Uniform, non-asymptotic confidence sequences for sequential estimation of average treatment effect

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

June 25, 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%

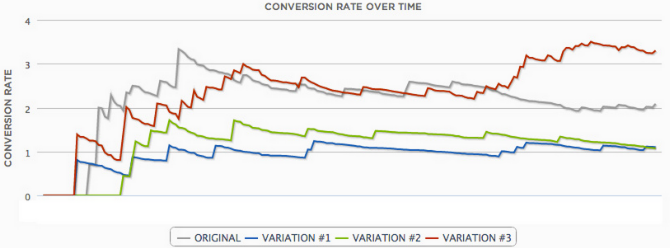

The promise and peril of sequential monitoring

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.

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]

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]

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]

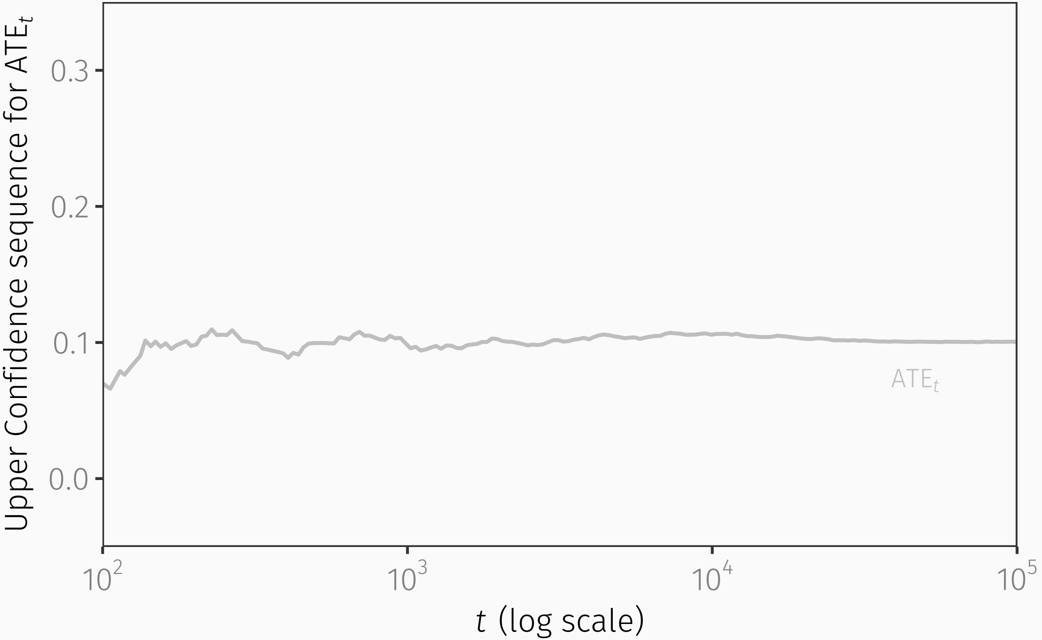
Upper Confidence sequence for ATE_t
0.3
0.2
0.1
0.0
ATE_t
10^2
10^3
10^4
10^5
t (log scale)

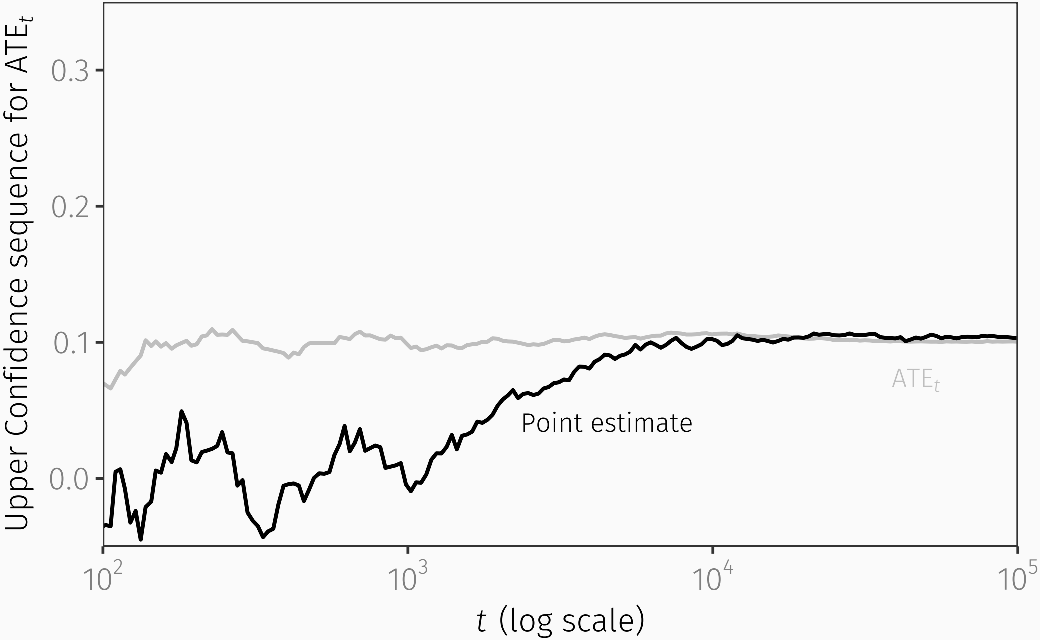

Upper Confidence sequence for ATE_t
0.3
0.2
0.1
0.0
ATE_t
Point estimate
10^2
10^3
10^4
10^5
t (log scale)

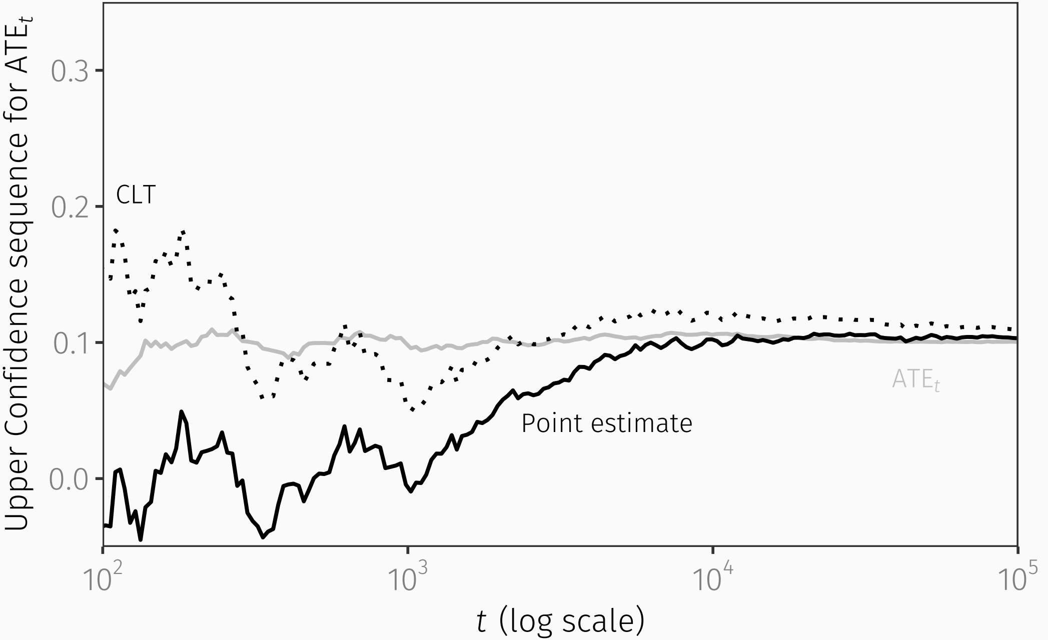

Upper Confidence sequence for ATE$_t$
0.3
0.2
0.1
0.0
CLT
Point estimate
ATE$_t$
10^2
10^3
10^4
10^5
t (log scale)

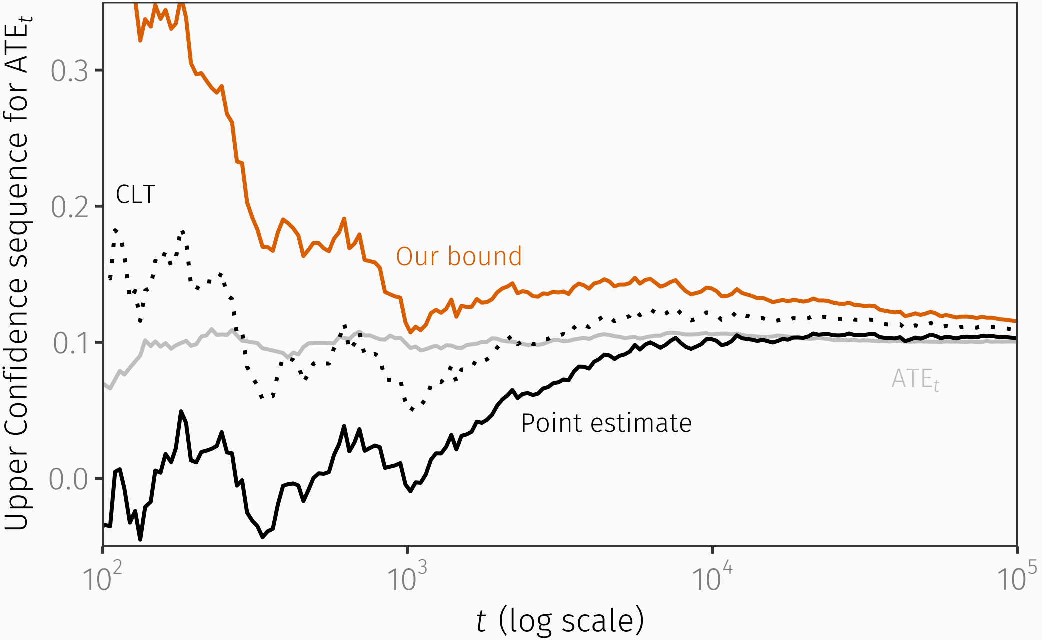
Upper Confidence sequence for ATE$_t$
0.3
0.2
0.1
0.0
CLT
Our bound
ATE$_t$
Point estimate
10^2
10^3
10^4
10^5
t (log scale)

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 $Y_t(1) - Y_t(0)$.

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

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

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

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 $\mathcal{O}(\sqrt{n})$

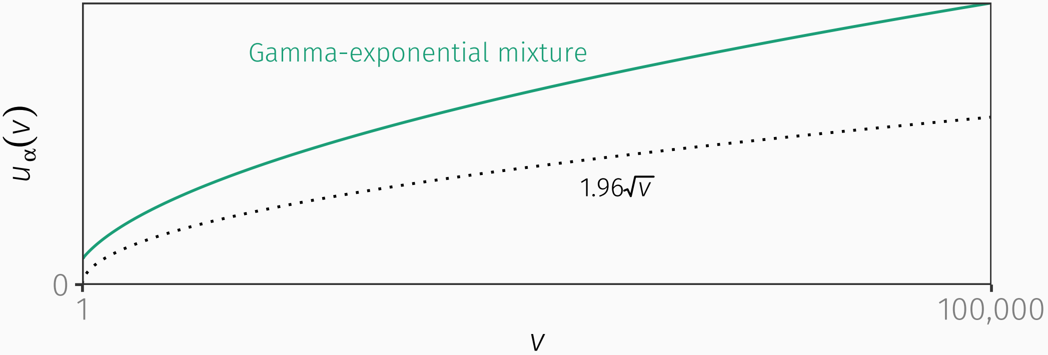

Theorem

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.

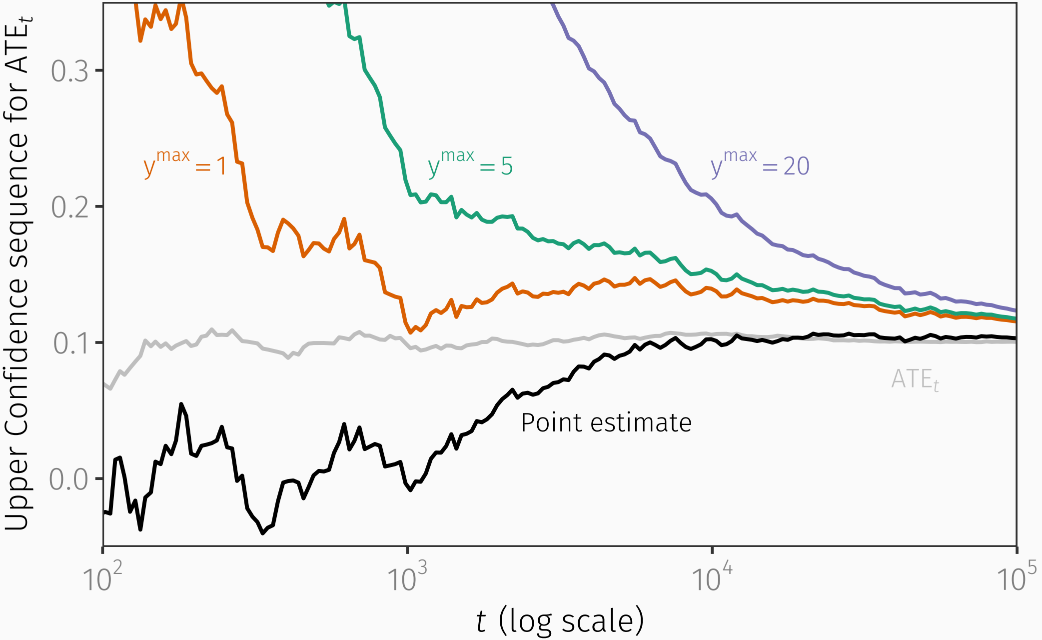

Upper Confidence sequence for ATE_t
0.3
0.2
0.1
0.0
$y^{max} = 1$
$y^{max} = 5$
$y^{max} = 20$
ATE_t
Point estimate
10^2
10^3
10^4
10^5
t (log scale)

Recap

- Non-asymptotic 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

Not covered today:

- Valid under superpopulation model as well
- Seamlessly handles biased coin designs or adaptive allocation
- Can estimate quantiles with similar tools (in i.i.d. framework)

Thank you!

stevehoward@berkeley.edu

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