Nonasymptotic confidence sequences for sequential estimation of treatment effects in randomized trials

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

June 30, 2021

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%

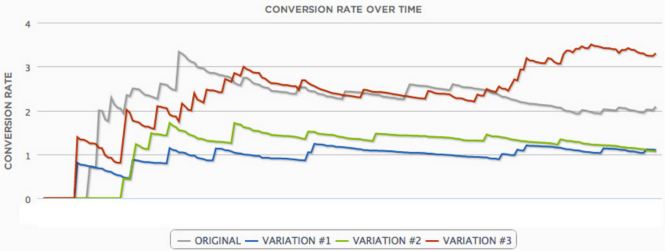

The promise and peril of sequential monitoring

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.

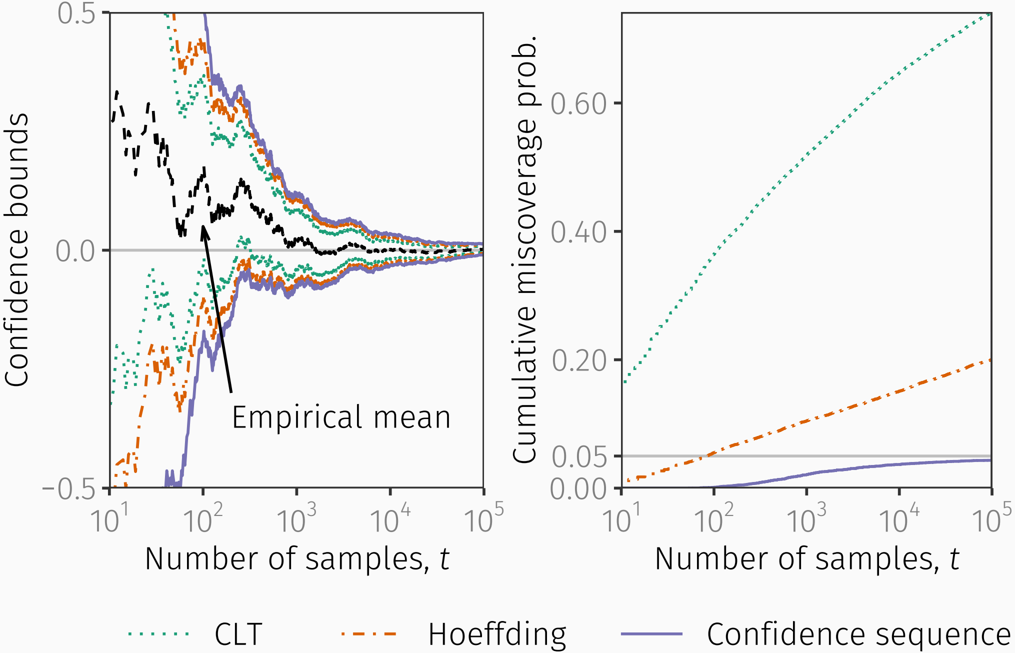

0.5
0.0
−0.5
Confidence bounds
Empirical mean
10^1 10^2 10^3 10^4 10^5
Number of samples, t
Cumulative miscoverage prob.
0.60
0.40
0.20
0.05
0.00
10^1 10^2 10^3 10^4 10^5
Number of samples, t
CLT
Hoeffding
Confidence sequence

Outline

Sequential estimation of average treatment effect

A taste of the underlying framework

Self-normalized bounds, matrix bounds, quantile estimation

Sequential estimation of average treatment effect

We give a method to

- sequentially gather data from a randomized experiment,

We give a method to

- sequentially gather data from a randomized experiment,
- repeatedly compute a confidence interval for average treatment effect,

We give a method to

- sequentially gather data from a randomized experiment,
- repeatedly compute a confidence interval for average treatment effect,
- and stop the experiment at any time,

We give a method to

- sequentially gather data from a randomized experiment,
- repeatedly compute a confidence interval for average treatment effect,
- and stop the experiment at any time,

with the following guarantee:

No matter how you adaptively stop, you have a valid confidence interval for the average treatment effect.

We give a method to

- sequentially gather data from a randomized experiment,
- repeatedly compute a confidence interval for average treatment effect,
- and stop the experiment at any time,

with the following guarantee:

No matter how you adaptively stop, you have a valid confidence interval for the average treatment effect.

(Can compute always-valid p-values instead, if desired.)

We avoid some unrealistic assumptions and problematic restrictions.

- No superpopulation or stationarity assumptions.

We avoid some unrealistic assumptions and problematic restrictions.

- No superpopulation or stationarity assumptions.
- No knowledge of the stopping rule.

We avoid some unrealistic assumptions and problematic restrictions.

- No superpopulation or stationarity assumptions.
- No knowledge of the stopping rule.
- No bound on sample size.

We avoid some unrealistic assumptions and problematic restrictions.

- No superpopulation or stationarity assumptions.
- No knowledge of the stopping rule.
- No bound on sample size.
- No asymptotic approximations or sharp null hypothesis.

We assume fixed, bounded potential outcomes.

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

We assume fixed, bounded potential outcomes.

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

We assign unit t randomly to treatment or control and observe $Y_t(1)$ or $Y_t(0)$ accordingly [Neyman 1923, Rubin 1974].

We assume fixed, bounded potential outcomes.

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

We assign unit t randomly to treatment or control and observe $Y_t(1)$ or $Y_t(0)$ accordingly [Neyman 1923, Rubin 1974]. This is the only source of randomness.

We assume fixed, bounded potential outcomes.

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

We assign unit t randomly to treatment or control and observe $Y_t(1)$ or $Y_t(0)$ accordingly [Neyman 1923, Rubin 1974]. This is the only source of randomness.

Assumption: no interference

We assume fixed, bounded potential outcomes.

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

We assign unit t randomly to treatment or control and observe $Y_t(1)$ or $Y_t(0)$ accordingly [Neyman 1923, Rubin 1974]. This is the only source of randomness.

Assumption: no interference

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

- More on this later

We define a sequence of average treatment effect estimands.

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

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

We define a sequence of average treatment effect estimands.

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

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

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

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

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

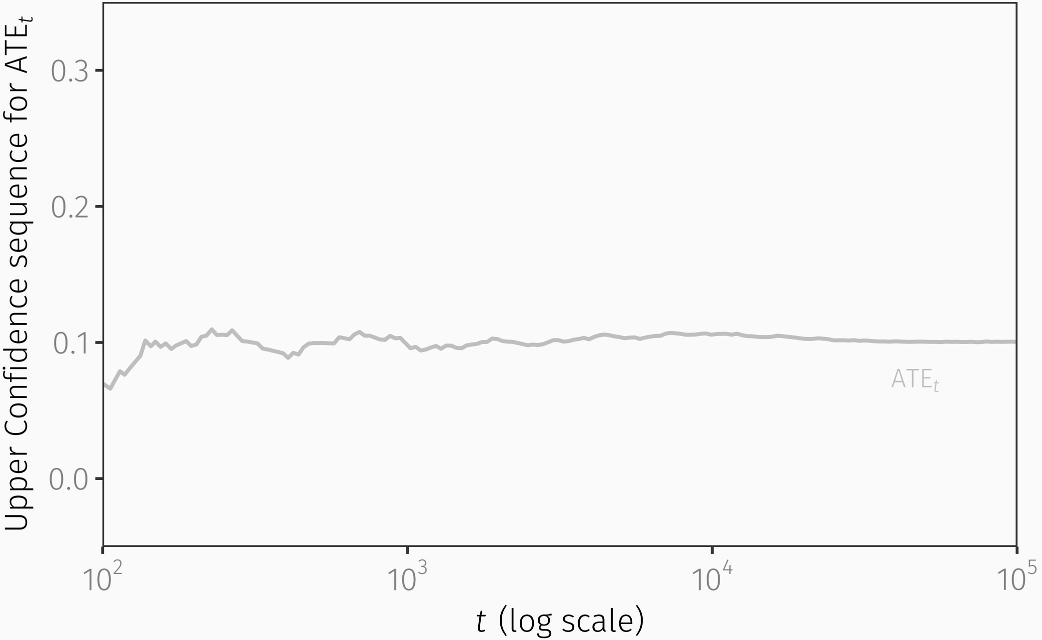
Upper Confidence sequence for ATE_t
0.3
0.2
0.1
0.0
ATE_t
10^2
10^3
10^4
10^5
t (log scale)

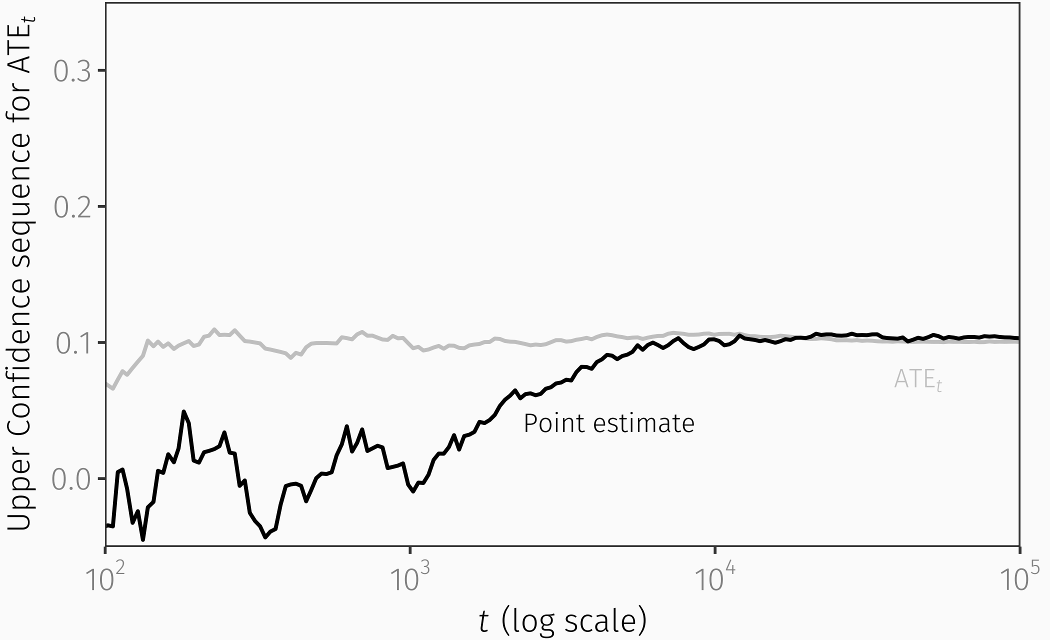

Upper Confidence sequence for ATE_t
0.3
0.2
0.1
0.0
ATE_t
Point estimate
10^2
10^3
10^4
10^5
t (log scale)

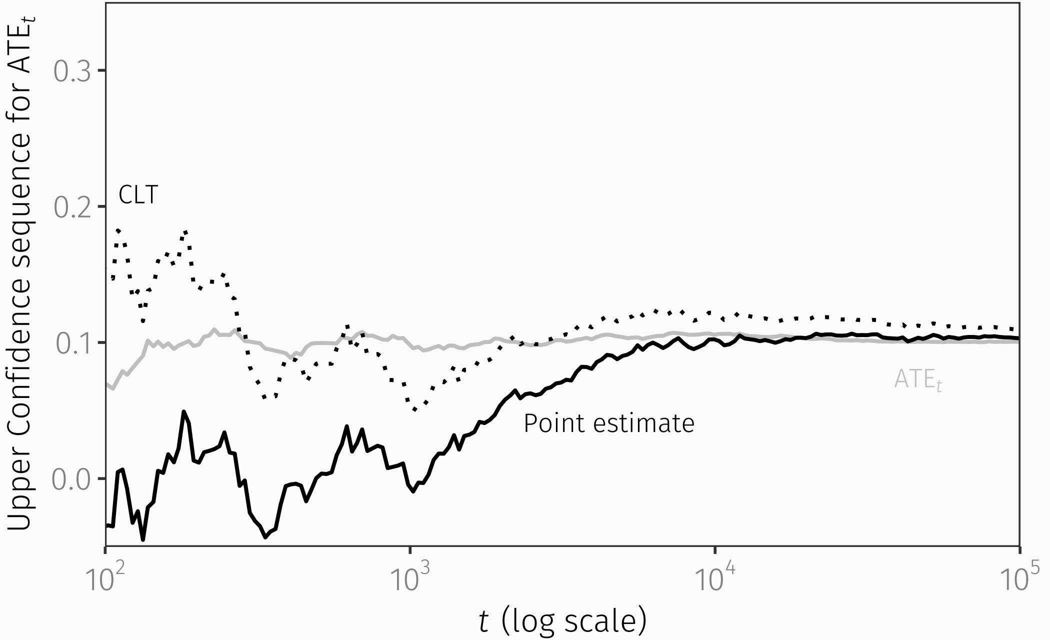
Upper Confidence sequence for ATE$_t$
0.3
0.2
0.1
0.0
CLT
Point estimate
ATE$_t$
10^2
10^3
10^4
10^5
t (log scale)

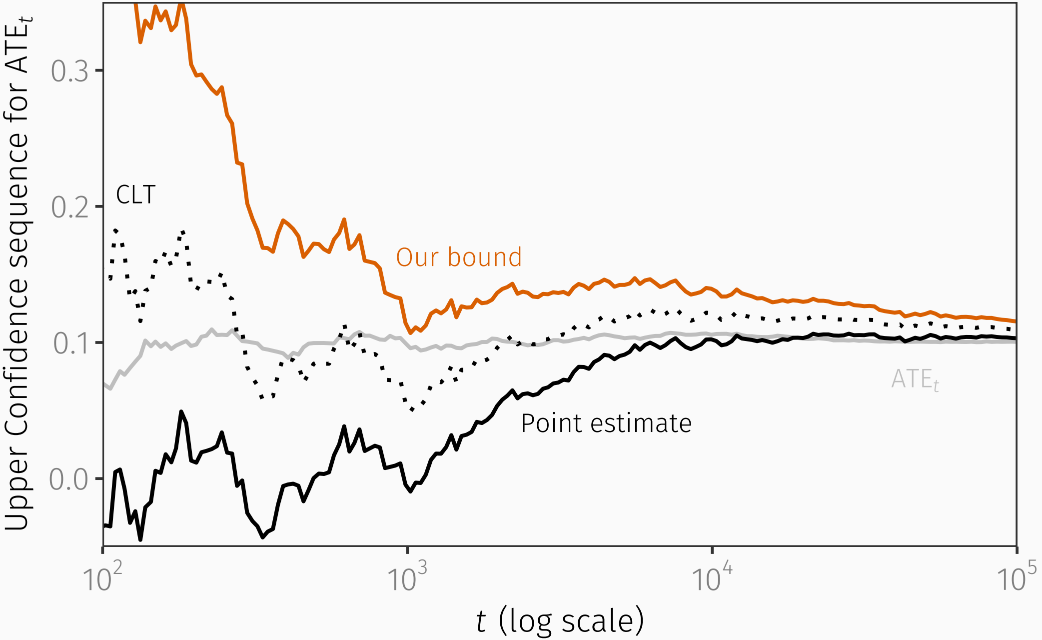

Upper Confidence sequence for ATE_t
0.3
0.2
0.1
0.0
CLT
Our bound
ATE_t
Point estimate
10^2
10^3
10^4
10^5
t (log scale)

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 $Y_t(1) - Y_t(0)$. (An augmented inverse propensity weighted estimator.)

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 $Y_t(1) - Y_t(0)$. (An augmented inverse propensity weighted estimator.)

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

Two 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 $Y_t(1) - Y_t(0)$. (An augmented inverse propensity weighted estimator.)

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

Two 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

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

CLT confidence bounds

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

Uniform, nonasymptotic 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

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

Uniform, nonasymptotic 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

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

Uniform, nonasymptotic 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 $\mathcal{O}(\sqrt{n})$

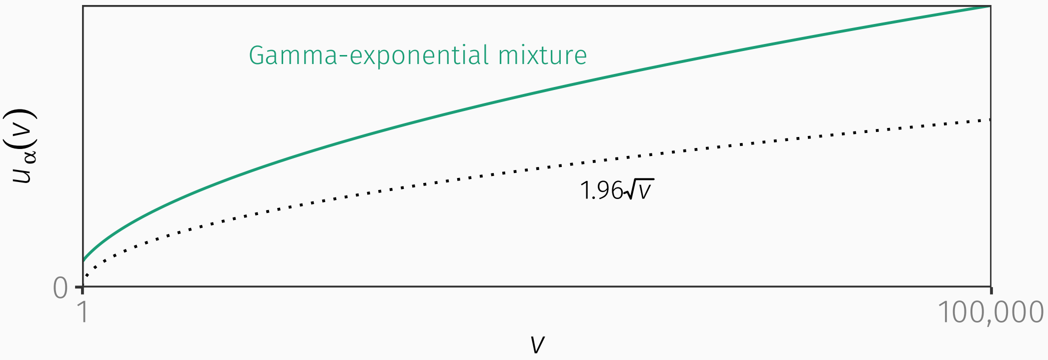

Theorem

Assuming no interference, if $Y_t(k) \in [0,1]$ for all k, t, then

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

Theorem

Assuming no interference, if $Y_t(k) \in [0,1]$ for all k, t, then

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

This implies

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

gives a $(1-\alpha)$-confidence sequence for (ATE_t).

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.

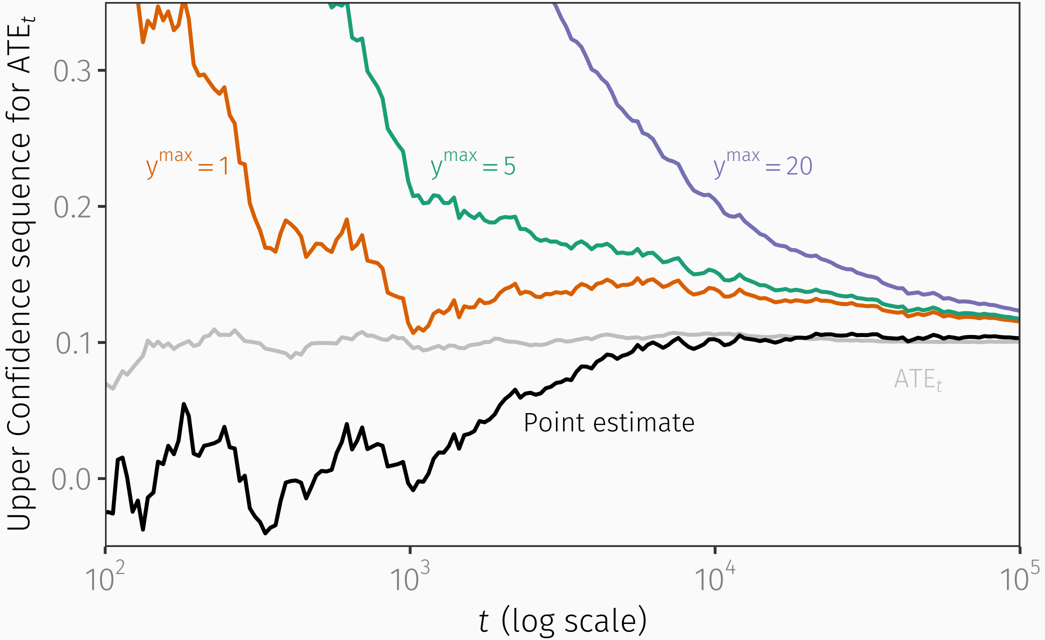

Upper Confidence sequence for ATE$_t$
0.3
0.2
0.1
0.0
$y^{max} = 1$
$y^{max} = 5$
$y^{max} = 20$
ATE$_t$
Point estimate
10^2
10^3
10^4
10^5
t (log scale)

Recap of ATE application

- Nonasymptotic confidence sequences for ATE_t
- Flexible inferential tool for sequential experiments
- Provable coverage under the assumption of bounded potential outcomes
- Replace central limit theorem argument with uniform concentration bounds
- Seamlessly handles “biased coin” or other adaptive allocation designs (not covered today)

A taste of the proof techniques

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.

To test $H_0 : p \leq p_0$ for **fixed** t, we need a result like

$$P_0(S_t - tp > u_\alpha(t)) \leq \alpha \qquad (\text{e.g, } u_\alpha(t) = z_{1-\alpha}\sqrt{p(1-p)t}).$$

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.

To test $H_0 : p \leq p_0$ for **fixed** t, we need a result like

$$P_0(S_t - tp > u_\alpha(t)) \leq \alpha \qquad \text{(e.g, } u_\alpha(t) = z_{1-\alpha}\sqrt{p(1-p)t}\text{)}.$$

For a **sequential** test, we need a result like

$$P_0(S_t - tp > u_\alpha(t) \text{ for some } t \in \mathbb{N}) \leq \alpha.$$

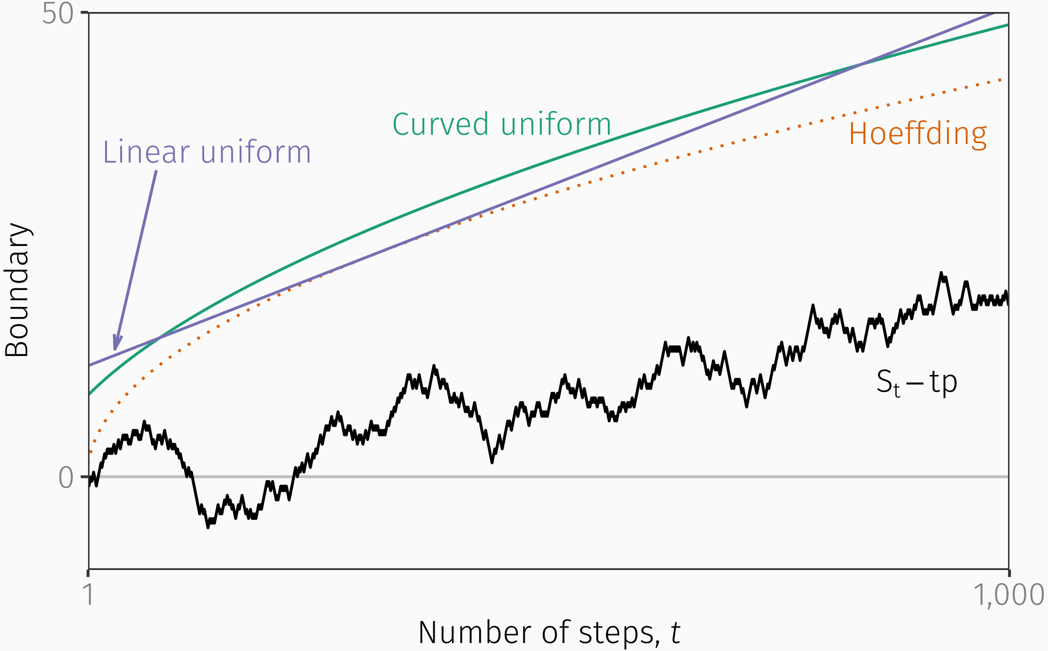
50
0
Boundary
Linear uniform
Curved uniform
Hoeffding
$S_t - tp$
1
1,000
Number of steps, t

A simple confidence sequence

For a **sequential** test, we need a result like

$$P_0(S_t - tp > u_\alpha(t) \text{ for some } t \in \mathbb{N}) \leq \alpha.$$

A simple confidence sequence

For a **sequential** test, we need a result like

$$P_0(S_t - tp > u_\alpha(t) \text{ for some } t \in \mathbb{N}) \leq \alpha.$$

Here's one such result: for any $\lambda > 0$,

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

A simple confidence sequence

For a **sequential** test, we need a result like

$$P_0(S_t - tp > u_\alpha(t) \text{ for some } t \in \mathbb{N}) \leq \alpha.$$

Here's one such result: 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

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

Fixed-sample Hoeffding bound

$S_t - tp$ is a *sub-Gaussian random variable* with variance parameter $t/4$:

$$\mathbb{E} \exp\{\lambda(S_t - tp)\} \leq \exp\left\{\frac{\lambda^2}{2} \cdot \frac{t}{4}\right\}.$$

Fixed-sample Hoeffding bound

$S_t - tp$ is a *sub-Gaussian random variable* with variance parameter $t/4$:

$$\mathbb{E} \exp \{\lambda(S_t - tp)\} \le \exp\left\{\frac{\lambda^2}{2} \cdot \frac{t}{4}\right\}.$$

Markov's inequality:

$$\mathbb{P}\left(\exp\left\{\lambda(S_t - tp) - \frac{\lambda^2}{2} \cdot \frac{t}{4}\right\} \ge x\right) \le \frac{1}{x}.$$

Time-uniform Hoeffding bound

The process $(S_t - tp)_{t \in \mathbb{N}}$ is a *sub-Gaussian process* with *variance process* $V_t = t/4$:

$$\exp\left\{\lambda(S_t - tp) - \frac{\lambda^2}{2} \cdot \frac{t}{4}\right\} \text{ is a supermartingale.}$$

The process $(S_t - tp)_{t\in\mathbb{N}}$ is a *sub-Gaussian process* with *variance process* $V_t = t/4$:

$$\exp\left\{\lambda(S_t - tp) - \frac{\lambda^2}{2}\cdot\frac{t}{4}\right\} \text{ is a supermartingale.}$$

Ville's inequality:

$$\mathbb{P}\left(\exp\left\{\lambda(S_t - tp) - \frac{\lambda^2}{2}\cdot\frac{t}{4}\right\} \geq x \text{ for some } t\in\mathbb{N}\right) \leq \frac{1}{x}.$$

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

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

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

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

No assumption on the distribution F.

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 underlying framework extends the Cramér-Chernoff method, unifying many existing results and yielding new confidence sequences in diverse settings.

Thank you!

- `steve@stevehoward.org`
- Time-uniform, nonparametric, nonasymptotic confidence sequences. Annals of Statistics, 2021. With A. Ramdas, J. McAuliffe, J. Sekhon.
- Time-uniform Chernoff bounds via nonnegative supermartingales. Probability Surveys, 2020. With A. Ramdas, J. McAuliffe, J. Sekhon.
- Sequential estimation of quantiles with applications best-arm identification and A/B-testing. Bernoulli, to appear. With A. Ramdas.
- Implementations of many uniform boundaries and confidence sequences: `https://github.com/gostevehoward/confseq`
- Slides: `stevehoward.org`