Confidence sequences for sequential experimentation and best-arm identification

Steve Howard
Joint work with Aaditya Ramdas, Jon McAuliffe, and Jasjeet Sekhon

September 5, 2019

Sequential monitoring of experiment results is problematic.

Click on the button

The percentage of visitors who clicked on a tracked element.

Variation #3 is beating Original by +58.0%.

VARIATIONS	VISITORS	CONVERSIONS	CONVERSION RATE	IMPROVEMENT	CHANCE TO BEAT BASELINE
Variation #3	970	32	3.3% (±1.12%)	+58.0%	95.2%
Original BASELINE	1,006	21	2.1% (±0.88%)	---	---
Variation #1	999	11	1.1% (±0.65%)	-47.3%	3.9%
Variation #2	1,027	11	1.1% (±0.63%)	-48.7%	3.3%

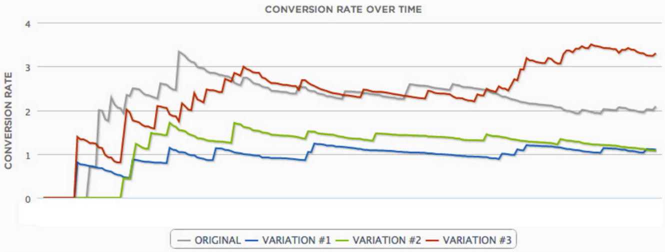

One path of p-values from a fair coin

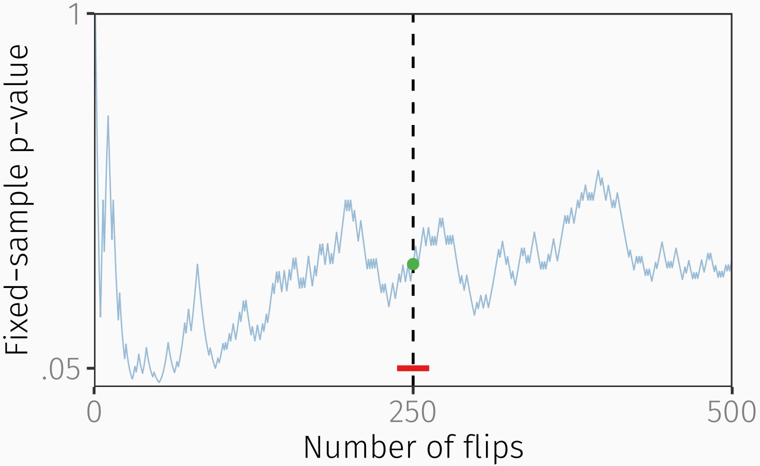

Let’s look at many such paths…

With no bias, we only rarely conclude the coin is biased.

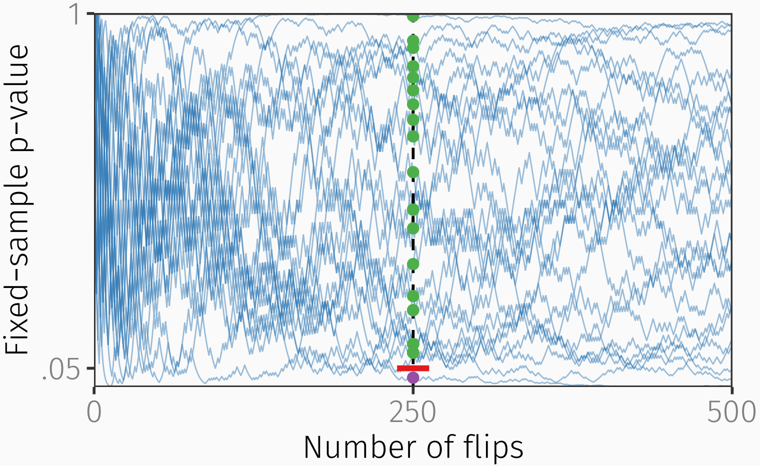

Just one out of 25 p-values is below 0.05.

Continuous monitoring of fixed-sample *p*-values inflates the false positive rate.

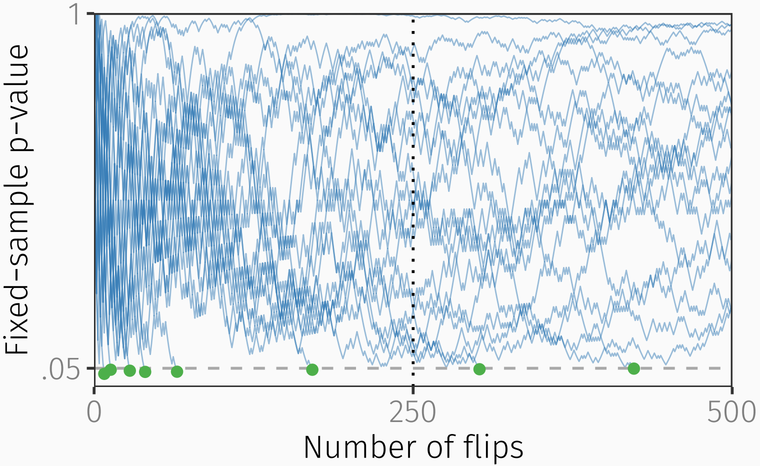

Here, with a fair coin, eight out of 25 paths reach significance.

The false positive rate grows arbitrarily large with enough flips.

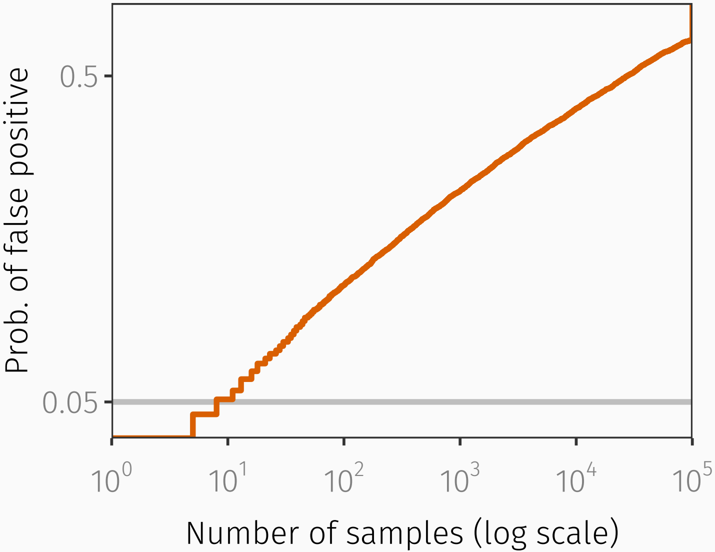

The promise and peril of sequential monitoring

Sequential monitoring can substantially inflate false positive rates.

The promise and peril of sequential monitoring

Sequential monitoring can substantially inflate false positive rates.

But there are good reasons to do it:

- If the treatment effect is stronger than expected, we can stop early.

The promise and peril of sequential monitoring

Sequential monitoring can substantially inflate false positive rates.

But there are good reasons to do it:

- If the treatment effect is stronger than expected, we can stop early.
- If the treatment effect is weaker than expected (or the budget has increased), we can extend the experiment.

The promise and peril of sequential monitoring

Sequential monitoring can substantially inflate false positive rates.

But there are good reasons to do it:

- If the treatment effect is stronger than expected, we can stop early.
- If the treatment effect is weaker than expected (or the budget has increased), we can extend the experiment.

How should we enable sequential monitoring without inflating false positive rates?

Confidence sequences solve the problem of continuous monitoring.

A **confidence sequence** for $(\theta_t)_{t=1}^{\infty}$ is a sequence of intervals $(\mathrm{CI}_t)_{t=1}^{\infty}$ satisfying

$$\mathbb{P}(\theta_t \in \mathrm{CI}_t \text{ for all } t \in \mathbb{N}) \geq 1 - \alpha.$$

[Darling and Robbins 1967, Lai 1984, Jennison and Turnbull 1989, Johari et al. 2015, H. et al. 2018]

Confidence sequences solve the problem of continuous monitoring.

A **confidence sequence** for $(\theta_t)_{t=1}^{\infty}$ is a sequence of intervals $(\mathrm{CI}_t)_{t=1}^{\infty}$ satisfying

$$\mathbb{P}(\theta_t \in \mathrm{CI}_t \text{ for all } t \in \mathbb{N}) \geq 1 - \alpha.$$

[Darling and Robbins 1967, Lai 1984, Jennison and Turnbull 1989, Johari et al. 2015, H. et al. 2018]

Much stronger than the fixed-sample guarantee:

$$\text{For all } t \in \mathbb{N}, \mathbb{P}(\theta_t \in \mathrm{CI}_t) \geq 1 - \alpha.$$

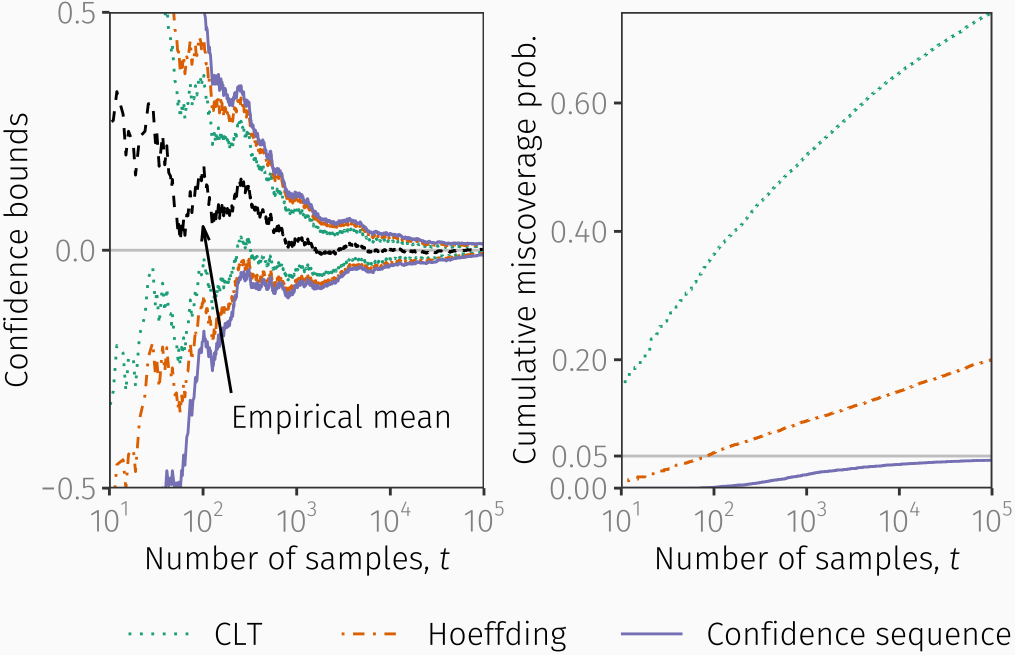
Confidence bounds
0.5
0.0
−0.5
Empirical mean
Number of samples, t
10^1 10^2 10^3 10^4 10^5
Cumulative miscoverage prob.
0.60
0.40
0.20
0.05
0.00
CLT
Hoeffding
Confidence sequence

Outline

Some key results

Application to best-arm identification

A taste of the underlying framework

Some key results

Empirical-Bernstein confidence sequence

Theorem (H., Ramdas, McAuliffe, Sekhon 2018)

Suppose X_i are independent and $[a, b]$-valued for all i. Let $\widehat{X}_i$ be any predictable sequence and u_α be any sub-exponential uniform boundary with scale $b - a$. Then

$$\mathbb{P}\left(\left|\bar{X}_t - \mathbb{E}\bar{X}_t\right| < \frac{u_\alpha\left(\sum_{i=1}^t (X_i - \widehat{X}_i)^2\right)}{t} \text{ for all } t \in \mathbb{N} \right) \geq 1 - 2\alpha.$$

Empirical-Bernstein confidence sequence

Theorem (H., Ramdas, McAuliffe, Sekhon 2018)

Suppose X_i are independent and $[a, b]$-valued for all i. Let $\widehat{X}_i$ be any predictable sequence and u_α be any sub-exponential uniform boundary with scale $b - a$. Then

$$\mathbb{P}\left(\left|\bar{X}_t - \mathbb{E}\bar{X}_t\right| < \frac{u_\alpha\left(\sum_{i=1}^t (X_i - \widehat{X}_i)^2\right)}{t} \text{ for all } t \in \mathbb{N} \right) \geq 1 - 2\alpha.$$

Empirical-Bernstein confidence sequence

Theorem (H., Ramdas, McAuliffe, Sekhon 2018)

Suppose X_i are independent and $[a, b]$-valued for all i. Let $\widehat{X}_i$ be any predictable sequence and u_α be any sub-exponential uniform boundary with scale $b - a$. Then

$$\mathbb{P}\left(|\bar{X}_t - \mathbb{E}\bar{X}_t| < \frac{u_\alpha\left(\sum_{i=1}^t (X_i - \widehat{X}_i)^2\right)}{t} \text{ for all } t \in \mathbb{N} \right) \geq 1 - 2\alpha.$$

Empirical-Bernstein confidence sequence

Theorem (H., Ramdas, McAuliffe, Sekhon 2018)

Suppose X_i are independent and $[a, b]$-valued for all i. Let $\widehat{X}_i$ be any predictable sequence and u_α be any sub-exponential uniform boundary with scale $b - a$. Then

$$\mathbb{P}\left(\left|\bar{X}_t - \mathbb{E}\bar{X}_t\right| < \frac{u_\alpha\left(\sum_{i=1}^t (X_i - \widehat{X}_i)^2\right)}{t} \text{ for all } t \in \mathbb{N} \right) \geq 1 - 2\alpha.$$

Here $u_\alpha(v)$ is $\mathcal{O}(\sqrt{v \log v})$ or $\mathcal{O}(\sqrt{v \log \log v})$.

Empirical-Bernstein confidence sequence

Theorem (H., Ramdas, McAuliffe, Sekhon 2018)

Suppose X_i are independent and $[a, b]$-valued for all i. Let $\widehat{X}_i$ be any predictable sequence and u_α be any sub-exponential uniform boundary with scale $b - a$. Then

$$\mathbb{P}\left(\left| \bar{X}_t - \mathbb{E}\bar{X}_t \right| < \frac{u_\alpha\left(\sum_{i=1}^t (X_i - \widehat{X}_i)^2\right)}{t} \text{ for all } t \in \mathbb{N} \right) \geq 1 - 2\alpha.$$

Using package `confseq`

```
u(v) = confseq.boundaries.gamma_exponential_mixture_bound(
    v, alpha, c = b - a, v_opt = ...)
```

Average treatment effect: setup

Unit i has fixed potential outcomes $Y_i(0), Y_i(1)$, for $i = 1, 2, \ldots$

Average treatment effect: setup

Unit i has fixed potential outcomes $Y_i(0), Y_i(1)$, for $i = 1, 2, \ldots$

We assign unit i randomly to treatment with probability p or control with probability $1 - p$, and observe $Y_i(1)$ or $Y_i(0)$ accordingly. [Neyman 1923, Rubin 1974]

Average treatment effect: setup

Unit i has fixed potential outcomes $Y_i(0), Y_i(1)$, for $i = 1, 2, \ldots$

We assign unit i randomly to treatment with probability p or control with probability $1 - p$, and observe $Y_i(1)$ or $Y_i(0)$ accordingly. [Neyman 1923, Rubin 1974]

Assumption: no interference

Average treatment effect: setup

Unit i has fixed potential outcomes $Y_i(0), Y_i(1)$, for $i = 1, 2, \ldots$

We assign unit i randomly to treatment with probability p or control with probability $1 - p$, and observe $Y_i(1)$ or $Y_i(0)$ accordingly. [Neyman 1923, Rubin 1974]

Assumption: no interference

Our goal: after observing units $1, \ldots, t$, we'd like to estimate

$$\mathrm{ATE}_t := \frac{1}{t} \sum_{i=1}^{t} [Y_i(1) - Y_i(0)].$$

Average treatment effect: setup

Unit i has fixed potential outcomes $Y_i(0), Y_i(1)$, for $i = 1, 2, \ldots$

We assign unit i randomly to treatment with probability p or control with probability $1 - p$, and observe $Y_i(1)$ or $Y_i(0)$ accordingly. [Neyman 1923, Rubin 1974]

Assumption: no interference

Our goal: after observing units $1, \ldots, t$, we'd like to estimate

$$\mathrm{ATE}_t := \frac{1}{t} \sum_{i=1}^{t} [Y_i(1) - Y_i(0)].$$

Assumption: $Y_i(k) \in [0, 1]$ for $k = 0, 1$, all i.

Average treatment effect: theorem

For each unit i, we construct an estimator X_i of the individual treatment effect $Y_i(1) - Y_i(0)$ with two key properties:

1. **Unbiased**: $\mathbb{E}X_i = Y_i(1) - Y_i(0)$
2. **Variance** of X_i depends on **prediction errors** $(Y_i(1) - \widehat{Y}_i(1))^2$ and $(Y_i(0) - \widehat{Y}_i(0))^2$.

Average treatment effect: theorem

For each unit i, we construct an estimator X_i of the individual treatment effect $Y_i(1) - Y_i(0)$ with two key properties:

1. **Unbiased**: $\mathbb{E}X_i = Y_i(1) - Y_i(0)$
2. **Variance** of X_i depends on **prediction errors** $(Y_i(1) - \widehat{Y}_i(1))^2$ and $(Y_i(0) - \widehat{Y}_i(0))^2$.

Theorem (H., Ramdas, McAuliffe, Sekhon 2018)

Assume no interference and $Y_t(k) \in [0,1]$ for all k, t. Let u_α be any sub-exponential uniform boundary with scale $2/\min\{p, 1-p\}$. Then

$$\mathbb{P}\left(\left|\bar{X}_t - \mathrm{ATE}_t\right| < \frac{u_\alpha\left(\sum_{i=1}^t (X_i - \widehat{X}_t)^2\right)}{t} \text{ for all } t \in \mathbb{N} \right) \geq 1 - \alpha.$$

Average treatment effect: illustration

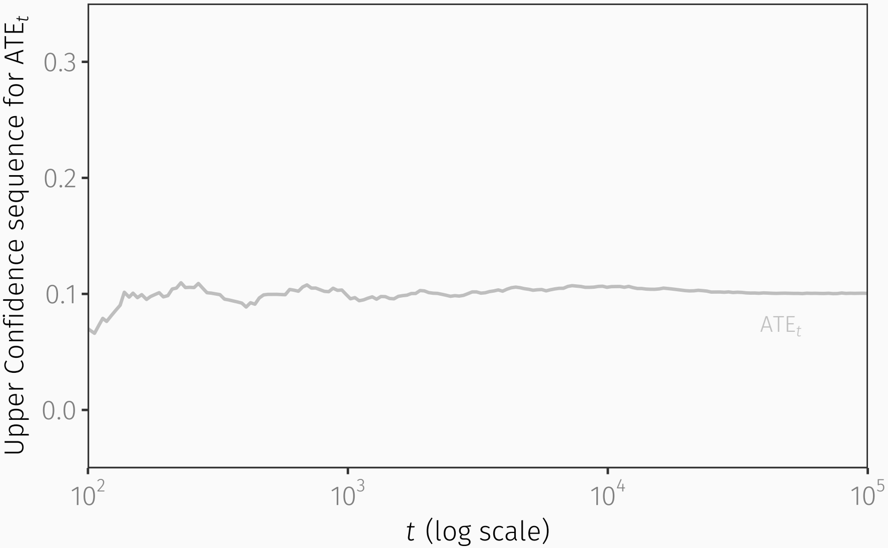

Average treatment effect: illustration

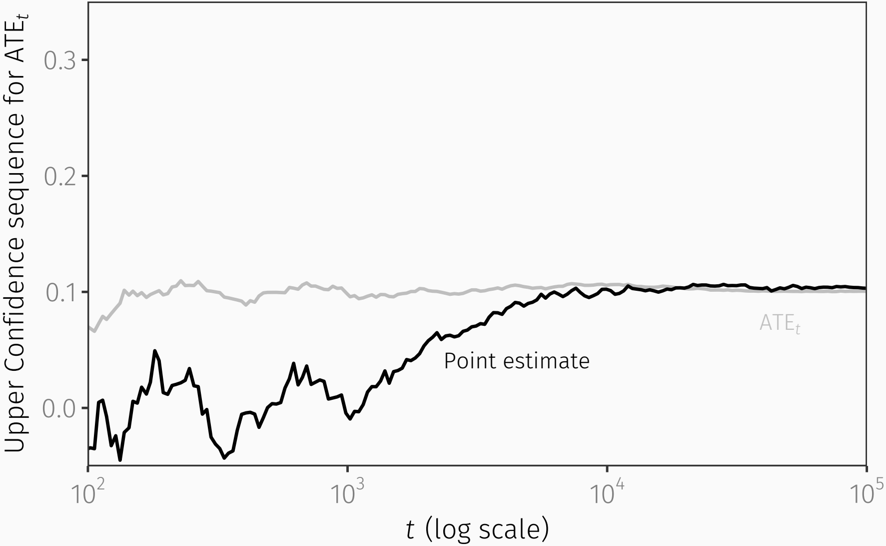

Average treatment effect: illustration

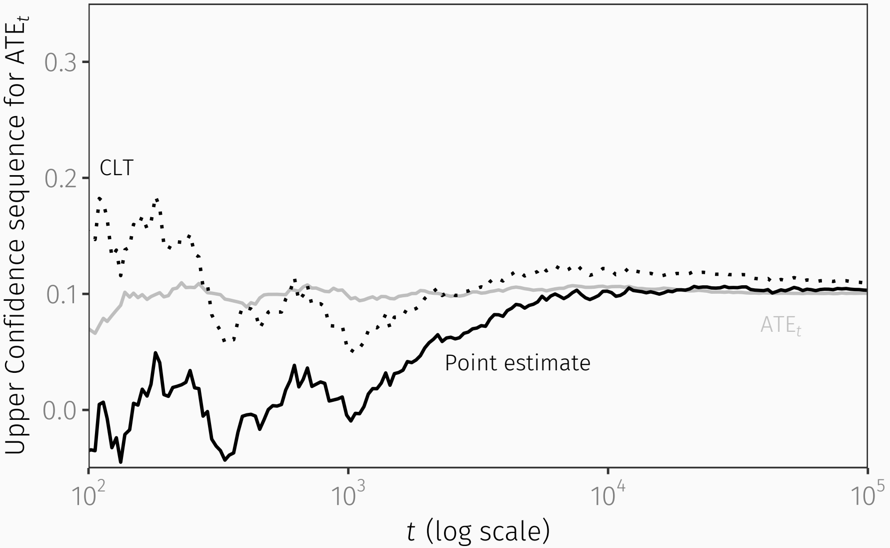

Average treatment effect: illustration

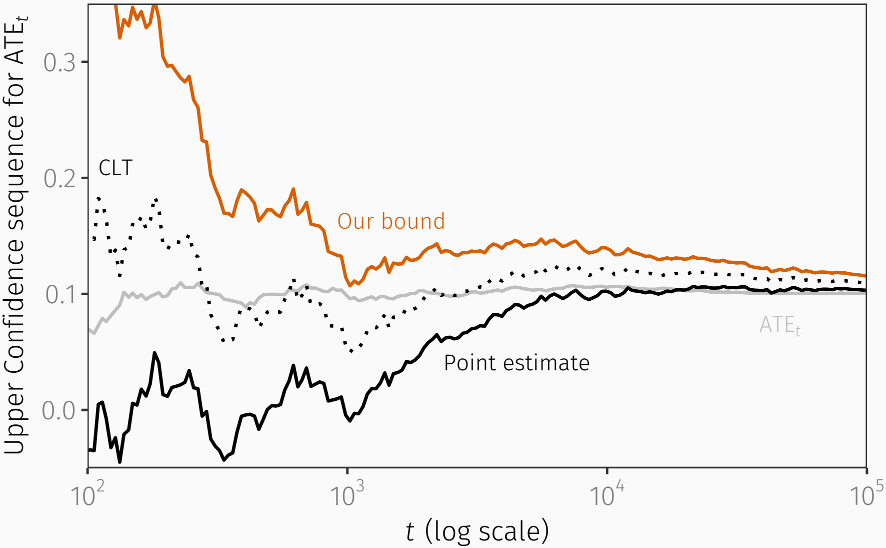

Quantile estimation

$X_1, X_2, \ldots$ i.i.d. from any distribution F. Let q be the p^{th} quantile of F, let $\widehat{Q}_t(p)$ denote the p^{th} sample quantile at time t.

Quantile estimation

$X_1, X_2, \ldots$ i.i.d. from any distribution F. Let q be the p^{th} quantile of F, let $\widehat{Q}_t(p)$ denote the p^{th} sample quantile at time t.

Theorem

Suppose X_i are i.i.d. from any distribution F. Let $u_{\alpha,p}$ be an appropriately scaled sub-Bernoulli uniform boundary. Then

$$\mathbb{P}\left(\widehat{Q}_t\left(p - \frac{u_{\alpha,1-p}(t)}{t}\right) \leq q \leq \widehat{Q}_t\left(p + \frac{u_{\alpha,p}(t)}{t}\right) \text{ for all } t \in \mathbb{N}\right) \geq 1 - \alpha.$$

Quantile estimation

$X_1, X_2, \ldots$ i.i.d. from any distribution F. Let q be the p^{th} quantile of F, let $\widehat{Q}_t(p)$ denote the p^{th} sample quantile at time t.

Theorem

Suppose X_i are i.i.d. from any distribution F. Let $u_{\alpha,p}$ be an appropriately scaled sub-Bernoulli uniform boundary. Then

$$\mathbb{P}\left(\widehat{Q}_t\left(p - \frac{u_{\alpha,1-p}(t)}{t}\right) \leq q \leq \widehat{Q}_t\left(p + \frac{u_{\alpha,p}(t)}{t}\right) \text{ for all } t \in \mathbb{N}\right) \geq 1 - \alpha.$$

No assumption on the distribution F.

Quantile estimation

Theorem

Suppose X_i are i.i.d. from any distribution F. Let $u_{\alpha,p}$ be an appropriately scaled sub-Bernoulli uniform boundary. Then

$$\mathbb{P}\left(\widehat{Q}_t\left(p - \frac{u_{\alpha,1-p}(t)}{t}\right) \leq q \leq \widehat{Q}_t\left(p + \frac{u_{\alpha,p}(t)}{t}\right) \text{ for all } t \in \mathbb{N}\right) \geq 1 - \alpha.$$

Using package `confseq`

```
u(t) = confseq.boundaries.beta_binomial_mixture_bound(
    p * (1 - p) * t, alpha,
    g = 1 - p, h = p,
    v_opt = ...)
```

Estimation of a cumulative distribution function

Theorem

Suppose X_i are i.i.d. from any distribution F. Let $\widehat{F}_t$ denote the empirical cumulative distribution function at time t. Then

$$\mathbb{P}\left(\left\| \widehat{F}_t - F \right\|_\infty \leq A\sqrt{\frac{\log\log(et) + C}{t}} \text{ for all } t \in \mathbb{N} \right) \geq 1 - e^{-\mathcal{O}(A^2 C)}.$$

Estimation of a cumulative distribution function

Theorem

Suppose X_i are i.i.d. from any distribution F. Let $\widehat{F}_t$ denote the empirical cumulative distribution function at time t. Then

$$\mathbb{P}\left(\left\|\widehat{F}_t - F\right\|_\infty \leq A\sqrt{\frac{\log\log(et) + C}{t}} \text{ for all } t \in \mathbb{N}\right) \geq 1 - e^{-\mathcal{O}(A^2C)}.$$

Using package `confseq`

```
u(t) = confseq.quantiles.empirical_process_lil_bound(
  t, alpha, t_min=1)
```

Best-arm identification

Confidence sequences are instrumental to many best-arm identification algorithms.

Multi-armed bandit: set of distributions indexed by $k = 1, \ldots, K$. Distribution k has mean μ_k.

Confidence sequences are instrumental to many best-arm identification algorithms.

Multi-armed bandit: set of distributions indexed by $k = 1, \ldots, K$. Distribution k has mean μ_k.

ϵ-best-arm identification with fixed confidence $1 - \delta$: choose an arm $k_\star$ such that, with probability at least $1 - \delta$, we have $\mu_{k_\star} \geq \max_k \mu_k - \epsilon$. [Even-Dar et al. 2002]

Confidence sequences are instrumental to many best-arm identification algorithms.

Multi-armed bandit: set of distributions indexed by $k = 1, \ldots, K$. Distribution k has mean μ_k.

ϵ-best-arm identification with fixed confidence $1-\delta$: choose an arm $k_\star$ such that, with probability at least $1-\delta$, we have $\mu_{k_\star} \geq \max_k \mu_k - \epsilon$. [Even-Dar et al. 2002]

Common strategy:

- Construct confidence sequence $(L_{kt}, U_{kt})_{t=1}^{\infty}$ for each arm k, so that $L_{kt} \leq \mu_k \leq U_{kt}$ for all k, t with probability at least $1-\delta$.
- Stop the first time there exists some $k_\star$ such that

$$L_{k_\star t} \geq U_{kt} - \epsilon \text{ for all } k \neq k_\star.$$

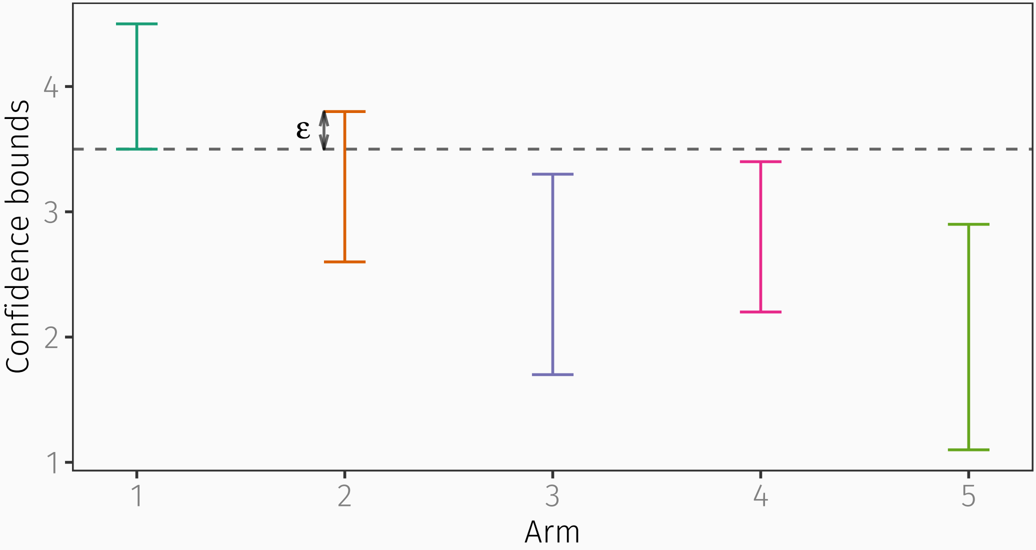
ε
Confidence bounds
4
3
2
1
1
2
3
4
5
Arm

Arm differences can be examined directly.

Comparing confidence intervals for *individual arm means* is wasteful.

Better to compute confidence intervals on *pairwise differences of arm means* directly.

- Related to asymptotically optimal stopping rule of Garivier and Kaufmann (2016).

Arm differences can be examined directly.

Comparing confidence intervals for *individual arm means* is wasteful.

Better to compute confidence intervals on *pairwise differences of arm means* directly.

- Related to asymptotically optimal stopping rule of Garivier and Kaufmann (2016).

Can be done in our framework, allowing more efficient stopping in nonparametric settings.

- Complete theory is work in progress

Quantile best-arm identification

Let $Q_k(p)$ denote the p^{th} quantile of arm k.

Quantile ϵ-best-arm identification with fixed confidence $1-\delta$: choose an arm $k_\star$ such that, with probability at least $1-\delta$, we have $Q_{k_\star}(p+\epsilon) \geq \max_k Q_k(p)$. [Szörényi et al. 2015]

Quantile best-arm identification

Let $Q_k(p)$ denote the p^{th} quantile of arm k.

Quantile ϵ-best-arm identification with fixed confidence $1 - \delta$: choose an arm $k_\star$ such that, with probability at least $1 - \delta$, we have $Q_{k_\star}(p + \epsilon) \geq \max_k Q_k(p)$. [Szörényi et al. 2015]

QLUCB algorithm [H. and Ramdas 2019]: at each round,

- sample arm h_t with highest LCB for $Q_k(p + \epsilon)$;
- sample arm with highest UCB for $Q_k(p)$, excluding h_t; and
- stop when LCB for $Q_k(p + \epsilon)$ is above UCB for $Q_j(p)$ for all $j \neq k$, for some k.

[cf. Kalyanakrishnan et al. 2012]

Quantile best-arm simulations

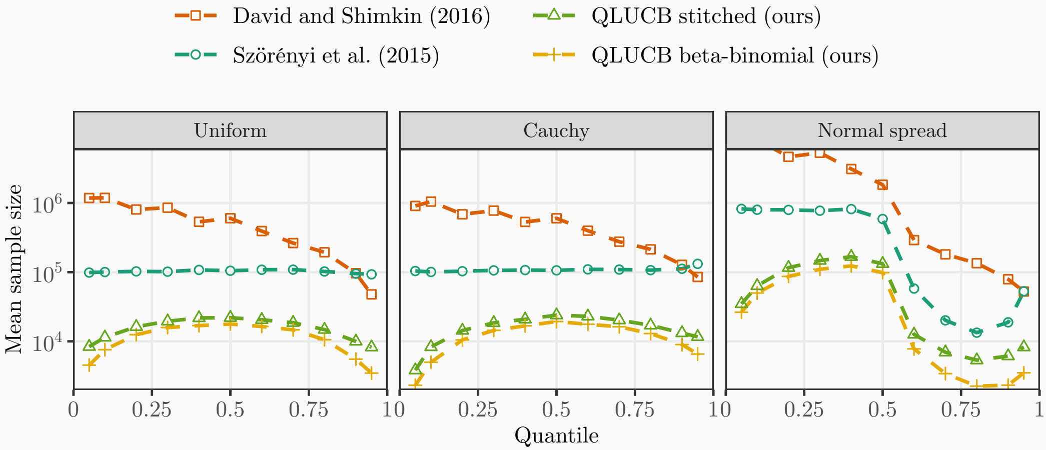

A taste of the underlying framework

Reminder: sub-Gaussianity and Hoeffding bound

A random variable X is *sub-Gaussian* with variance parameter σ^2 if

$$\log \mathbb{E} e^{\lambda X} \leq \frac{\lambda^2 \sigma^2}{2} \quad \text{for all } \lambda \in \mathbb{R}. \tag{1}$$

Reminder: sub-Gaussianity and Hoeffding bound

A random variable X is *sub-Gaussian* with variance parameter σ^2 if

$$\log \mathbb{E} e^{\lambda X} \leq \frac{\lambda^2 \sigma^2}{2} \quad \text{for all } \lambda \in \mathbb{R}. \tag{1}$$

Hoeffding bound (1963): if $X_i \in [0, 1]$ independent, $i = 1, \ldots, t$, then

$$\mathbb{P}\left(\sum_{i=1}^{t} (X_i - \mathbb{E} X_i) \geq \sqrt{\frac{t \log \alpha^{-1}}{2}} \right) \leq \alpha. \tag{2}$$

Reminder: sub-Gaussianity and Hoeffding bound

A random variable X is *sub-Gaussian* with variance parameter σ^2 if

$$\log \mathbb{E} e^{\lambda X} \leq \frac{\lambda^2 \sigma^2}{2} \quad \text{for all } \lambda \in \mathbb{R}. \tag{1}$$

Hoeffding bound (1963): if $X_i \in [0,1]$ independent, $i = 1, \ldots, t$, then

$$\mathbb{P}\left(\sum_{i=1}^{t} (X_i - \mathbb{E} X_i) \geq \sqrt{\frac{t \log \alpha^{-1}}{2}} \right) \leq \alpha. \tag{2}$$

Proof involves two main pieces:

1. Show X_i is sub-Gaussian with variance parameter 1/4, and
2. Use Cramér-Chernoff method to obtain (2) from (1).

Reminder: sub-Gaussianity and Hoeffding bound

A random variable X is *sub-Gaussian* with variance parameter σ^2 if

$$\log \mathbb{E} e^{\lambda X} \leq \psi(\lambda)\sigma^2 \quad \text{for all } \lambda \in \mathbb{R}. \qquad \text{Here } \psi(\lambda) = \frac{\lambda^2}{2} \tag{1}$$

Hoeffding bound (1963): if $X_i \in [0, 1]$ independent, $i = 1, \ldots, t$, then

$$\mathbb{P}\left(\sum_{i=1}^{t}(X_i - \mathbb{E}X_i) \geq u_\psi(t\sigma^2)\right) \leq \alpha. \qquad \text{Here } u_\psi(v) = \sqrt{2v \log \alpha^{-1}}, \;\; \sigma^2 = \frac{1}{4} \tag{2}$$

Proof involves two main pieces:

1. Show X_i is sub-Gaussian with variance parameter $1/4$, and
2. Use Cramér-Chernoff method to obtain (2) from (1).

Our recipe for confidence sequences

1. Assume a null hypothesis $H_{\theta_0} : \theta = \theta_0$.

Our recipe for confidence sequences

1. Assume a null hypothesis $H_{\theta_0} : \theta = \theta_0$.
2. Show process $S_t^{\theta_0}$ is *sub-ψ with variance process* $V_t^{\theta_0}$ under H_{θ_0}.

Our recipe for confidence sequences

1. Assume a null hypothesis $H_{\theta_0} : \theta = \theta_0$.
2. Show process $S_t^{\theta_0}$ is *sub-ψ with variance process* $V_t^{\theta_0}$ under H_{θ_0}.
3. Choose any *sub-ψ uniform boundary* $u_\alpha : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$. Then, under H_{θ_0},

$$\mathbb{P}\left(S_t^{\theta_0} \geq u_\alpha(V_t^{\theta_0}) \text{ for some } t \in \mathbb{N}\right) \leq \alpha.$$

Our recipe for confidence sequences

1. Assume a null hypothesis $H_{\theta_0} : \theta = \theta_0$.
2. Show process $S_t^{\theta_0}$ is *sub-ψ with variance process* $V_t^{\theta_0}$ under H_{θ_0}.
3. Choose any *sub-ψ uniform boundary* $u_\alpha : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$. Then, under H_{θ_0},

$$\mathbb{P}\left(S_t^{\theta_0} \geq u_\alpha(V_t^{\theta_0}) \text{ for some } t \in \mathbb{N}\right) \leq \alpha.$$

4. At time t, a confidence set for θ is

$$\mathrm{CI}_t = \left\{\theta_0 \in \mathbb{R} : S_t^{\theta_0} < u_\alpha(V_t^{\theta_0})\right\}.$$

Example: testing the bias of a coin

Suppose a coin has bias p. Let S_t be the number of heads observed up to time t.

Example: testing the bias of a coin

Suppose a coin has bias p. Let S_t be the number of heads observed up to time t.

The process $S_t - tp$ is sub-Gaussian with variance process $V_t = t/4$.

Example: testing the bias of a coin

Suppose a coin has bias p. Let S_t be the number of heads observed up to time t.

The process $S_t - tp$ is sub-Gaussian with variance process $V_t = t/4$.

Then, for any $\lambda > 0$,

$$\mathbb{P}\left(S_t - tp \geq \underbrace{\frac{\log \alpha^{-1}}{\lambda} + \frac{\lambda}{2} \cdot \frac{t}{4}}_{\text{A linear boundary}} \text{ for some } t \in \mathbb{N} \right) \leq \alpha.$$

Example: testing the bias of a coin

Suppose a coin has bias p. Let S_t be the number of heads observed up to time t.

The process $S_t - tp$ is sub-Gaussian with variance process $V_t = t/4$.

Then, for any $\lambda > 0$,

$$\mathbb{P}\left(S_t - tp \geq \underbrace{\frac{\log \alpha^{-1}}{\lambda} + \frac{\lambda}{2} \cdot \frac{t}{4}}_{\text{A linear boundary}} \text{ for some } t \in \mathbb{N} \right) \leq \alpha.$$

This yields the confidence sequence

$$\left| \frac{S_t}{t} - p \right| < \frac{\log \alpha^{-1}}{\lambda t} + \frac{\lambda}{8}, \quad \text{for all } t, \text{ with probability at least } 1 - 2\alpha.$$

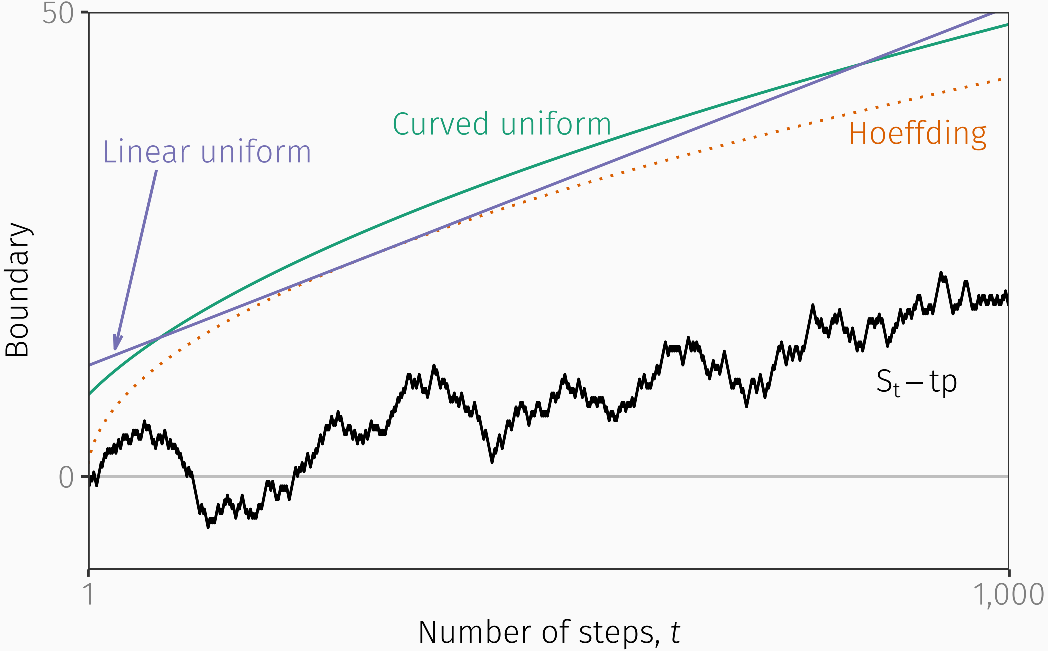

50
0
Curved uniform
Linear uniform
Hoeffding
$S_t - tp$
Boundary
1
1,000
Number of steps, t

Our framework includes self-normalized bounds.

Suppose X_i has a symmetric distribution conditional on $X_1, \ldots, X_{i-1}$.

Our framework includes self-normalized bounds.

Suppose X_i has a symmetric distribution conditional on $X_1, \ldots, X_{i-1}$.

Then $S_t = \sum_{i=1}^t X_i$ is sub-Gaussian with variance process $V_t = \sum_{i=1}^t X_i^2$.

[de la Peña 1999]

Our framework includes self-normalized bounds.

Suppose X_i has a symmetric distribution conditional on $X_1, \ldots, X_{i-1}$.

Then $S_t = \sum_{i=1}^t X_i$ is sub-Gaussian with variance process $V_t = \sum_{i=1}^t X_i^2$.
[de la Peña 1999]

So

$$\mathbb{P}\left(\sum_{i=1}^t X_i \geq u_\alpha\left(\sum_{i=1}^t X_i^2\right) \text{ for some } t\right) \leq \alpha.$$

Our framework includes matrix bounds.

Suppose X_i is a Hermitian, $d \times d$ random matrix with a conditionally symmetric distribution, $X_i \sim -X_i \mid X_1, \ldots, X_{i-1}$.

Our framework includes matrix bounds.

Suppose X_i is a Hermitian, $d \times d$ random matrix with a conditionally symmetric distribution, $X_i \sim -X_i \mid X_1, \ldots, X_{i-1}$.

Then $S_t = \gamma_{\max}\left(\sum_{i=1}^t X_i\right)$ is sub-Gaussian with variance process $V_t = \gamma_{\max}\left(\sum_{i=1}^t X_i^2\right)$. [de la Peña 1999, Tropp 2011]

Our framework includes matrix bounds.

Suppose X_i is a Hermitian, $d \times d$ random matrix with a conditionally symmetric distribution, $X_i \sim -X_i \mid X_1, \ldots, X_{i-1}$.

Then $S_t = \gamma_{\max}\left(\sum_{i=1}^t X_i\right)$ is sub-Gaussian with variance process $V_t = \gamma_{\max}\left(\sum_{i=1}^t X_i^2\right)$. [de la Peña 1999, Tropp 2011]

So

$$\mathbb{P}\left(\gamma_{\max}\left(\sum_{i=1}^t X_i\right) \geq u_{\alpha,d}\left(\gamma_{\max}\left(\sum_{i=1}^t X_i^2\right)\right) \text{ for some } t\right) \leq \alpha.$$

Summary

- Peeking is a serious problem, and confidence sequences are a flexible solution.

Summary

- Peeking is a serious problem, and confidence sequences are a flexible solution.
- We derive useful confidence sequences in a variety of nonparametric settings, including for estimating average treatment effect and quantiles.

Summary

- Peeking is a serious problem, and confidence sequences are a flexible solution.
- We derive useful confidence sequences in a variety of nonparametric settings, including for estimating average treatment effect and quantiles.
- Our confidence sequences immediately improve best-arm identification algorithms and extend validity to nonparametric settings.

Summary

- Peeking is a serious problem, and confidence sequences are a flexible solution.
- We derive useful confidence sequences in a variety of nonparametric settings, including for estimating average treatment effect and quantiles.
- Our confidence sequences immediately improve best-arm identification algorithms and extend validity to nonparametric settings.
- Our underlying framework extends the Cramér-Chernoff method, unifying many existing results and yielding new confidence sequences in diverse settings.

Thank you!

- stevehoward@berkeley.edu
- Exponential line-crossing inequalities: https://arxiv.org/abs/1808.03204
- Uniform, nonparametric, non-asymptotic confidence sequences: https://arxiv.org/abs/1810.08240
- Sequential estimation of quantiles with applications to A/B-testing and best-arm identification: https://arxiv.org/abs/1906.09712
- Implementations of many uniform boundaries and confidence sequences: https://github.com/gostevehoward/confseq
- Slides: stevehoward.org

Appendix

We construct an unbiased estimator of each individual treatment effect.

We require predictions $\widehat{Y}_t(1)$ and $\widehat{Y}_t(0)$ for $Y_t(1)$ and $Y_t(0)$.

We construct an unbiased estimator of each individual treatment effect.

We require predictions $\widehat{Y}_t(1)$ and $\widehat{Y}_t(0)$ for $Y_t(1)$ and $Y_t(0)$.

For each unit t, we construct an estimator X_t of the individual treatment effect $Y_t(1) - Y_t(0)$.

X_t depends on our predictions as well as the observed treatment and outcome.

We construct an unbiased estimator of each individual treatment effect.

We require predictions $\widehat{Y}_t(1)$ and $\widehat{Y}_t(0)$ for $Y_t(1)$ and $Y_t(0)$.

For each unit t, we construct an estimator X_t of the individual treatment effect $Y_t(1) - Y_t(0)$.

X_t depends on our predictions as well as the observed treatment and outcome.

Two key properties of X_t:

1. **Unbiased**: $\mathbb{E}X_t = Y_t(1) - Y_t(0)$
2. **Variance** of X_t depends on **prediction errors** $(Y_t(1) - \widehat{Y}_t(1))^2$ and $(Y_t(0) - \widehat{Y}_t(0))^2$.

We construct an unbiased estimator of each individual treatment effect.

We require predictions $\widehat{Y}_t(1)$ and $\widehat{Y}_t(0)$ for $Y_t(1)$ and $Y_t(0)$.

For each unit t, we construct an estimator X_t of the individual treatment effect $Y_t(1) - Y_t(0)$.

X_t depends on our predictions as well as the observed treatment and outcome.

Two key properties of X_t:

1. **Unbiased**: $\mathbb{E}X_t = Y_t(1) - Y_t(0)$
2. **Variance** of X_t depends on **prediction errors** $(Y_t(1) - \widehat{Y}_t(1))^2$ and $(Y_t(0) - \widehat{Y}_t(0))^2$.

Let $S_t = \sum_{i=1}^{t} X_i$. Then S_t/t is unbiased for ATE_t.

CLT confidence bounds

$$\frac{S_t}{t} \pm \frac{1.96\sqrt{\sum_{i=1}^{t}(X_i - \bar{X}_t)^2}}{t}$$

CLT confidence bounds

$$\frac{S_t}{t} \pm \frac{1.96\sqrt{\sum_{i=1}^{t}(X_i - \bar{X}_t)^2}}{t}$$

Uniform, non-asymptotic confidence bounds

$$\frac{S_t}{t} \pm \frac{u_\alpha\left(\sum_{i=1}^{t}(X_i - \widehat{X}_i)^2\right)}{t} \quad \text{where } \widehat{X}_i = \widehat{Y}_i(1) - \widehat{Y}_i(0).$$

CLT confidence bounds

$$\frac{S_t}{t} \pm \frac{1.96\sqrt{\sum_{i=1}^{t}(X_i - \bar{X}_t)^2}}{t}$$

Uniform, non-asymptotic confidence bounds

$$\frac{S_t}{t} \pm \frac{u_\alpha\left(\sum_{i=1}^{t}(X_i - \widehat{X}_i)^2\right)}{t} \quad \text{where } \widehat{X}_i = \widehat{Y}_i(1) - \widehat{Y}_i(0).$$

- $\sum_{i=1}^{t}(X_i - \widehat{X}_i)^2$ is an “online” estimate of $\mathrm{Var}(S_t)$.
- Estimation precision depends on prediction accuracy.

CLT confidence bounds

$$\frac{S_t}{t} \pm \frac{1.96\sqrt{\sum_{i=1}^{t}(X_i - \bar{X}_t)^2}}{t}$$

Uniform, non-asymptotic confidence bounds

$$\frac{S_t}{t} \pm \frac{u_\alpha\left(\sum_{i=1}^{t}(X_i - \widehat{X}_i)^2\right)}{t} \qquad \text{where } \widehat{X}_i = \widehat{Y}_i(1) - \widehat{Y}_i(0).$$

- $\sum_{i=1}^{t}(X_i - \widehat{X}_i)^2$ is an "online" estimate of $\mathrm{Var}(S_t)$.
- Estimation precision depends on prediction accuracy.
- $u_\alpha(v) = \mathcal{O}(\sqrt{v \log v})$, so $u_\alpha(v)$ is like $z_{1-\alpha}\sqrt{v}$, but the "z-factor" grows over time (slowly).

The uniform boundary grows only slightly faster than $\mathcal{O}(\sqrt{n})$

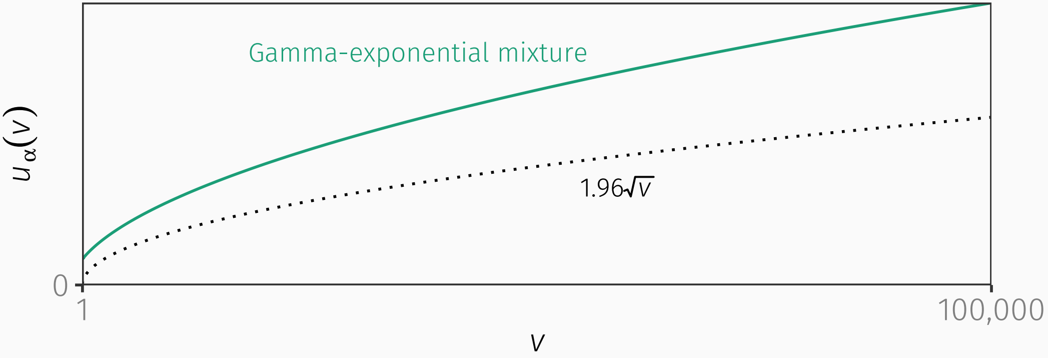

What about that boundedness assumption?

Without some similar assumption, it is impossible to construct a confidence interval.

What about that boundedness assumption?

Without some similar assumption, it is impossible to construct a confidence interval.

The problem: one outlier can have arbitrarily large influence.

- Consider $Y_1(1) = t$, $Y_1(0) = 0$, and $Y_i(k) = 0$ for all $k = 0, 1$ and $i = 2, \ldots, t$.

What about that boundedness assumption?

Without some similar assumption, it is impossible to construct a confidence interval.

The problem: one outlier can have arbitrarily large influence.

- Consider $Y_1(1) = t$, $Y_1(0) = 0$, and $Y_i(k) = 0$ for all $k = 0, 1$ and $i = 2, \ldots, t$.

Asymptotic arguments often sweep this issue under the rug.

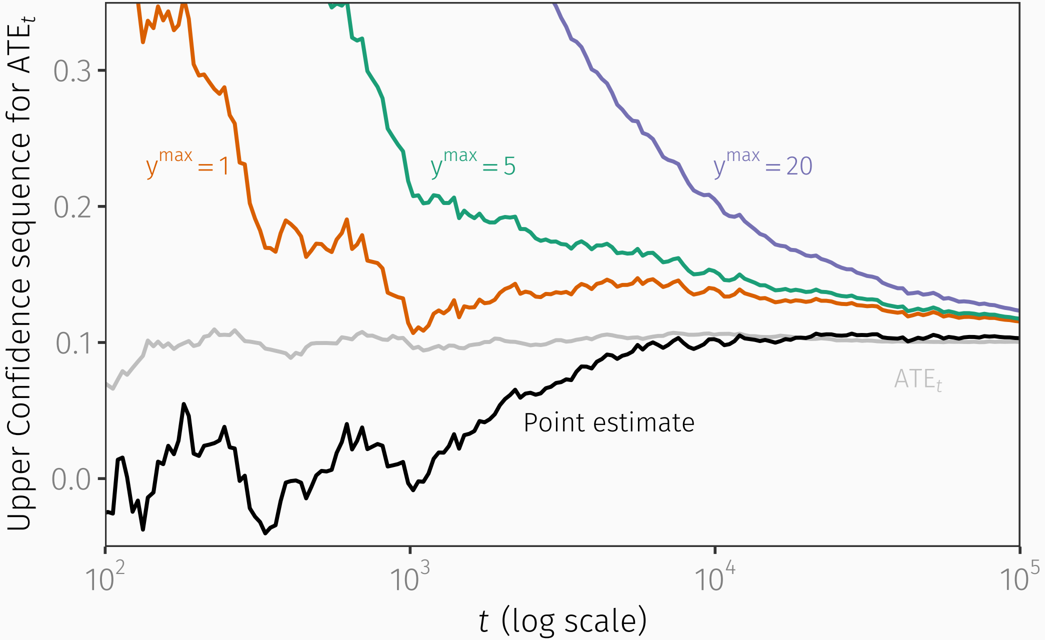
Upper Confidence sequence for ATE_t
0.3
0.2
0.1
0.0
$y^{max} = 1$
$y^{max} = 5$
$y^{max} = 20$
ATE_t
Point estimate
10^2
10^3
10^4
10^5
t (log scale)