Bayes-Assisted Confidence Sequences for CDFs

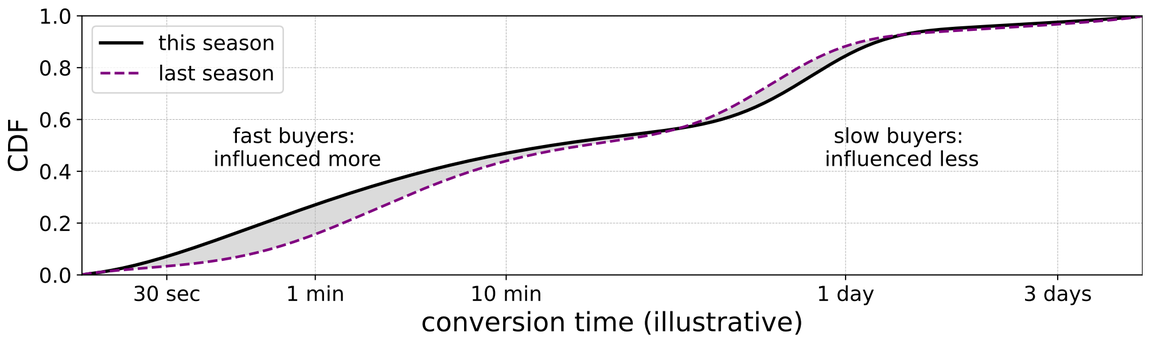
Valentin Kilian*, Shirley Xiaoqi Liu*, Judith Rousseau, Francesco Orabona, Patrick Rebeschini

University of Oxford · Dauphine Université PSL Paris Sciences & Lettres · King Abdullah University of Science and Technology

Ongoing work: feedback and related ideas very welcome! Correspondence to: shirley.liu@stats.ox.ac.uk

A motivating application

- Amazon adds an express-checkout button. Monitor sequentially what fraction of users convert in 30s, 1min, 1 day, ...? This is about **whole CDF** not mean.
- Instead of starting from scratch, can use CDF from last season as a **prior**.



A primer: single-mean inference by betting

Consider i.i.d. random variables $X_1, X_2, \cdots \in [0,1]$ with mean $\mu := \mathbb{E}[X_t]$.
Goal: Confidence sequence (CS) via intervals $I_t = [\ell_t(X_{1:t}), u_t(X_{1:t})]$ such that

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

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

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

  Under $H_0\colon \mu = m$, i.e., $\mathbb{E}[c_t(m)] = 0$.
- Bet $\beta_t \in [-\frac{1}{1-m}, \frac{1}{m}]$ fraction of current wealth on the coin's outcome, compounding wealth

$$W_t(\beta_{1:t}, m) = W_{t-1}(\beta_{1:t-1}, m)\big(1+\beta_t\, c_t(m)\big)\,, \qquad W_0(m) = 1\,.$$

- By Ville's inequality, a valid I_t is the set of candidate means not yet rejected:

$$I_t = \Big\{m \in [0,1] : \underbrace{\max_{\beta_{1:i}=\beta} \ln W_i(\beta_{1:i}, m) - \ln R_i(m)}_{\text{by definition of regret}} < \ln \tfrac{1}{\delta} \quad \forall i \leq t\Big\}\,.$$

Recipe: Smart betting algorithm $\implies$ small regret $\ln R_t(m)$ $\implies$ tight CS.

Known effects of the prior on the bet

Betting can be written as 2-stock portfolio selection:

$$\overbrace{1+\beta c}^{\text{per-round wealth multiplier}} = 1 + \underbrace{\Big[\Big(\tfrac{1}{1-m}+\tfrac{1}{m}\Big)b - \tfrac{1}{1-m}\Big]}_{=\beta} c = \overbrace{\begin{bmatrix} b \\ 1-b \end{bmatrix}}^{\text{wealth allocations to 2 stocks}} \cdot \overbrace{\begin{bmatrix} 1+\frac{c}{m} \\ 1-\frac{c}{1-m} \end{bmatrix}}^{\text{market gains of 2 stocks}}, \quad b \in [0,1]\,.$$

State-of-the-art CS is obtained through universal portfolio [Cover and Ordentlich, 1996, Orabona and Jun, 2024], which mixes constant-b betting strategies over a fixed prior F.

Regret $\ln R_t(m)$ and the resulting CS are governed entirely by the prior F:

$$\ln R_t(m) = \max_{b\in[0,1]} \ln W_t(b,m) - \ln \textstyle\int_0^1 W_t(b,m)\mathrm{d}F(b)$$

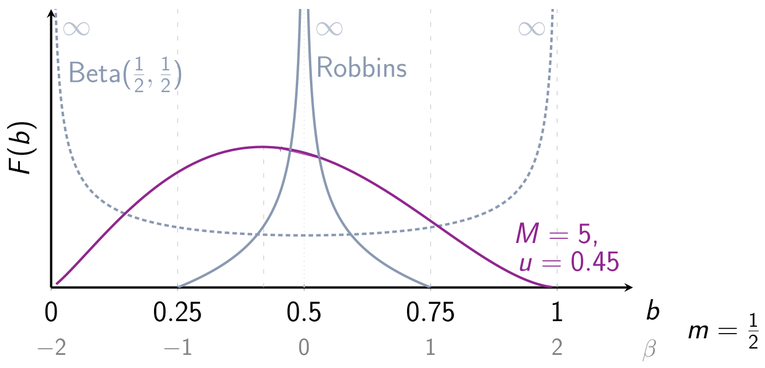
- Beta$(\frac{1}{2},\frac{1}{2})$ prior $F(b) = 1/\big[\pi\sqrt{b(1-b)}\big] \Rightarrow$ never-vacuous CS at *small* t.
- Robbins [1970] prior $F(\beta) \asymp 1/\Big[|\beta| \ln\big(\frac{1}{|\beta|}\big)\Big(\ln\ln\big(\frac{1}{|\beta|}\big)\Big)^2\Big] \Rightarrow \ln\ln t$ rate as $t\to\infty$

Both priors are tuned for worst-case regret, ignoring prior knowledge of μ.

Question. How to incorporate prior knowledge of μ to tighten the CS?

A Bayes-assisted 2-stock universal portfolio algorithm

As a starting point, consider $X_1, X_2, \ldots \overset{\text{i.i.d.}}{\sim} \text{Bern}(\mu)$.
Encode a guess u of μ through a **Beta prior on the bet** b with mean $u \in (0,1)$ and concentration $M > 0$: $F = \text{Beta}\big(Mu,\ M(1-u)\big)$.



Use this prior to form wealth process and CS.

Well-specified prior yields constant-factor gains

Theorem. Define $\hat{\mu} := \frac{1}{t}\sum_{i=1}^t \mathbb{1}\{X_i = 1\}$. If $\min\{t\hat{\mu}, t(1-\hat{\mu}), Mu, M(1-u)\} \geq 1$, then the regret of the universal portfolio algorithm with $F = \text{Beta}\big(Mu, M(1-u)\big)$ satisfies

$$\ln R_t(m) = \underbrace{\tfrac{1}{2}\ln t}_{\text{fixed rate}} + \underbrace{C(\hat{\mu}, M, u)}_{\mathcal{O}(1)\text{ term}}\,,$$

(data-dependent: $\hat{\mu}$; prior-dependent: M, u)

$$\text{where} \quad C(\hat{\mu}, M, u) \approx -\tfrac{1}{2}\ln M + M\cdot \mathrm{kl}(u\,\|\,\hat{\mu}) + \tfrac{1}{2}\ln\frac{\hat{\mu}(1-\hat{\mu})}{u(1-u)}\,,$$

with $\mathrm{kl}(u\|\hat{\mu}) = u\ln\frac{u}{\hat{\mu}} + (1-u)\ln\frac{1-u}{1-\hat{\mu}}$.

Effects of the prior:

- A good guess $u \approx \mu$ reduces $M\,\mathrm{kl}(u\,\|\,\hat{\mu})$ ☺.
- A larger M reduces $-\frac{1}{2}\ln M$ ☺ but amplifies misspecification via $M\,\mathrm{kl}(u\|\hat{\mu})$ ☹.

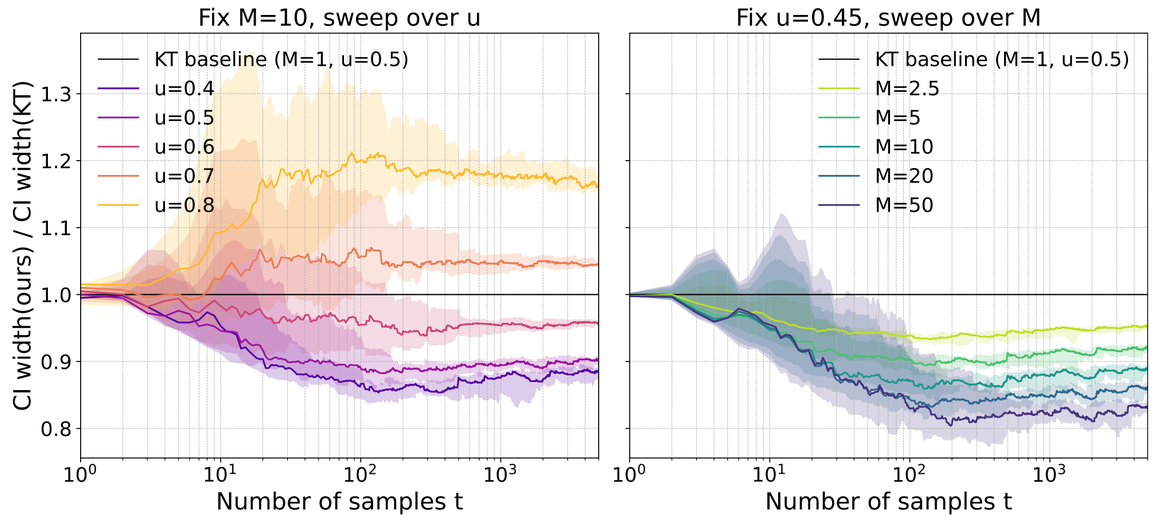


Figure: Solid curves and shaded bands show median and interquartile range over 20 trials, compared to the Krichevsky-Trofimov (KT) bettor as baseline. $X_1, X_2, \ldots \overset{\text{i.i.d.}}{\sim} \text{Bern}(0.4)$, $\delta = 0.05$.

Bayes-assisted confidence sequences for CDFs

Consider i.i.d. random variables $Z_1, Z_2, \cdots \in \mathcal{Y} \subseteq \mathbb{R}$ with CDF

$$F(y) := \mathbb{P}(Z_t \leq y) \quad \text{for } y \in \mathcal{Y}\,.$$

Goal: Anytime-valid confidence bands via intervals $I_t(y)$ such that

$$\mathbb{P}\big(F(y) \in I_t(y),\ \forall t \geq 1, y \in \mathcal{Y}\big) \geq 1-\delta\,.$$

- Run the same universal portfolio machinery per threshold $y \in \mathcal{Y}$
- **Use PAC-Bayes to extract pointwise-in-y guarantees** [Clerico et al., 2026]. For reference measure π and localised measure ρ, both on $\mathcal{Y}$ with $\rho \ll \pi$,

$$\int \max_{\beta_{1:t}=\beta} \ln W_t\big(\beta_{1:t}, F(y), y\big)\mathrm{d}\rho(y) - \underbrace{\mathrm{KL}(\rho\,\|\,\pi)}_{\text{localisation slack}} - \underbrace{\int \ln R_t\big(F(y), y\big)\mathrm{d}\rho(y)}_{\text{regret integral}} \leq \ln\tfrac{1}{\delta}\,.$$

- Well-specified prior on F reduces regret integral ☺.

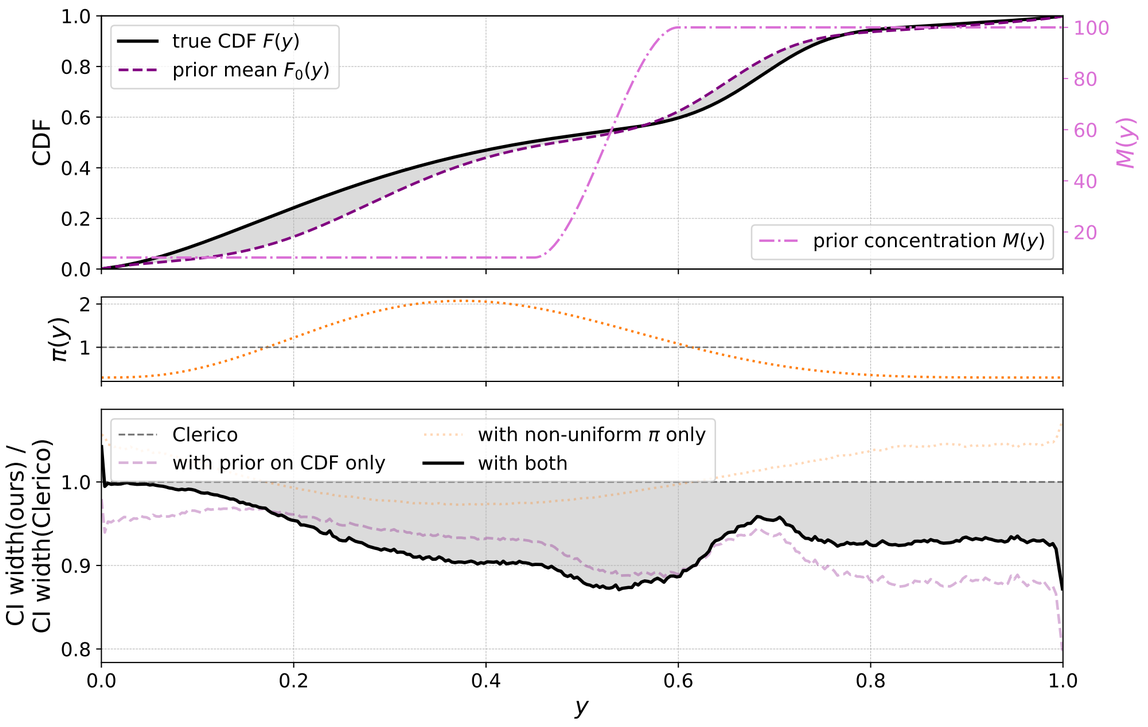


Figure: Per-y prior used is $F = \text{Beta}(M(y)F_0(y), M(y)(1-F_0(y)))$. $t = 2000, \delta = 0.05$

Many interesting next steps: Better choices of π, ρ? Improve robustness? ...

References

Eugenio Clerico, Hamish E. Flynn, and Patrick Rebeschini. Uniform mean estimation for monotonic processes. *Statistics & Probability Letters*, 228:110558, 2026. ISSN 0167-7152.

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

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

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