Estimating means of bounded random variables by betting

Ian Waudby-Smith & Aaditya Ramdas
Carnegie Mellon University
Royal Statistical Society meeting, 2023

Given bounded rvs: $X_1, X_2, \ldots, X_n \in [0,1]$
with mean $\mathbb{E}(X_i) = \mu$,

Goal: produce a confidence interval
$C_n \equiv C(X_1, \ldots, X_n)$ for μ:

$$\mathbb{P}(\mu \in C_n) \geq 1 - \alpha\,.$$

Hoeffding's inequality (1963) provides one solution:

$$C_n^H := \frac{1}{n}\sum_{i=1}^{n} X_i \pm \sqrt{\frac{\log(2/\alpha)}{2n}}$$

(No asymptotics or parametric assumptions.)

The downside? Not very sharp,
especially for small variance $\sigma^2 := \mathrm{Var}(X_i)$.

Given bounded rvs: $X_1, X_2, \ldots, X_n \in [0,1]$
with mean $\mathbb{E}(X_i) = \mu$,

Goal: produce a confidence interval
$C_n \equiv C(X_1, \ldots, X_n)$ for μ:

$$\mathbb{P}(\mu \in C_n) \geq 1 - \alpha,$$

so that C_n adapts to the underlying variance σ^2.

Hoeffding:

$$C_n^H := \left\{ m \in [0,1] : \prod_{i=1}^{n} \exp\left\{\lambda(X_i - m) - \lambda^2/8\right\} < \frac{1}{\alpha} \right\}, \qquad \lambda \leftarrow \sqrt{\frac{8\log(1/\alpha)}{n}}.$$

Our bound:

$$C_n := \left\{ m \in [0,1] : \prod_{i=1}^{n} (1 + \lambda_i(X_i - m)) < \frac{1}{\alpha} \right\}. \qquad \text{(design } \lambda_i \text{ later)}$$

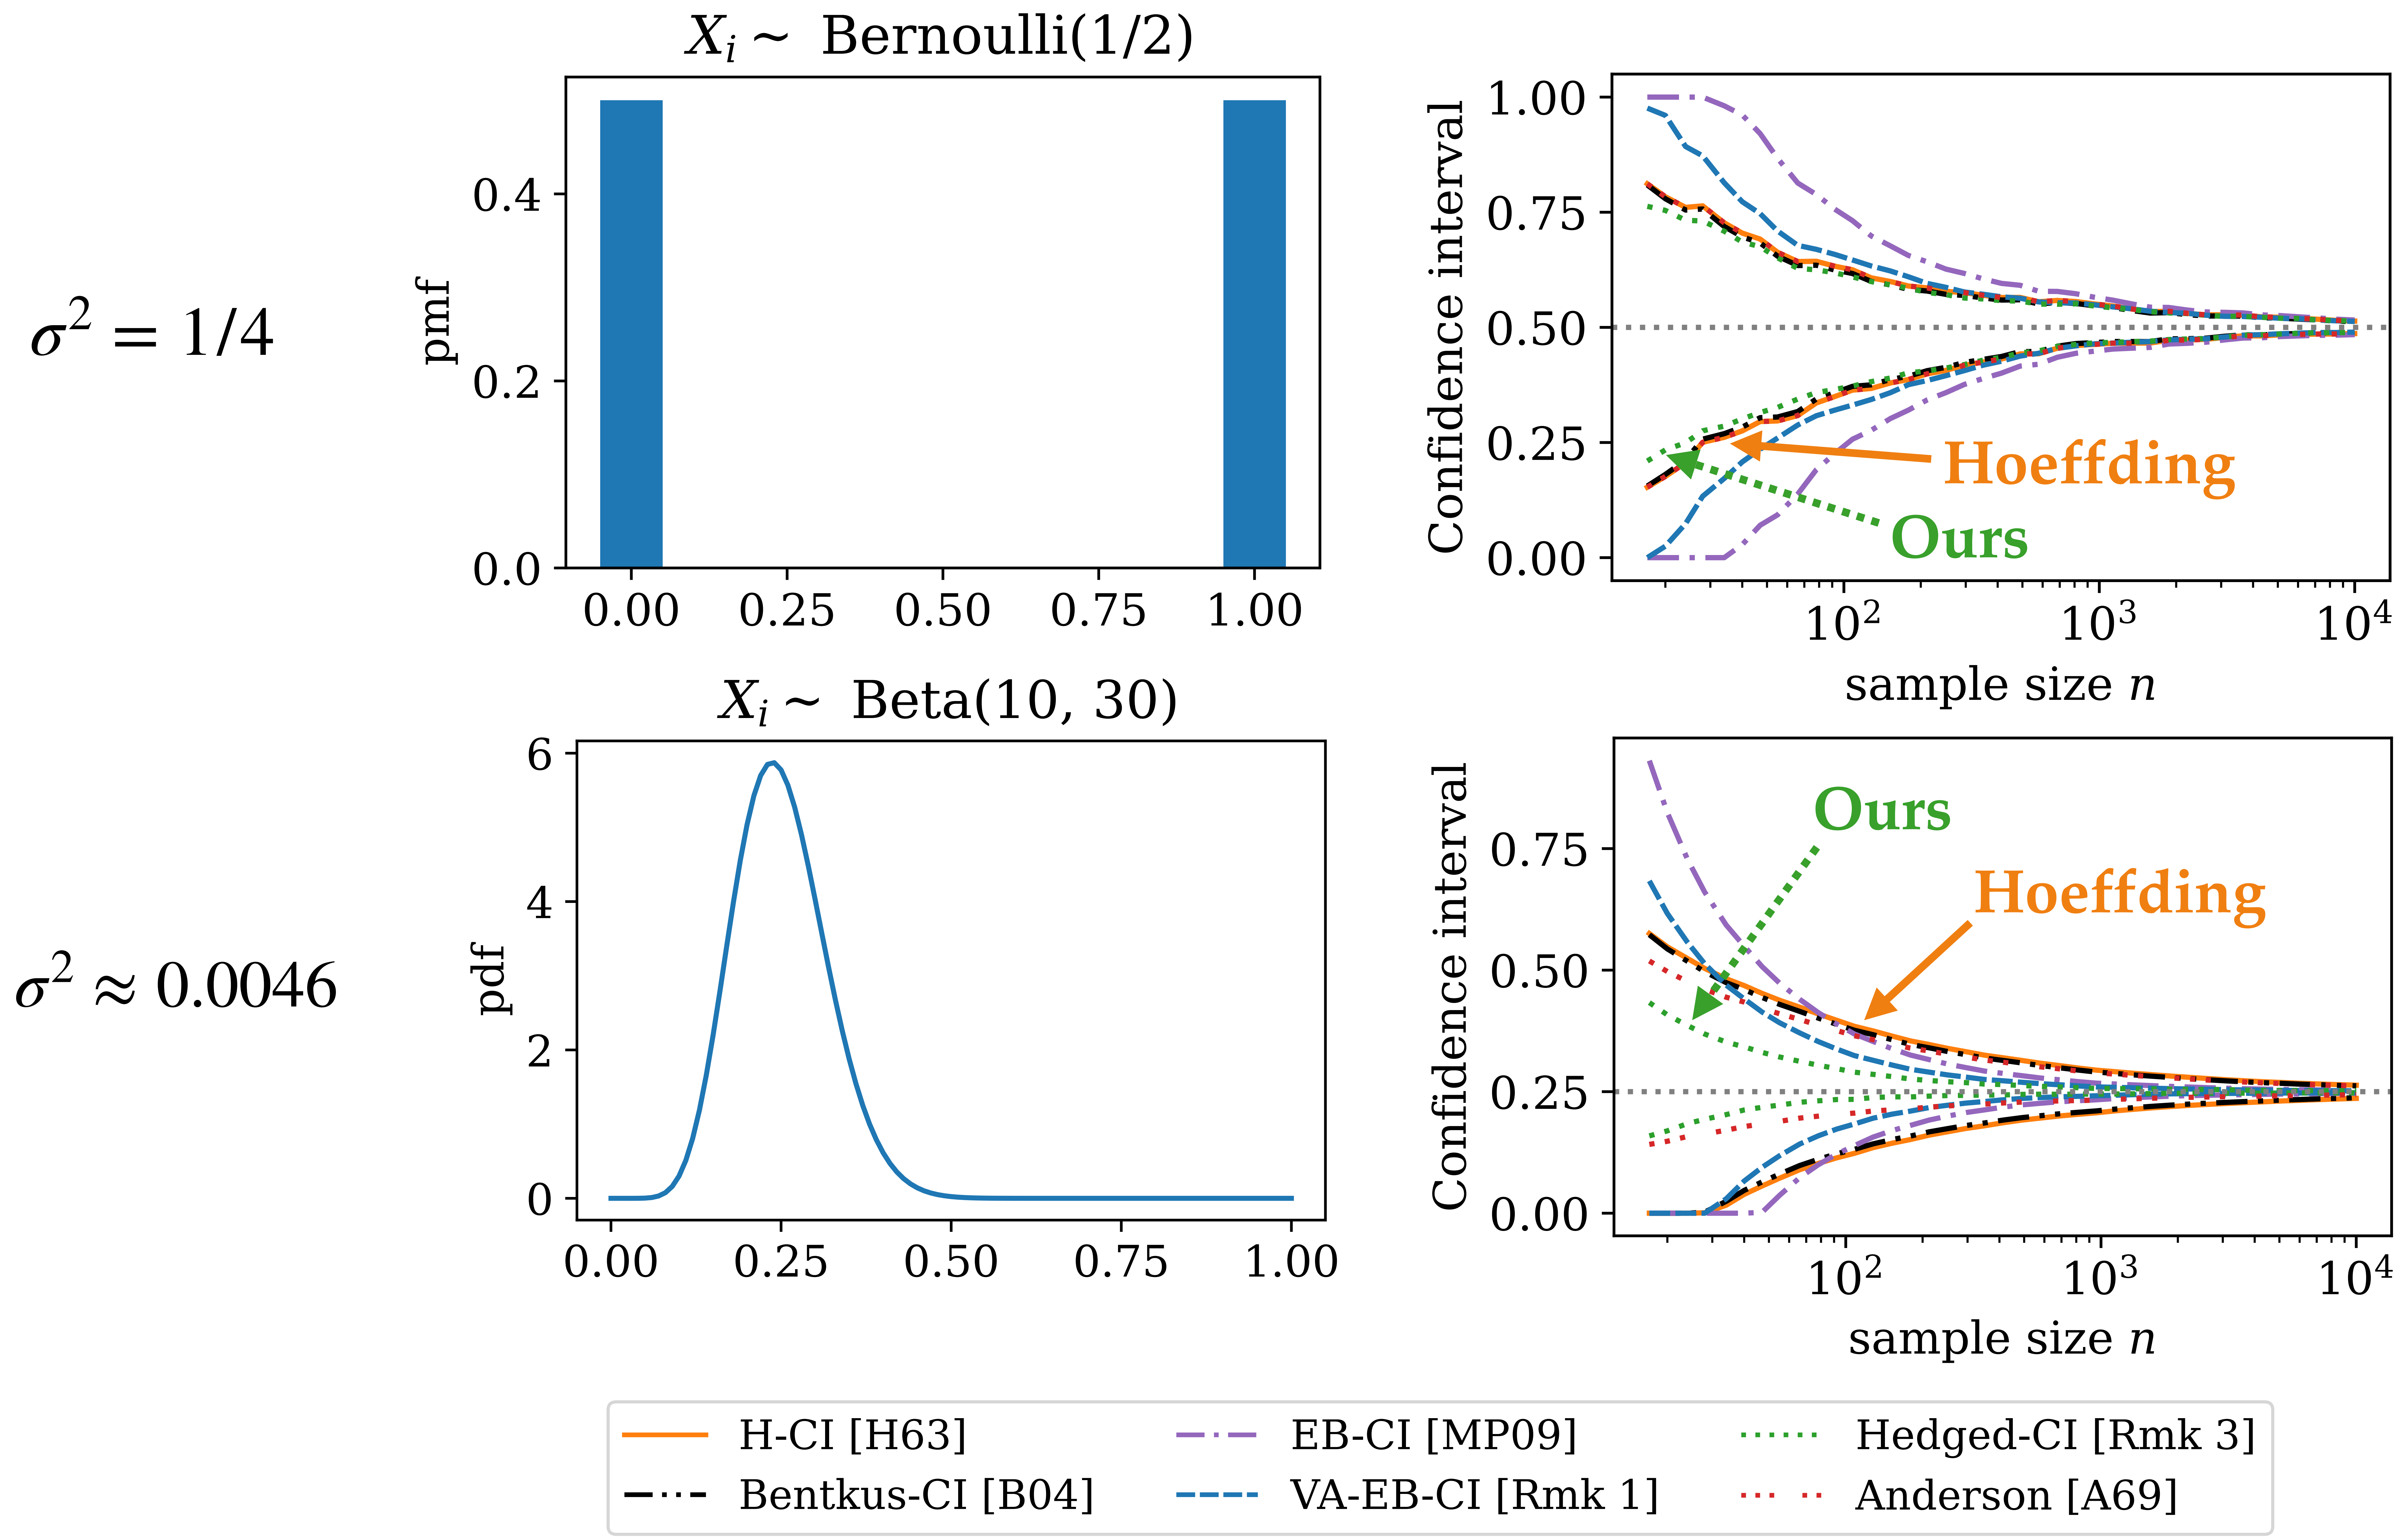
$\sigma^2 = 1/4$
$X_i \sim$ Bernoulli(1/2)
pmf
0.0
0.2
0.4
0.00
0.25
0.50
0.75
1.00
Confidence interval
1.00
0.75
0.50
0.25
0.00
Hoeffding
Ours
10^2
10^3
10^4
sample size n
$\sigma^2 \approx 0.0046$
$X_i \sim$ Beta(10, 30)
pdf
0
2
4
6
0.00
0.25
0.50
0.75
1.00
Confidence interval
0.75
0.50
0.25
0.00
Ours
Hoeffding
10^2
10^3
10^4
sample size n
H-CI [H63]
Bentkus-CI [B04]
EB-CI [MP09]
VA-EB-CI [Rmk 1]
Hedged-CI [Rmk 3]
Anderson [A69]

Quick detour: motivation

Motivation 1 of 2:
Risk-controlling prediction sets.

Bates, Angelopoulos, Lei, Malik, Jordan (2021)

Their goal: prediction sets for tumors while controlling risk.

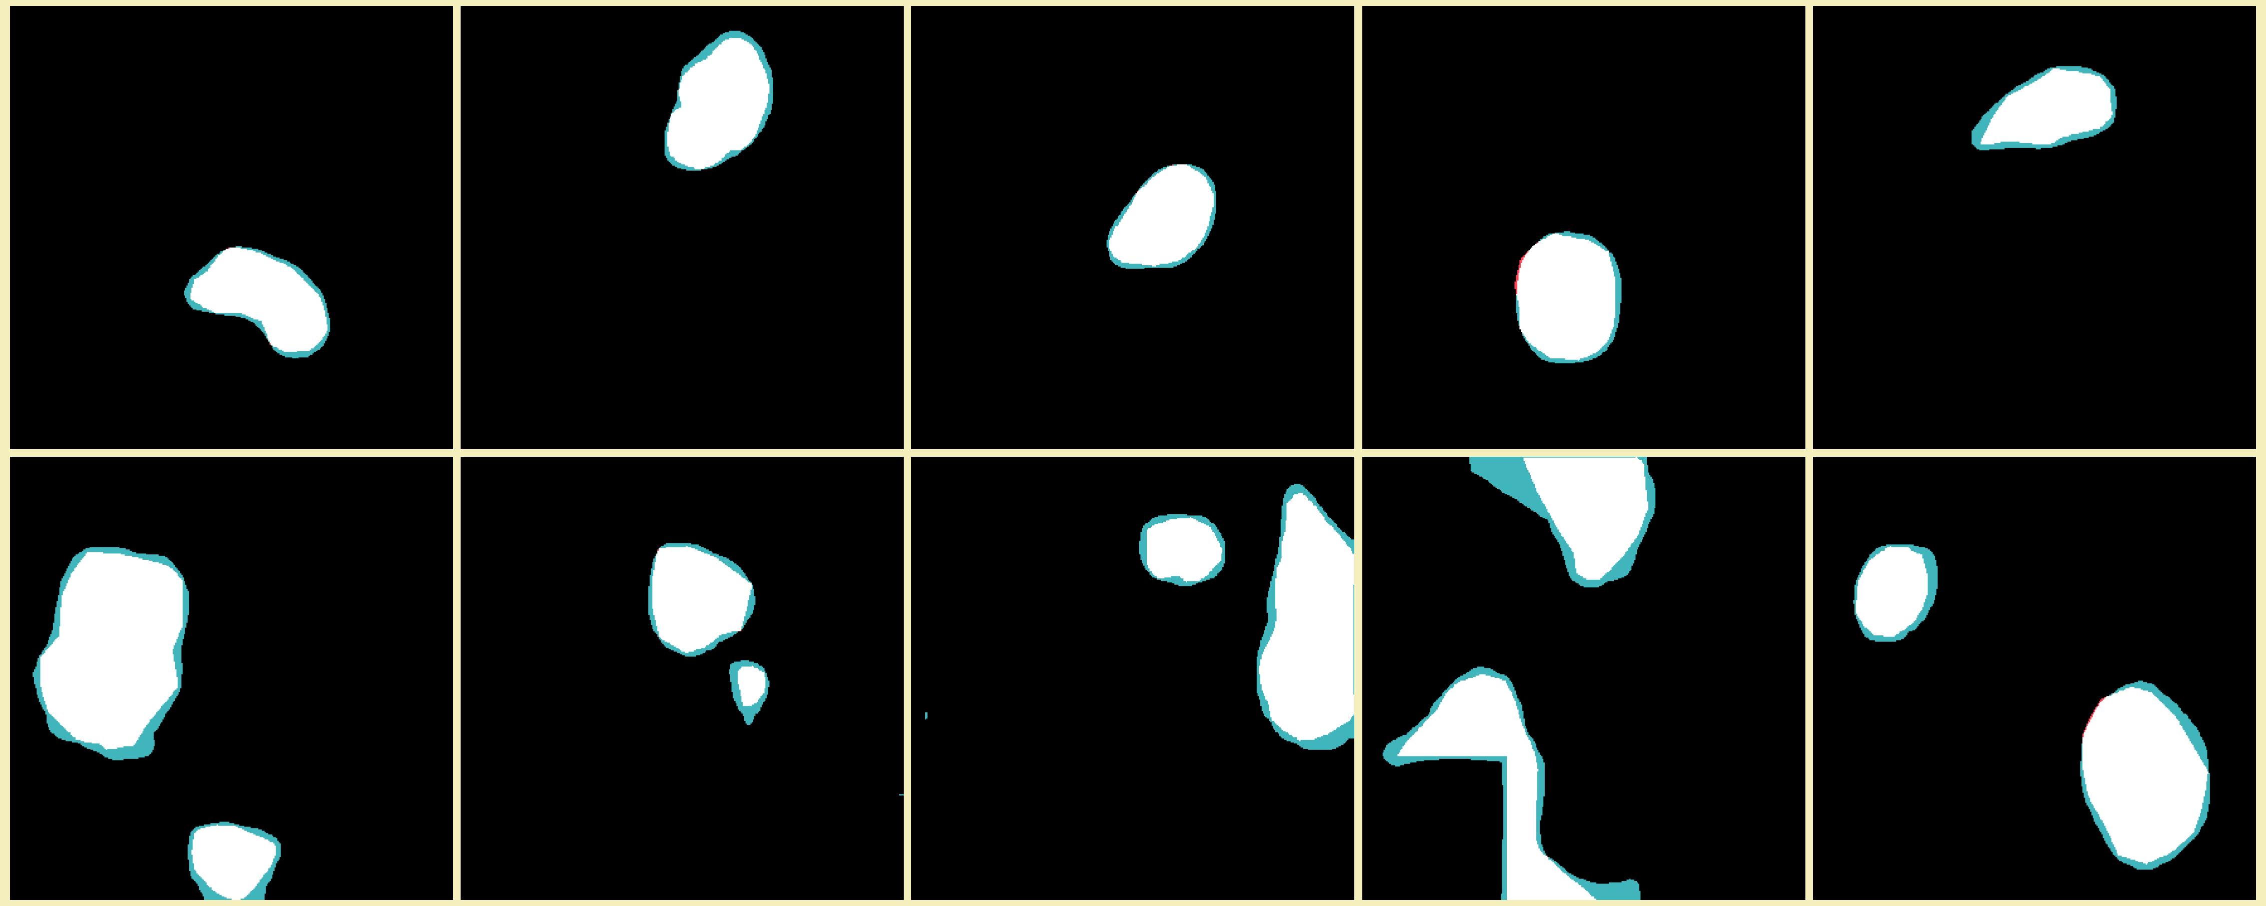

Tight confidence interval
$\implies$ sharper prediction sets.

Waudby-Smith, Stark, and Ramdas (2021)

Motivation 2 of 2:
Risk-limiting election audits

The goal: audit the outcome of an election using random samples of ballots.

Tight confidence sequence $\Longrightarrow$ faster audit.

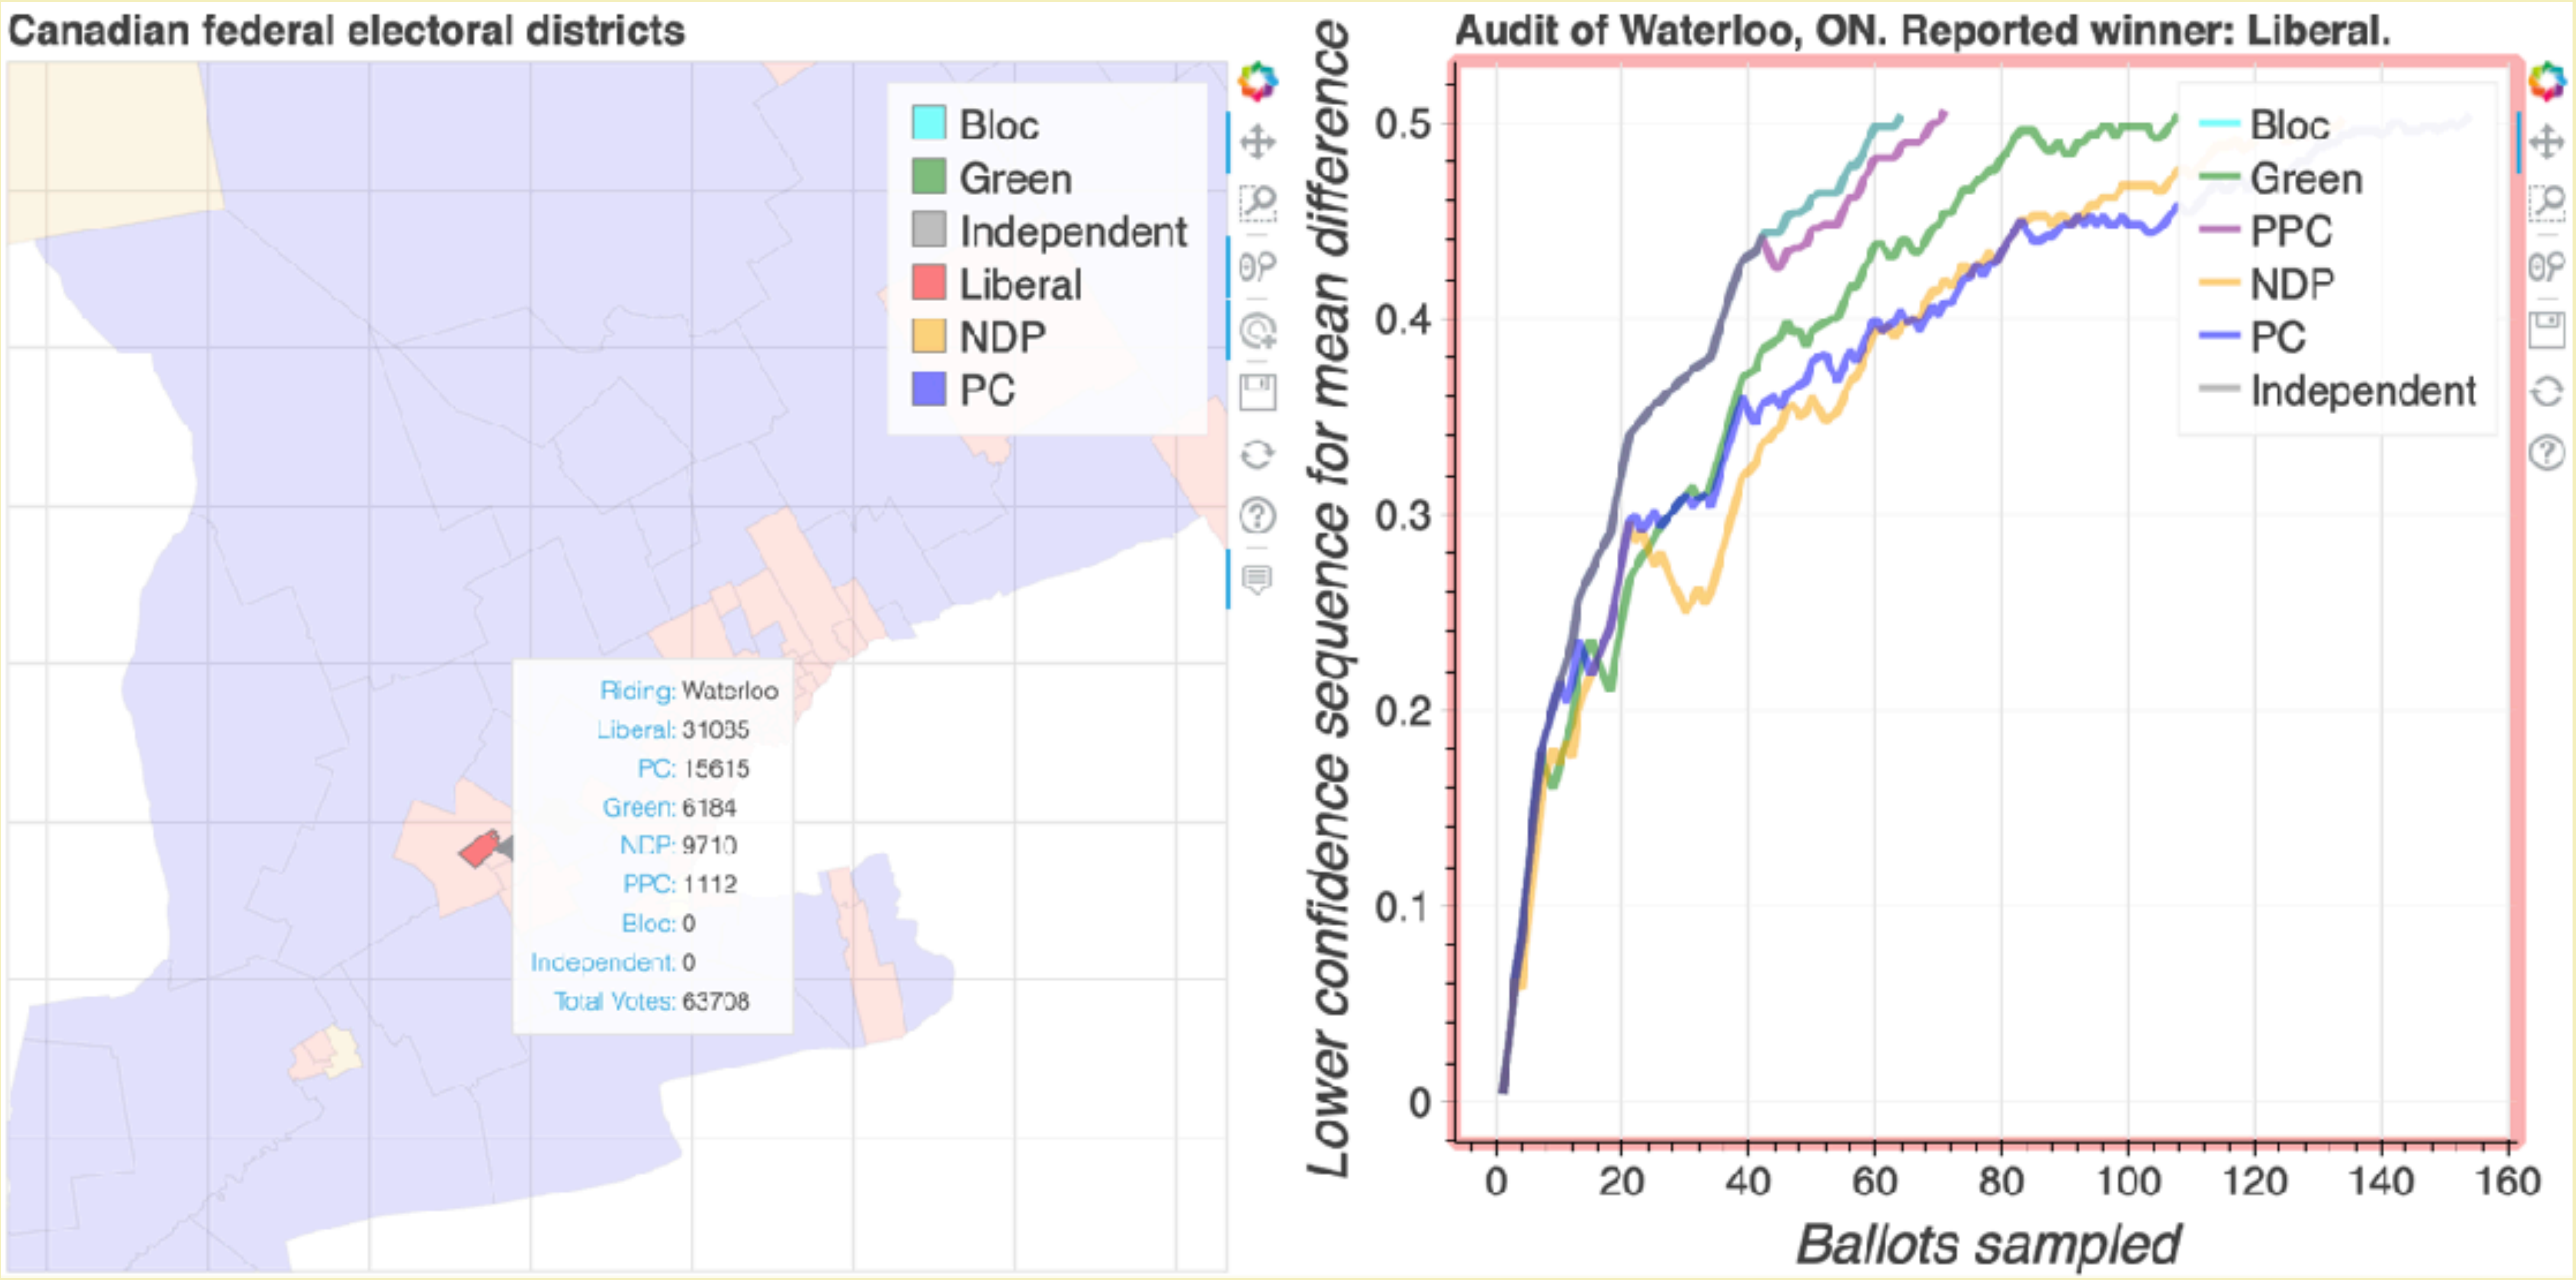

Formal setup:

1. Observe $X_1, X_2, X_3, \ldots$
2. $X_t \in [0,1]$ almost surely.
3. $\mathbb{E}(X_t \mid X_1, \ldots, X_{t-1}) = \mu$.

Familiar special case: $X_1, X_2, \ldots \stackrel{iid}{\sim} \mathbb{P}$, with $\mathbb{E}_{\mathbb{P}}(X_1) = \mu$.

Forgetting about confidence sets for a moment,
consider the following game for each m

$K_0 \leftarrow \$1$

For $t = 1,2,3,\ldots$:

Gambler chooses bet $\lambda_t \in (-1/(1-m),\ 1/m)$ (based on $X_1, \ldots, X_{t-1}$)

Observe X_t

$K_t \leftarrow K_{t-1} + K_{t-1} \cdot \lambda_t \cdot (X_t - m)$

EndFor

$$K_t \leftarrow K_{t-1} + K_{t-1} \cdot \lambda_t \cdot (X_t - m)$$

- What if $\mu \gg m$? $K_t \to \infty$ by cleverly choosing $\lambda_t > 0$.
- What if $\mu \ll m$? $K_t \to \infty$ by cleverly choosing $\lambda_t < 0$.
- What if $\mu = m$? Then the gambler can *never* make much money, no matter how λ_t is chosen!

$$C_t := \left\{ m \in [0,1] : K_t(m) < \frac{1}{\alpha} \right\}.$$

So C_t is "the set of all $m \in [0,1]$ for which $K_t(m)$ is small."

The "capital process" $(K_t(\mu))_{t=0}^{\infty}$

is a nonnegative martingale starting at one.

Proof:

$$\mathbb{E}(K_t(\mu) \mid X_1^{t-1}) = K_{t-1}(\mu) + K_{t-1}(\mu) \cdot \underbrace{\mathbb{E}(\lambda_t(X_t - \mu) \mid X_1^{t-1})}$$

$$= \lambda_t \left(\mathbb{E}(X_t \mid X_1^{t-1}) - \mu \right)$$

$$= 0$$

$$\therefore \mathbb{E}(K_t(\mu) \mid X_1^{t-1}) = K_{t-1}(\mu)$$

Ideas also in
Robbins et al.
Shafer & Vovk,
& many others

Jean Ville (1939): $\mathbb{P}\left(\exists t \geq 1 : K_t(\mu) \geq \frac{1}{\alpha}\right) \leq \alpha$

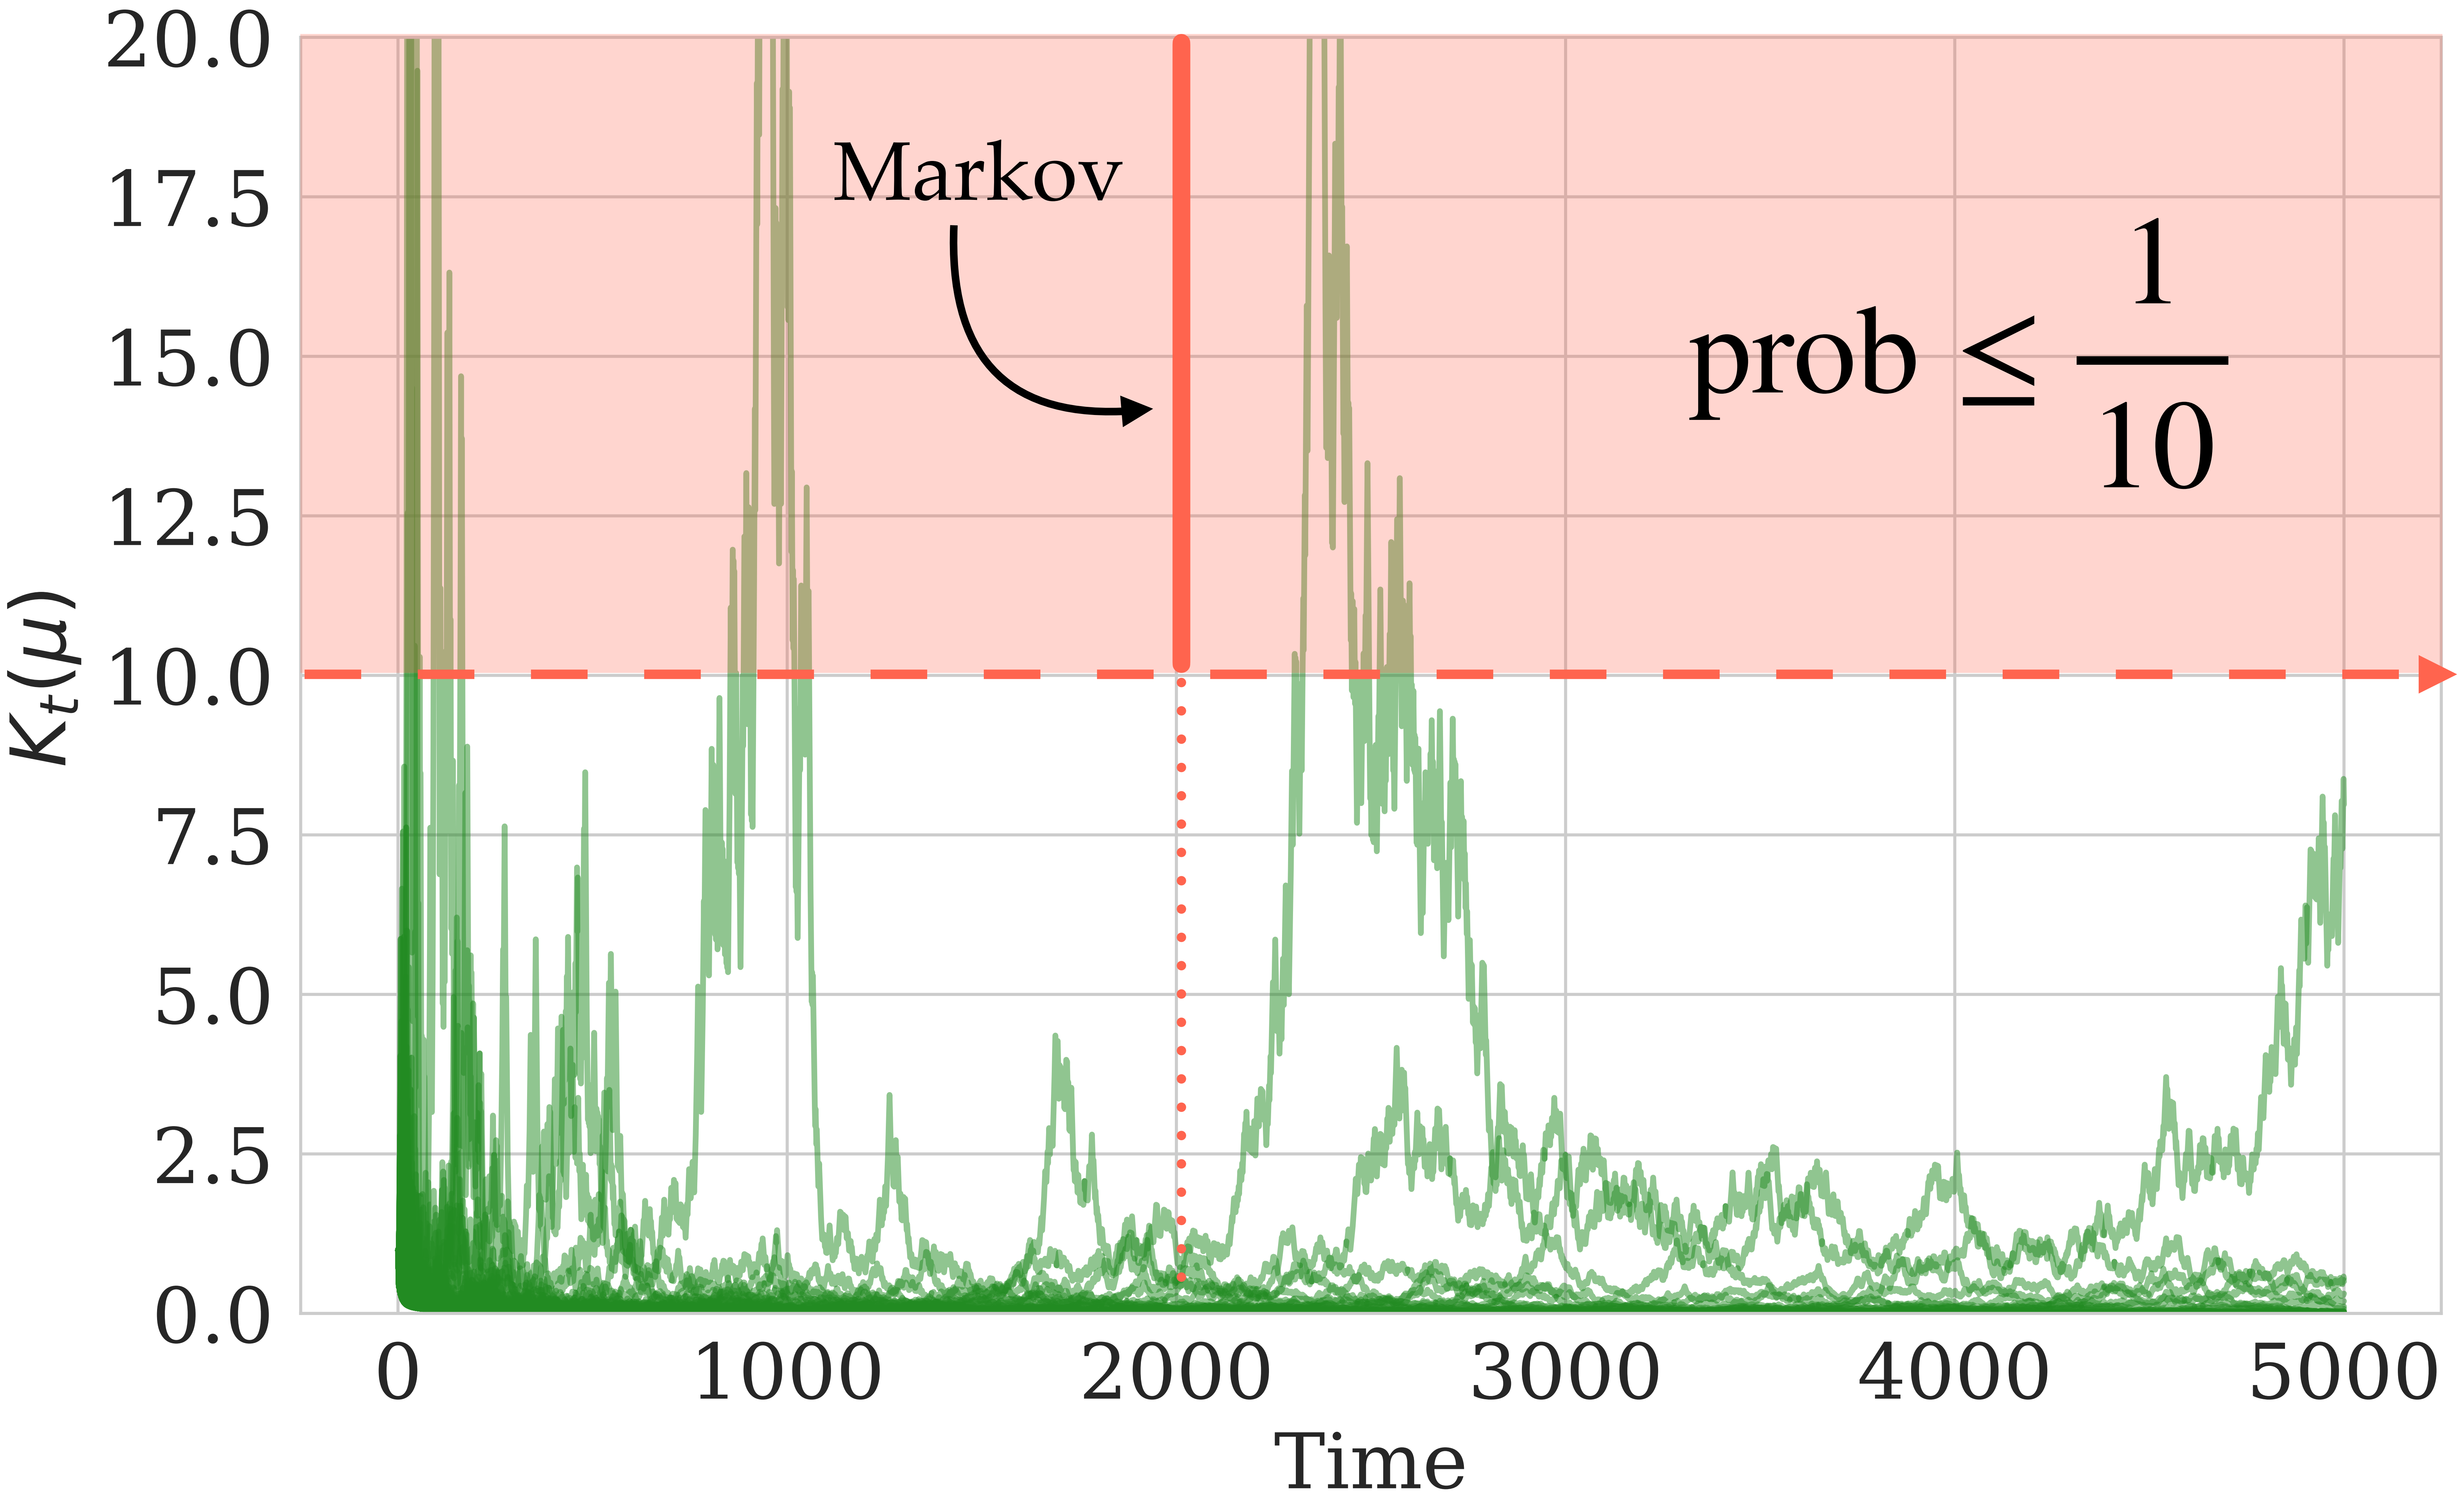

"Invert" Ville's inequality:

Robbins et al. (1960s-1970s),
Shafer & Vovk (2001, 2019)
Johari et al. (2015)
Jun & Orabona (2019),
Howard et al. (2020).
Grunwald et al. (2019)

$$C_t := \left\{ m \in [0,1] : K_t(m) < \frac{1}{\alpha} \right\}.$$

"The set of m for which the gambler didn't make much money".

$$\mathbb{P}\left(\exists t \geq 1 : \mu \notin C_t\right) \leq \alpha.$$

$(C_t)_{t=1}^{\infty}$ forms a $(1-\alpha)$-***confidence sequence.***

Detour: confidence sequences

Herbert Robbins, 1960s/70s

+ Siegmund, Darling, & Lai

Confidence sequence

$$\mathbb{P}\left(\exists t \geq 1 : \mu \notin C_t\right) \leq \alpha$$

$$\mathbb{P}\left(\forall t \geq 1, \mu \in C_t\right) \geq 1 - \alpha$$

Confidence interval

$$\forall n, \mathbb{P}\left(\mu \notin C_n\right) \leq \alpha$$

$$\forall n, \mathbb{P}\left(\mu \in C_n\right) \geq 1 - \alpha$$

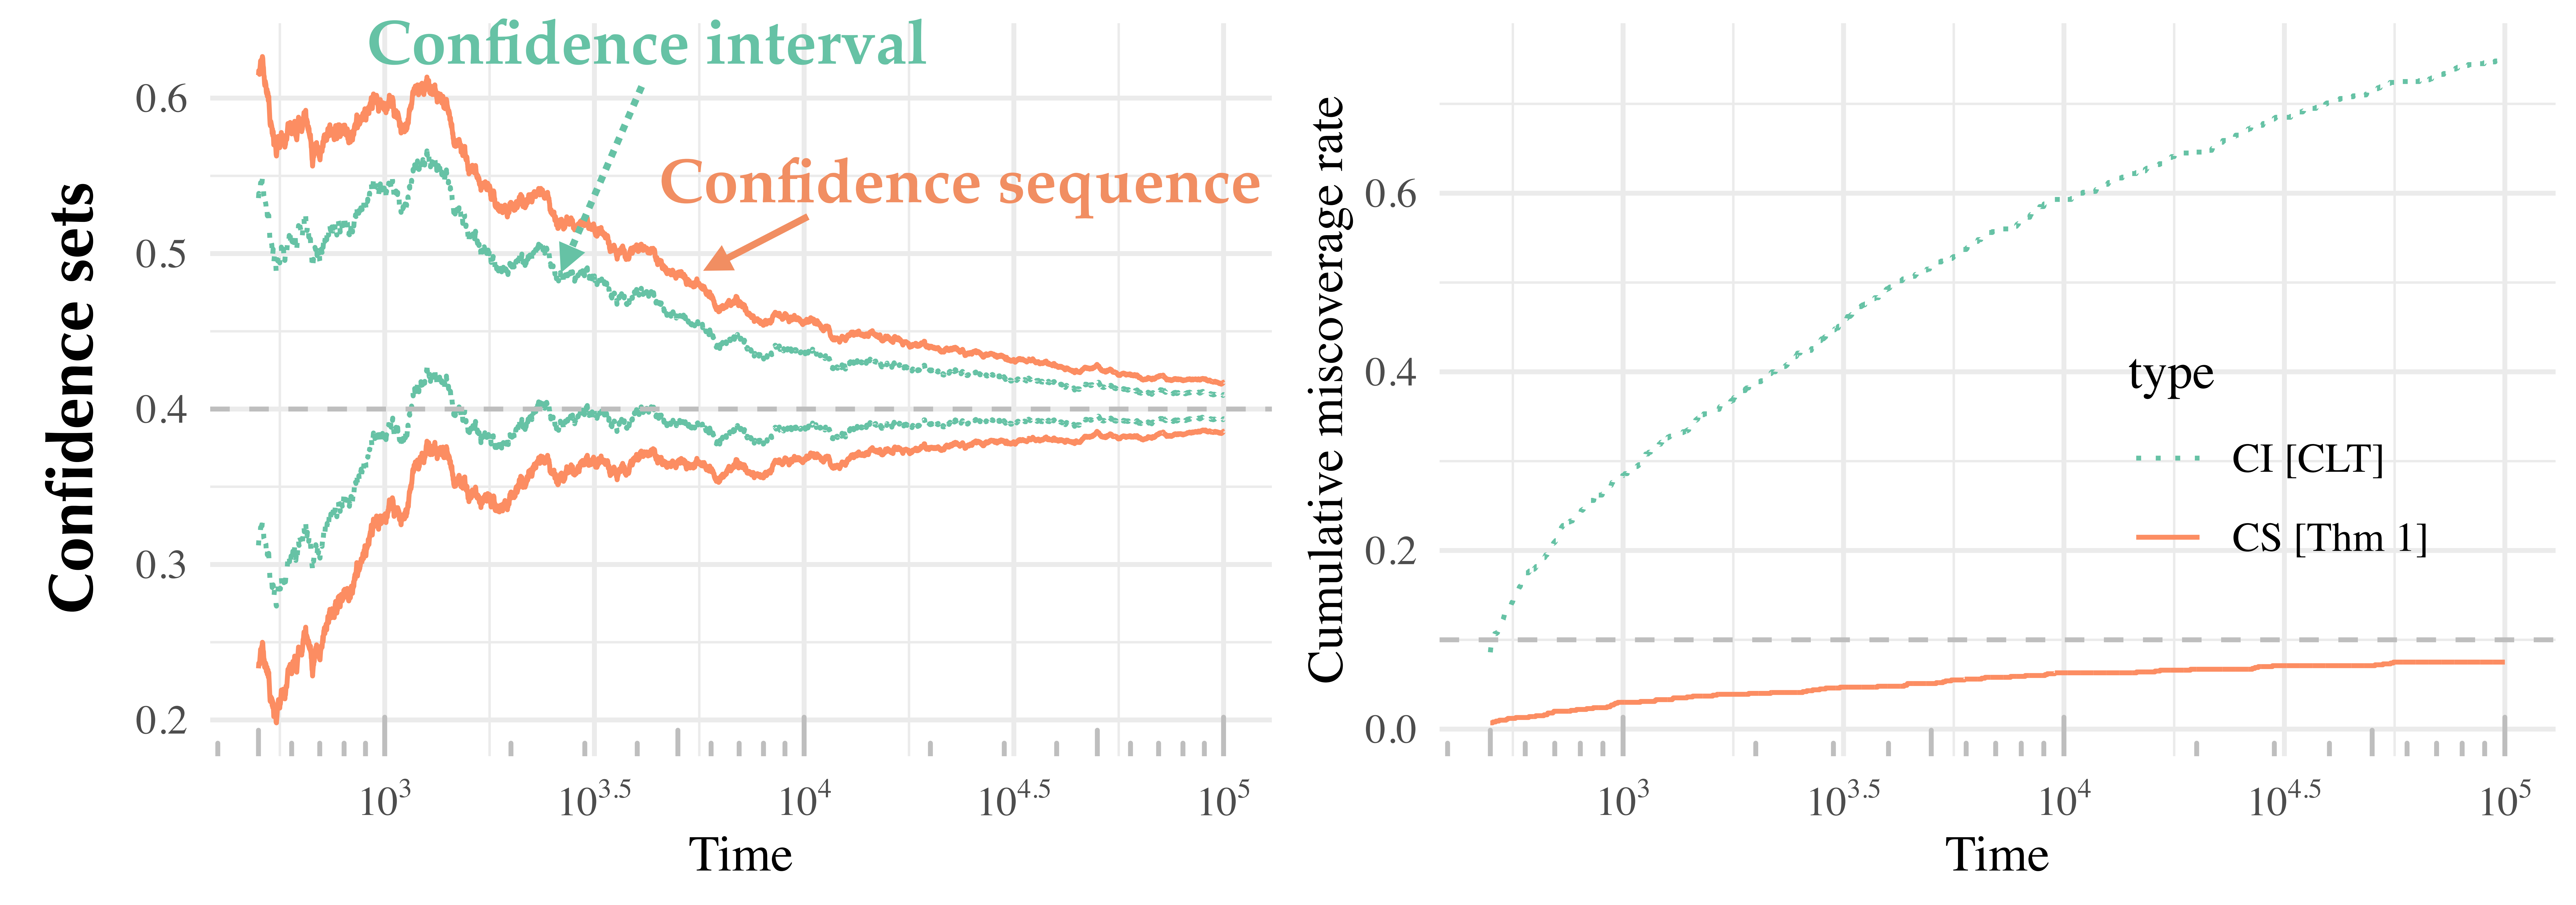

Confidence intervals are valid *at a single sample size.*
Confidence sequences are valid *at all sample sizes simultaneously.*

Back to confidence sequences for means of bounded random variables

Our confidence sequence:

$$C_t := \left\{ m \in [0,1] : \prod_{i=1}^{t} (1+\lambda_i(m) \cdot (X_i - m)) < \frac{1}{\alpha} \right\}$$

is valid for any λ_i but what is a smart choice?

Maximize the **G**rowth **R**ate **A**dapted to the **P**articular **A**lternative (**GRAPA**).

$$\text{Choose } \lambda_t(m) = \underset{\lambda}{\operatorname{argmax}} \frac{1}{t-1} \sum_{i=1}^{t-1} \log\left(1 + \lambda \cdot (X_i - m)\right)$$

$$\text{Or } \lambda_t(m) \approx \frac{\widehat{\mu}_{t-1} - m}{\widehat{\sigma}_{t-1}^2 + (\widehat{\mu}_{t-1} - m)^2}$$

(*approximate* GRAPA)

⋮

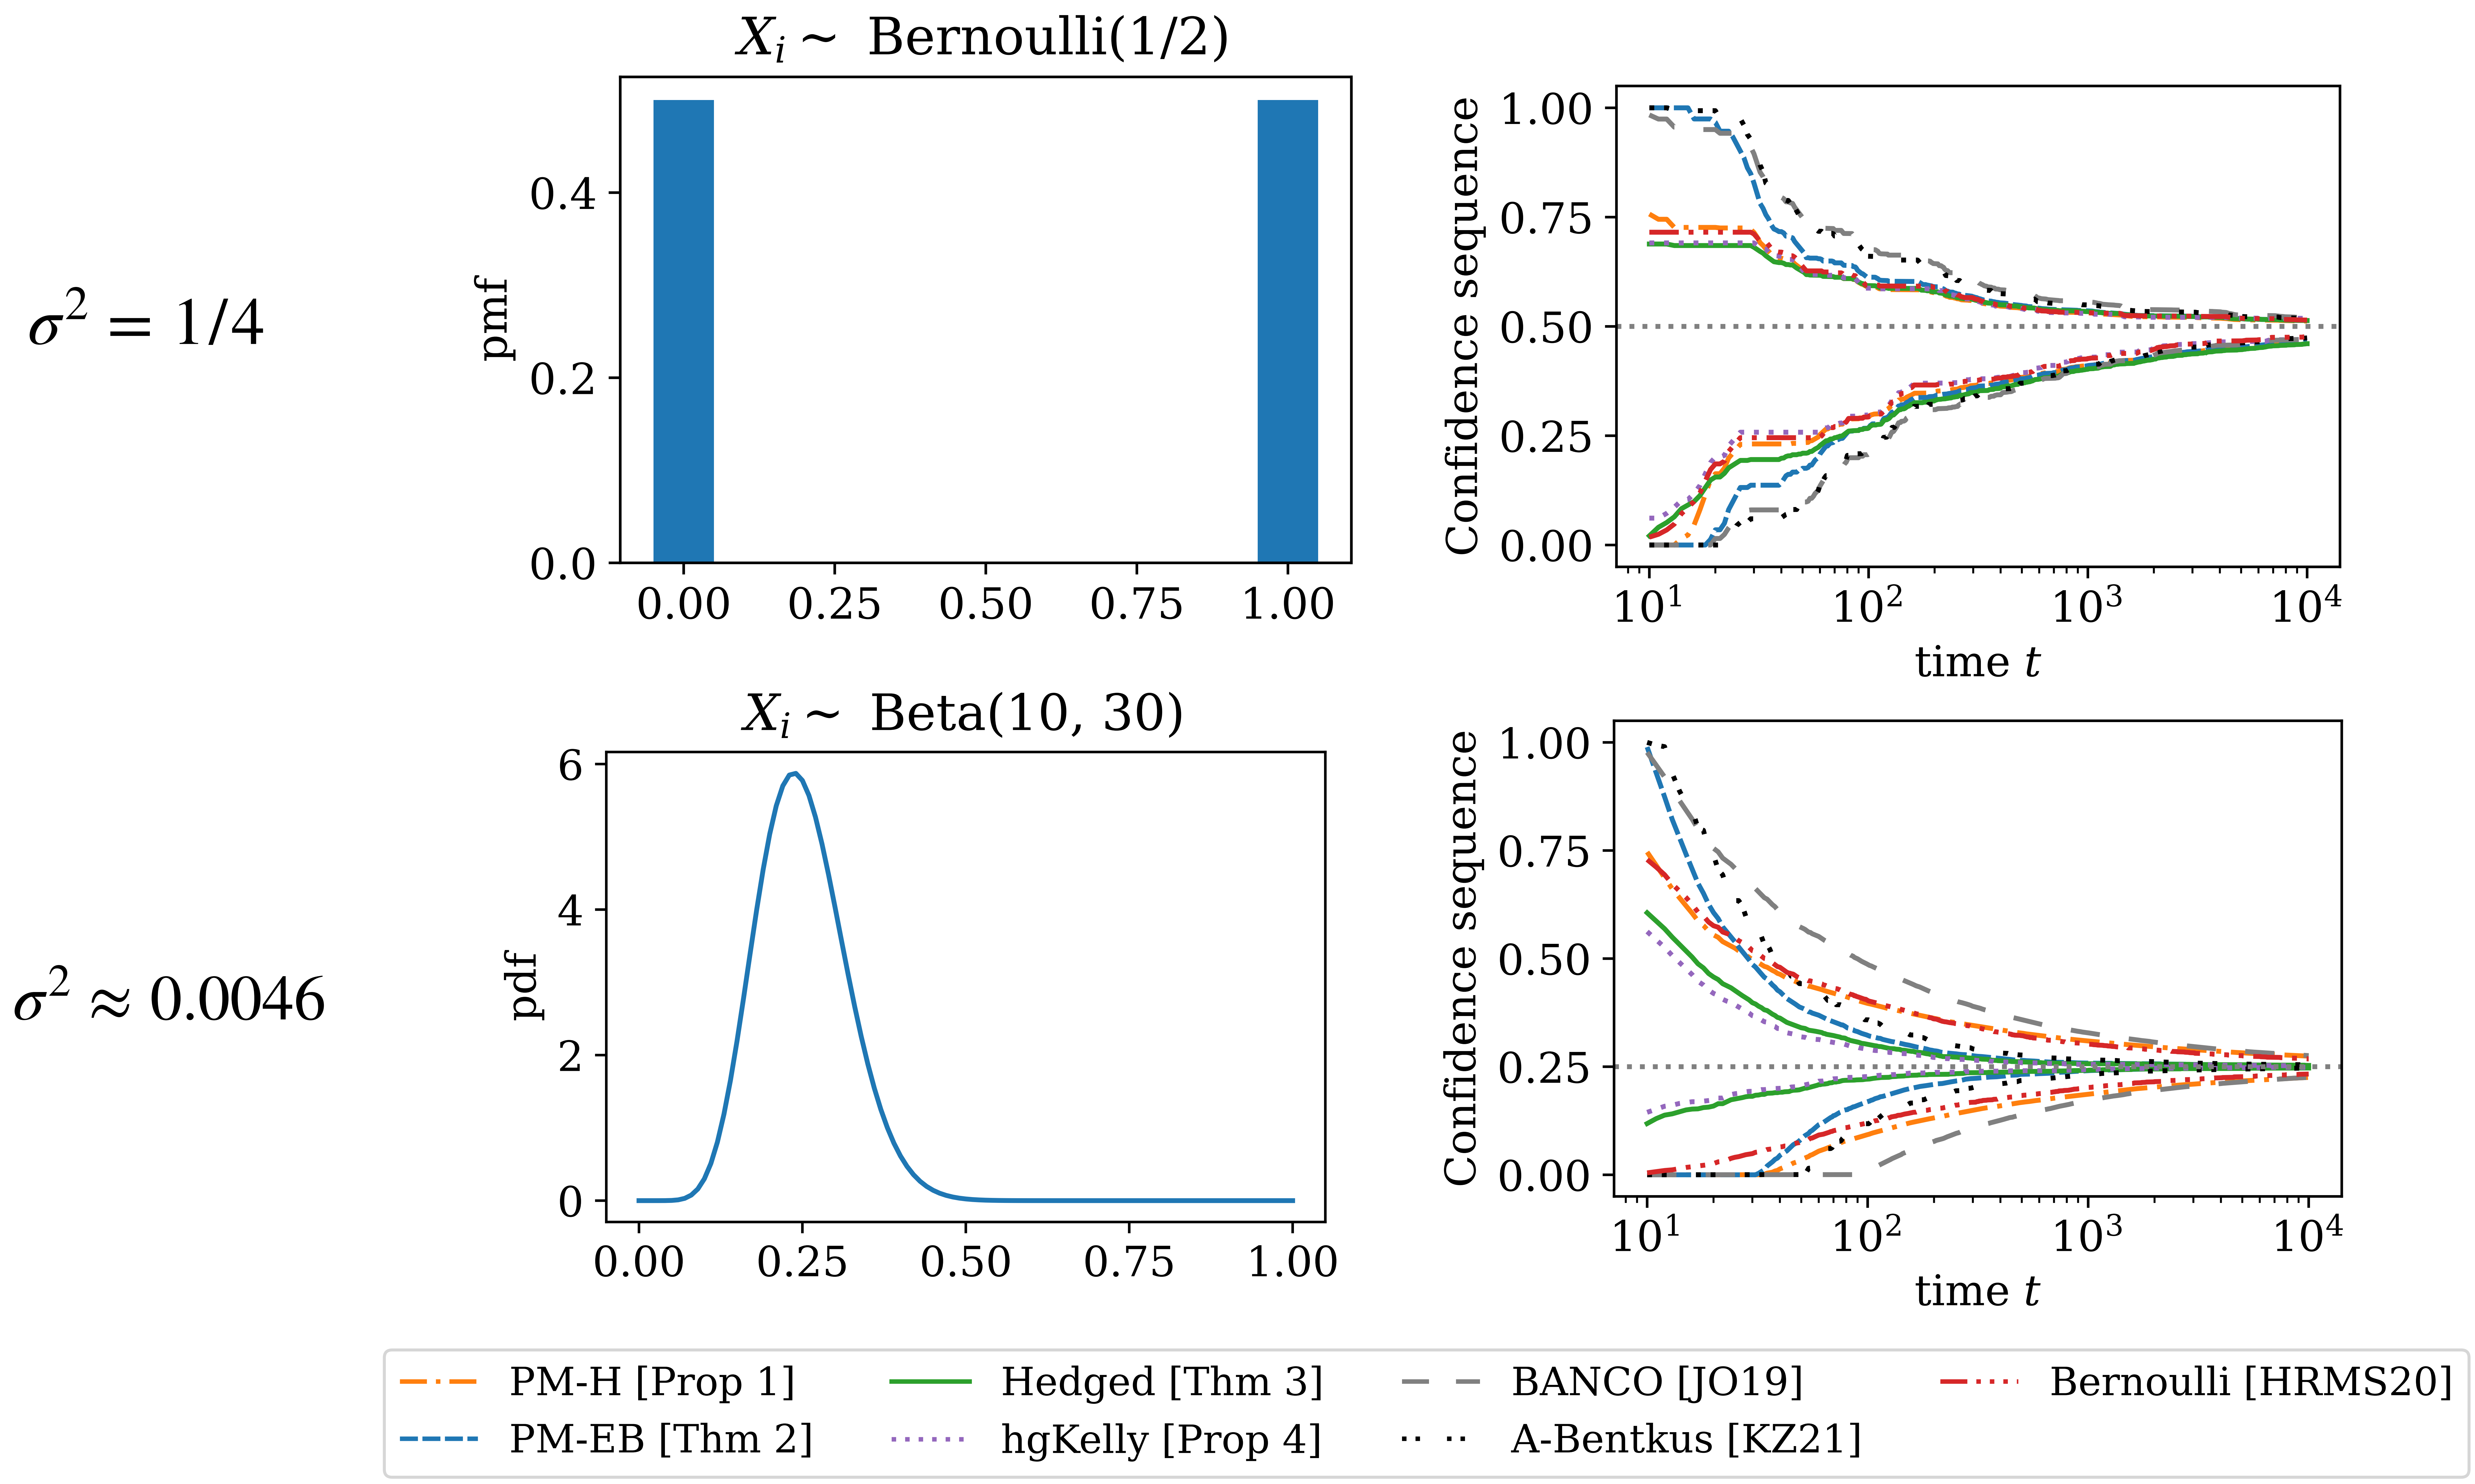

$\sigma^2 = 1/4$
$X_i \sim$ Bernoulli(1/2)
pmf
0.0
0.2
0.4
0.00
0.25
0.50
0.75
1.00
Confidence sequence
time t
10^1
10^2
10^3
10^4
$\sigma^2 \approx 0.0046$
$X_i \sim$ Beta(10, 30)
pdf
0
2
4
6
PM-H [Prop 1]
PM-EB [Thm 2]
Hedged [Thm 3]
hgKelly [Prop 4]
BANCO [JO19]
A-Bentkus [KZ21]
Bernoulli [HRMS20]

Indeed, the same bound

$$C_t := \left\{ m \in [0,1] : \prod_{i=1}^{t} (1+\lambda_i(m) \cdot (X_i - m)) < \frac{1}{\alpha} \right\}$$

can be used to derive state-of-the-art confidence intervals *and* sequences for both fixed-*n* *and* sequential regimes.

Even if we only care about fixed-n confidence intervals, deriving a confidence sequence first can be beneficial.

If $(C_t)_{t=1}^{\infty}$ is a confidence sequence, then C_n is a confidence interval for any fixed n.

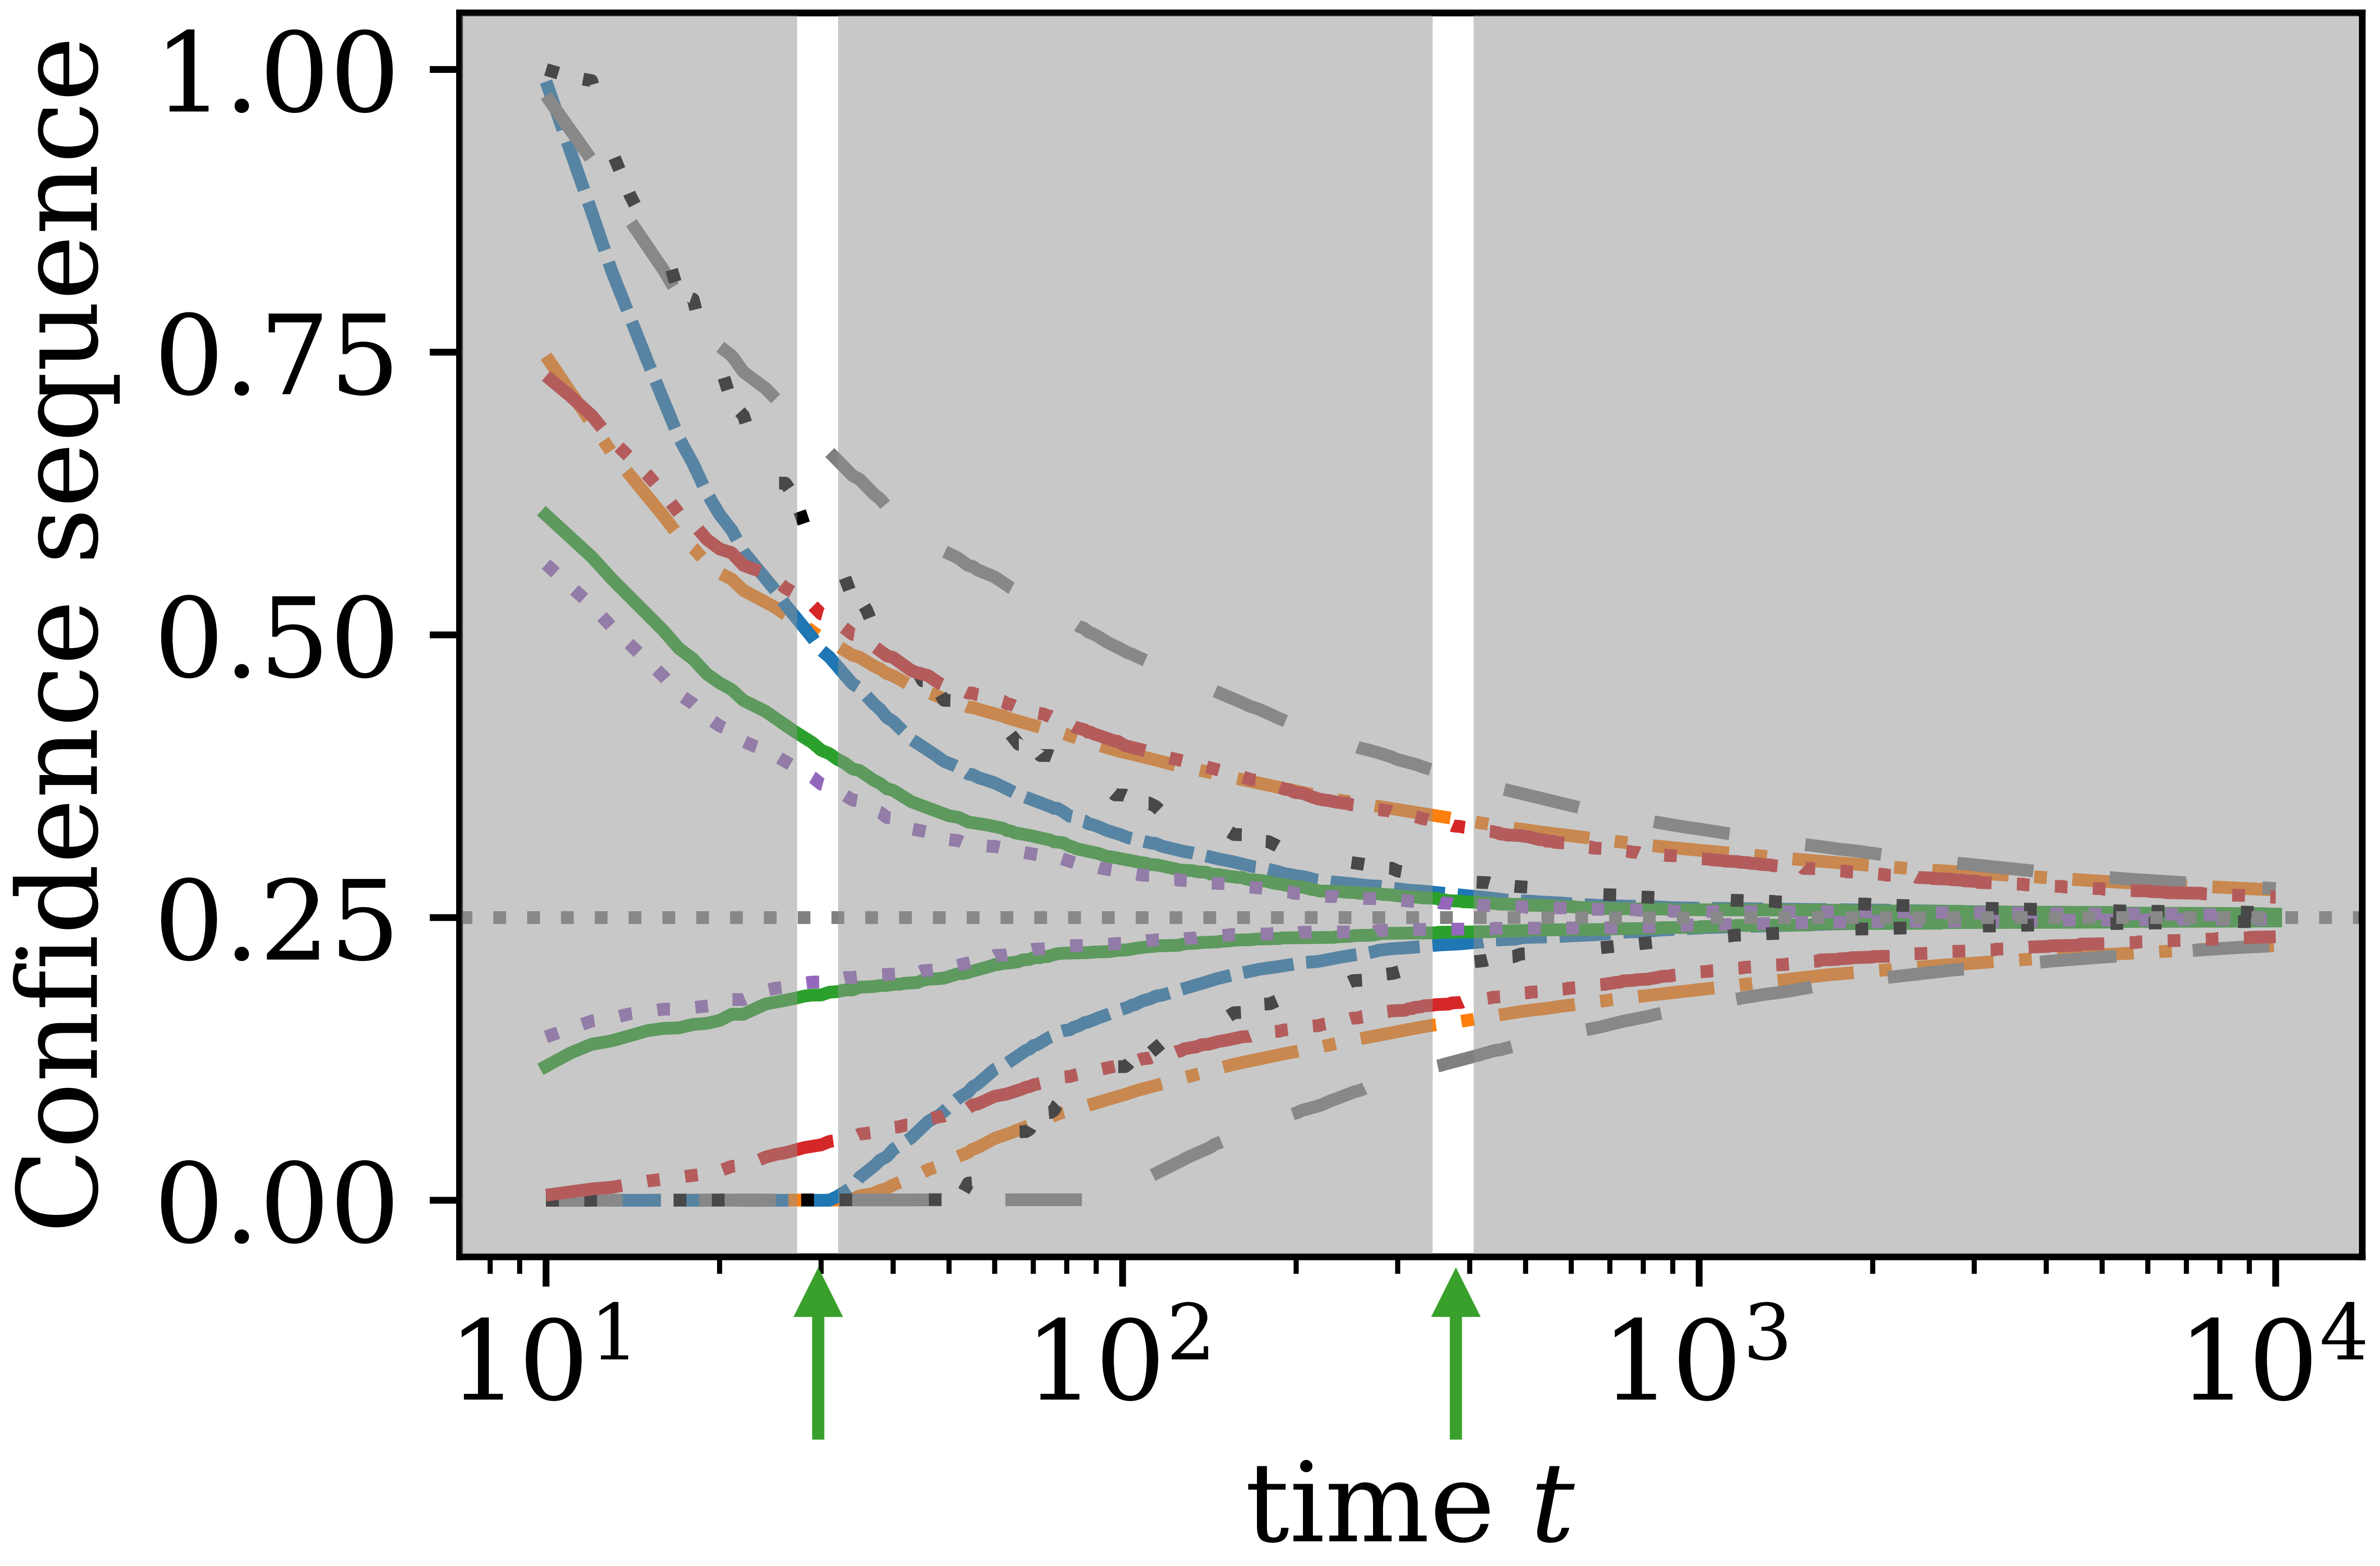

$$C_t := \left\{ m \in [0,1] : K_t(m) < \frac{1}{\alpha} \right\}$$

After some heuristic calculations, a good choice is

$$\lambda_i^+(m) := \sqrt{\frac{2\log(2/\alpha)}{n\hat{\sigma}_{i-1}^2}} \wedge \frac{1/2}{m}$$

$$\lambda_i^-(m) := \sqrt{\frac{2\log(2/\alpha)}{n\hat{\sigma}_{i-1}^2}} \wedge \frac{1/2}{1-m}$$

$$K_n^{\pm}(m) := \max\left\{ \frac{1}{2}\prod_{i=1}^{n}(1 + \lambda_i^+ \cdot (X_i - m)), \frac{1}{2}\prod_{i=1}^{n}(1 - \lambda_i^- \cdot (X_i - m)) \right\}$$

data-dependent tuning parameters without sample splitting!

So, given $X_1, \ldots, X_n$ bounded with mean μ,

$$C_n^{\pm} := \left\{ m \in [0,1] : K_n^{\pm}(m) < \frac{1}{\alpha} \right\}$$

is a sharp confidence interval for μ.

There's one more modification we can make to get **strictly** tighter confidence intervals!

If $C_1, C_2, \ldots, C_n, \ldots$ forms a $(1-\alpha)$-confidence sequence, then so does $\bigcap_{i \le n} C_i$.

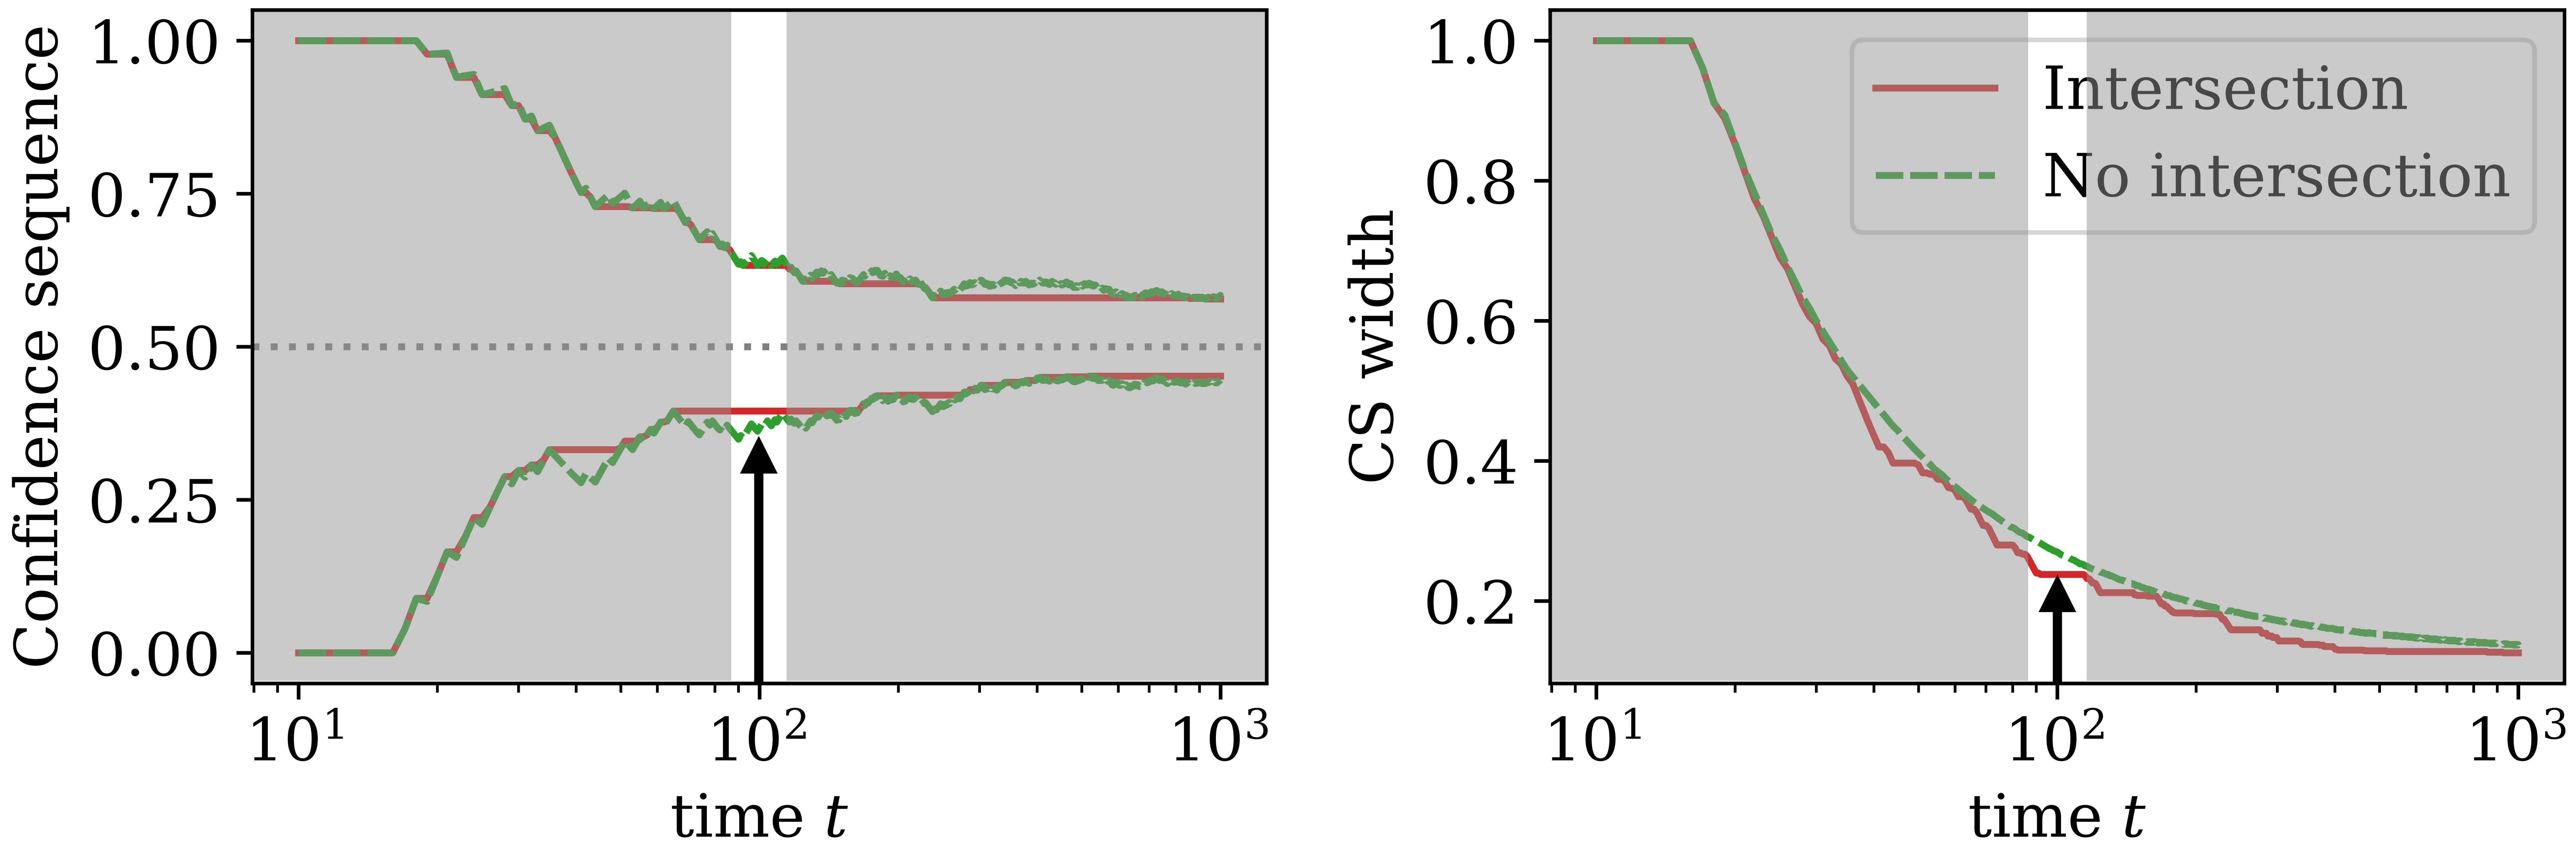

So, $\bigcap_{i \le n} C_i^{\pm}$ is a strict improvement over $C_n^{\pm}$ **for free.**

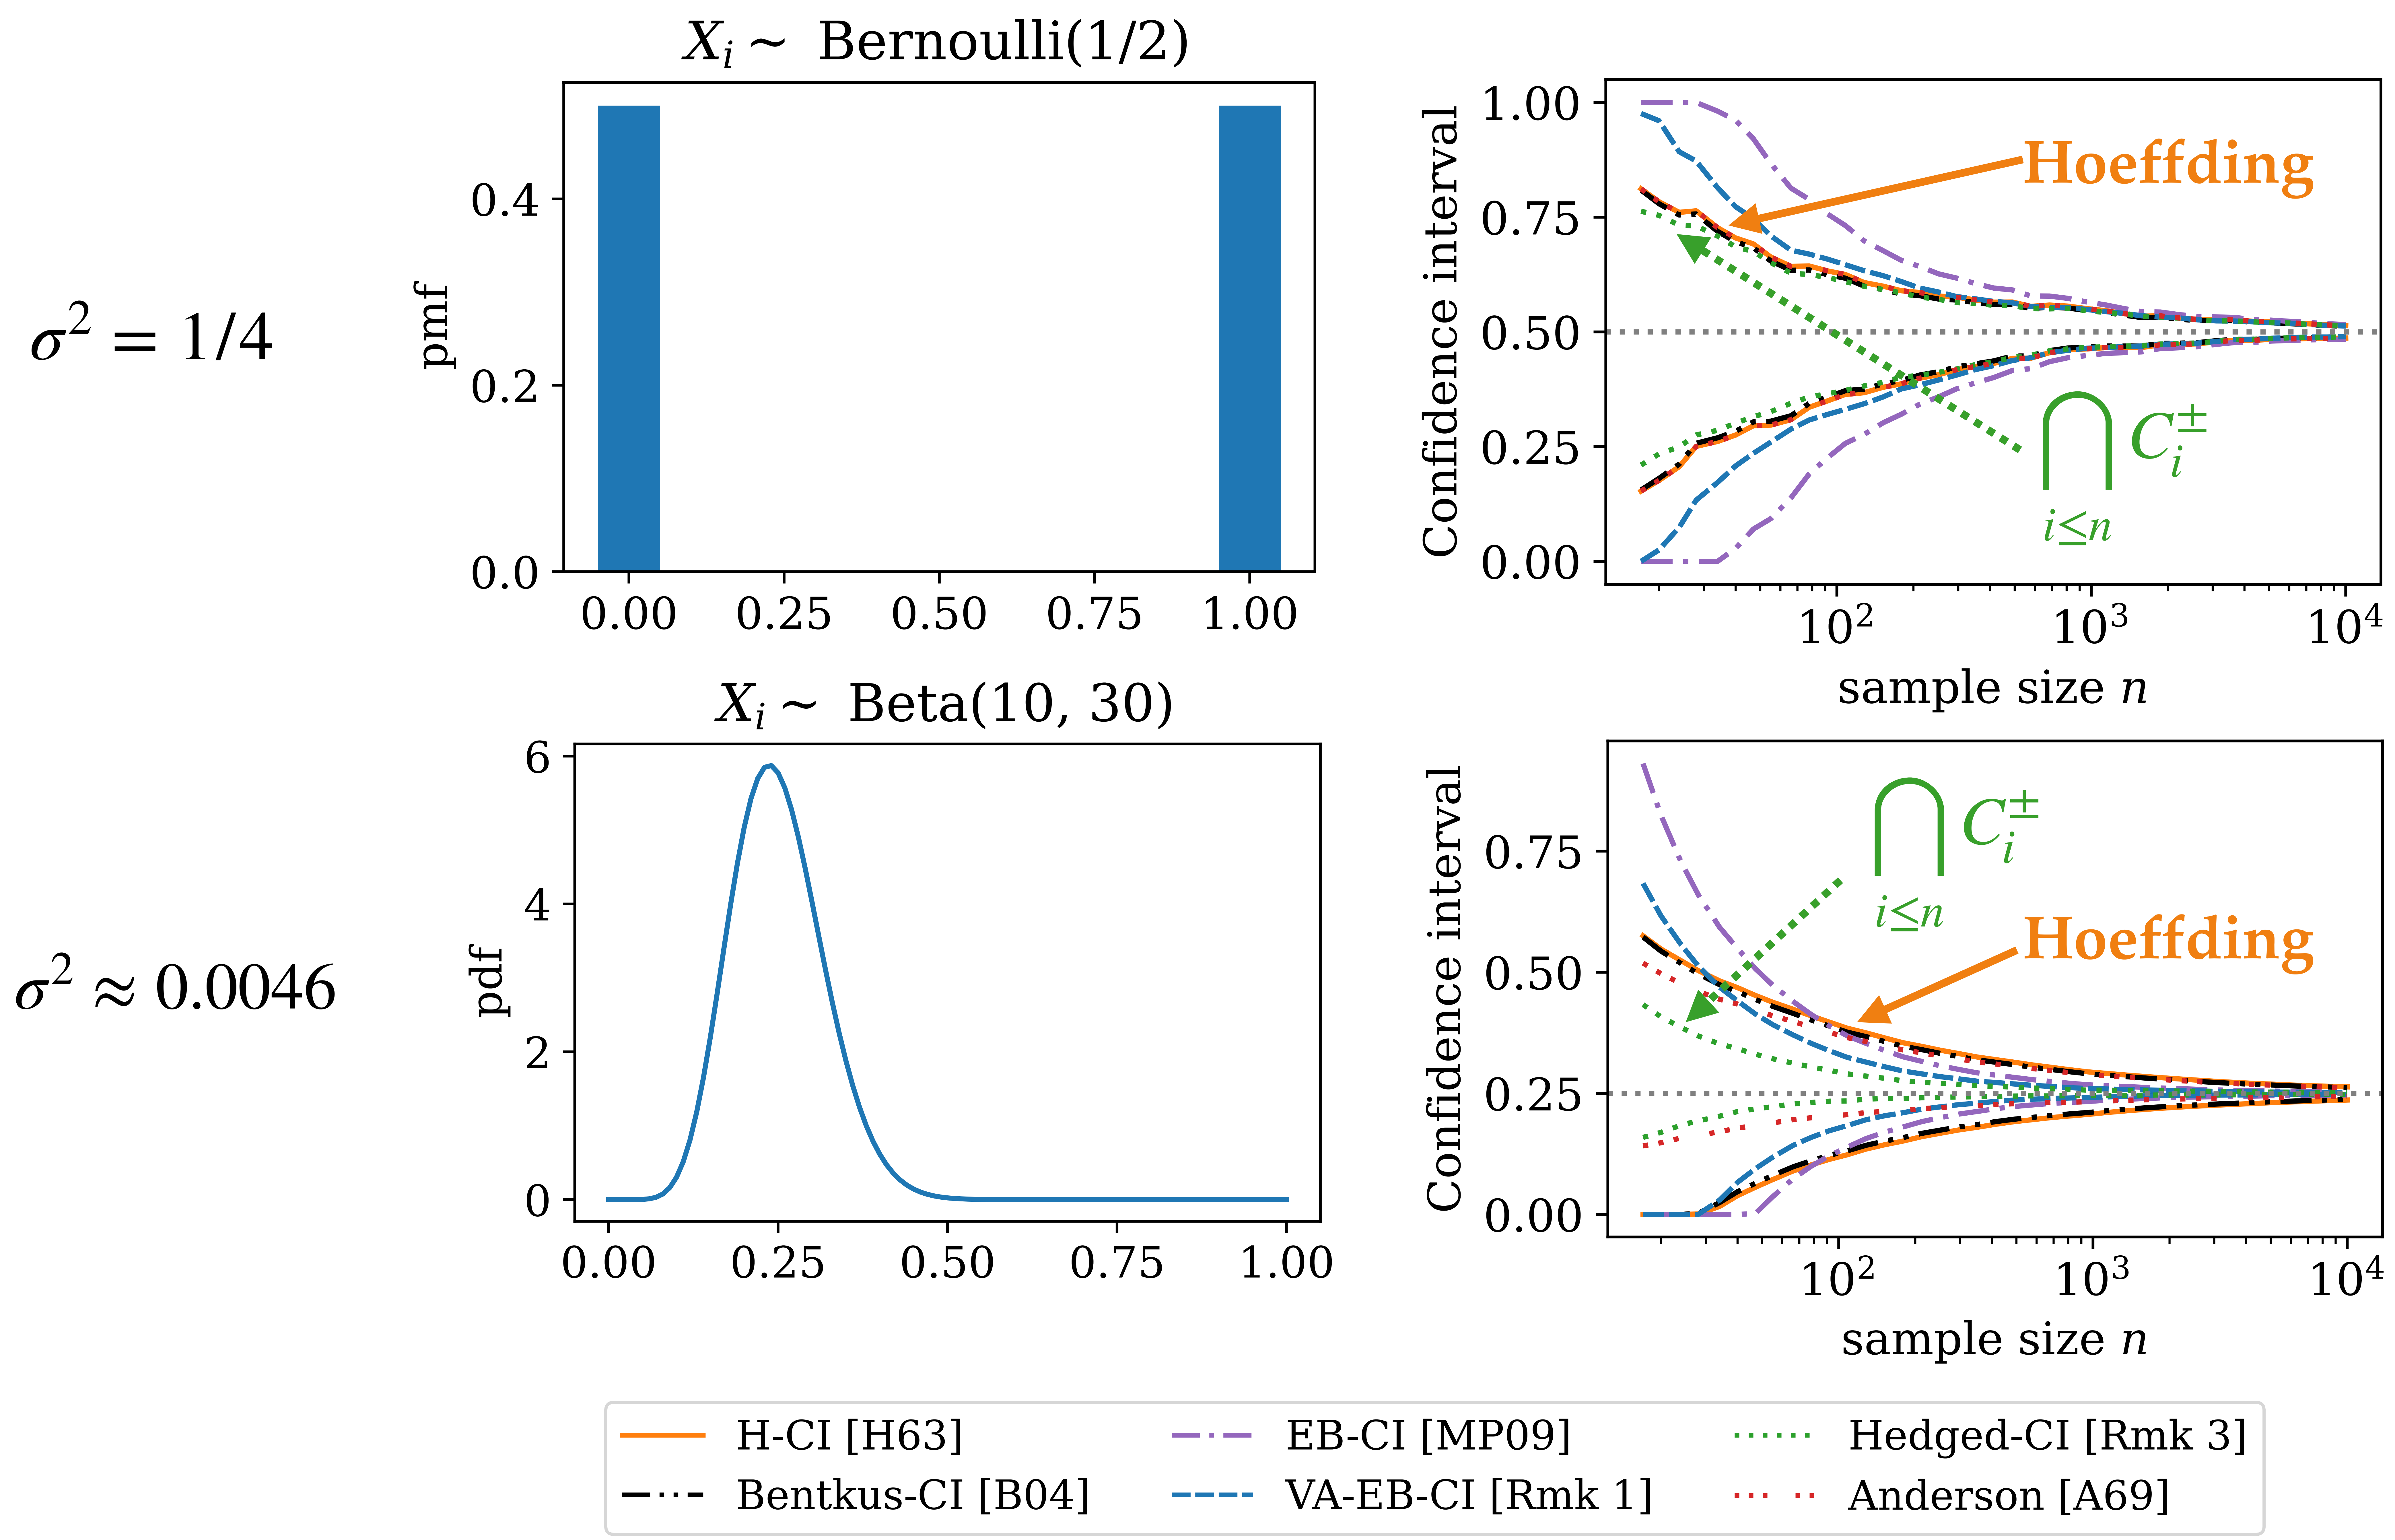

$\sigma^2 = 1/4$
$X_i \sim$ Bernoulli(1/2)
pmf
Confidence interval
sample size n
Hoeffding
$\bigcap_{i\le n} C_i^{\pm}$
$\sigma^2 \approx 0.0046$
$X_i \sim$ Beta(10, 30)
pdf
Confidence interval
sample size n
Hoeffding
$\bigcap_{i\le n} C_i^{\pm}$
H-CI [H63]
Bentkus-CI [B04]
EB-CI [MP09]
VA-EB-CI [Rmk 1]
Hedged-CI [Rmk 3]
Anderson [A69]

Summary

1. Developed nonparametric, nonasymptotic confidence sets for means of bounded random variables.
2. Valid at arbitrary stopping times, w/ no penalties for peeking at data early.
3. Substantially outperform prior work on this problem.

\+ Closed-form empirical Bernstein confidence sets and extensions to sampling without replacement in the full paper

Thank you.

ian.waudbysmith.com

stat.cmu.edu/~aramdas/

Confidence sets for sampling without replacement

Without-replacement (WoR) sampling:

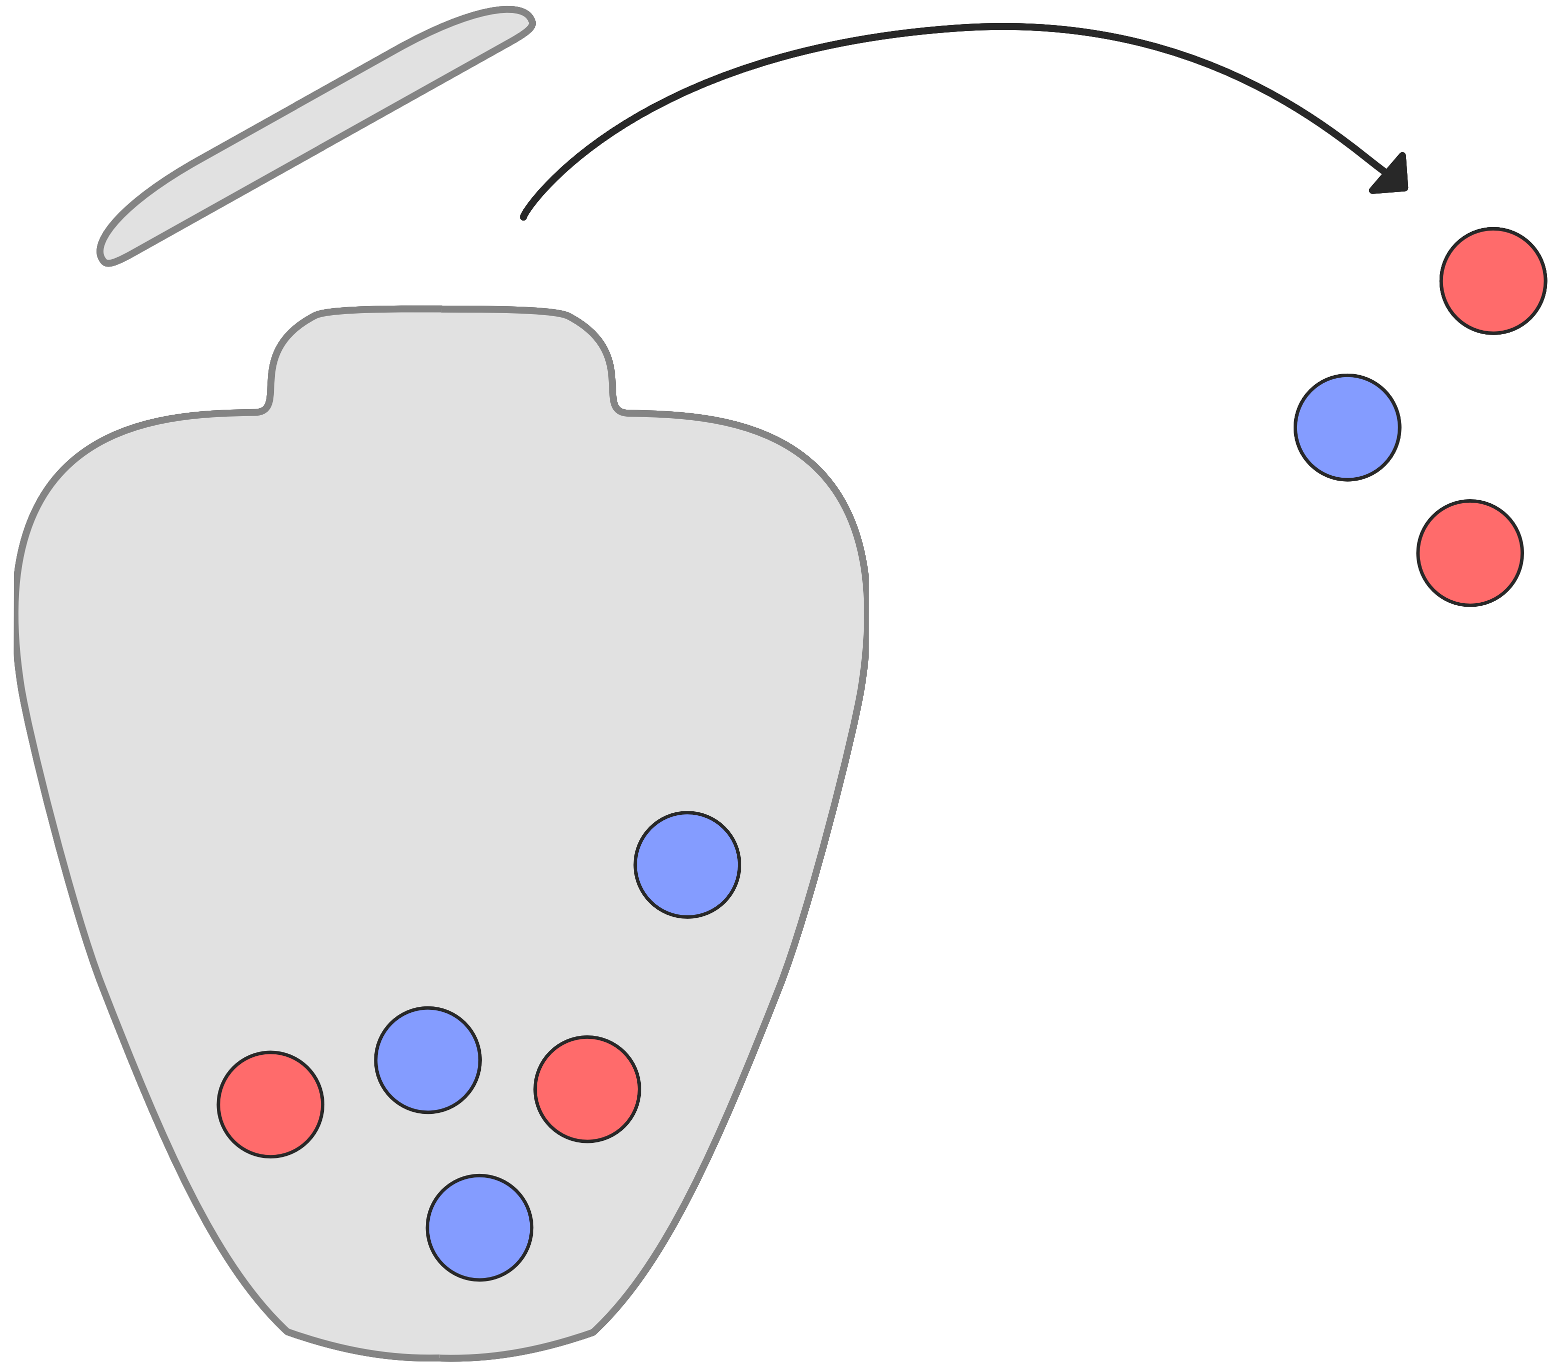

$$(x_1, \ldots, x_N) \in [0,1]^N, \ \mu := \frac{1}{N}\sum_{i=1}^{N} x_i$$

$$X_1 \sim \text{Unif}\left((x_1, \ldots, x_N)\right)$$

$$X_2 \sim \text{Unif}\left((x_1, \ldots, x_N) \setminus X_1\right)$$

$$\vdots$$

$$X_t \sim \text{Unif}\left((x_1, \ldots, x_N) \setminus X_1^{t-1}\right)$$

Want to estimate $\mu := \frac{1}{N} \sum_{i=1}^{N} x_i$

Goal: construct a game so that $\left(K_t(\mu)\right)_{t=0}^{N}$ is a martingale under WoR sampling.

$$X_t \sim \text{Unif}\left((x_1, \ldots, x_N) \backslash X_1^{t-1}\right) \implies \mathbb{E}(X_t \mid X_1^{t-1}) = \underbrace{\frac{N\mu - \sum_{i=1}^{t-1} X_i}{N - t + 1}}_{=: \, \mu_t^{\text{WoR}}}$$

Consider a "candidate mean" $m \in [0,1]$

$K_0 \leftarrow \$1$

For $t = 1,2,3,\ldots$:

Gambler chooses bet $\lambda_t \in (-1/(1-m_t^{\text{WoR}}),\ 1/m_t^{\text{WoR}})$

Observe X_t

$K_t \leftarrow K_{t-1} + K_{t-1} \cdot \lambda_t \cdot (X_t - m_t^{\text{WoR}})$

EndFor

$$m_t^{\text{WoR}} = \frac{Nm - \sum_{i=1}^{t-1} X_i}{N - t + 1}$$

Consider a "candidate mean" $m \in [0,1]$

$K_0 \leftarrow \$1$

For $t = 1,2,3,\ldots$:

Gambler chooses bet $\lambda_t \in (-1/(1-m),\ 1/m)$

Observe X_t

$K_t \leftarrow K_{t-1} \ + \ K_{t-1} \cdot \lambda_t \cdot (X_t - m)$

EndFor

Consider a "candidate mean" $m \in [0,1]$

$K_0 \leftarrow \$1$

For $t = 1,2,3,\ldots$:

Gambler chooses bet $\lambda_t \in (-1/(1-m_t^{\mathrm{WoR}}),\ 1/m_t^{\mathrm{WoR}})$

Observe X_t

$K_t \leftarrow K_{t-1} + K_{t-1} \cdot \lambda_t \cdot (X_t - m_t^{\mathrm{WoR}})$

EndFor

$$m_t^{\mathrm{WoR}} = \frac{Nm - \sum_{i=1}^{t-1} X_i}{N - t + 1}$$

Then,

$$K_t^{\text{WoR}}(\mu) := \prod_{i=1}^{t} (1 + \lambda_i \cdot (X_i - \mu_t^{\text{WoR}}))$$

forms a nonnegative martingale, and

$$C_t^{\text{WoR}} := \left\{ m \in [0,1] : K_t^{\text{WoR}}(m) < \frac{1}{\alpha} \right\}$$

forms a $(1-\alpha)$-confidence sequence.

Confidence sequences for sampling WoR

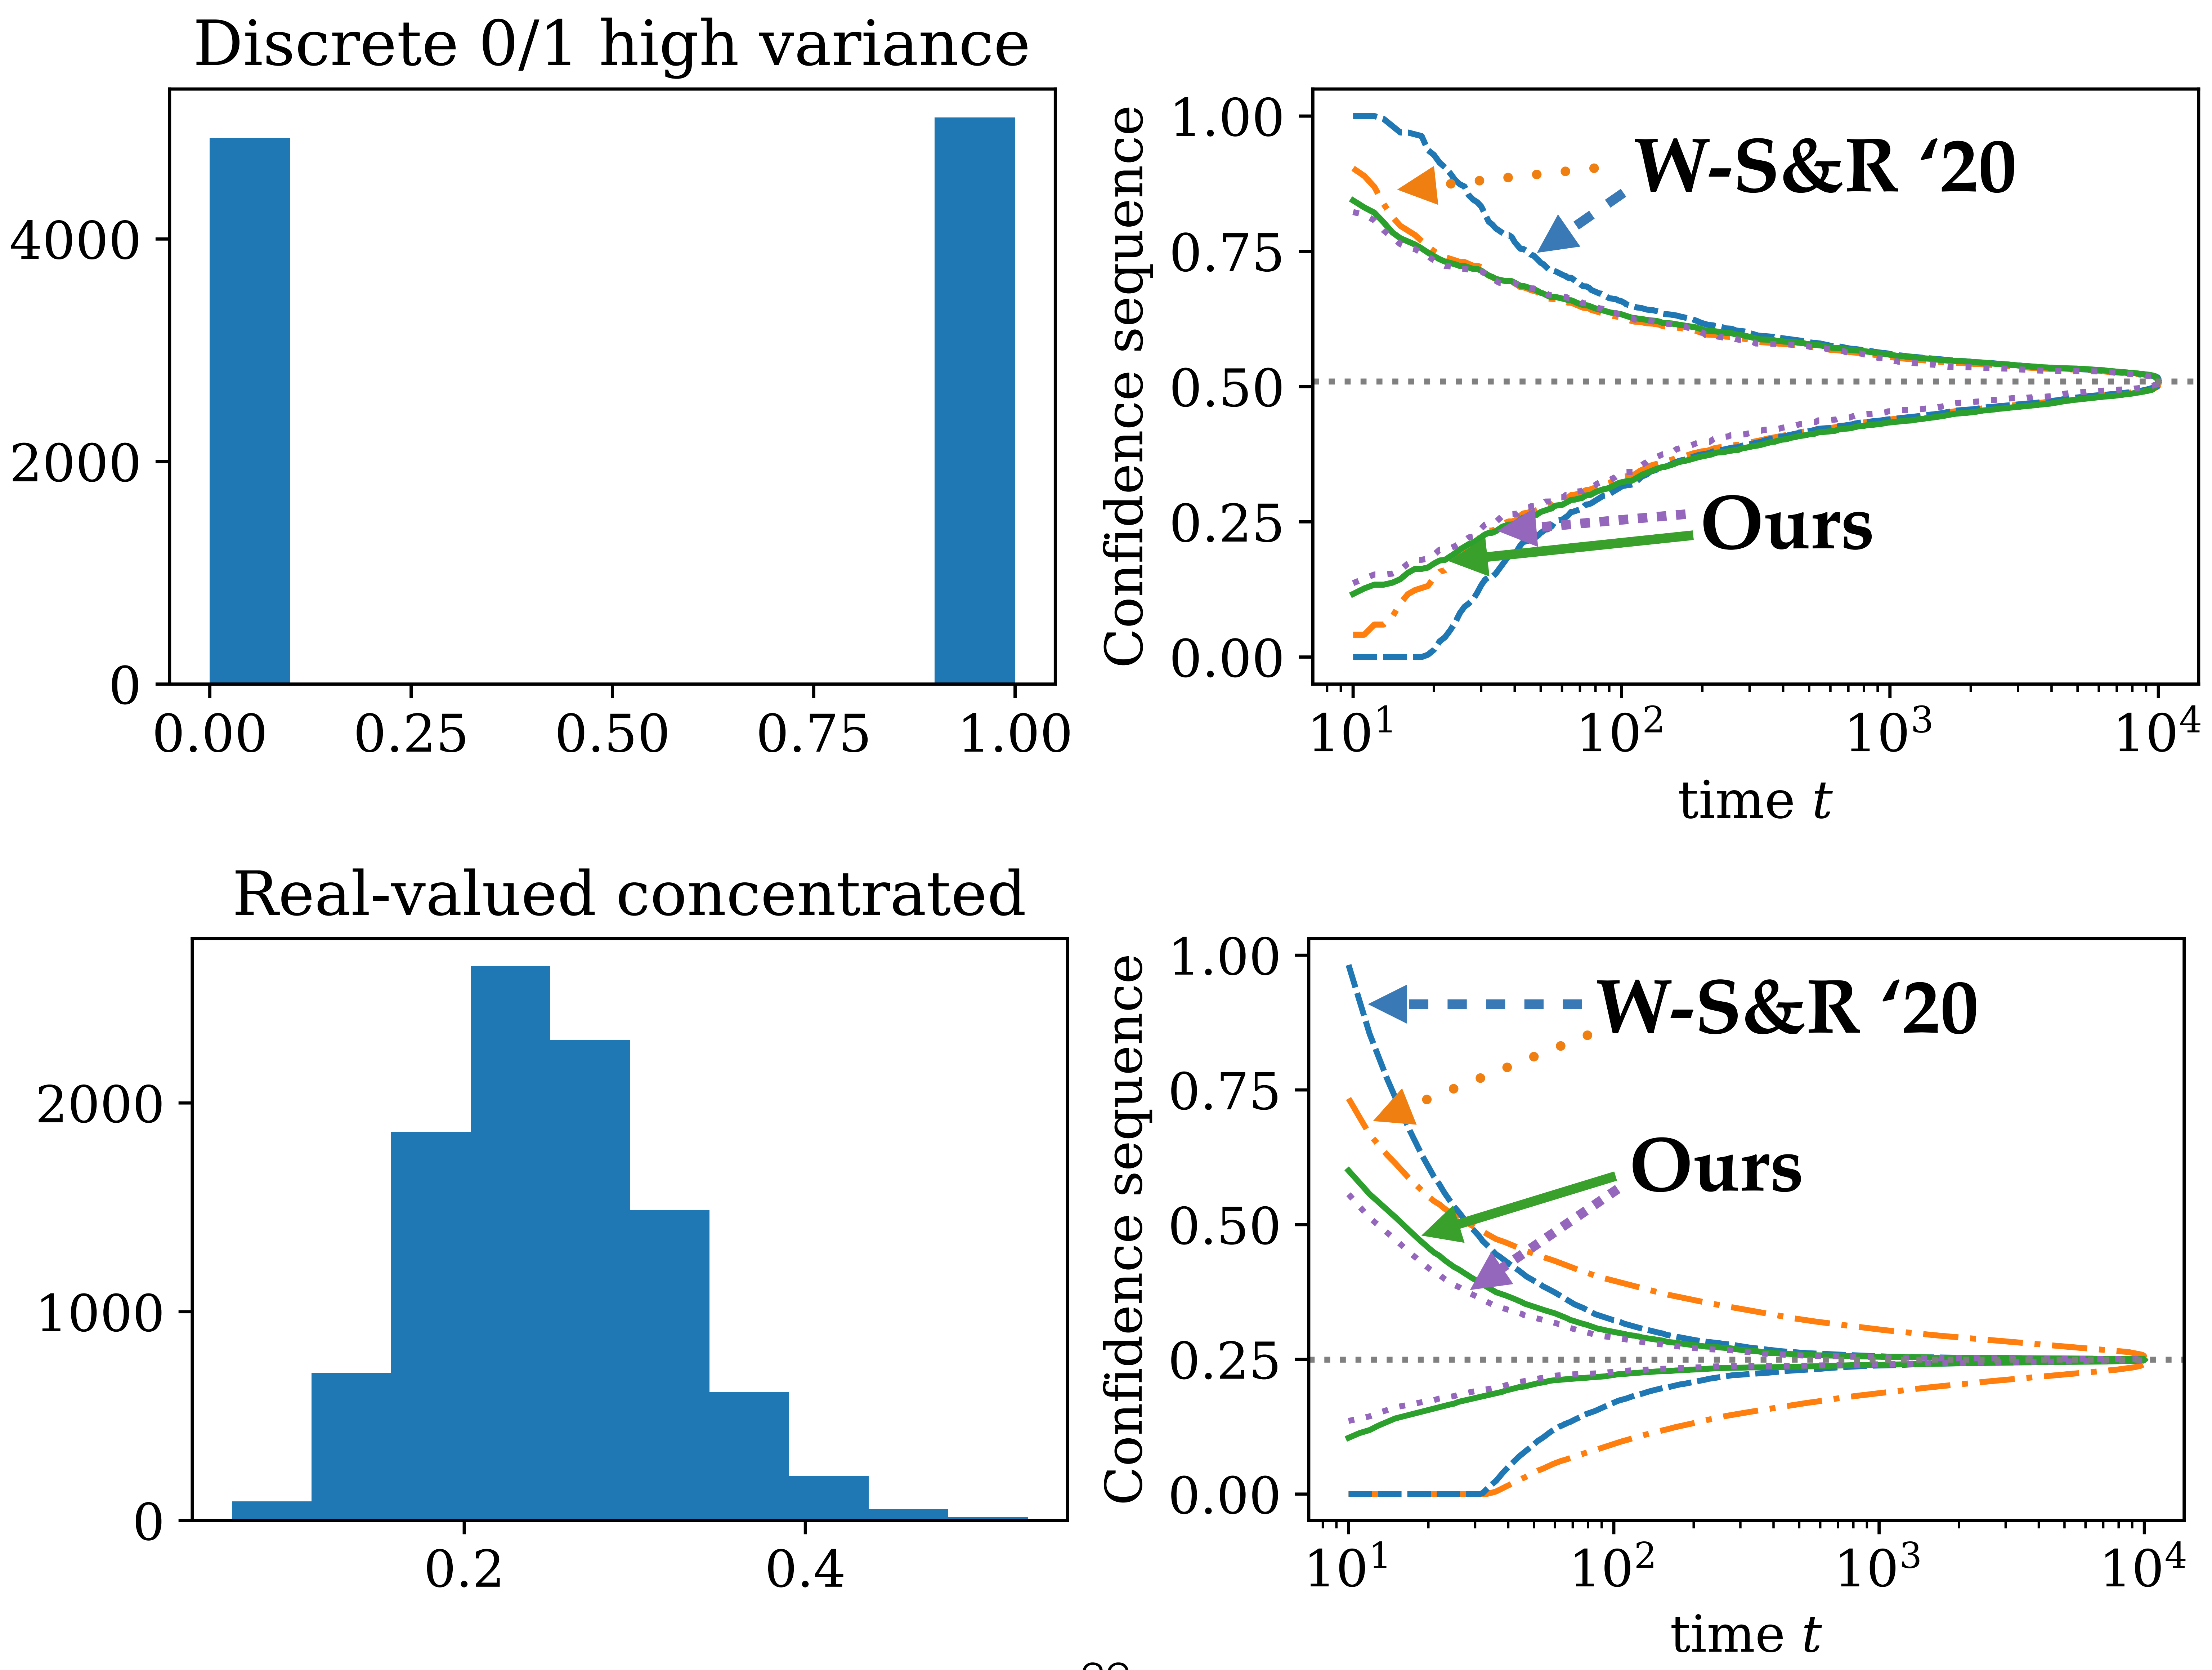

Confidence intervals for sampling WoR

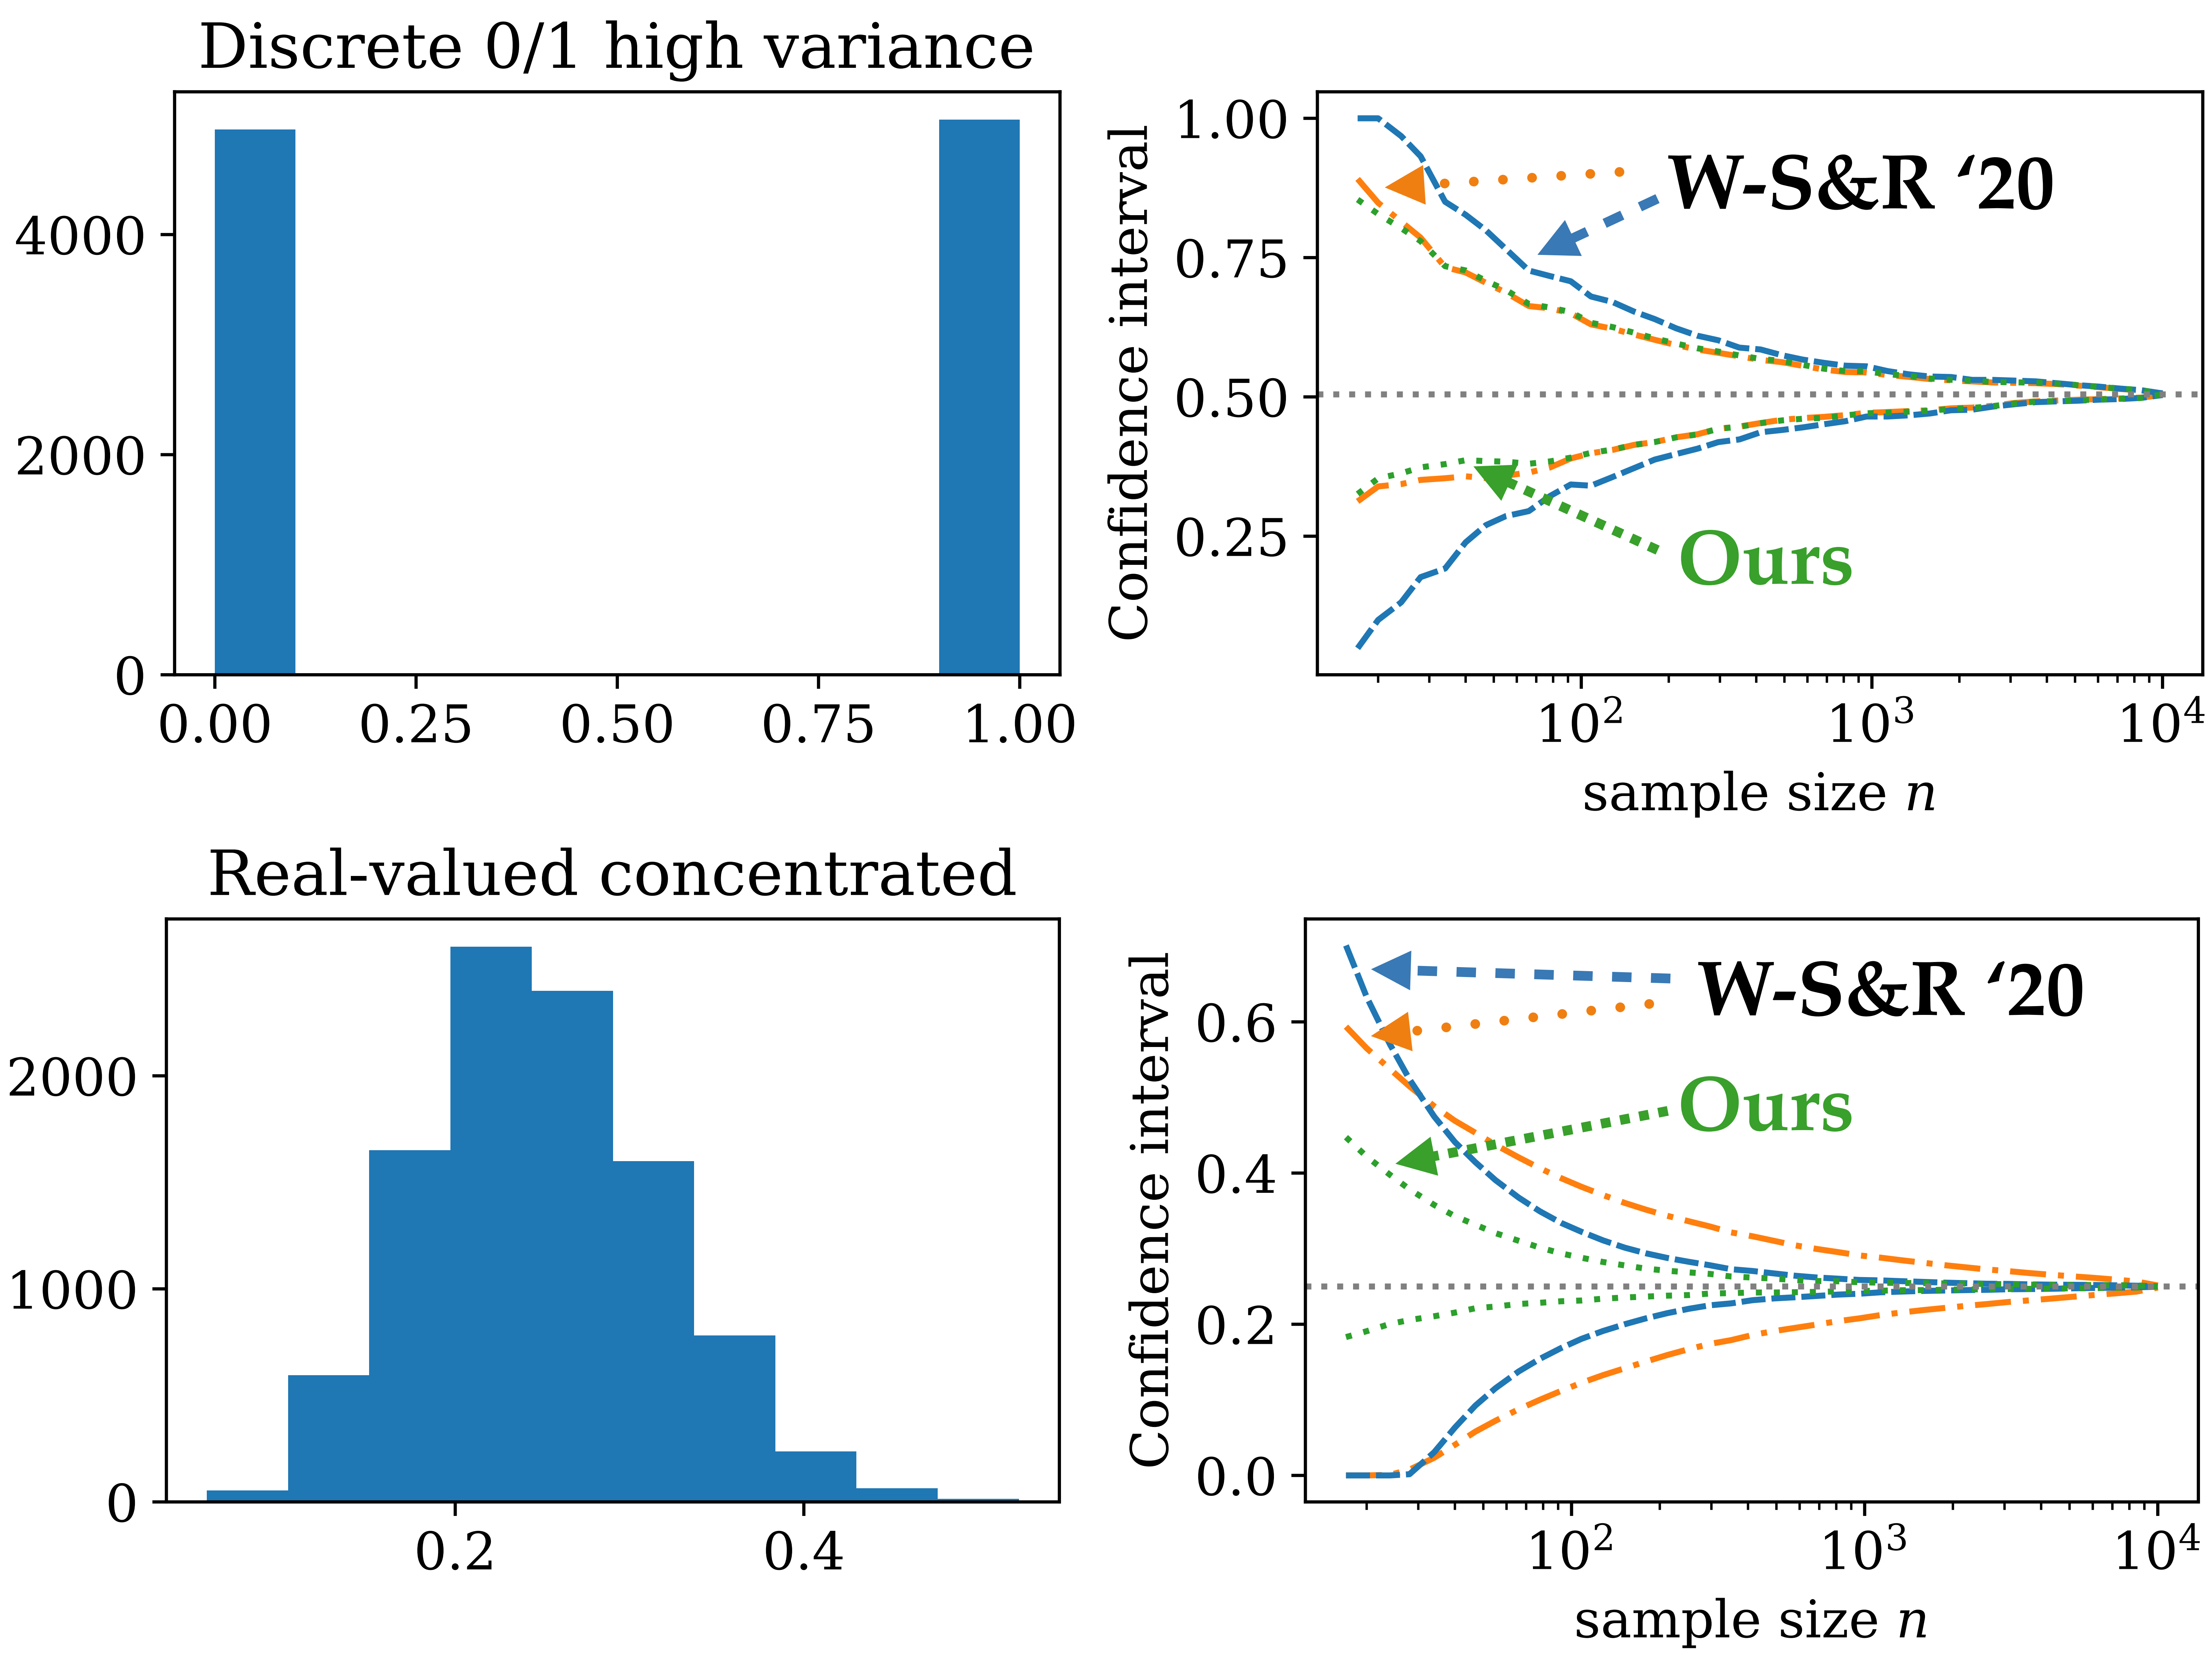

Closed-form empirical Bernstein confidence sequences & confidence intervals

$$C_t := \left\{ m \in [0,1] : \prod_{i=1}^{t} (1 + \lambda_i(m) \cdot (X_i - m) \right\}$$

While C_t is easy to compute, it is not closed-form.

However,

$$C_t^{\text{PMEB}} := \left\{ m \in [0,1] : \prod_{i=1}^{t} \exp \left\{ \lambda_i(X_i - m) - 4(X_i - \widehat{\mu}_{i-1})^2 \psi_E(\lambda_i) \right\} \right\}$$ **is!**

$$C_t^{\mathrm{PMEB}} := \left(\frac{\sum_{i=1}^t \lambda_i X_i}{\sum_{i=1}^t \lambda_i} \pm \frac{\log(2/\alpha) + 4\sum_{i=1}^t (X_i - \widehat{\mu}_{i-1})^2 \psi_E(\lambda_i)}{\sum_{i=1}^t \lambda_i} \right)$$

where $$\psi_E(\lambda) := -(\log(1-\lambda) - \lambda)/4,$$

$$\lambda_t := \sqrt{\frac{2\log(2/\alpha)}{\widehat{\sigma}_{t-1} t \log(1+t)}} \wedge \frac{1}{2},$$

$$\widehat{\sigma}_t^2 := \frac{1/4 + \sum_{i=1}^t (X_i - \widehat{\mu}_i^2)^2}{t+1}, \quad \text{and} \quad \widehat{\mu}_t := \frac{1/2 + \sum_{i=1}^t X_i}{t+1}.$$

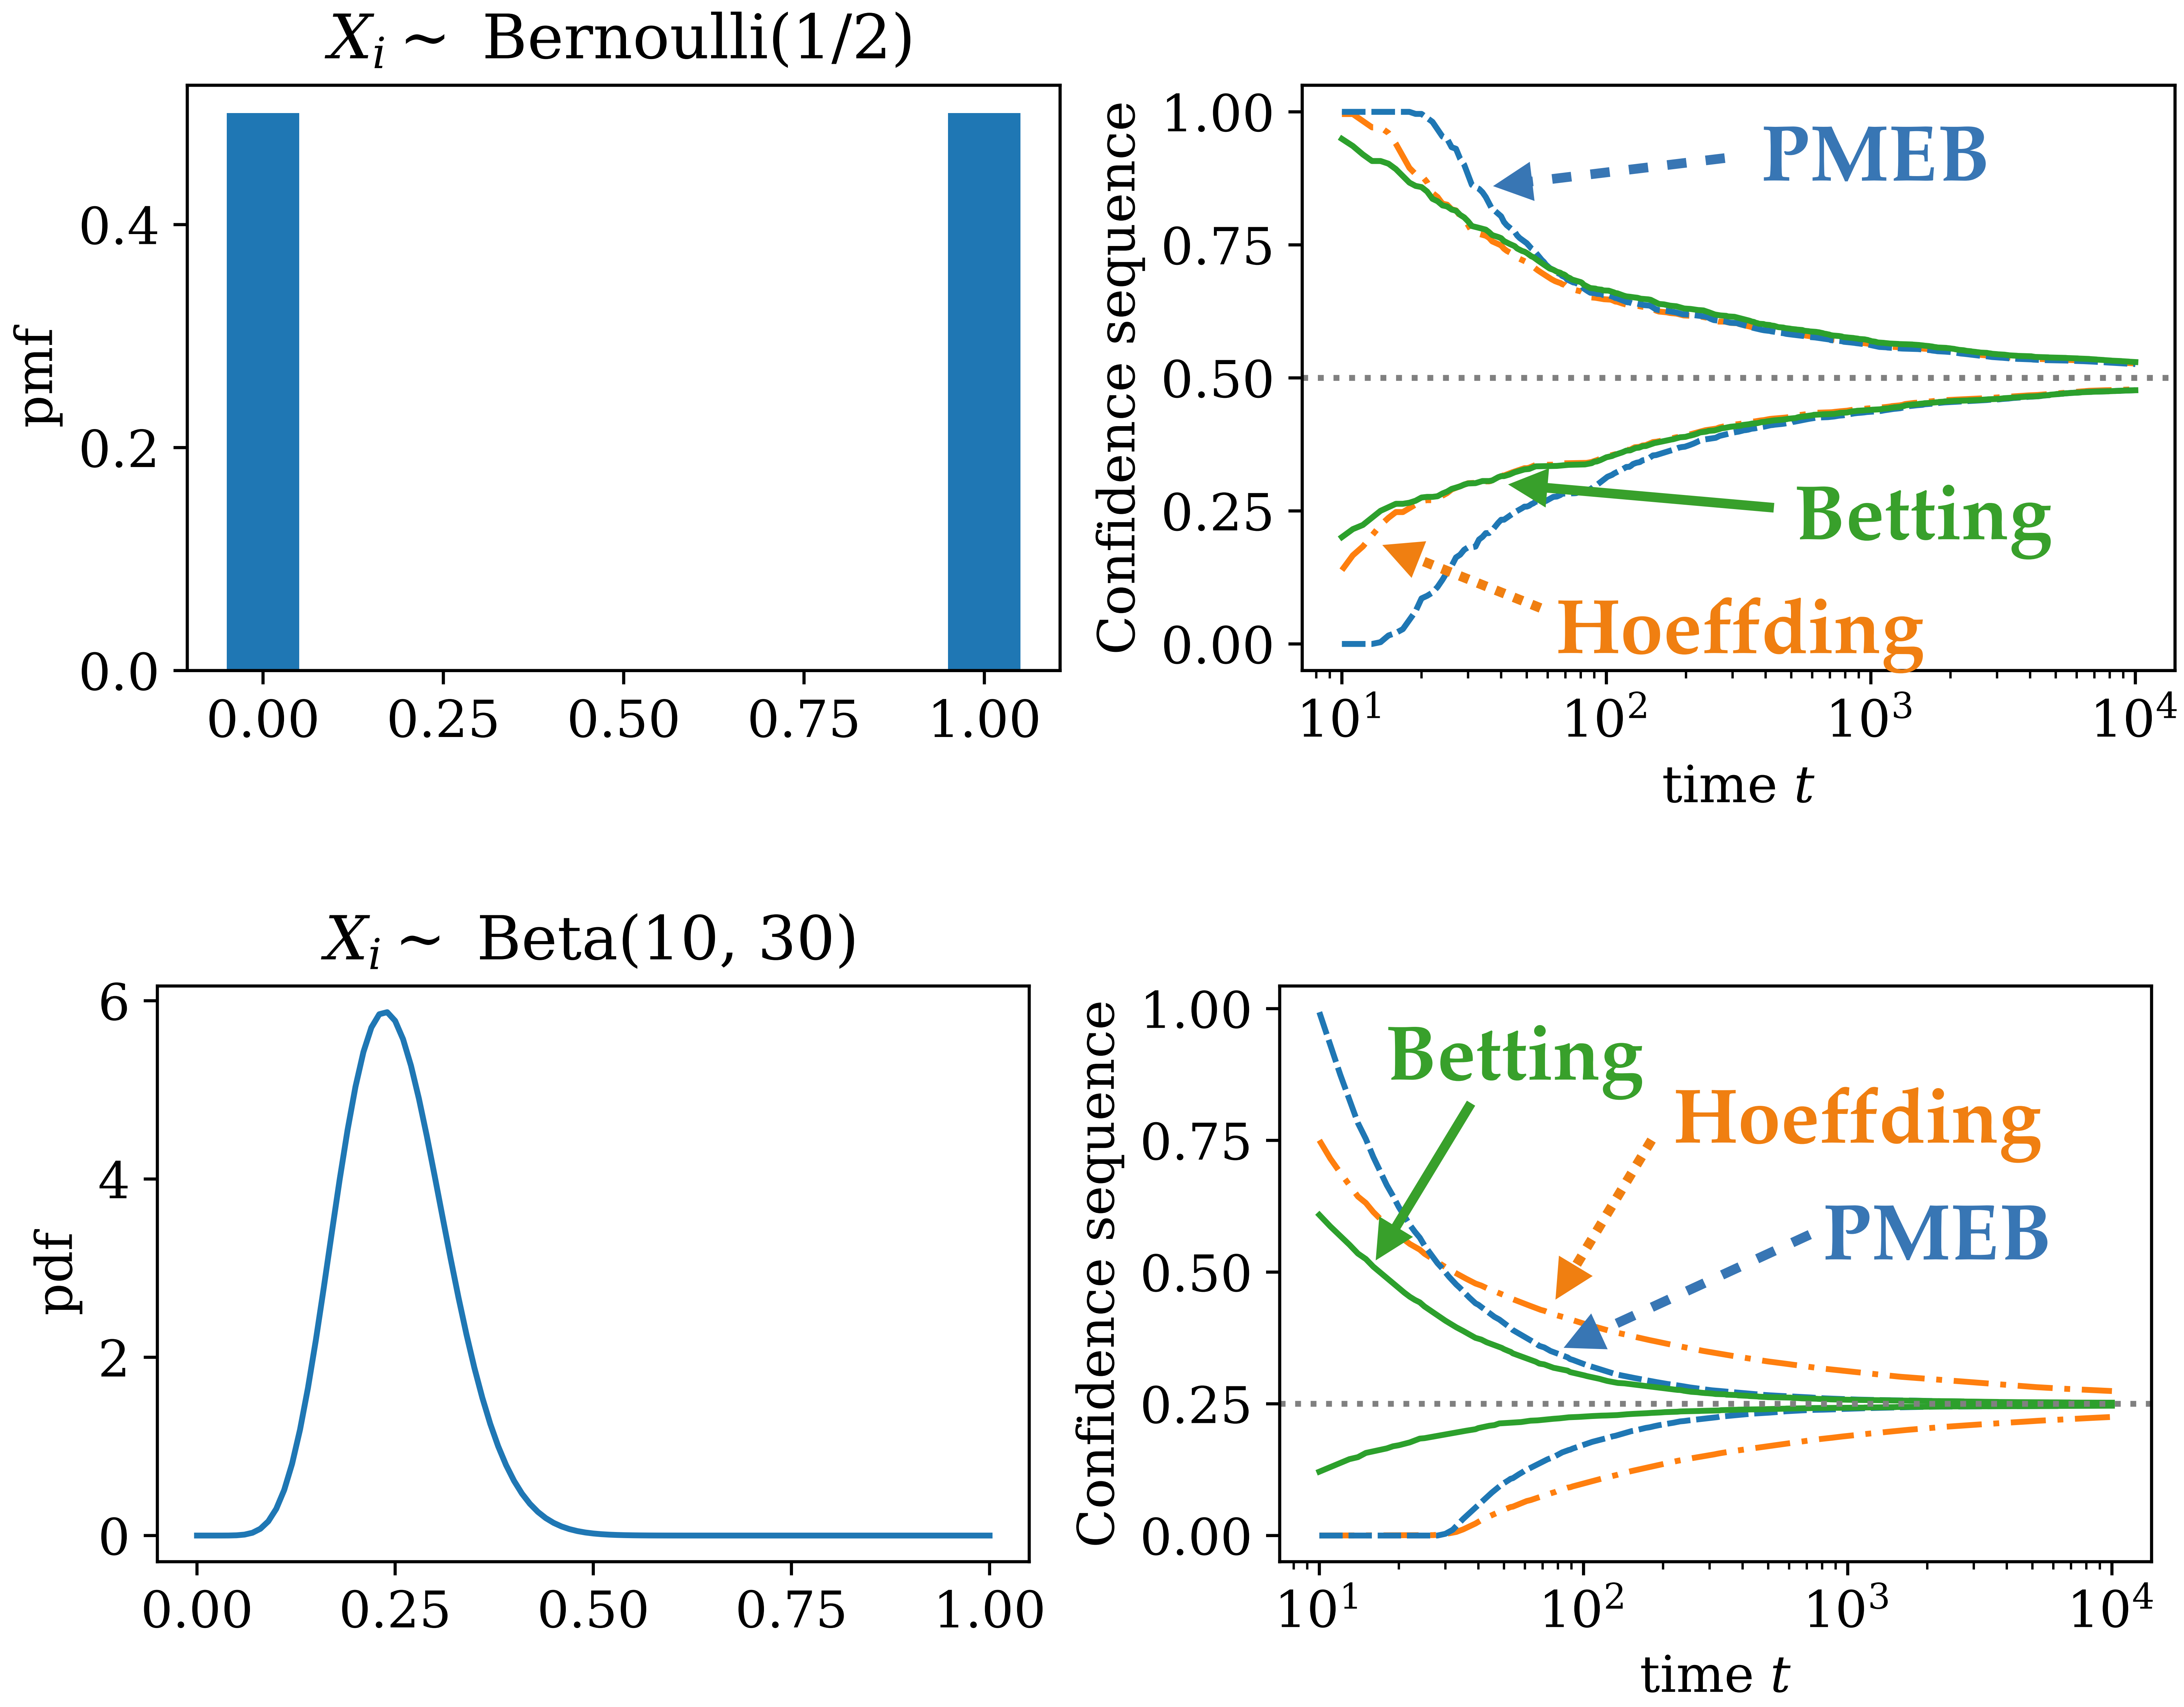

$X_i \sim$ Bernoulli(1/2)
pmf
Confidence sequence
PMEB
Betting
Hoeffding
time t
$X_i \sim$ Beta(10, 30)
pdf
Confidence sequence
Betting
Hoeffding
PMEB
time t

Similarly for fixed-time confidence intervals:

$$C_n^{\text{PMEB}} := \left(\frac{\sum_{i=1}^n \lambda_i X_i}{\sum_{i=1}^n \lambda_i} \pm \frac{\log(2/\alpha) + 4\sum_{i=1}^n (X_i - \widehat{\mu}_{i-1})^2 \psi_E(\lambda_i)}{\sum_{i=1}^n \lambda_i} \right)$$

but here, $\lambda_i := \sqrt{\frac{2\log(2/\alpha)}{n\widehat{\sigma}_{i-1}}} \wedge \frac{1}{2},$

Final bound: $\bigcap_{i \le n} C_i^{\text{PMEB}}$

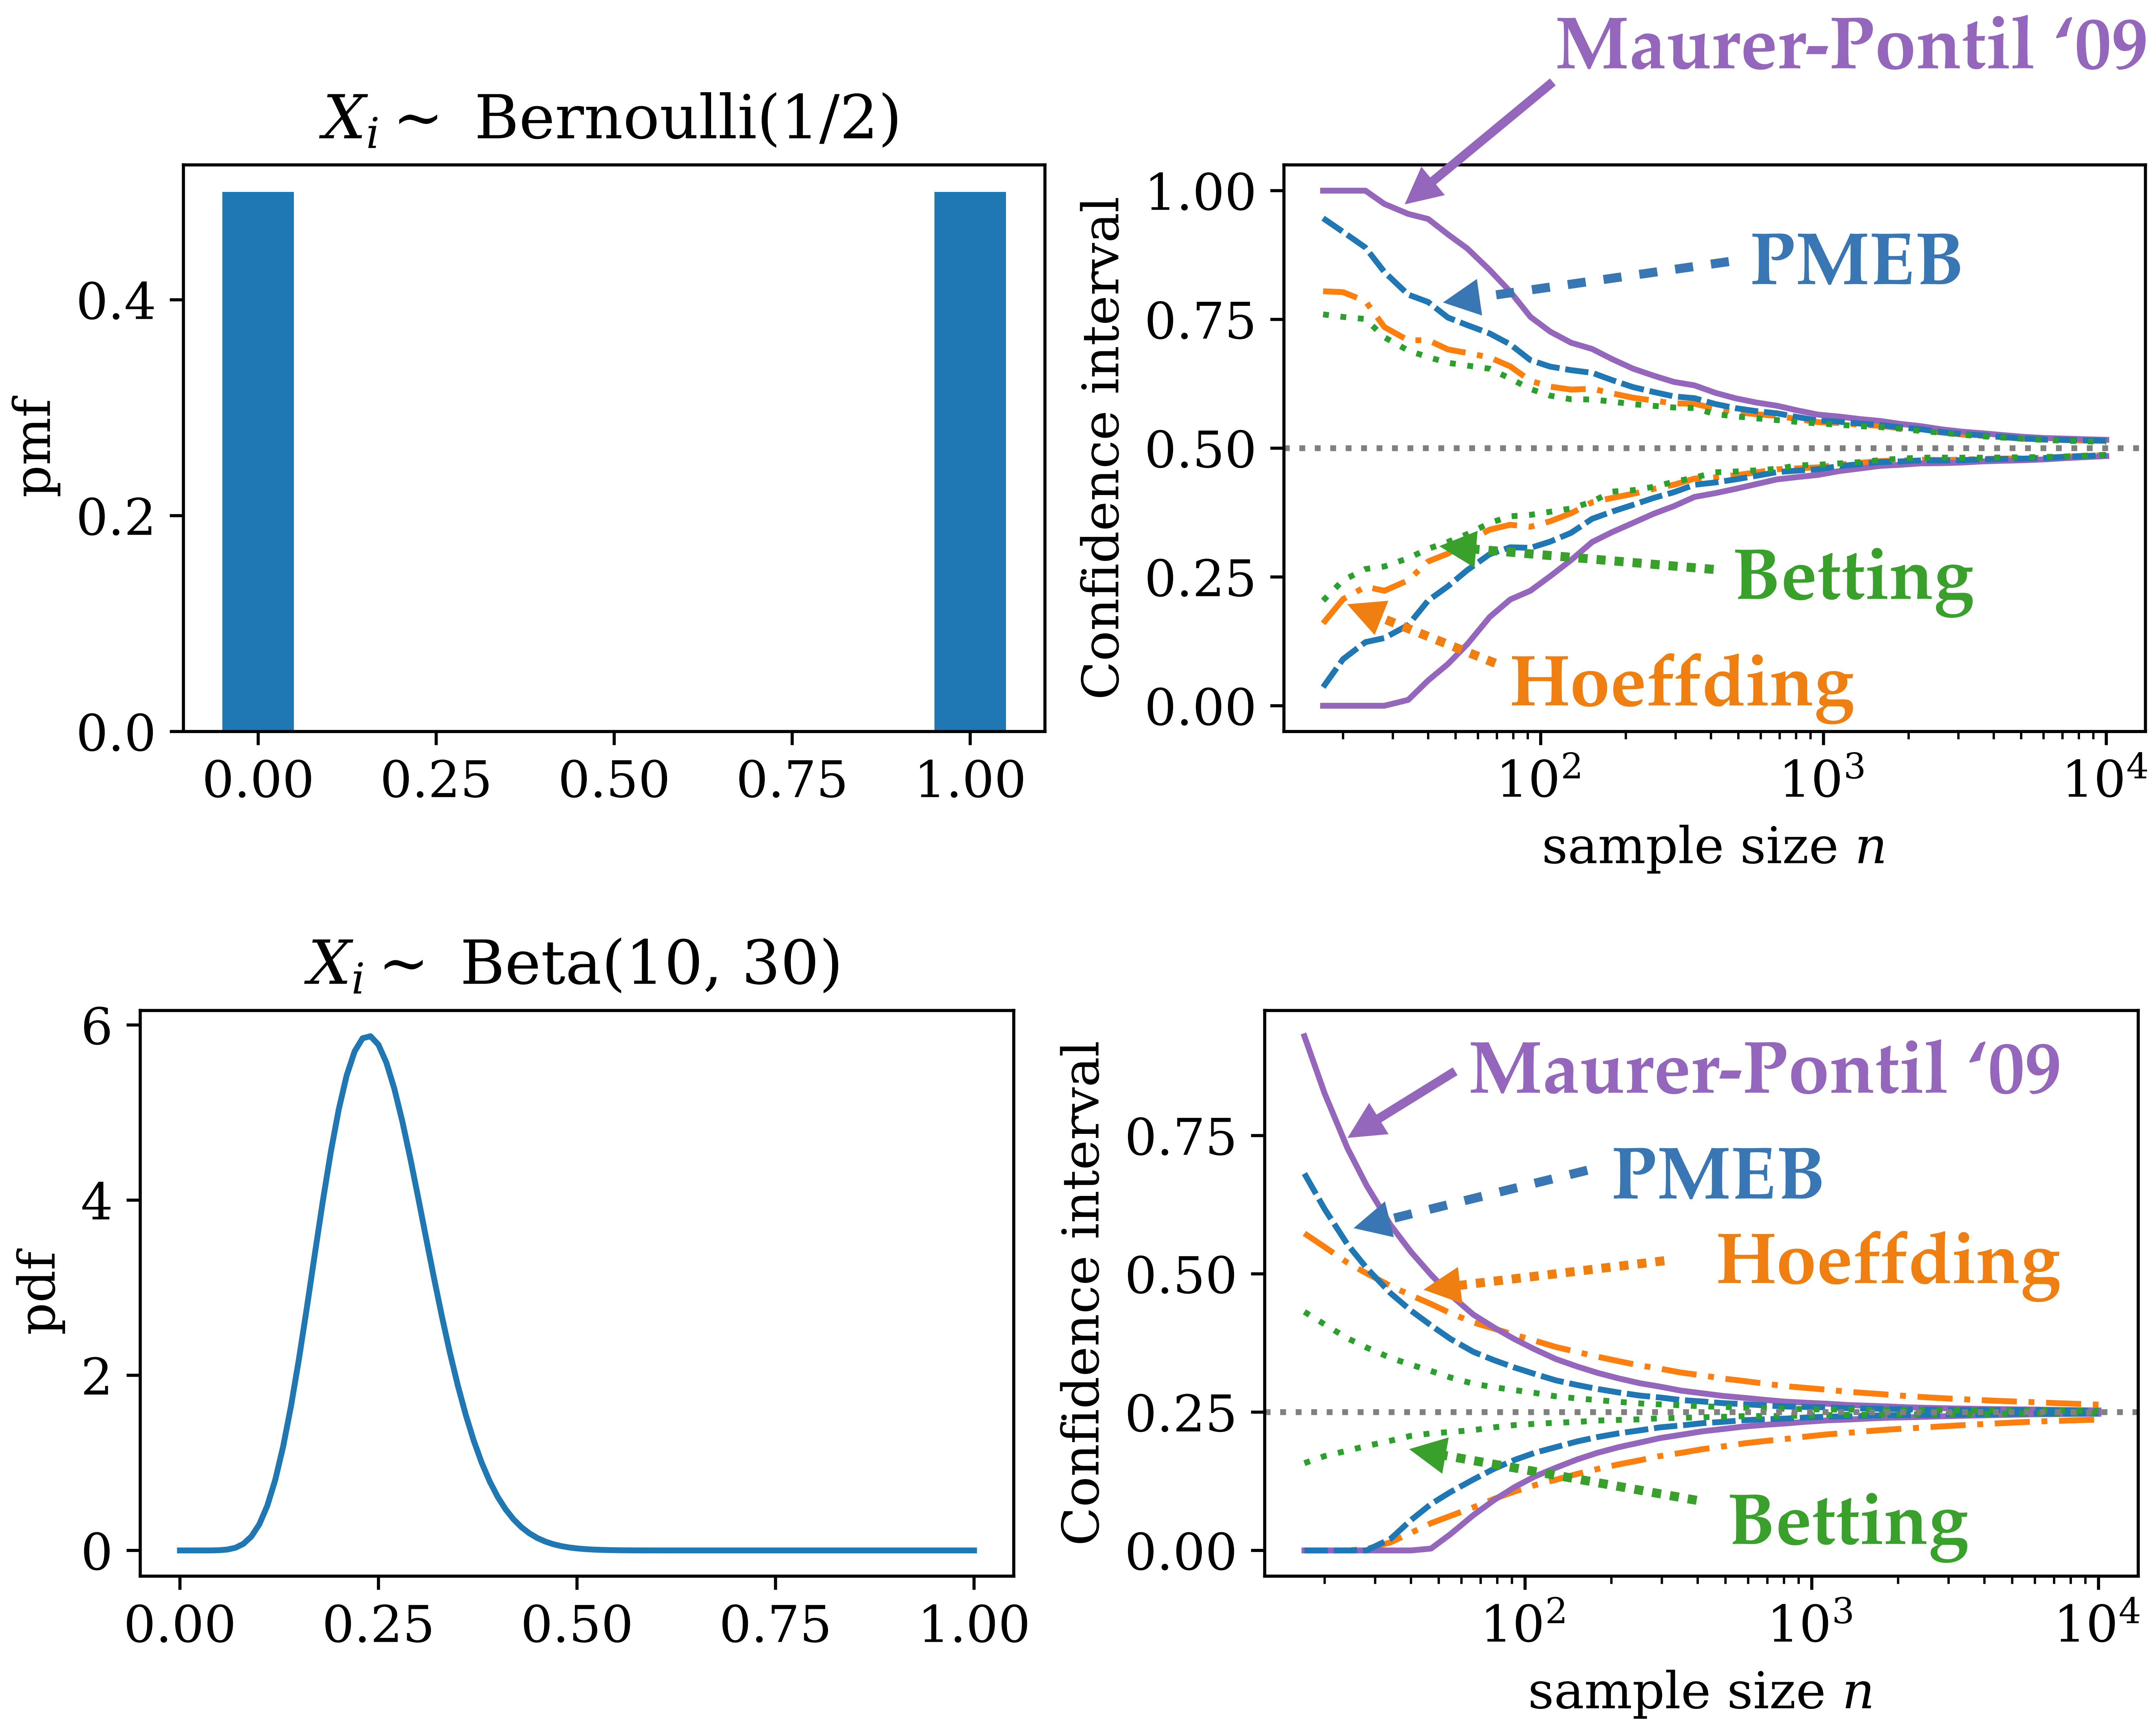

$X_i \sim$ Bernoulli(1/2)
pmf
Maurer-Pontil '09
PMEB
Betting
Hoeffding
Confidence interval
sample size n
$X_i \sim$ Beta(10, 30)
pdf
Maurer-Pontil '09
PMEB
Hoeffding
Betting
Confidence interval
sample size n

Choice of $(\lambda_t^+)_{t=1}^n$ and $(\lambda_t^-)_{t=1}^n$ for fixed-time confidence intervals

Why does $\lambda_i^+(m) := \sqrt{\frac{2\log(2/\alpha)}{n\widehat{\sigma}_{i-1}^2}} \wedge \frac{1}{m}$ perform so well?

$$K_n^+(\mu) := \prod_{i=1}^{n} \left(1 + \lambda \cdot (X_i - \mu)\right)$$

$$\gtrsim \underbrace{\prod_{i=1}^{n} \exp\left\{\lambda \cdot (X_i - \mu) - (X_i - \widehat{\mu}_{i-1})^2 \lambda^2/2\right\}}$$

$$\implies \text{Width}_n := \frac{\log(2/\alpha) + \sum_{i=1}^{n} (X_i - \widehat{\mu}_{i-1})^2/2}{t\lambda}$$

$$\approx \frac{\log(2/\alpha) + n\sigma^2/2}{t\lambda}$$

$$\underset{\lambda}{\operatorname{argmin}}\ \mathrm{Width}_n = \sqrt{\frac{2\log(2/\alpha)}{n\sigma^2}}$$

$$\lambda_t^+(m) := \sqrt{\frac{2\log(2/\alpha)}{n\,\widehat{\sigma}_{t-1}^2}} \wedge \frac{1}{m} \qquad \lambda_t^-(m) := \sqrt{\frac{2\log(2/\alpha)}{n\,\widehat{\sigma}_{t-1}^2}} \wedge \frac{1}{1-m}$$

Brief selective history of betting ideas