}{\delta},\ \forall t \geq 0\Big) \geq 1 - \delta.$$

# A reformulation using online learning tools

To build CS, for each $t \geq 0$, collect $m \in [0,1]$ such that

$$\ln \mathrm{Wealth}_t(\beta_t^*, m) - \underbrace{\Big( \ln \mathrm{Wealth}_t(\beta_t^*, m) - \ln \mathrm{Wealth}_t(\beta, m) \Big)}_{=:\mathrm{Regret}_t(m) \text{ (associated with log-wealth)}} < \ln \tfrac{1}{\delta} \,,$$

where $\beta_t^*$ is best fixed bet in hindsight:

$$\beta_t^\star(m) \in \operatorname*{arg\,max}_{\beta \in \left[-\frac{1}{1-m}, \frac{1}{m}\right]} \ln \mathrm{Wealth}_t(\beta, m), \quad \ln \mathrm{Wealth}_t(\beta, m) = \sum_{i=1}^{t} \ln\big(1 + \beta(X_i - m)\big).$$

The paper proves inequalities of the schematic form

$$\mathbb{P}\Big( \ln \mathrm{Wealth}_t(\beta_t^*, \mu) \;\leq\; \mathrm{Regret}_t(\mu) + \ln \tfrac{1}{\delta},\ \forall t \geq 0 \Big) \geq 1 - \delta.$$

**Issue:** near-optimal betting algorithms are often computationally infeasible.

# A reformulation using online learning tools

To build CS, for each $t \geq 0$, collect $m \in [0,1]$ such that

$$\ln \mathrm{Wealth}_t(\beta_t^*, m) - \underbrace{\Big(\ln \mathrm{Wealth}_t(\beta_t^*, m) - \ln \mathrm{Wealth}_t(\beta, m)\Big)}_{=:\mathrm{Regret}_t(m) \text{ (associated with log-wealth)}} < \ln \tfrac{1}{\delta},$$

where $\beta_t^*$ is best fixed bet in hindsight:

$$\beta_t^\star(m) \in \operatorname*{arg\,max}_{\beta \in \left[-\frac{1}{1-m}, \frac{1}{m}\right]} \ln \mathrm{Wealth}_t(\beta, m), \quad \ln \mathrm{Wealth}_t(\beta, m) = \sum_{i=1}^{t} \ln\big(1 + \beta(X_i - m)\big).$$

The paper proves inequalities of the schematic form

$$\mathbb{P}\Big(\ln \mathrm{Wealth}_t(\beta_t^*, \mu) \ \leq \ \mathrm{Regret}_t(\mu) + \ln \tfrac{1}{\delta}, \ \forall t \geq 0\Big) \geq 1 - \delta.$$

**Issue:** near-optimal betting algorithms are often computationally infeasible.
**Remedy:** Instead of regret, plug in a regret upper bound $R_t(m)$, so no need to run the algorithm!

# A reformulation using online learning tools

To build CS, for each $t \geq 0$, collect $m \in [0,1]$ such that

$$\ln \mathrm{Wealth}_t(\beta_t^*, m) - \underbrace{\Big(\ln \mathrm{Wealth}_t(\beta_t^*, m) - \ln \mathrm{Wealth}_t(\beta, m)\Big)}_{=:\mathrm{Regret}_t(m) \text{ (associated with log-wealth)}} < \ln \tfrac{1}{\delta}\,,$$

where $\beta_t^*$ is best fixed bet in hindsight:

$$\beta_t^\star(m) \in \operatorname*{arg\,max}_{\beta \in \left[-\frac{1}{1-m}, \frac{1}{m}\right]} \ln \mathrm{Wealth}_t(\beta, m), \quad \ln \mathrm{Wealth}_t(\beta, m) = \sum_{i=1}^{t} \ln\big(1 + \beta(X_i - m)\big).$$

The paper proves inequalities of the schematic form

$$\mathbb{P}\Big(\ln \mathrm{Wealth}_t(\beta_t^*, \mu) \;\leq\; R_t(\mu) + \ln \tfrac{1}{\delta},\ \forall t \geq 0\Big) \geq 1 - \delta.$$

**Issue:** near-optimal betting algorithms are often computationally infeasible.
**Remedy:** Instead of regret, plug in a regret upper bound $R_t(m)$, so no need to run the algorithm!

**Take-away message:** Any betting algorithm with a guaranteed regret (whether or not computationally feasible) yields a CS.

**Take-away message:** Any betting algorithm with a guaranteed regret (whether or not computationally feasible) yields a CS.

$\rightarrow$ Paper focuses on betting algorithms via a reduction to **online portfolio selection**.

# Outline

# From coin betting to 2-stock online portfolio selection

# From coin betting to 2-stock online portfolio selection

**Lemma 1.** Let $m \in [0,1]$ and $c \in [-m, 1-m]$. Set two market gains $w_1 = 1 + \frac{c}{m}$ and $w_2 = 1 - \frac{c}{1-m}$ (note $w_1, w_2 \geq 0$ so they are valid market gains). Define $\boldsymbol{b} = [b, 1-b]^\top$ to be the play of a 2-stocks portfolio algorithm, where $b \in [0,1]$. Then, by taking

$$\beta = -\frac{1}{1-m} + \left(\frac{1}{1-m} + \frac{1}{m}\right) b \in \left[-\frac{1}{1-m}, \frac{1}{m}\right]$$

as the signed betting fraction, a continuous coin-betting algorithm on $c$ ensures that the gain in the coin-betting problem coincides with the gain in the portfolio selection problem.

# From coin betting to 2-stock online portfolio selection

**Lemma 1.** Let $m \in [0,1]$ and $c \in [-m, 1-m]$. Set two market gains $w_1 = 1 + \frac{c}{m}$ and $w_2 = 1 - \frac{c}{1-m}$ (note $w_1, w_2 \geq 0$ so they are valid market gains). Define $\boldsymbol{b} = [b, 1-b]^\top$ to be the play of a 2-stocks portfolio algorithm, where $b \in [0,1]$. Then, by taking

$$\beta = -\frac{1}{1-m} + \left(\frac{1}{1-m} + \frac{1}{m}\right) b \in \left[-\frac{1}{1-m}, \frac{1}{m}\right]$$

as the signed betting fraction, a continuous coin-betting algorithm on $c$ ensures that the gain in the coin-betting problem coincides with the gain in the portfolio selection problem.

**Proof sketch.**

$$\begin{bmatrix} b \\ 1-b \end{bmatrix} \cdot \begin{bmatrix} w_1 \\ w_2 \end{bmatrix} = b\left(1 + \frac{c}{m}\right) + (1-b)\left(1 - \frac{c}{1-m}\right) = 1 + c\left(\frac{b}{m} - \frac{1-b}{1-m}\right) = 1 + c\beta.$$

# From coin betting to 2-stock online portfolio selection

**Lemma 1.** Let $m \in [0,1]$ and $c \in [-m, 1-m]$. Set two market gains $w_1 = 1 + \frac{c}{m}$ and $w_2 = 1 - \frac{c}{1-m}$ (note $w_1, w_2 \geq 0$ so they are valid market gains). Define $\boldsymbol{b} = [b, 1-b]^\top$ to be the play of a 2-stocks portfolio algorithm, where $b \in [0,1]$. Then, by taking

$$\beta = -\frac{1}{1-m} + \left(\frac{1}{1-m} + \frac{1}{m}\right) b \in \left[-\frac{1}{1-m}, \frac{1}{m}\right]$$

as the signed betting fraction, a continuous coin-betting algorithm on $c$ ensures that the gain in the coin-betting problem coincides with the gain in the portfolio selection problem.

**Interpretation:** Betting on an asymmetric bounded coin is equivalent to choosing a portfolio weight in a 2-stock online universal portfolio.

$\Rightarrow \ln \mathrm{Wealth}^{\mathrm{coin}}(\beta) = \ln \mathrm{Wealth}^{\mathrm{portfolio}}(\boldsymbol{b})$ and $\mathrm{Regret}_t^{\mathrm{coin}}(\beta) = \mathrm{Regret}_t^{\mathrm{portfolio}}(\boldsymbol{b})$

# Universal portfolio algorithms

# Universal portfolio algorithms

A portfolio algorithm is universal if the regret against *any* sequence of market gain vectors $[w_1, w_2]^\top$ is sublinear in $T$ (Cover and Ordentlich, 1996).

# Universal portfolio algorithms

A portfolio algorithm is universal if the regret against *any* sequence of market gain vectors $[w_1, w_2]^\top$ is sublinear in $T$ (Cover and Ordentlich, 1996).

Paper focuses on a mixture-based universal portfolio algorithm called the $F$-weighted portfolio algorithm (Cover, 1991)

# Universal portfolio algorithms

A portfolio algorithm is universal if the regret against *any* sequence of market gain vectors $[w_1, w_2]^\top$ is sublinear in $T$ (Cover and Ordentlich, 1996).

Paper focuses on a mixture-based universal portfolio algorithm called the $F$-weighted portfolio algorithm (Cover, 1991) that outputs at each step

$$\boldsymbol{b}_t = \frac{\int_B \boldsymbol{b}\mathrm{Wealth}_{t-1}(\boldsymbol{b})\mathrm{d}F(\boldsymbol{b})}{\int_B \mathrm{Wealth}_{t-1}(\boldsymbol{b})\mathrm{d}F(\boldsymbol{b})}, \quad B := \left\{[b, 1-b]^\top \in \mathbb{R}^2 : b \in [0,1]\right\},$$

# Universal portfolio algorithms

A portfolio algorithm is universal if the regret against *any* sequence of market gain vectors $[w_1, w_2]^\top$ is sublinear in $T$ (Cover and Ordentlich, 1996).

Paper focuses on a mixture-based universal portfolio algorithm called the $F$-weighted portfolio algorithm (Cover, 1991) that outputs at each step

$$\boldsymbol{b}_t = \frac{\int_B \boldsymbol{b} \mathrm{Wealth}_{t-1}(\boldsymbol{b}) \mathrm{d}F(\boldsymbol{b})}{\int_B \mathrm{Wealth}_{t-1}(\boldsymbol{b}) \mathrm{d}F(\boldsymbol{b})}, \quad B := \left\{ [b, 1-b]^\top \in \mathbb{R}^2 : b \in [0,1] \right\},$$

with regret guarantee

$$\mathrm{Regret}_T = \max_{\boldsymbol{b} \in B} \ln \mathrm{Wealth}_T(\boldsymbol{b}) - \ln \underbrace{\int_B \mathrm{Wealth}_t(\boldsymbol{b}) \mathrm{d}F(\boldsymbol{b})}_{\substack{\text{wealth of the } F\text{-weighted} \\ \text{portfolio algorithm}}} = o(T).$$

# Outline

# PRECiSE-CO96:

## Portfolio REgret for Confidence SEquences Using Cover and Ordentlich (1996)

# PRECiSE-CO96:

## Portfolio REgret for Confidence SEquences Using Cover and Ordentlich (1996)

For the $F$-weighted universal portfolio algorithm with $F = \text{Beta}(\frac{1}{2}, \frac{1}{2})$, its regret is upper bounded by $\text{Regret}_t \leq \ln \frac{\sqrt{\pi}\Gamma(t+1)}{\Gamma(t+\frac{1}{2})} \leq \frac{1}{2}\ln(4t)$.

# PRECiSE-CO96:

## Portfolio REgret for Confidence SEquences Using Cover and Ordentlich (1996)

For the $F$-weighted universal portfolio algorithm with $F = \text{Beta}(\frac{1}{2}, \frac{1}{2})$, its regret is upper bounded by $\text{Regret}_t \leq \ln \frac{\sqrt{\pi}\Gamma(t+1)}{\Gamma(t+\frac{1}{2})} \leq \frac{1}{2}\ln(4t)$. This yields the following CS:

**Theorem 1.** Let $\delta \in (0,1)$. Let $X_1, X_2, \dots \in [0,1]$ such that $\mathbb{E}[X_i \mid X_{1:i-1}] = \mu$ a.s.. Then, PRECiSE-CO96 yields

$$\mathbb{P}\left(\exists t \geq 1 : \max_{\beta \in \left[-\frac{1}{1-\mu}, \frac{1}{\mu}\right]} \sum_{i=1}^{t} \ln\left(1 + \beta(X_i - \mu)\right) \geq \ln(\tfrac{1}{\delta}) + \tfrac{1}{2}\ln(4t)\right) \leq \delta.$$

# PRECiSE-CO96:

## Portfolio REgret for Confidence SEquences Using Cover and Ordentlich (1996)

For the $F$-weighted universal portfolio algorithm with $F = \text{Beta}(\frac{1}{2}, \frac{1}{2})$, its regret is upper bounded by $\text{Regret}_t \le \ln \frac{\sqrt{\pi}\Gamma(t+1)}{\Gamma(t+\frac{1}{2})} \le \frac{1}{2}\ln(4t)$. This yields the following CS:

**Theorem 1.** Let $\delta \in (0,1)$. Let $X_1, X_2, \dots \in [0,1]$ such that $\mathbb{E}[X_i \mid X_{1:i-1}] = \mu$ a.s.. Then, PRECiSE-CO96 yields

$$\mathbb{P}\left(\exists t \ge 1 : \max_{\beta \in \left[-\frac{1}{1-\mu}, \frac{1}{\mu}\right]} \sum_{i=1}^{t} \ln\left(1 + \beta(X_i - \mu)\right) \ge \ln(\tfrac{1}{\delta}) + \tfrac{1}{2}\ln(4t)\right) \le \delta.$$

Compare with KT guarantee: $\mathbb{P}\left(\exists t \ge 1 : \frac{(\sum_{i=1}^{t}(X_i - \mu))^2}{2t} \ge \ln(\frac{1}{\delta}) + \frac{1}{2}\ln(4t)\right) \le \delta.$

- ✗ No closed-form concentration, need numerical inversion through bisection
- ✓ Never-vacuous for any sample size

# PRECiSE-CO96: why never-vacuous?

# PRECiSE-CO96: why never-vacuous?

Prior placed on portfolio weight $b \in [0, 1]$ is Beta$(\frac{1}{2}, \frac{1}{2})$: $F(b) = \frac{1}{\pi\sqrt{b(1-b)}}$.

This has several advantages:

# PRECiSE-CO96: why never-vacuous?

Prior placed on portfolio weight $b \in [0, 1]$ is Beta$(\frac{1}{2}, \frac{1}{2})$: $F(b) = \frac{1}{\pi\sqrt{b(1-b)}}$.

This has several advantages:

- Maps to full admissible betting range $\beta = -\frac{1}{1-m} + \left(\frac{1}{1-m} + \frac{1}{m}\right) b \in \left[-\frac{1}{1-m}, \frac{1}{m}\right]$ which depends on $m$

# PRECiSE-CO96: why never-vacuous?

Prior placed on portfolio weight $b \in [0,1]$ is Beta$(\frac{1}{2}, \frac{1}{2})$: $F(b) = \frac{1}{\pi\sqrt{b(1-b)}}$.

This has several advantages:

- Maps to full admissible betting range $\beta = -\frac{1}{1-m} + \left(\frac{1}{1-m} + \frac{1}{m}\right) b \in \left[-\frac{1}{1-m}, \frac{1}{m}\right]$ which depends on $m$
- Allocates sufficient mass near extreme bets $\Rightarrow$ prevents early-time regret explosion $\Rightarrow$ CS non-vacuous for all $t$

# PRECiSE-CO96: why never-vacuous?

Prior placed on portfolio weight $b \in [0,1]$ is Beta$(\frac{1}{2},\frac{1}{2})$: $F(b) = \frac{1}{\pi\sqrt{b(1-b)}}$.

This has several advantages:

- Maps to full admissible betting range $\beta = -\frac{1}{1-m} + \left(\frac{1}{1-m} + \frac{1}{m}\right) b \in \left[-\frac{1}{1-m}, \frac{1}{m}\right]$ which depends on $m$
- Allocates sufficient mass near extreme bets $\Rightarrow$ prevents early-time regret explosion $\Rightarrow$ CS non-vacuous for all $t$
- Deeper connections to universal coding, Jefferys prior

# PRECiSE-CO96 yields time-uniform empirical Bernstein bound

# PRECiSE-CO96 yields time-uniform empirical Bernstein bound

The implicit concentration in Theorem 1 implies an empirical-Bernstein-type CS, in contrast to the Hoeffding-type CS based on KT.

# PRECiSE-CO96 yields time-uniform empirical Bernstein bound

The implicit concentration in Theorem 1 implies an empirical-Bernstein-type CS, in contrast to the Hoeffding-type CS based on KT. In detail:

**Theorem 2.** Under the assumptions of Theorem 1, denote by $\hat{\mu}_i = \frac{1}{i}\sum_{j=1}^{i} X_j$, $V_i = \sum_{j=1}^{i}(X_j - \hat{\mu}_i)^2$. Then,

$$\mathbb{P}\left(\mu \in \left[\max_{1\le i\le t} \hat{\mu}_i - \varepsilon_i,\ \min_{1\le i\le t} \hat{\mu}_i + \varepsilon_i\right],\ \forall t \ge 1\right) \ge 1-\delta,$$

where for large $t$, $\varepsilon_t \approx \frac{1}{t}\sqrt{2V_t(\frac{1}{2}\ln t + \ln\frac{1}{\delta})} + \frac{4}{3t}(\frac{1}{2}\ln t + \ln\frac{1}{\delta})$.

# PRECiSE-CO96 yields time-uniform empirical Bernstein bound

The implicit concentration in Theorem 1 implies an empirical-Bernstein-type CS, in contrast to the Hoeffding-type CS based on KT. In detail:

**Theorem 2.** Under the assumptions of Theorem 1, denote by $\hat{\mu}_i = \frac{1}{i}\sum_{j=1}^{i} X_j$, $V_i = \sum_{j=1}^{i}(X_j - \hat{\mu}_i)^2$. Then,

$$\mathbb{P}\left(\mu \in \left[\max_{1\le i\le t} \hat{\mu}_i - \varepsilon_i,\ \min_{1\le i\le t} \hat{\mu}_i + \varepsilon_i\right],\ \forall t \ge 1\right) \ge 1 - \delta,$$

where for large $t$, $\varepsilon_t \approx \frac{1}{t}\sqrt{2V_t(\frac{1}{2}\ln t + \ln\frac{1}{\delta})} + \frac{4}{3t}(\frac{1}{2}\ln t + \ln\frac{1}{\delta})$.

**Remarks:**

# PRECiSE-CO96 yields time-uniform empirical Bernstein bound

The implicit concentration in Theorem 1 implies an empirical-Bernstein-type CS, in contrast to the Hoeffding-type CS based on KT. In detail:

**Theorem 2.** Under the assumptions of Theorem 1, denote by $\hat{\mu}_i = \frac{1}{i}\sum_{j=1}^{i} X_j$, $V_i = \sum_{j=1}^{i}(X_j - \hat{\mu}_i)^2$. Then,

$$\mathbb{P}\left(\mu \in \left[\max_{1\le i\le t} \hat{\mu}_i - \varepsilon_i,\ \min_{1\le i\le t} \hat{\mu}_i + \varepsilon_i\right],\ \forall t \ge 1\right) \ge 1 - \delta,$$

where for large $t$, $\varepsilon_t \approx \frac{1}{t}\sqrt{2V_t(\frac{1}{2}\ln t + \ln\frac{1}{\delta})} + \frac{4}{3t}(\frac{1}{2}\ln t + \ln\frac{1}{\delta})$.

**Remarks:**

- Note $\frac{1}{t}\sqrt{V_t \ln\frac{1}{\delta}}$ is the empirical Bernstein rate for fixed $t$, and $\frac{1}{t}\sqrt{V_t \ln t}$ is price for time-uniformity

# PRECiSE-CO96 yields time-uniform empirical Bernstein bound

The implicit concentration in Theorem 1 implies an empirical-Bernstein-type CS, in contrast to the Hoeffding-type CS based on KT. In detail:

**Theorem 2.** Under the assumptions of Theorem 1, denote by $\hat{\mu}_i = \frac{1}{i}\sum_{j=1}^{i} X_j$, $V_i = \sum_{j=1}^{i}(X_j - \hat{\mu}_i)^2$. Then,

$$\mathbb{P}\left(\mu \in \left[\max_{1\le i\le t} \hat{\mu}_i - \varepsilon_i,\ \min_{1\le i\le t} \hat{\mu}_i + \varepsilon_i\right],\ \forall t \ge 1\right) \ge 1 - \delta,$$

where for large $t$, $\varepsilon_t \approx \frac{1}{t}\sqrt{2V_t(\frac{1}{2}\ln t + \ln\frac{1}{\delta})} + \frac{4}{3t}(\frac{1}{2}\ln t + \ln\frac{1}{\delta})$.

**Remarks:**

- Note $\frac{1}{t}\sqrt{V_t \ln\frac{1}{\delta}}$ is the empirical Bernstein rate for fixed $t$, and $\frac{1}{t}\sqrt{V_t \ln t}$ is price for time-uniformity
- **Next:** Going from $\frac{1}{t}\sqrt{V_t \ln t}$ (i.e., CLT rate $\times\sqrt{\ln t}$) to $\frac{1}{t}\sqrt{V_t \ln\ln t}$ (i.e., law of the iterated logarithm (LIL) rate, asymptotically optimal)

# Outline

# From PRECiSE-CO96 to -R70: motivation for a different prior

# From PRECiSE-CO96 to -R70: motivation for a different prior

By the first-order condition, we have

$$\beta_t^\star(m) \in \operatorname*{arg\,max}_{\beta \in \left[-\frac{1}{1-m}, \frac{1}{m}\right]} \sum_{i=1}^{t} \ln\big(1 + \beta(X_i - m)\big) \approx \frac{\sum_{i=1}^{t}(X_i - m)}{\sum_{i=1}^{t}(X_i - m)^2}.$$

# From PRECiSE-CO96 to -R70: motivation for a different prior

By the first-order condition, we have

$$\beta_t^\star(m) \in \underset{\beta \in \left[-\frac{1}{1-m}, \frac{1}{m}\right]}{\arg\max} \sum_{i=1}^{t} \ln\big(1 + \beta(X_i - m)\big) \approx \frac{\sum_{i=1}^{t}(X_i - m)}{\sum_{i=1}^{t}(X_i - m)^2}.$$

Under $m = \mu$, the LIL gives $\sum_{i=1}^{t}(X_i - \mu) = \mathcal{O}\big(\sqrt{V_t \ln\ln t}\big)$, which implies

$$\beta_t^\star(\mu) = \mathcal{O}\left(\sqrt{\frac{\ln\ln t}{V_t}}\right) \;\to\; 0.$$

# From PRECiSE-CO96 to -R70: motivation for a different prior

By the first-order condition, we have

$$\beta_t^\star(m) \in \operatorname*{arg\,max}_{\beta \in \left[-\frac{1}{1-m}, \frac{1}{m}\right]} \sum_{i=1}^{t} \ln\big(1 + \beta(X_i - m)\big) \approx \frac{\sum_{i=1}^t (X_i - m)}{\sum_{i=1}^t (X_i - m)^2}.$$

Under $m = \mu$, the LIL gives $\sum_{i=1}^t (X_i - \mu) = \mathcal{O}\big(\sqrt{V_t \ln\ln t}\big)$, which implies

$$\beta_t^\star(\mu) = \mathcal{O}\left(\sqrt{\frac{\ln\ln t}{V_t}}\right) \to 0.$$

**Take-away message:** The optimal bet $\to 0$ as $t \to \infty$, so a prior must maintain mass near 0 at all times; one that places mass away from 0 pays extra regret.

# From PRECiSE-CO96 to -R70: motivation for a different prior

By the first-order condition, we have

$$\beta_t^\star(m) \in \underset{\beta \in \left[-\frac{1}{1-m}, \frac{1}{m}\right]}{\arg\max} \sum_{i=1}^{t} \ln\big(1 + \beta(X_i - m)\big) \approx \frac{\sum_{i=1}^t (X_i - m)}{\sum_{i=1}^t (X_i - m)^2}.$$

Under $m = \mu$, the LIL gives $\sum_{i=1}^t (X_i - \mu) = \mathcal{O}\big(\sqrt{V_t \ln\ln t}\big)$, which implies

$$\beta_t^\star(\mu) = \mathcal{O}\left(\sqrt{\frac{\ln\ln t}{V_t}}\right) \;\to\; 0.$$

**Take-away message:** The optimal bet $\to 0$ as $t \to \infty$, so a prior must maintain mass near 0 at all times; one that places mass away from 0 pays extra regret.

$\Rightarrow$ Instead of $F(b) = \text{Beta}(\frac{1}{2}, \frac{1}{2})$, PRECiSE-R70 uses the following $F(\beta)$ from Robbins (1970)

$$F(\beta) \asymp \frac{1}{|\beta| \ln\left(\frac{1}{|\beta|}\right) \left(\ln\ln\left(\frac{1}{|\beta|}\right)\right)^2}.$$

Note $F(\beta)$ is a proper prior with $F(\beta) \to \infty$ as $\beta \to 0$.

# PRECiSE-R70:

## Portfolio REgret for Confidence SEquences Using Robbins (1970)

# PRECiSE-R70:

## Portfolio REgret for Confidence SEquences Using Robbins (1970)

**Theorem 3.** Let $\delta \in (0,1)$. Let $X_1, X_2, \dots \in [0,1]$ such that $\mathbb{E}[X_i \mid X_{1:i-1}] = \mu$ a.s.. Denote by $\hat{\mu}_t = \frac{1}{t}\sum_{i=1}^t X_i$, $V_t = \sum_{i=1}^t (X_i - \hat{\mu}_t)^2$. Then, PRECiSE-R70 yields a CS such that

$$\mathbb{P}\big(\mu \in [\ell_t, u_t],\ \forall t \geq 1\big) \geq 1 - \delta\,,$$

where the corresponding confidence width satisfies

$$\max\{u_t - \hat{\mu}_t, \hat{\mu}_t - \ell_t\} \;\lesssim\; \frac{1}{t}\sqrt{2 V_t \big(\ln \ln t + \ln \tfrac{1}{\delta}\big)}.$$

# PRECiSE-R70:

## Portfolio REgret for Confidence SEquences Using Robbins (1970)

**Theorem 3.** Let $\delta \in (0,1)$. Let $X_1, X_2, \dots \in [0,1]$ such that $\mathbb{E}[X_i \mid X_{1:i-1}] = \mu$ a.s.. Denote by $\hat{\mu}_t = \frac{1}{t}\sum_{i=1}^t X_i$, $V_t = \sum_{i=1}^t (X_i - \hat{\mu}_t)^2$. Then, PRECiSE-R70 yields a CS such that

$$\mathbb{P}\big(\mu \in [\ell_t, u_t],\ \forall t \geq 1\big) \geq 1 - \delta\,,$$

where the corresponding confidence width satisfies

$$\max\{u_t - \hat{\mu}_t, \hat{\mu}_t - \ell_t\} \lesssim \frac{1}{t}\sqrt{2V_t\big(\ln\ln t + \ln\tfrac{1}{\delta}\big)}.$$

Compare with

- Empirical Bernstein: $|\hat{\mu}_t - \mu| \lesssim \frac{1}{t}\sqrt{2V_t \ln\frac{1}{\delta}}$ for fixed $t$
- PRECiSE-CO96: $|\hat{\mu}_t - \mu| \lesssim \frac{1}{t}\sqrt{2V_t(\ln t + \ln\frac{1}{\delta})}$ uniformly in $t$

# Outline

# Simulations

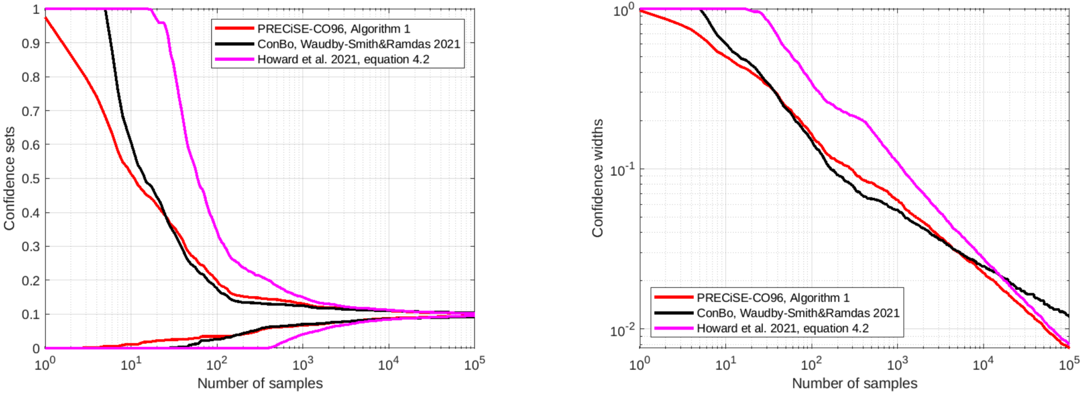


Figure: Time-uniform CSs (left) and their widths (right) for a sequence of i.i.d. Bernoulli(0.1) r.v.s.

# Simulations

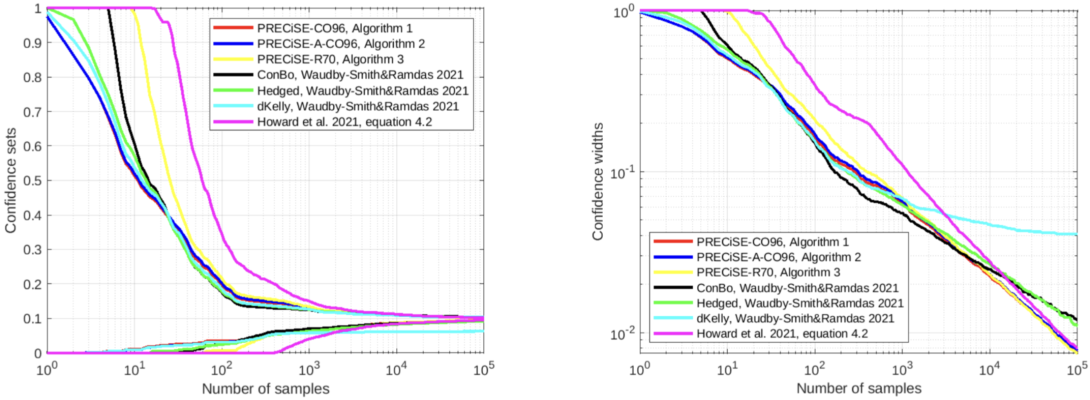


Figure: Time-uniform CSs (left) and their widths (right) for a sequence of i.i.d. Bernoulli(0.1) r.v.s.

# Simulations

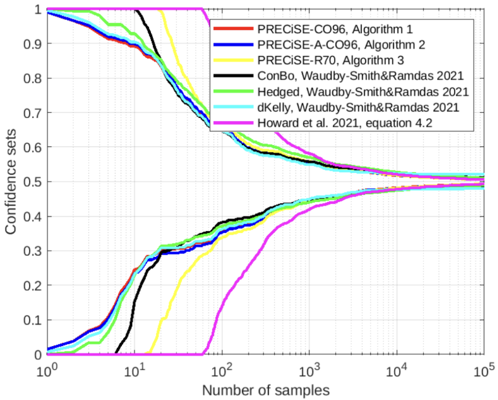


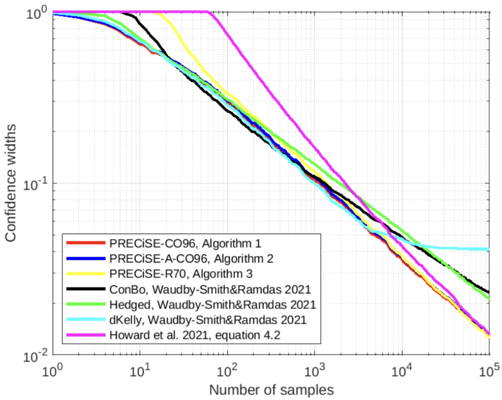


Figure: Time-uniform CSs (left) and their widths (right) for a sequence of i.i.d. Bernoulli(0.5) r.v.s.

# Simulations

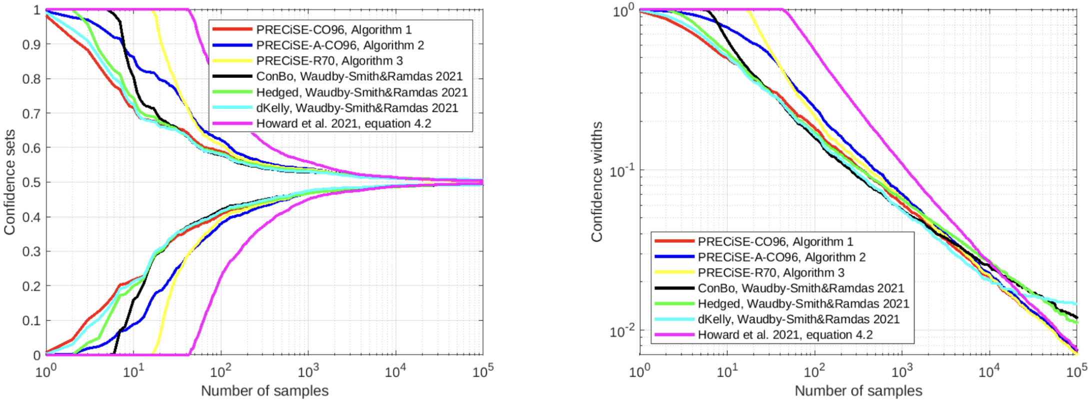


Figure: Time-uniform CSs (left) and their widths (right) for a sequence of i.i.d. Beta$(1,1)$ r.v.s.

# Outline

# Summary

# Summary

1. Betting $\rightarrow$ Ville $\rightarrow$ CS is mechanically simple once you have an e-process $\text{Wealth}_t(m)$

# Summary

1. Betting $\rightarrow$ Ville $\rightarrow$ CS is mechanically simple once you have an e-process $\text{Wealth}_t(m)$
2. The paper's novelty: center the concentration inequality at log-wealth at best bet in hindsight and control the gap by universal portfolio regret

# Summary

1. Betting $\rightarrow$ Ville $\rightarrow$ CS is mechanically simple once you have an e-process $\text{Wealth}_t(m)$
2. The paper's novelty: center the concentration inequality at log-wealth at best bet in hindsight and control the gap by universal portfolio regret
3. We do not need to run the ideal betting strategy: a regret bound implies the existence of a CS and enables inversion

# Summary

1. Betting $\to$ Ville $\to$ CS is mechanically simple once you have an e-process $\text{Wealth}_t(m)$
2. The paper's novelty: center the concentration inequality at log-wealth at best bet in hindsight and control the gap by universal portfolio regret
3. We do not need to run the ideal betting strategy: a regret bound implies the existence of a CS and enables inversion
4. **PRECiSE-CO96** uses $\text{Beta}(\frac{1}{2}, \frac{1}{2})$ prior on the portfolio weight $b \in [0, 1]$
   $\Rightarrow$ first never-vacuous CS for *any* sample size $T$

# Summary

1. Betting $\to$ Ville $\to$ CS is mechanically simple once you have an e-process $\text{Wealth}_t(m)$
2. The paper's novelty: center the concentration inequality at log-wealth at best bet in hindsight and control the gap by universal portfolio regret
3. We do not need to run the ideal betting strategy: a regret bound implies the existence of a CS and enables inversion
4. **PRECiSE-CO96** uses $\text{Beta}(\frac{1}{2}, \frac{1}{2})$ prior on the portfolio weight $b \in [0, 1]$

   $\Rightarrow$ first never-vacuous CS for *any* sample size $T$
5. **PRECiSE-R70** uses a zero-centered prior on the bet $\beta \in [-1, 1]$

   $\Rightarrow$ LIL rate asymptotically

# References


Cover, T. and Ordentlich, E. (1996). Universal portfolios with side information. *IEEE Transactions on Information Theory*, 42(2):348–363.

Cover, T. M. (1991). Universal portfolios. *Mathematical Finance*, 1(1):1–29.

Orabona, F. and Jun, K.-S. (2024). Tight concentrations and confidence sequences from the regret of universal portfolio. *IEEE Transactions on Information Theory*, 70(1):436–455.

Orabona, F. and Pal, D. (2016). Coin Betting and Parameter-Free Online Learning. In Lee, D., Sugiyama, M., Luxburg, U., Guyon, I., and Garnett, R., editors, *Advances in Neural Information Processing Systems*, volume 29. Curran Associates, Inc.

Ramdas, A. and Wang, R. (2025). Hypothesis testing with e-values. *Foundations and Trends® in Statistics*, 1(1-2):1–390.

Robbins, H. (1970). Statistical Methods Related to the Law of the Iterated Logarithm. *The Annals of Mathematical Statistics*, 41(5):1397 – 1409.