

Sequential estimation of quantiles with applications to A/B testing and bandits



Aaditya Ramdas

Assistant Professor
Dept. of Statistics and Data Science
Machine Learning Dept.
Carnegie Mellon University

Sequential estimation of quantiles with applications to *A/B* testing and bandits



Aaditya Ramdas

Assistant Professor
Dept. of Statistics and Data Science
Machine Learning Dept.
Carnegie Mellon University



Steven R. Howard

PhD student
Dept. of Statistics
UC Berkeley

Users of app or website



50%

50%

A



B



Users of app or website



50%

50%

A



B



Total time spent on website:
44 hours

Time spent: 77 hours

Users of app or website



50%

50%

A



B



Total time spent on website:
44 hours

Time spent: 77 hours

B wins!?

Users of app or website



50%

50%

A



B



Total time spent on website:
44 hours

Time spent: 77 hours

B wins!?

Summarize evidence as a p-value

Users of app or website



50%

50%

A



B



Total time spent on website:
44 hours

Time spent: 77 hours

B wins!?

Summarize evidence as a **p-value**

(Null hypothesis: A and B have equal means)

Calculate p-value

Calculate p-value




Possible
observations


Calculate p-value




Calculate p-value




Calculate p-value



Calculate p-value




Calculate p-value



Reject null if $P \leq \alpha$

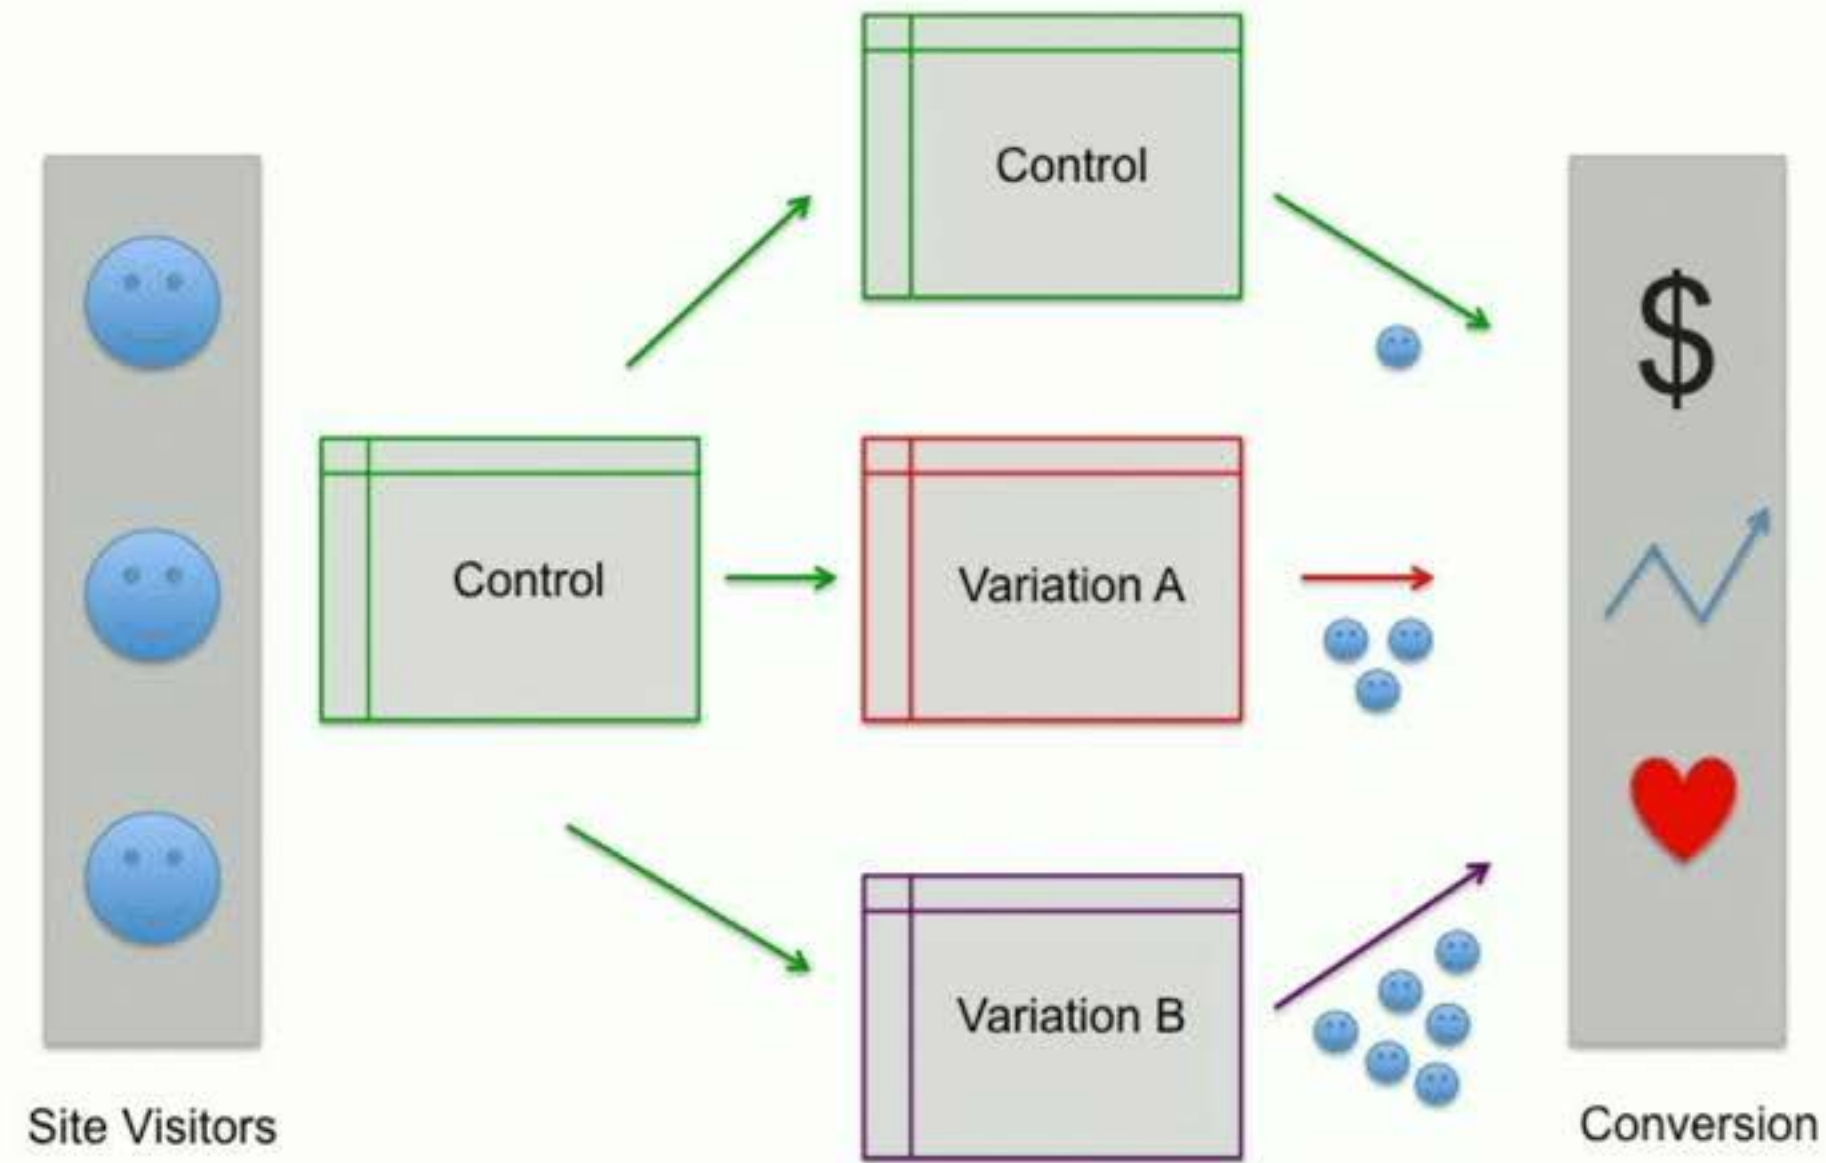
Calculate p-value



Reject null if $P \leq \alpha$

Then, $\Pr(\text{false positive}) \leq \alpha$.

What if we are testing more than one alternative?



Much more traffic needed by an A/B/n test

Multi-armed bandits



Multi-armed bandits



For example, we would like to test null hypothesis

$$H_0 : \mu_A \geq \max\{\mu_B, \mu_C\}.$$

Multi-armed bandits



For example, we would like to test null hypothesis

$$H_0 : \mu_A \geq \max\{\mu_B, \mu_C\}.$$

Or, maybe we want to identify the best arm with high probability.

Switch from estimating means to quantiles?

Let $X \sim F$. Define the p -quantile as $Q(p) := \sup\{x : F(x) \leq p\}$,
and $Q^-(p) := \sup\{x : F(x) < p\}$ (hence $q_{1/2}$ is the median)

Switch from estimating means to quantiles?

Let $X \sim F$. Define the p -quantile as $Q(p) := \sup\{x : F(x) \leq p\}$,
and $Q^-(p) := \sup\{x : F(x) < p\}$ (hence $q_{1/2}$ is the median)

Reasons to use quantiles include:

Switch from estimating means to quantiles?

Let $X \sim F$. Define the p -quantile as $Q(p) := \sup\{x : F(x) \leq p\}$,
and $Q^-(p) := \sup\{x : F(x) < p\}$ (hence $q_{1/2}$ is the median)

Reasons to use quantiles include:

- Quantiles always exist for any distribution,
while means (moments) do not always exist (eg: Cauchy).

Switch from estimating means to quantiles?

Let $X \sim F$. Define the p -quantile as $Q(p) := \sup\{x : F(x) \leq p\}$,
and $Q^-(p) := \sup\{x : F(x) < p\}$ (hence $q_{1/2}$ is the median)

Reasons to use quantiles include:

- Quantiles always exist for any distribution, while means (moments) do not always exist (eg: Cauchy).
- Quantiles can be defined for any totally ordered space, eg: ratings A-F, where “distance between ratings” undefined.

Switch from estimating means to quantiles?

Let $X \sim F$. Define the p -quantile as $Q(p) := \sup\{x : F(x) \leq p\}$,
and $Q^-(p) := \sup\{x : F(x) < p\}$ (hence $q_{1/2}$ is the median)

Reasons to use quantiles include:

- Quantiles always exist for any distribution, while means (moments) do not always exist (eg: Cauchy).
- Quantiles can be defined for any totally ordered space, eg: ratings A-F, where “distance between ratings” undefined.
- Estimating quantiles can be done sequentially.

Switch from estimating means to quantiles?

Let $X \sim F$. Define the p -quantile as $Q(p) := \sup\{x : F(x) \leq p\}$,
and $Q^-(p) := \sup\{x : F(x) < p\}$ (hence $q_{1/2}$ is the median)

Reasons to use quantiles include:

- Quantiles always exist for any distribution, while means (moments) do not always exist (eg: Cauchy).
- Quantiles can be defined for any totally ordered space, eg: ratings A-F, where “distance between ratings” undefined.
- Estimating quantiles can be done sequentially.
- Can get p-values for testing difference in quantiles.

Switch from estimating means to quantiles?

Let $X \sim F$. Define the p -quantile as $Q(p) := \sup\{x : F(x) \leq p\}$,
and $Q^-(p) := \sup\{x : F(x) < p\}$ (hence $q_{1/2}$ is the median)

Reasons to use quantiles include:

- Quantiles always exist for any distribution, while means (moments) do not always exist (eg: Cauchy).
- Quantiles can be defined for any totally ordered space, eg: ratings A-F, where “distance between ratings” undefined.
- Estimating quantiles can be done sequentially.
- Can get p-values for testing difference in quantiles.
- Can run bandit experiments, including best-arm identification.

Switch from estimating means to quantiles?

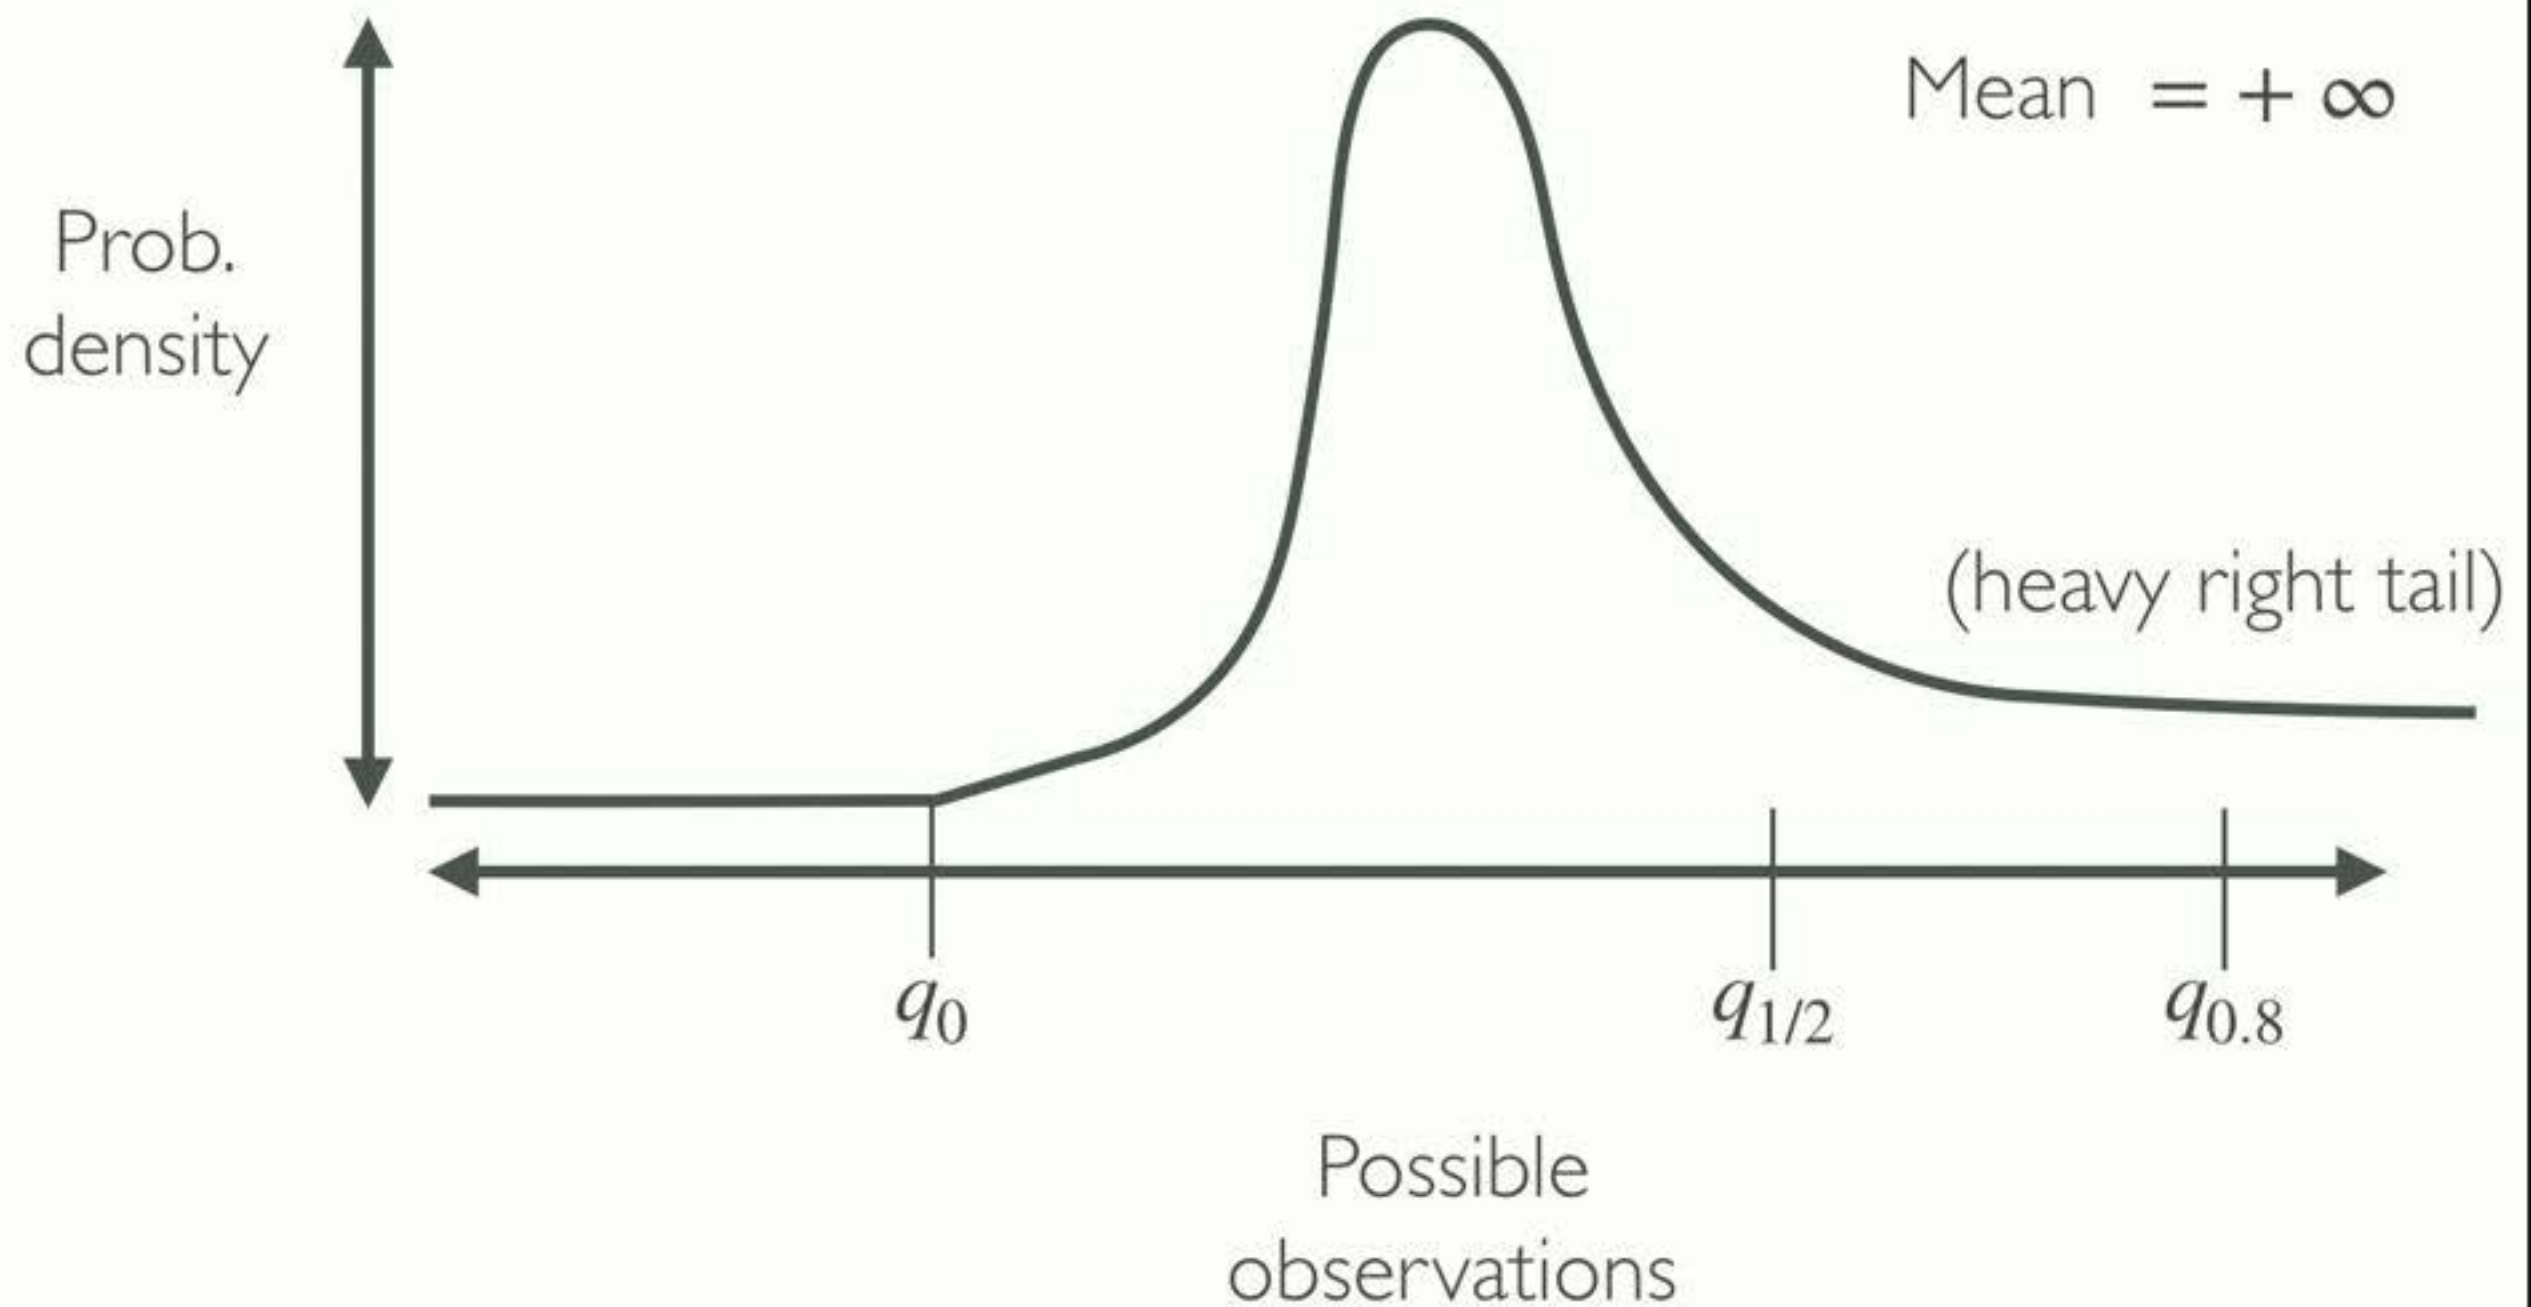
Let $X \sim F$. Define the p -quantile as $Q(p) := \sup\{x : F(x) \leq p\}$,
and $Q^-(p) := \sup\{x : F(x) < p\}$ (hence $q_{1/2}$ is the median)

Reasons to use quantiles include:

- Quantiles always exist for any distribution, while means (moments) do not always exist (eg: Cauchy).
- Quantiles can be defined for any totally ordered space, eg: ratings A-F, where “distance between ratings” undefined.
- Estimating quantiles can be done sequentially.
- Can get p-values for testing difference in quantiles.
- Can run bandit experiments, including best-arm identification.


Can also estimate all quantiles simultaneously.

Quantiles are informative for heavy tails




Eg: amount of time spent on Reddit

The mean need not even exist (eg: Cauchy)



Eg: amount of money won/lost in a casino


The mean need not even exist (eg: Cauchy)



Eg: amount of money won/lost in a casino

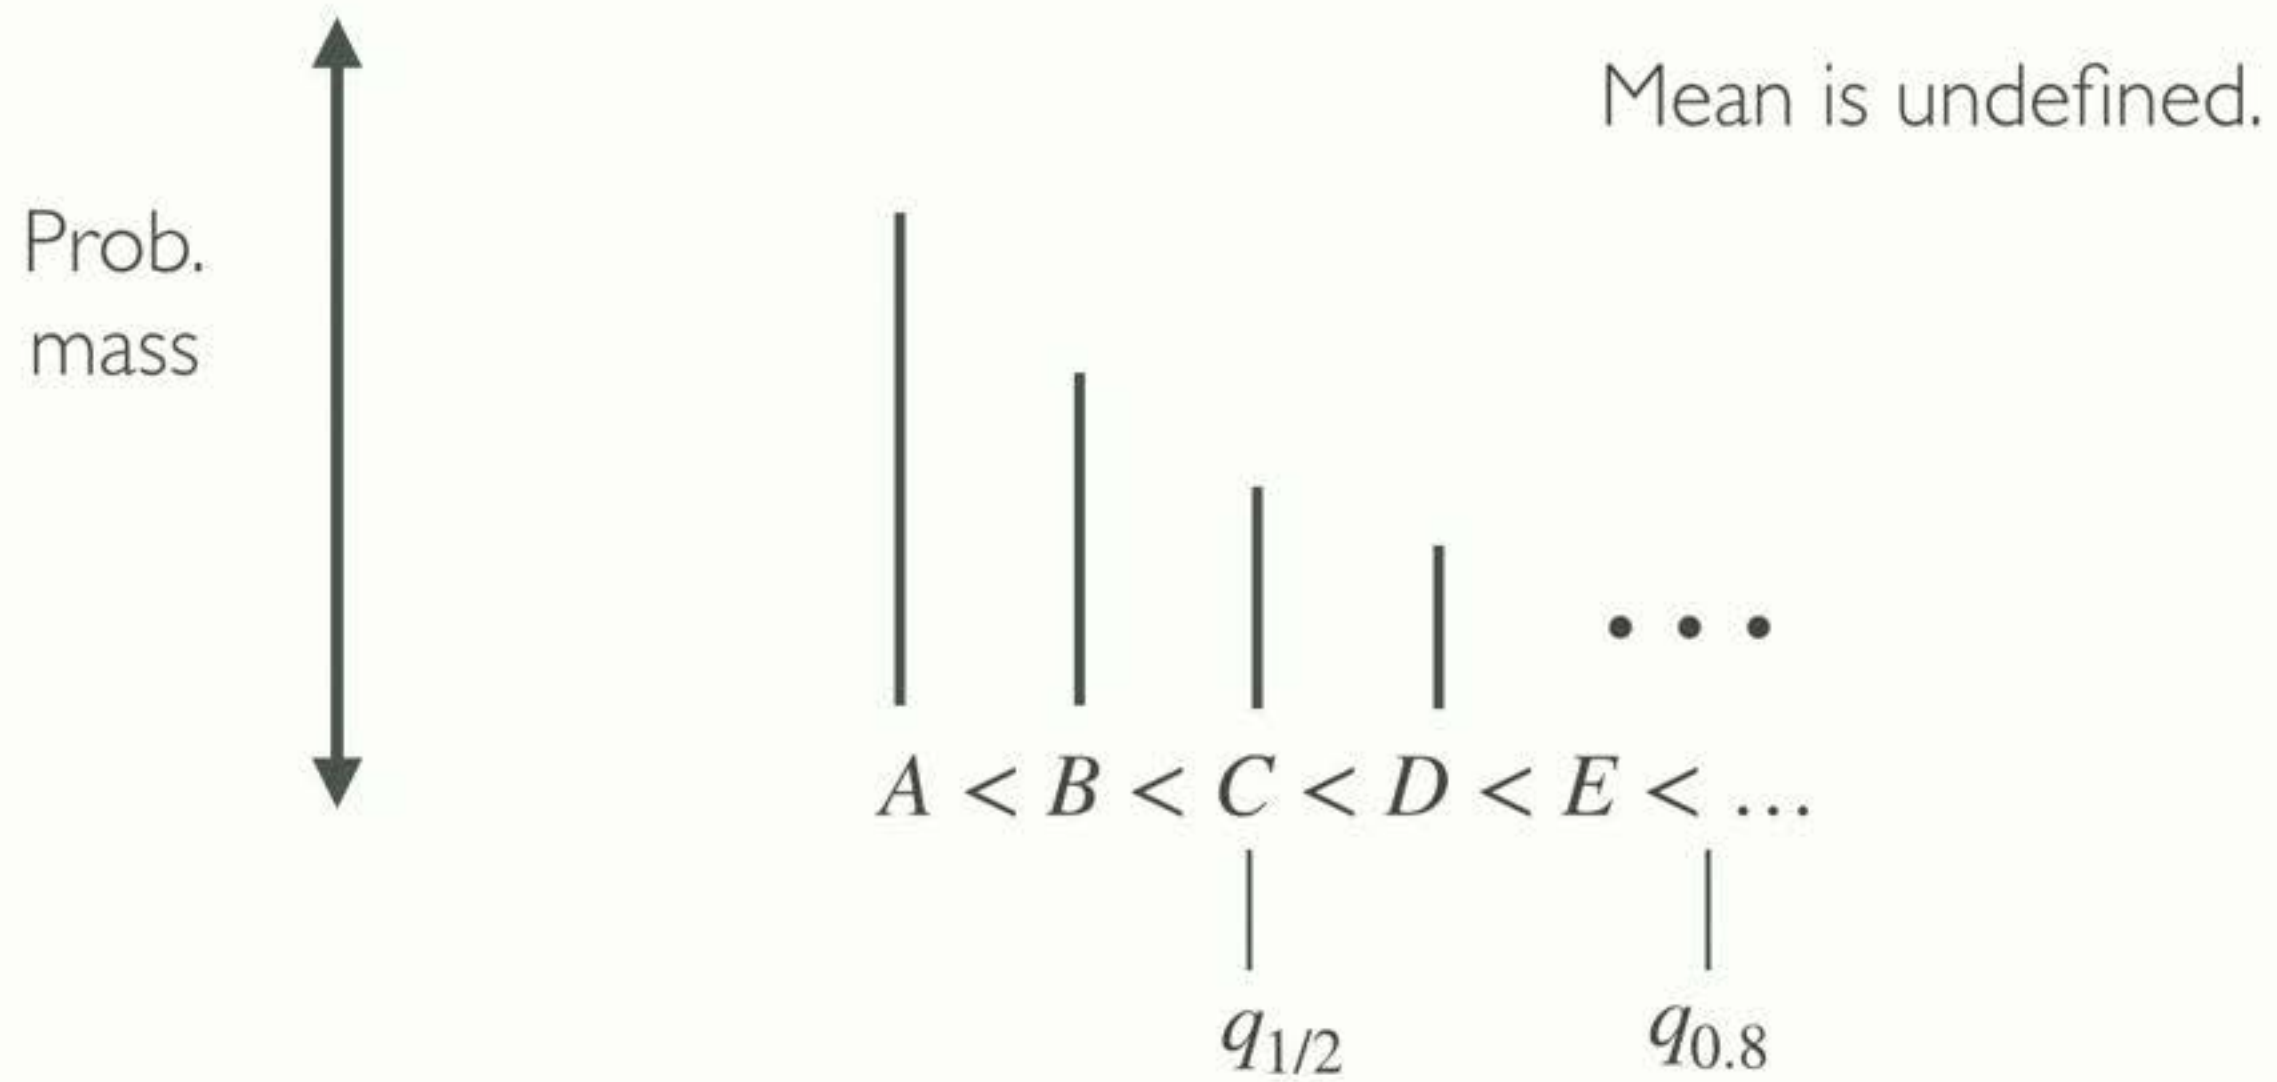
Do not need to resort to trimming "outliers".
(How to pick threshold? Throw away or cap?)

The same could arise in discrete settings




Eg: number of links clicked

Quantile sensible in totally ordered settings



Eg: grades or non-numerical ratings

Quantile sensible in totally ordered settings



Eg: grades or non-numerical ratings

Do not need to artificially assign numerical values.
(Are they equally spaced? Spacing and start point matter.)

Q I. A/B testing with quantiles

First pick target quantile α (say 0.9).

$$H_0 : q_{0.9}(A) = q_{0.9}(B)$$

$$H_1 : q_{0.9}(A) < q_{0.9}(B)$$

Q1.A/B testing with quantiles

First pick target quantile α (say 0.9).

$$H_0 : q_{0.9}(A) = q_{0.9}(B)$$

$$H_1 : q_{0.9}(A) < q_{0.9}(B)$$

Can we construct a p-value?

Q I. A/B testing with quantiles

First pick target quantile α (say 0.9).

$$H_0 : q_{0.9}(A) = q_{0.9}(B)$$

$$H_1 : q_{0.9}(A) < q_{0.9}(B)$$

Can we construct a p-value?

If numerical, can we get a confidence interval for $q_{0.9}(B) - q_{0.9}(A)$?

Q2. Best-arm identification with quantiles



Which arm has the highest 80% quantile?

Q2. Best-arm identification with quantiles



Which arm has the highest 80% quantile?

Can we design MAB algorithms to adaptively determine the “best” arm with a prescribed failure probability?

Q2. Best-arm identification with quantiles



Which arm has the highest 80% quantile?

Can we design MAB algorithms to adaptively determine the “best” arm with a prescribed failure probability?


If the first arm is “special”, can we design MAB algorithms to adaptively test the null hypothesis that A is best, and get a p-value?

Commonly taught confidence intervals
and p-values are only valid (correctly control error)
if the sample size is fixed in advance.

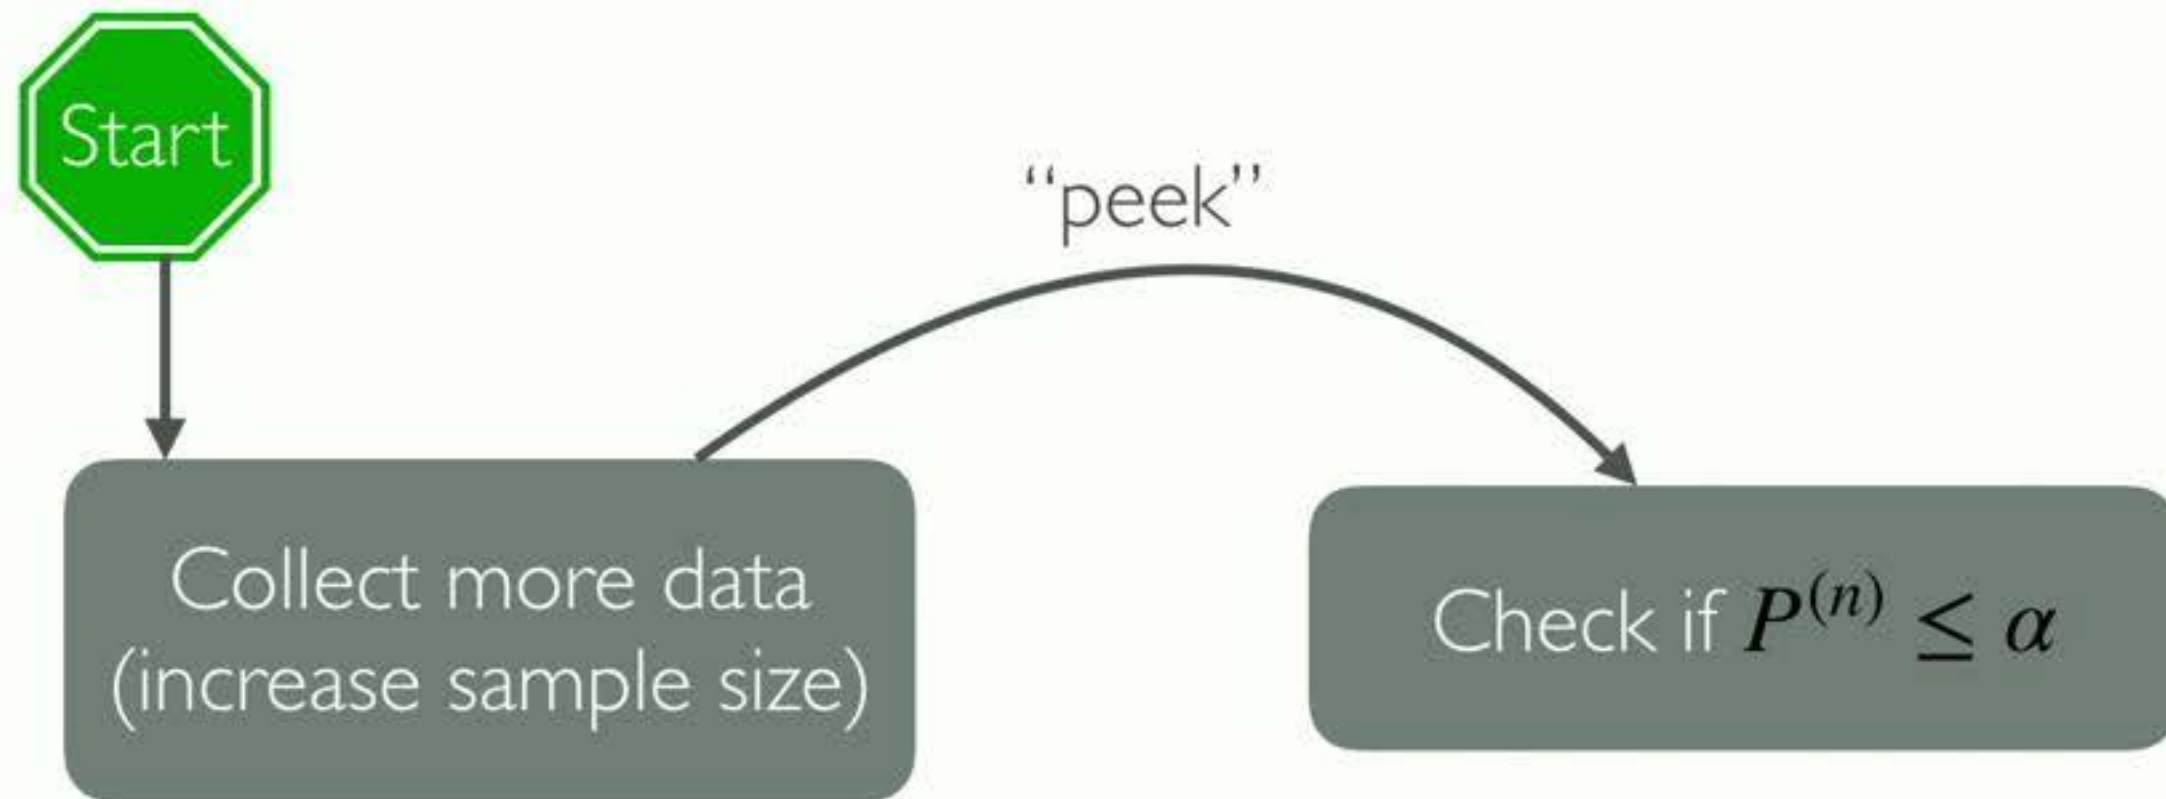
High-level caricature of an A/B-test



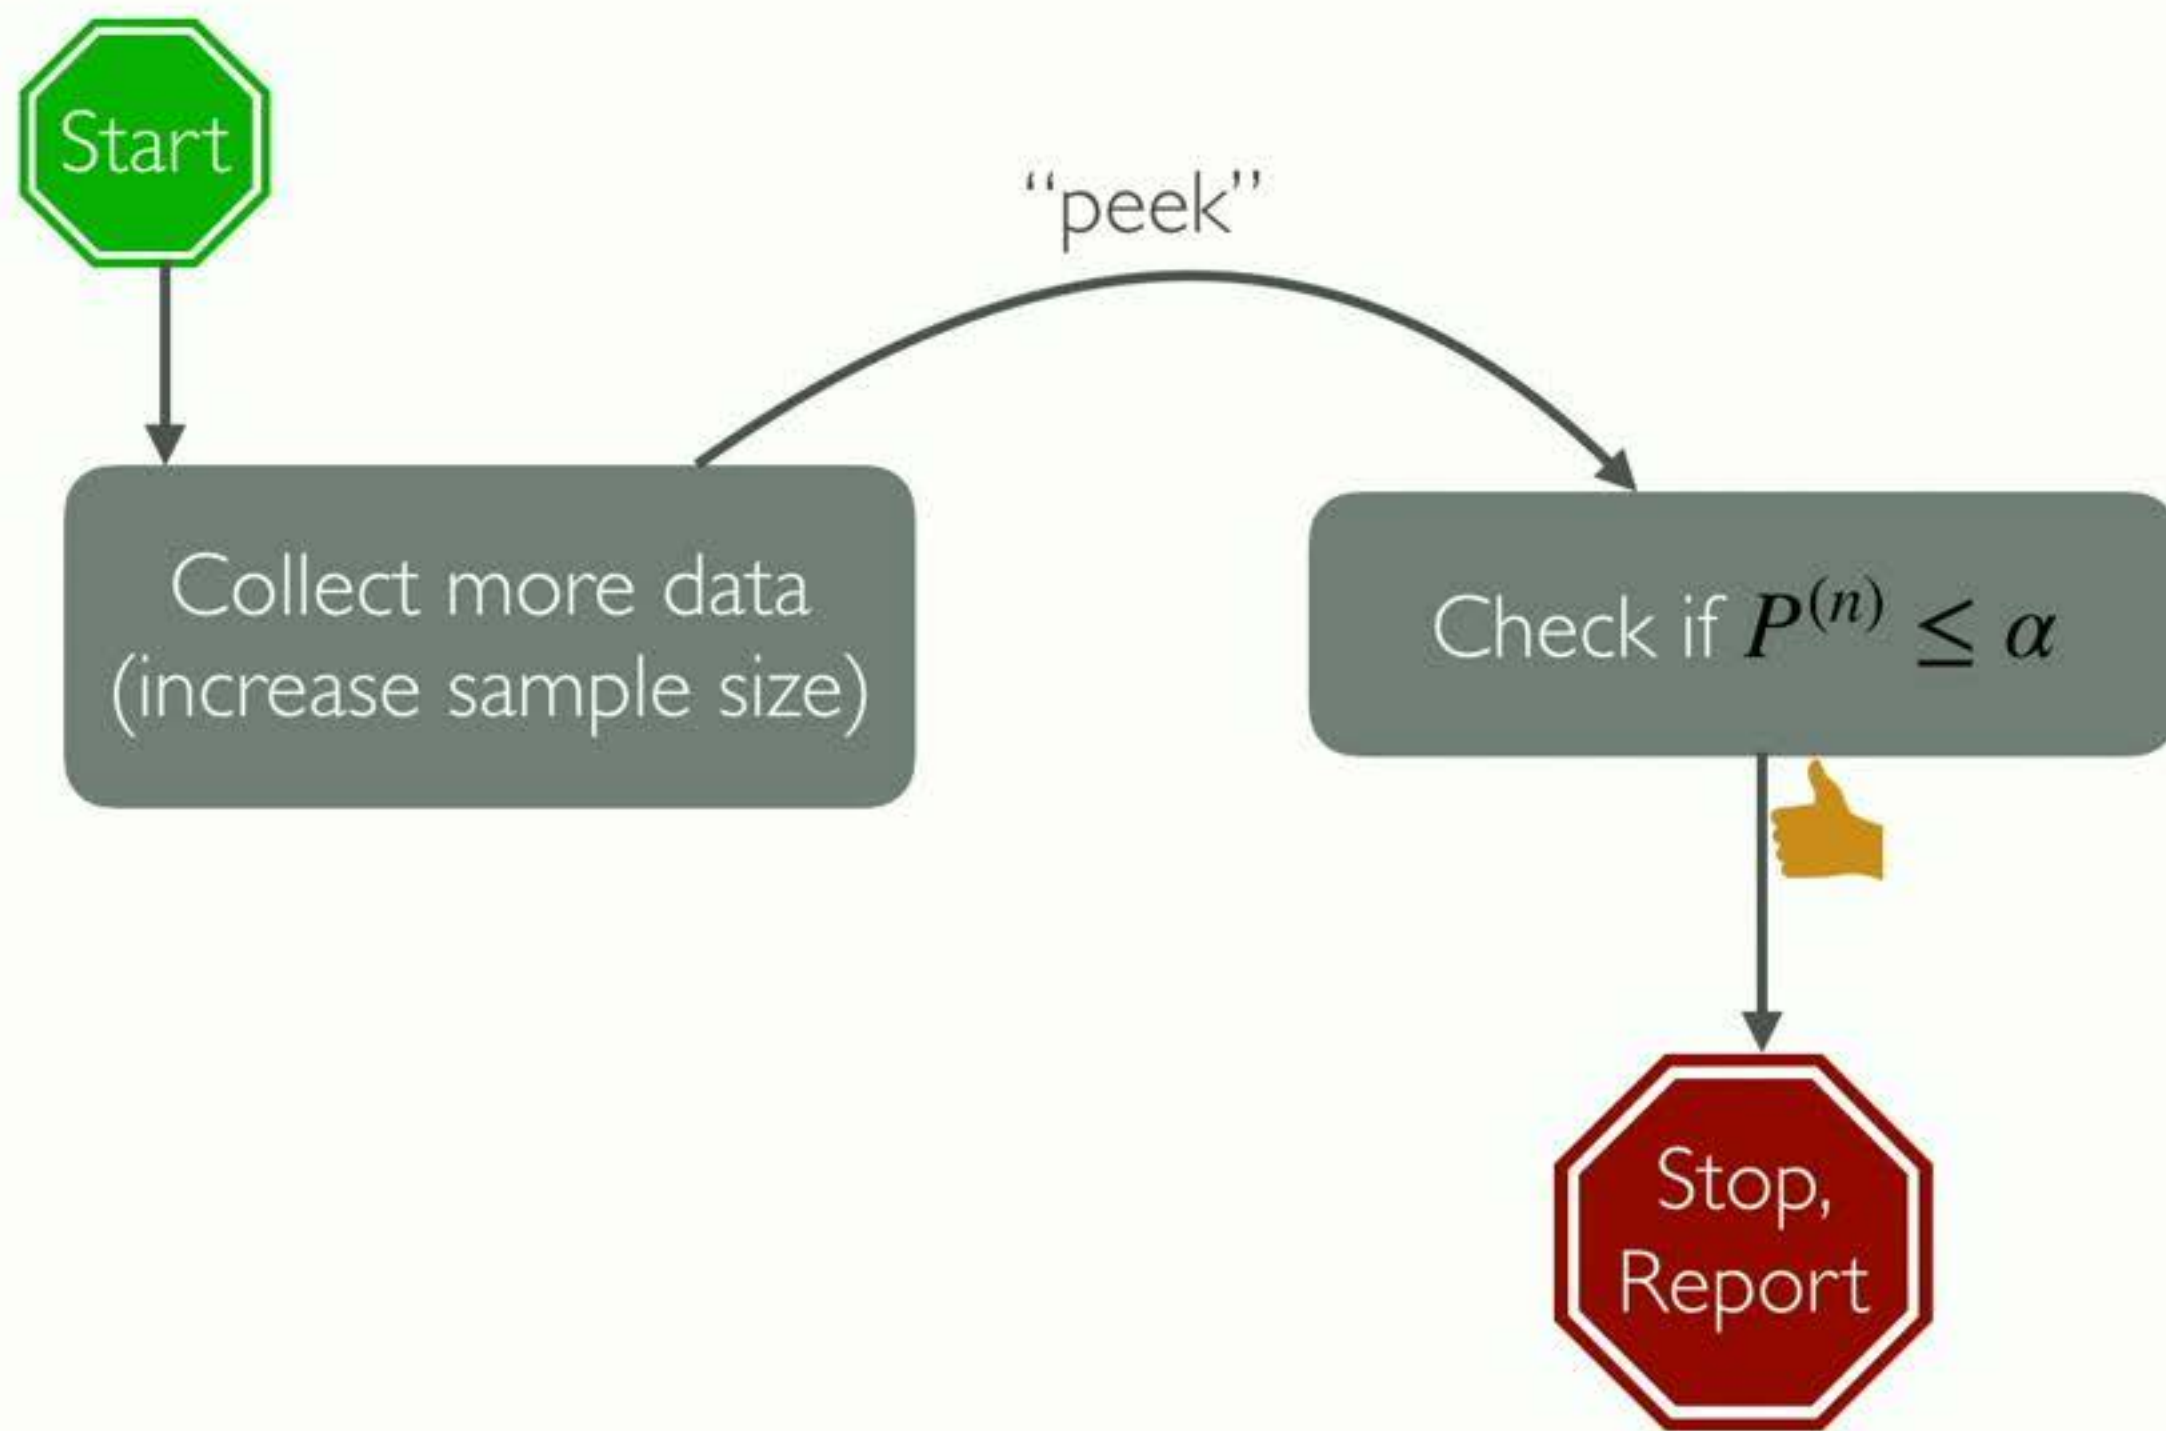
High-level caricature of an A/B-test



High-level caricature of an A/B-test

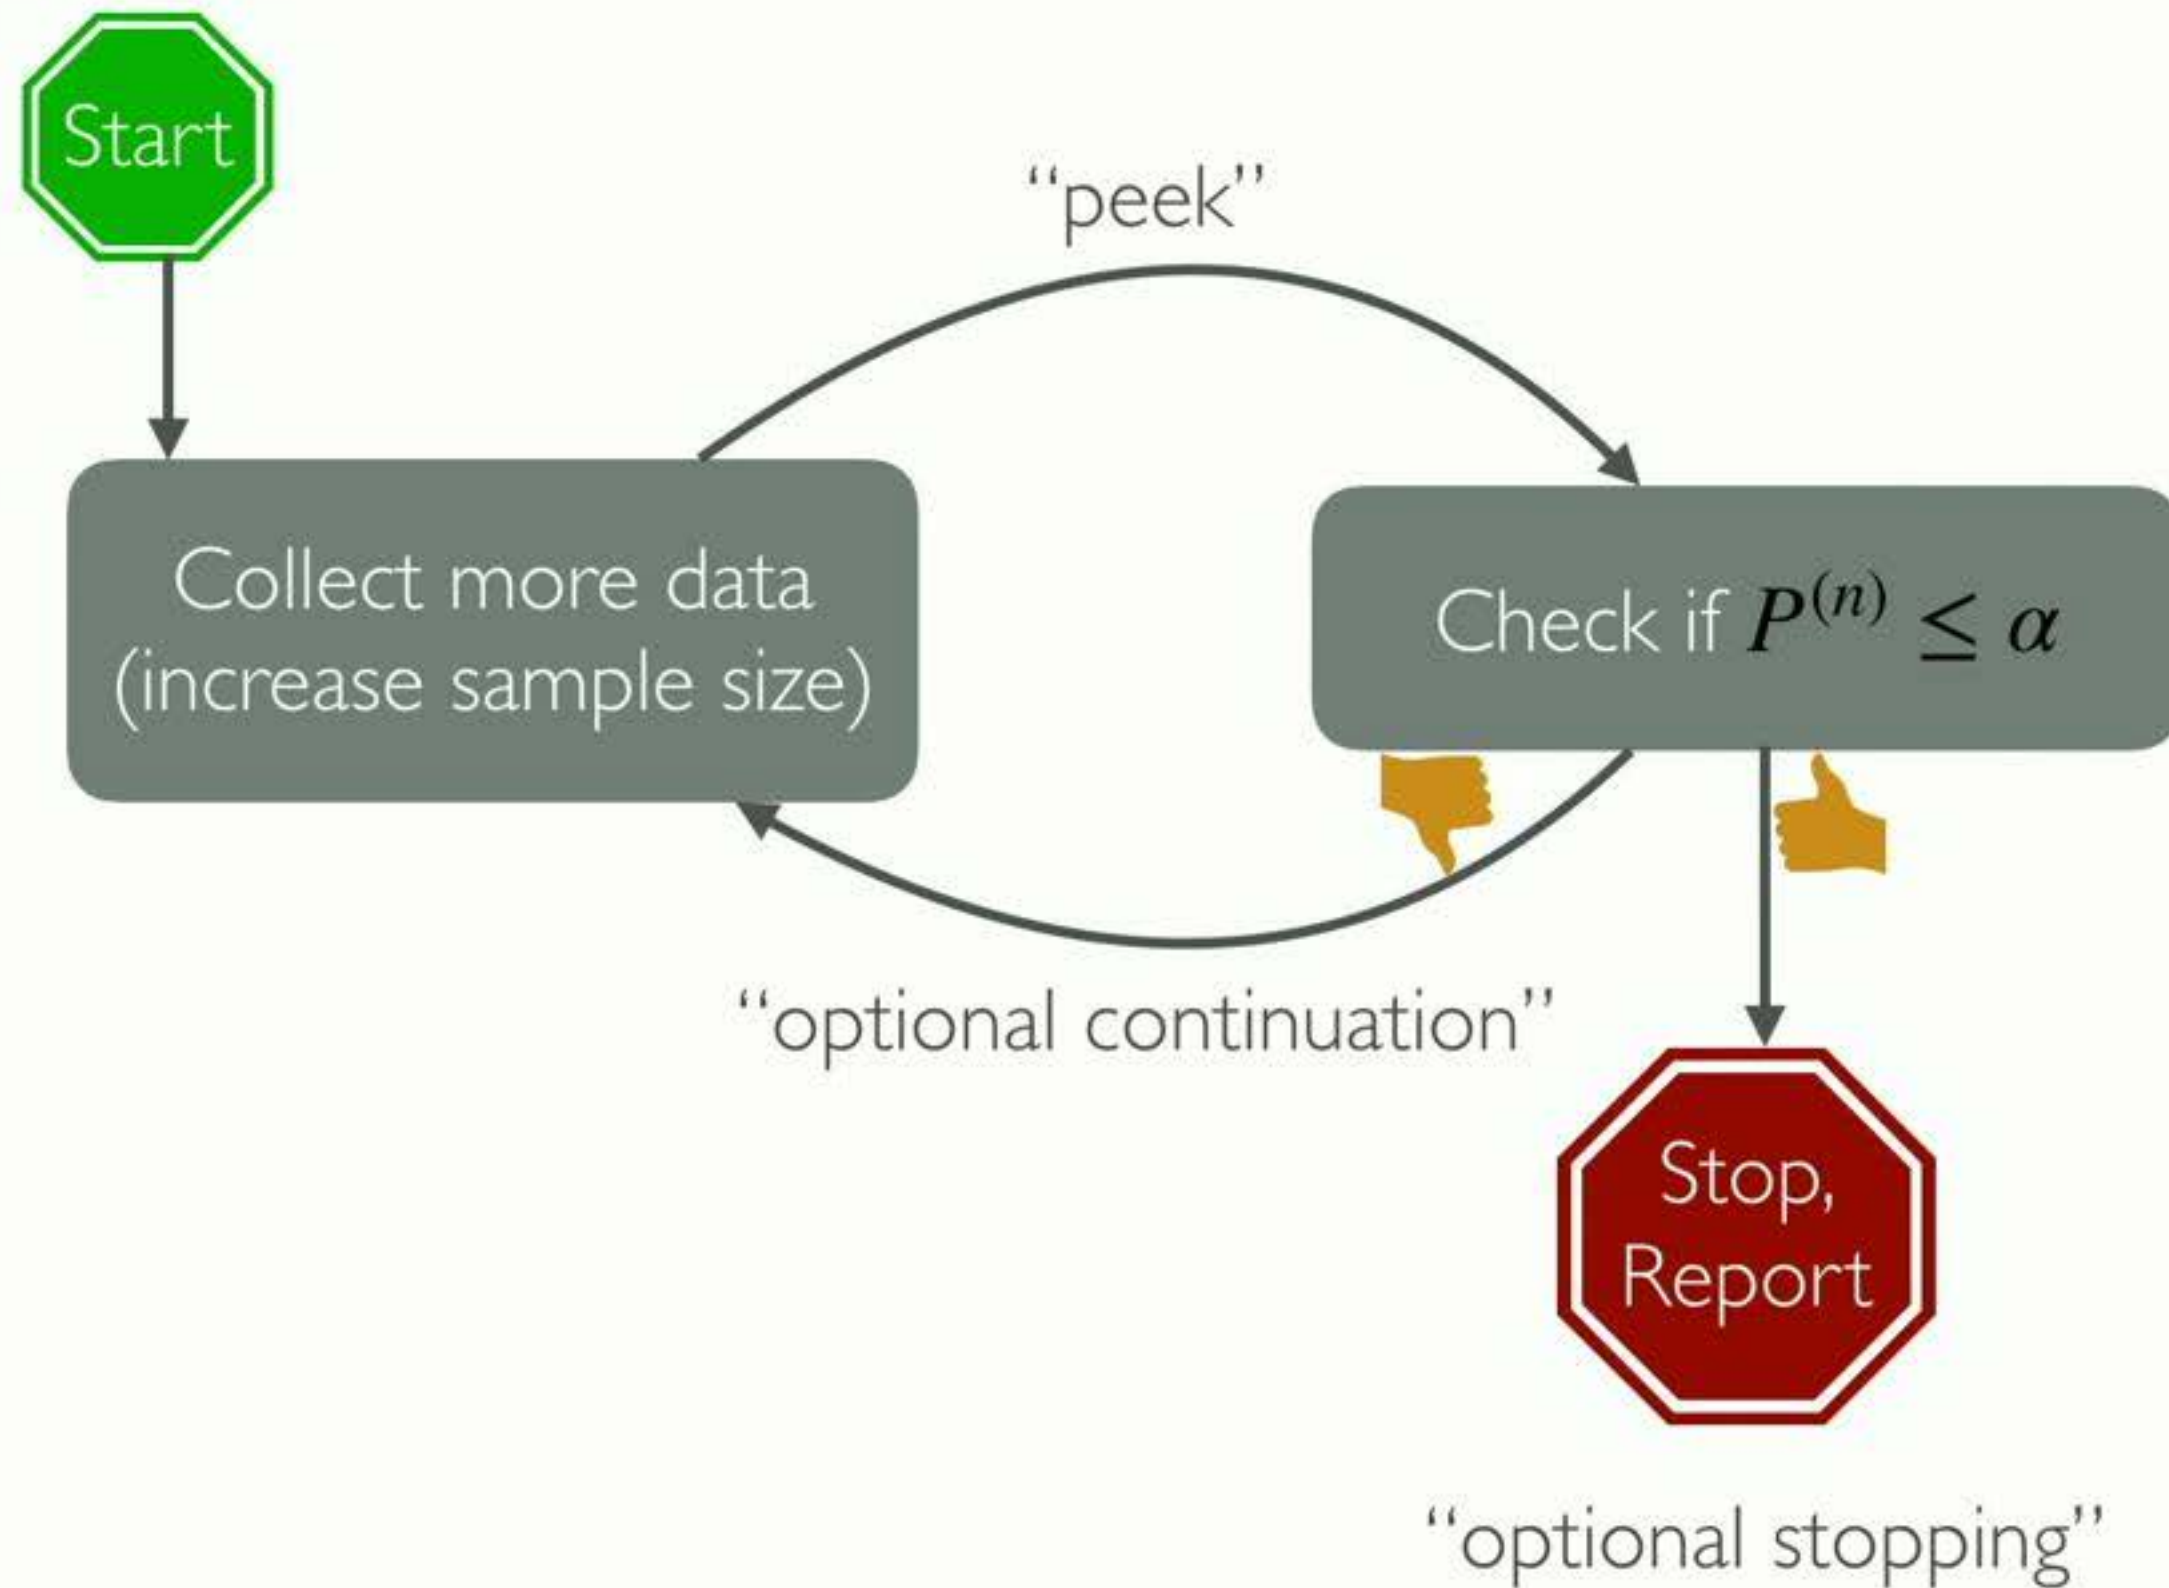


High-level caricature of an A/B-test

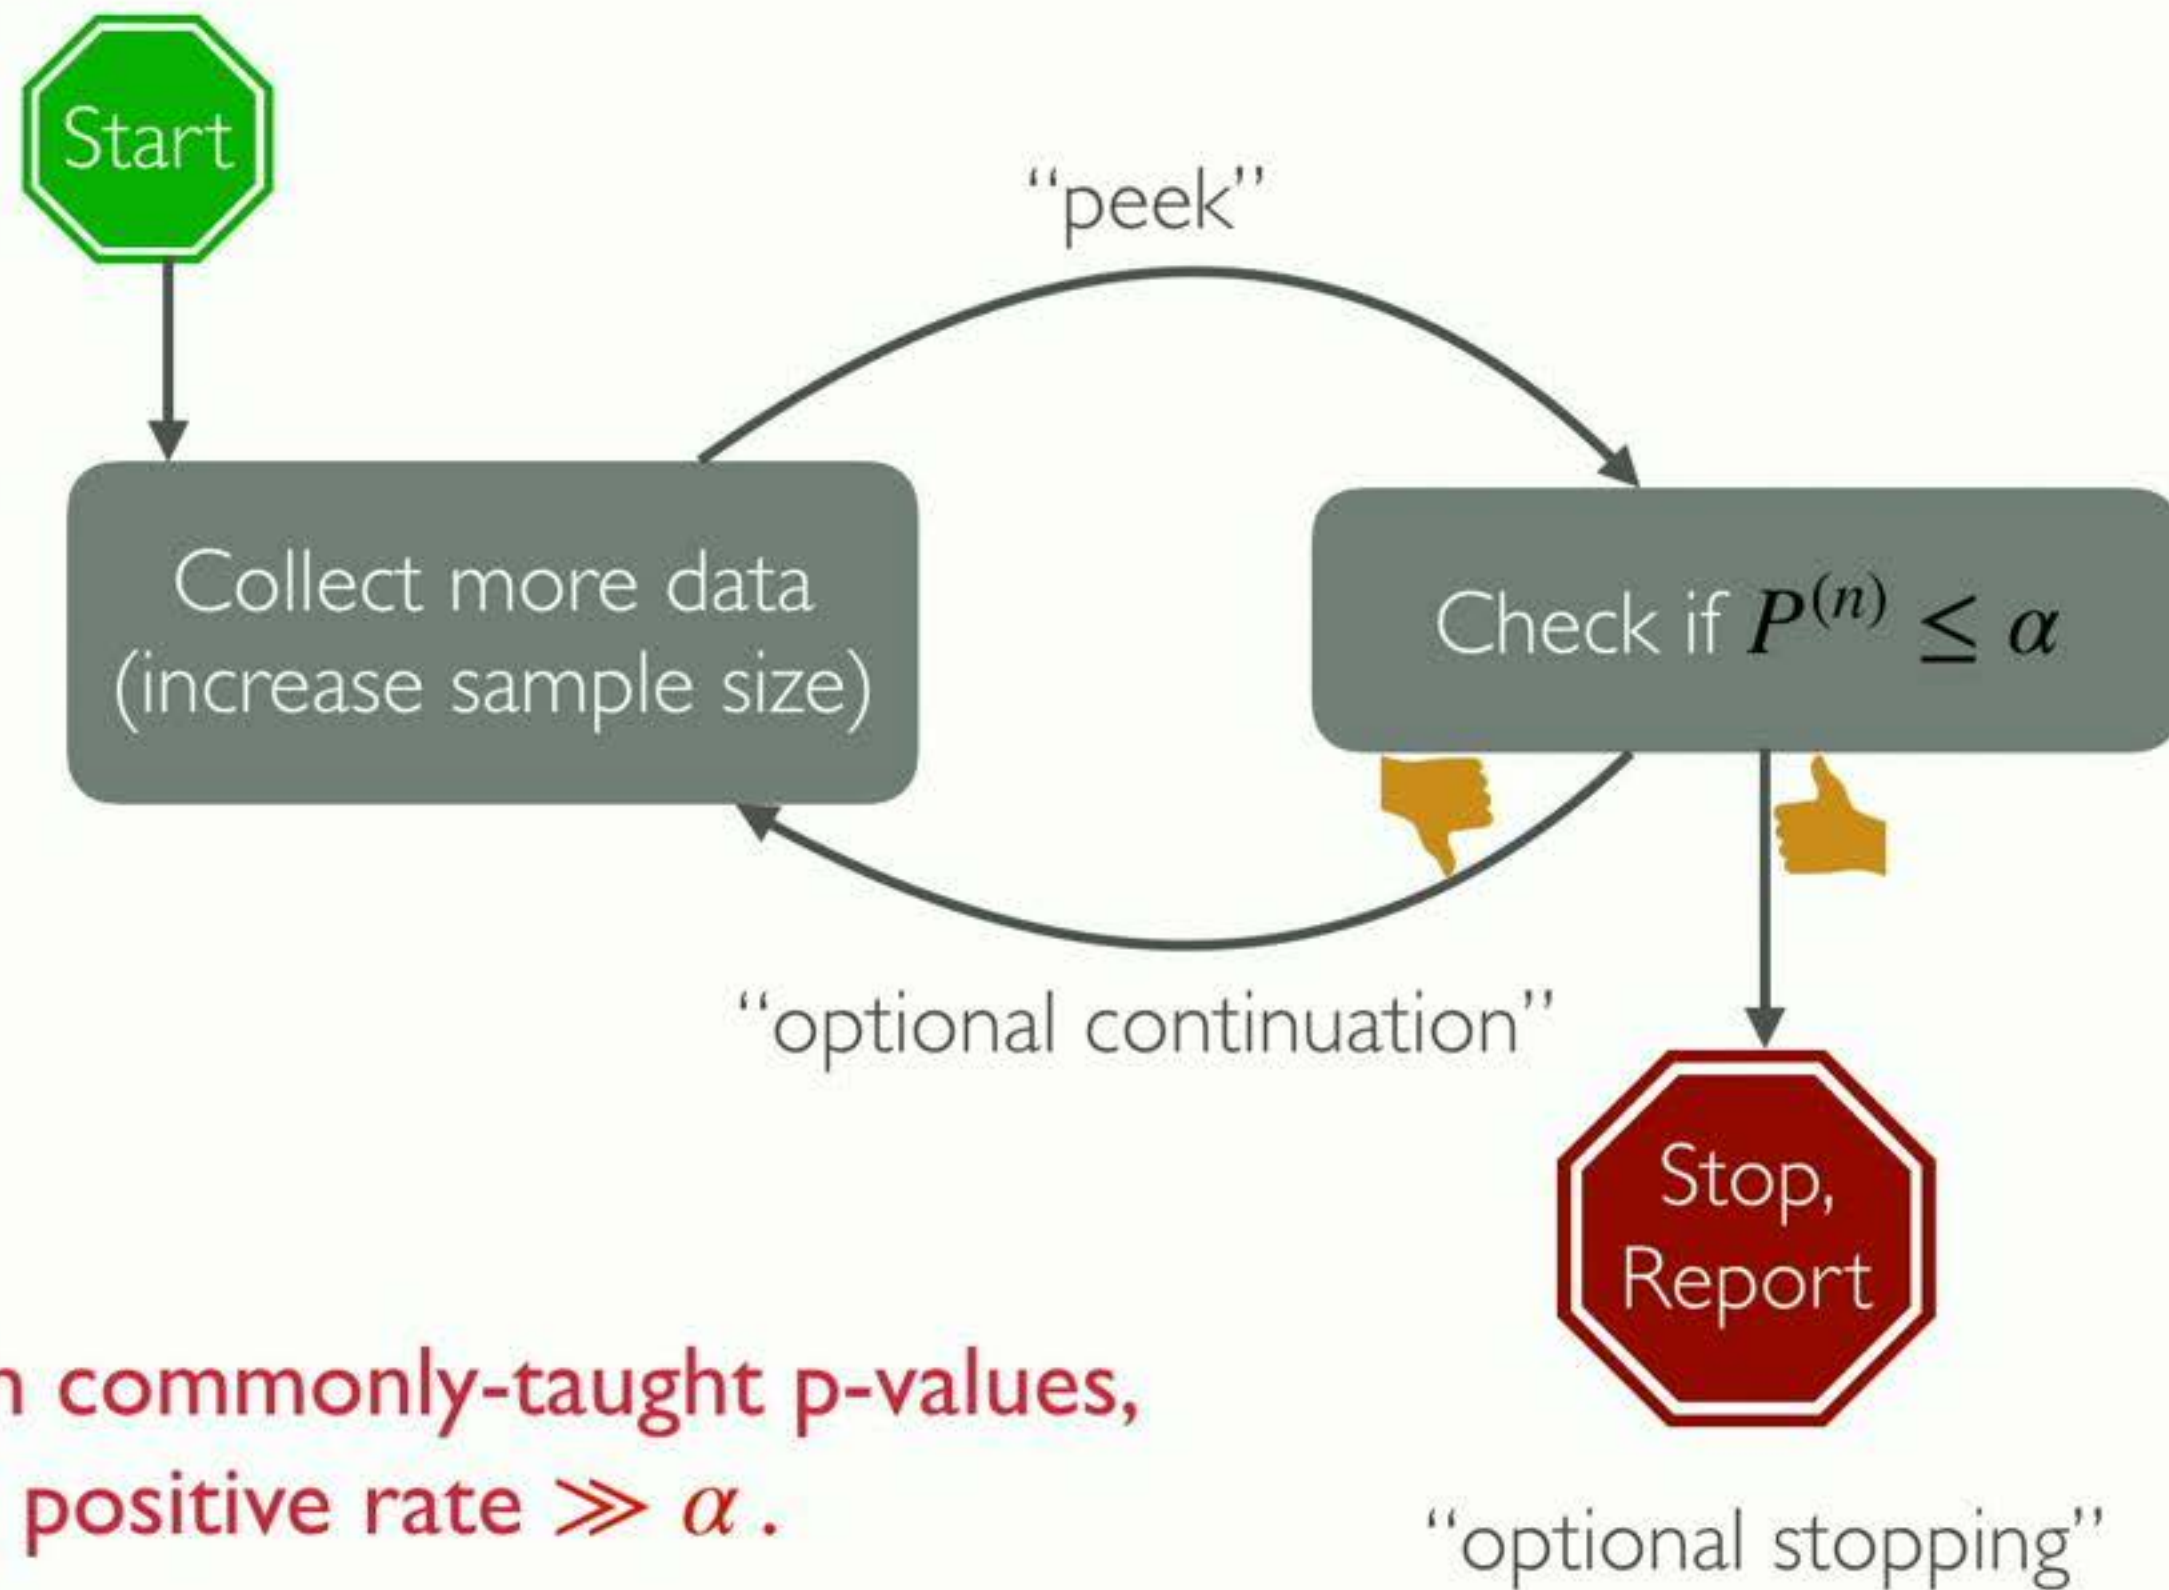


"optional stopping"

High-level caricature of an A/B-test



High-level caricature of an A/B-test



With commonly-taught p-values,
false positive rate $\gg \alpha$.



~~$P < 0.05$~~

After 10 people



~~$P < 0.05$~~

After 10 people

After 284 people



~~$P < 0.05$~~

~~$P < 0.05$~~

After 10 people

After 284 people





~~$P < 0.05$~~

~~$P < 0.05$~~

~~$P < 0.05$~~

After 10 people

After 284 people

After 1214 people

After 2398 people



~~$P < 0.05$~~

After 10 people

~~$P < 0.05$~~

After 284 people

~~$P < 0.05$~~

After 1214 people

~~$P < 0.05$~~

After 2398 people

After 7224 people



~~$P < 0.05$~~

After 10 people

~~$P < 0.05$~~

After 284 people

~~$P < 0.05$~~

After 1214 people

~~$P < 0.05$~~

After 2398 people

~~$P < 0.05$~~

After 7224 people



After 11,219 people, STOP!



~~$P < 0.05$~~

After 10 people

~~$P < 0.05$~~

After 284 people

~~$P < 0.05$~~


After 1214 people

~~$P < 0.05$~~

After 2398 people

~~$P < 0.05$~~

After 7224 people



After 11,219 people, STOP!

Let $P^{(n)}$ be a classical p-value (eg: t-test),
calculated using the first n samples.

Let $P^{(n)}$ be a classical p-value (eg: t-test),
calculated using the first n samples.

Under the null hypothesis (no treatment effect),

$$\forall n \geq 1, \quad \underbrace{\Pr(P^{(n)} \leq \alpha)}_{\text{prob. of false positive}} \leq \alpha.$$

Let $P^{(n)}$ be a classical p-value (eg: t-test),
calculated using the first n samples.

Under the null hypothesis (no treatment effect),

$$\forall n \geq 1, \quad \underbrace{\Pr(P^{(n)} \leq \alpha)}_{\text{prob. of false positive}} \leq \alpha.$$

Let τ be the stopping time of the experiment.

Often, τ depends on data, eg: $\tau := \min\{n \in \mathbb{N} : P_n \leq \alpha\}$.

Let $P^{(n)}$ be a classical p-value (eg: t-test),
calculated using the first n samples.

Under the null hypothesis (no treatment effect),

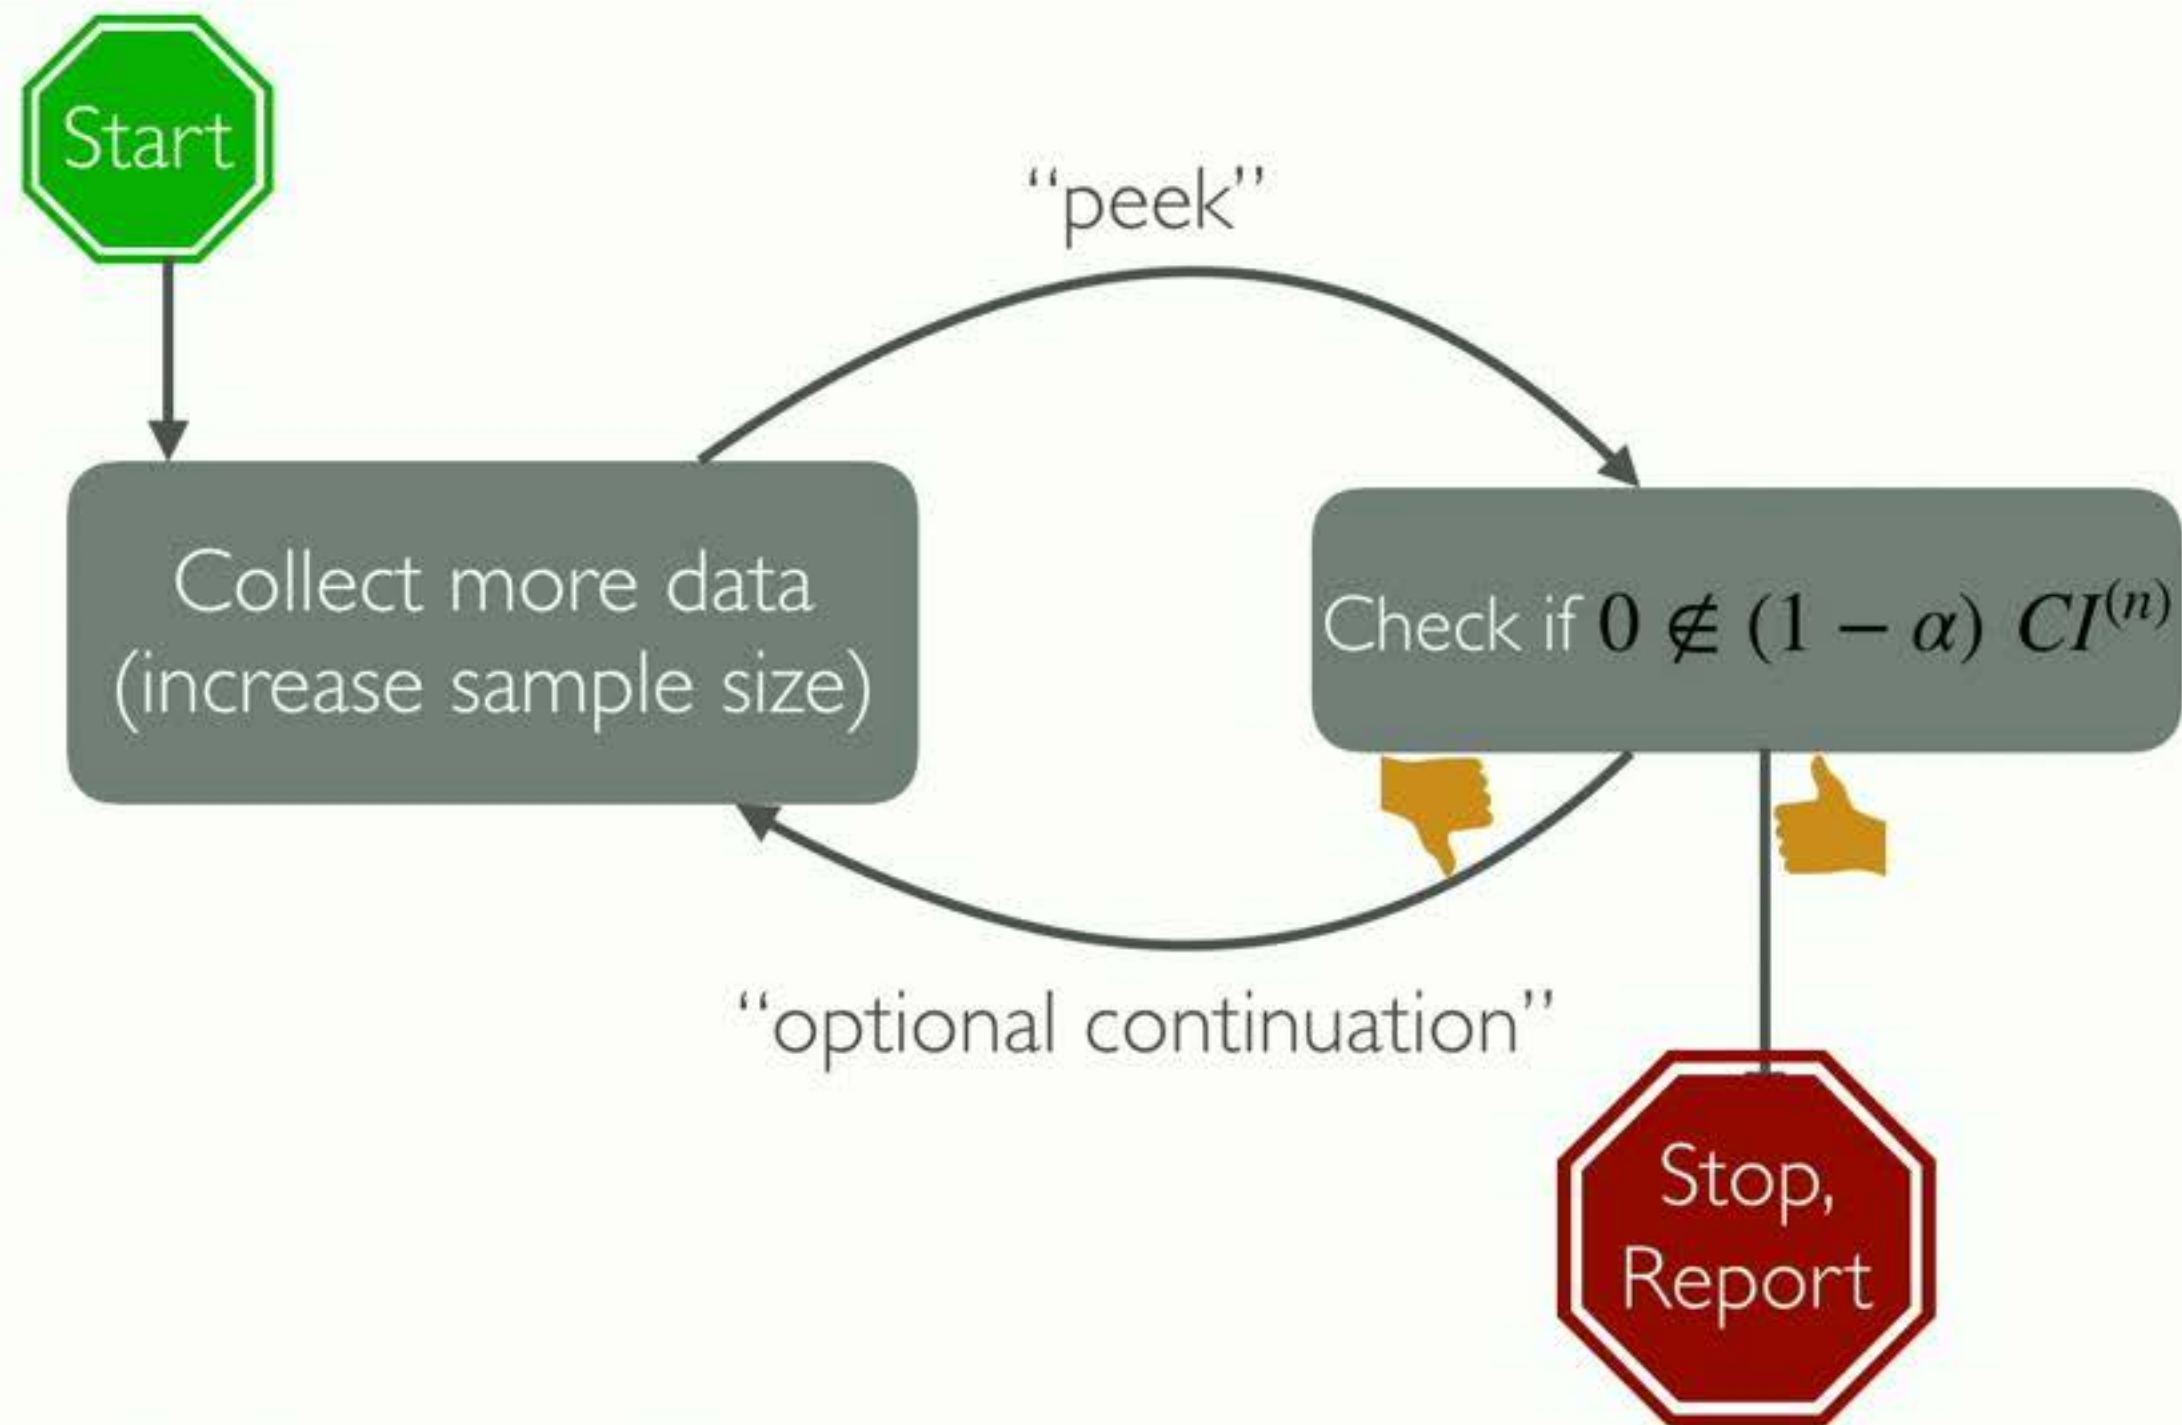
$$\forall n \geq 1, \quad \underbrace{\Pr(P^{(n)} \leq \alpha)}_{\text{prob. of false positive}} \leq \alpha.$$

Let τ be the stopping time of the experiment.

Often, τ depends on data, eg: $\tau := \min\{n \in \mathbb{N} : P_n \leq \alpha\}$.

Unfortunately, under the null, $\Pr(P^{(\tau)} \leq \alpha) \not\leq \alpha$.
In other words, $\Pr(\exists n \in \mathbb{N} : P^{(n)} \leq \alpha) \gg \alpha$.

Same problem with confidence interval (CI)



Again, false positive rate $\gg \alpha$.

"optional stopping"

Let $(L^{(n)}, U^{(n)})$ be any classical $(1 - \alpha)$ CI,
calculated using the first n samples (eg: CLT).

Let $(L^{(n)}, U^{(n)})$ be any classical $(1 - \alpha)$ CI, calculated using the first n samples (eg: CLT).

When trying to estimate the treatment effect θ ,

$$\forall n \geq 1, \underbrace{\Pr(\theta \in (L^{(n)}, U^{(n)}))}_{\text{prob. of coverage}} \geq 1 - \alpha.$$

Let $(L^{(n)}, U^{(n)})$ be any classical $(1 - \alpha)$ CI, calculated using the first n samples (eg: CLT).

When trying to estimate the treatment effect θ ,

$$\forall n \geq 1, \underbrace{\Pr(\theta \in (L^{(n)}, U^{(n)}))}_{\text{prob. of coverage}} \geq 1 - \alpha.$$

Let τ be the stopping time of the experiment.

Again, τ may depend on data, eg: $\tau := \min\{n \in \mathbb{N} : L^{(n)} > 0\}$.

$$\text{Unfortunately, } \Pr(\theta \in (L^{(\tau)}, U^{(\tau)})) \not\geq 1 - \alpha.$$

Let $(L^{(n)}, U^{(n)})$ be any classical $(1 - \alpha)$ CI, calculated using the first n samples (eg: CLT).

When trying to estimate the treatment effect θ ,

$$\forall n \geq 1, \underbrace{\Pr(\theta \in (L^{(n)}, U^{(n)}))}_{\text{prob. of coverage}} \geq 1 - \alpha.$$

Let τ be the stopping time of the experiment.

Again, τ may depend on data, eg: $\tau := \min\{n \in \mathbb{N} : L^{(n)} > 0\}$.

Unfortunately, $\Pr(\theta \in (L^{(\tau)}, U^{(\tau)})) \not\geq 1 - \alpha.$

In other words, $\Pr(\forall n \geq 1 : \theta \in (L^{(n)}, U^{(n)})) \ll 1 - \alpha.$
usually = 0.

Solution: “confidence sequence”
(aka “anytime confidence intervals”)

or “sequential p-values” for testing
(aka “always-valid p-values”)

A “**confidence sequence**” for a parameter θ is a sequence of confidence intervals (L_n, U_n) with a **uniform (simultaneous)** coverage guarantee.

$$\mathbb{P}(\underbrace{\forall n \geq 1}_{\text{Sample size}} : \theta \in (L_n, U_n)) \geq 1 - \alpha.$$

A “**confidence sequence**” for a parameter θ is a sequence of confidence intervals (L_n, U_n) with a **uniform (simultaneous)** coverage guarantee.

$$\mathbb{P}(\underbrace{\forall n \geq 1}_{\text{Sample size}} : \theta \in (L_n, U_n)) \geq 1 - \alpha.$$

Darling, Robbins '67, '68
Lai '76, '84
Howard, Ramdas, McAuliffe, Sekhon '18

Example: tracking the mean of a Gaussian or Bernoulli from i.i.d. observations.

$$X_1, X_2, \dots \sim N(\theta, 1) \text{ or } \text{Ber}(\theta)$$

Example: tracking the mean of a Gaussian or Bernoulli from i.i.d. observations.

$$X_1, X_2, \dots \sim N(\theta, 1) \text{ or } \text{Ber}(\theta)$$

Producing a confidence *interval* at a fixed time is elementary statistics (~100 years old).


Example: tracking the mean of a Gaussian or Bernoulli from i.i.d. observations.

$$X_1, X_2, \dots \sim N(\theta, 1) \text{ or } \text{Ber}(\theta)$$

Producing a confidence *interval* at a fixed time is elementary statistics (~ 100 years old).

How do we produce a confidence sequence?
(which is like a confidence band over time)

(Fair coin)



..... Pointwise CI (CLT)

— Anytime CI

Eg: If X_i is 1-subGaussian, then


$$\frac{\sum_{i=1}^n X_i}{n} \pm 1.71 \sqrt{\frac{\log \log(2n) + 0.72 \log(5.19/\alpha)}{n}}$$

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

Eg: If X_i is 1-subGaussian, then

$$\frac{\sum_{i=1}^n X_i}{n} \pm 1.71 \sqrt{\frac{\log \log(2n) + 0.72 \log(5.19/\alpha)}{n}}$$

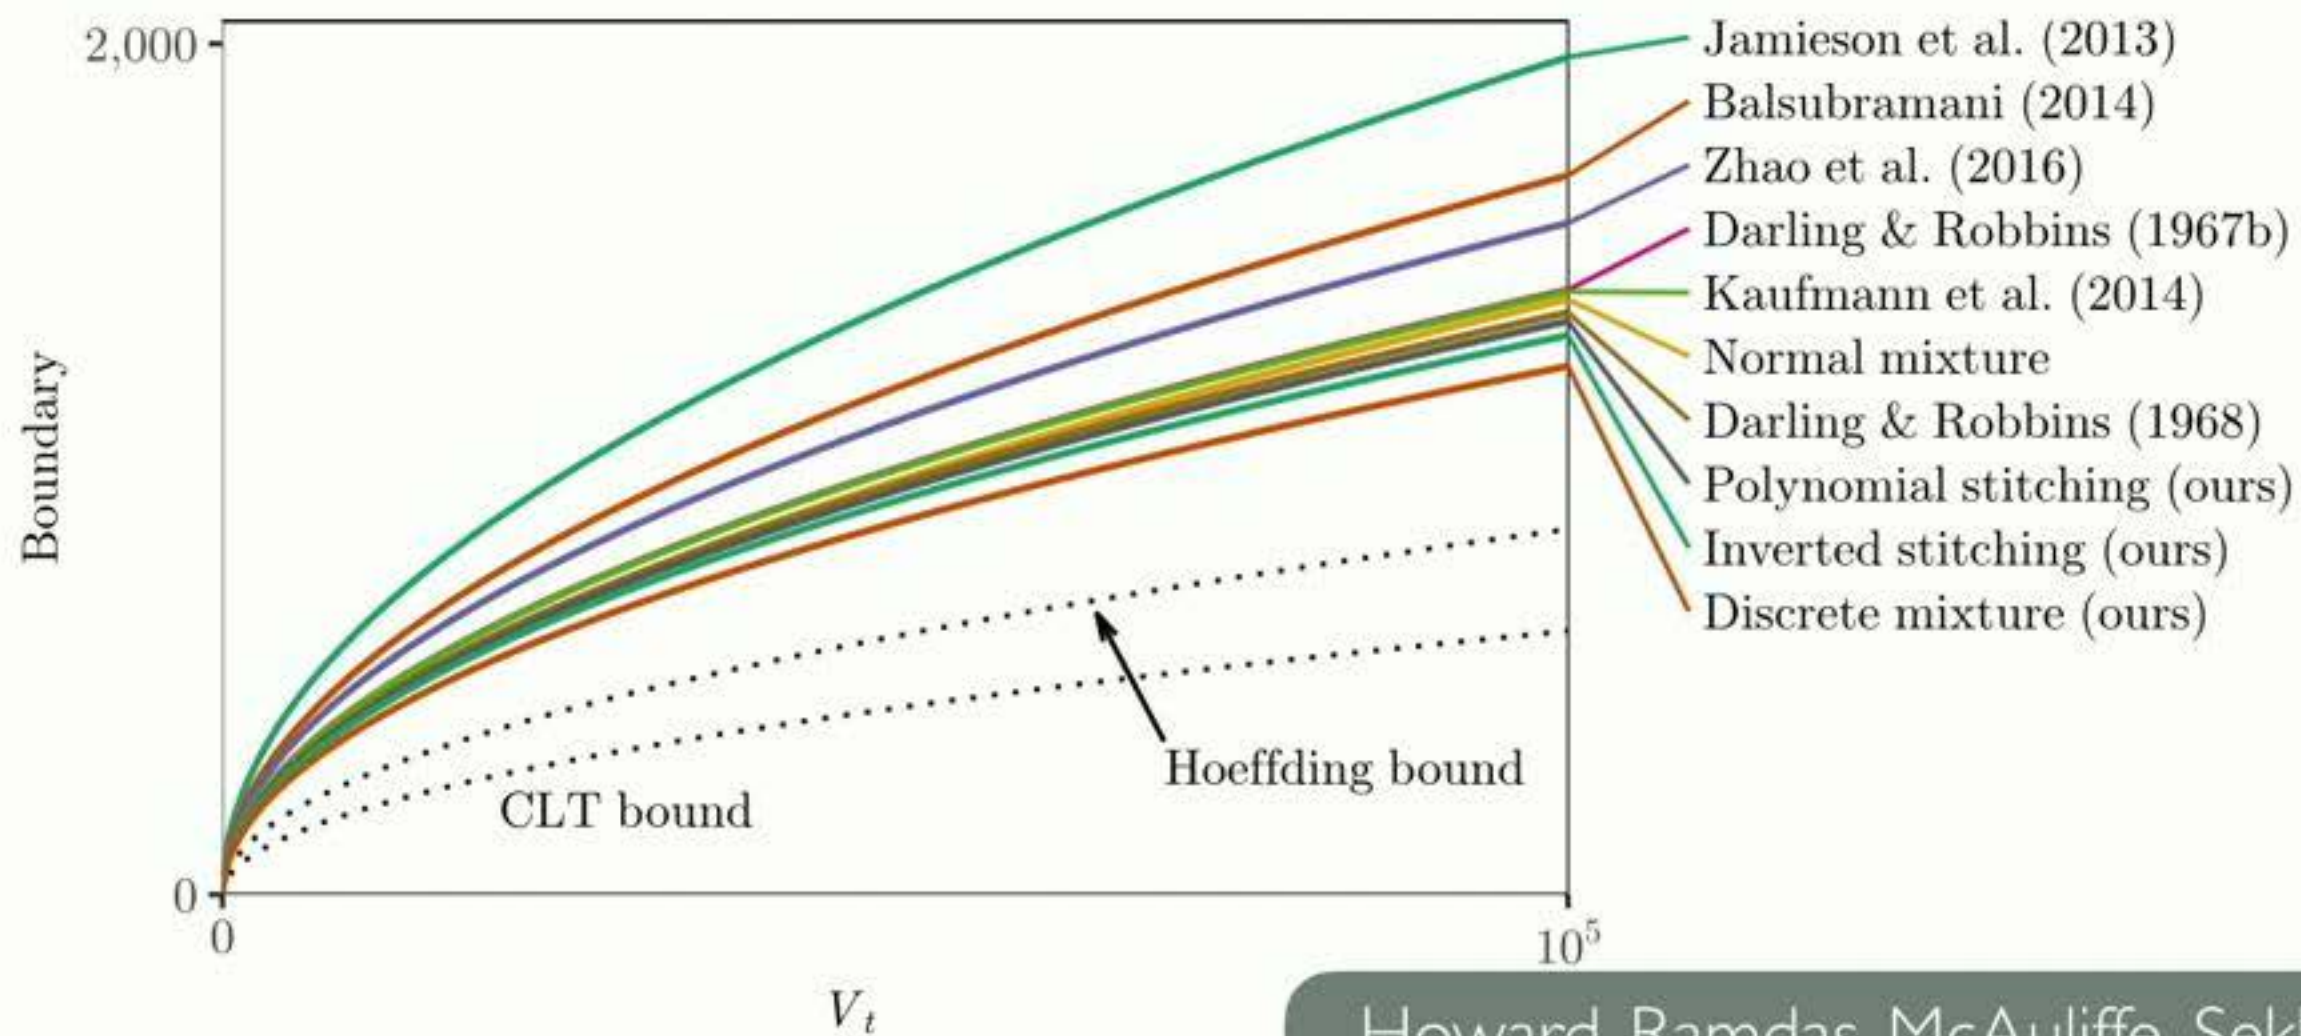
is a $(1 - \alpha)$ confidence sequence.



Eg: If X_i is 1-subGaussian, then

$$\frac{\sum_{i=1}^n X_i}{n} \pm 1.71 \sqrt{\frac{\log \log(2n) + 0.72 \log(5.19/\alpha)}{n}}$$

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



$$\mathbb{P}\left(\bigcup_{n \in \mathbb{N}} \{\theta \notin (L_n, U_n)\}\right) \leq \alpha.$$

Some implications:

$$\mathbb{P}\left(\bigcup_{n \in \mathbb{N}} \{\theta \notin (L_n, U_n)\}\right) \leq \alpha.$$

Some implications:

I. Valid inference at any time, even stopping times:

$$\mathbb{P}\left(\bigcup_{n \in \mathbb{N}} \{\theta \notin (L_n, U_n)\}\right) \leq \alpha.$$

Some implications:

I. Valid inference at any time, even stopping times:

For any stopping time τ : $\mathbb{P}(\theta \notin (L_\tau, U_\tau)) \leq \alpha$.

$$\mathbb{P}\left(\bigcup_{n \in \mathbb{N}} \{\theta \notin (L_n, U_n)\}\right) \leq \alpha.$$

Some implications:

1. Valid inference at any time, even stopping times:

For any stopping time τ : $\mathbb{P}(\theta \notin (L_\tau, U_\tau)) \leq \alpha$.

2. Valid post-hoc inference (in hindsight):

$$\mathbb{P}\left(\bigcup_{n \in \mathbb{N}} \{\theta \notin (L_n, U_n)\}\right) \leq \alpha.$$

Some implications:

1. Valid inference at any time, even stopping times:

For any stopping time τ : $\mathbb{P}(\theta \notin (L_\tau, U_\tau)) \leq \alpha$.

2. Valid post-hoc inference (in hindsight):

For any random time T : $\mathbb{P}(\theta \notin (L_T, U_T)) \leq \alpha$.

$$\mathbb{P}\left(\bigcup_{n \in \mathbb{N}} \{\theta \notin (L_n, U_n)\}\right) \leq \alpha.$$

Some implications:

1. Valid inference at any time, even stopping times:

For any stopping time τ : $\mathbb{P}(\theta \notin (L_\tau, U_\tau)) \leq \alpha$.

2. Valid post-hoc inference (in hindsight):

For any random time T : $\mathbb{P}(\theta \notin (L_T, U_T)) \leq \alpha$.

3. No pre-specified sample size:

can extend or stop experiments adaptively.

Eg: If X_i is 1-subGaussian, then

$$\frac{\sum_{i=1}^n X_i}{n} \pm 1.71 \sqrt{\frac{\log \log(2n) + 0.72 \log(5.19/\alpha)}{n}}$$

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

Duality between anytime p-value and CI

Duality between anytime p-value and CI

Define a set of null values \mathcal{H}_0 for θ .

Duality between anytime p-value and CI

Define a set of null values \mathcal{H}_0 for θ .

Let $P^{(n)} := \inf\{\alpha : \text{the } (1 - \alpha) \text{ CI}^{(n)} \text{ does not intersect } \mathcal{H}_0\}$

Duality between anytime p-value and CI

Define a set of null values \mathcal{H}_0 for θ .

Let $P^{(n)} := \inf\{\alpha : \text{the } (1 - \alpha) \text{ CI}^{(n)} \text{ does not intersect } \mathcal{H}_0\}$

If $CI^{(n)}$ is a pointwise CI then $P^{(n)}$ is a classical p-value.

For all fixed times n , $\underbrace{\Pr(P^{(n)} \leq \alpha)}_{\text{prob. of false positive}} \leq \alpha.$

Duality between anytime p-value and CI

Define a set of null values \mathcal{H}_0 for θ .

Let $P^{(n)} := \inf\{\alpha : \text{the } (1 - \alpha) \text{ CI}^{(n)} \text{ does not intersect } \mathcal{H}_0\}$

If $CI^{(n)}$ is a pointwise CI then $P^{(n)}$ is a classical p-value.

For all fixed times n , $\underbrace{\Pr(P^{(n)} \leq \alpha)}_{\text{prob. of false positive}} \leq \alpha.$

If $CI^{(n)}$ is an anytime CI then $P^{(n)}$ is an always-valid p-value.

For all stopping times τ , $\Pr(P^{(\tau)} \leq \alpha) \leq \alpha.$

For all data-dependent times T , $\Pr(P^{(T)} \leq \alpha) \leq \alpha.$

Can construct confidence sequences
(and hence always valid p-values)
in a wide variety of *nonparametric* settings
(eg: random variables that are
bounded, or subGaussian, or subexponential)

Confidence sequences for quantiles,
and for all quantiles simultaneously

Howard, Ramdas'19

Confidence sequence for fixed quantiles

Confidence sequence for fixed quantiles

Define $\ell_t := 1.4 \log \log(2.04t) + \log(9.97/\alpha)$ and $u_t := \sqrt{0.52t\ell(t)}$

Confidence sequence for fixed quantiles

Define $\ell_t := 1.4 \log \log(2.04t) + \log(9.97/\alpha)$ and $u_t := \sqrt{0.52t\ell(t)}$

Then $\Pr(\forall t \in \mathbb{N} : \widehat{Q}_t(1/2 - \frac{u_t}{t}) \leq Q(1/2) \leq \widehat{Q}_t(1/2 + \frac{u_t}{t})) \geq 1 - \alpha$.

Confidence sequence for fixed quantiles

Define $\ell_t := 1.4 \log \log(2.04t) + \log(9.97/\alpha)$ and $u_t := \sqrt{0.52t\ell(t)}$

Then $\Pr(\forall t \in \mathbb{N} : \widehat{Q}_t(1/2 - \frac{u_t}{t}) \leq Q(1/2) \leq \widehat{Q}_t(1/2 + \frac{u_t}{t})) \geq 1 - \alpha$.

For other fixed quantiles p , variance factor scales like $p(1-p)$

Confidence sequence for fixed quantiles

Define $\ell_t := 1.4 \log \log(2.04t) + \log(9.97/\alpha)$ and $u_t := \sqrt{0.52t\ell(t)}$

Then $\Pr(\forall t \in \mathbb{N} : \widehat{Q}_t(1/2 - \frac{u_t}{t}) \leq Q(1/2) \leq \widehat{Q}_t(1/2 + \frac{u_t}{t})) \geq 1 - \alpha$.

For other fixed quantiles p , variance factor scales like $p(1-p)$

Let $u_t(p) := \sqrt{2.06p(1-p)t\ell(t) + 0.16(1-2p)^2\ell^2(t) + 0.4(1-2p)\ell(t)}$

Confidence sequence for fixed quantiles

Define $\ell_t := 1.4 \log \log(2.04t) + \log(9.97/\alpha)$ and $u_t := \sqrt{0.52t\ell(t)}$

Then $\Pr(\forall t \in \mathbb{N} : \widehat{Q}_t(1/2 - \frac{u_t}{t}) \leq Q(1/2) \leq \widehat{Q}_t(1/2 + \frac{u_t}{t})) \geq 1 - \alpha$.

For other fixed quantiles p , variance factor scales like $p(1-p)$

Let $u_t(p) := \sqrt{2.06p(1-p)t\ell(t) + 0.16(1-2p)^2\ell^2(t) + 0.4(1-2p)\ell(t)}$

$\Pr(\forall t \in \mathbb{N} : \widehat{Q}_t(p - \frac{u_t(1-p)}{t}) \leq q \leq \widehat{Q}_t(p + \frac{u_t(p)}{t})) \geq 1 - \alpha$.

for any $q \in [Q^-(p), Q(p)]$.

Confidence sequence for fixed quantiles

Define $\ell_t := 1.4 \log \log(2.04t) + \log(9.97/\alpha)$ and $u_t := \sqrt{0.52t\ell(t)}$

Then $\Pr(\forall t \in \mathbb{N} : \widehat{Q}_t(1/2 - \frac{u_t}{t}) \leq Q(1/2) \leq \widehat{Q}_t(1/2 + \frac{u_t}{t})) \geq 1 - \alpha$.

For other fixed quantiles p , variance factor scales like $p(1-p)$

Let $u_t(p) := \sqrt{2.06p(1-p)t\ell(t) + 0.16(1-2p)^2\ell^2(t) + 0.4(1-2p)\ell(t)}$

$\Pr(\forall t \in \mathbb{N} : \widehat{Q}_t(p - \frac{u_t(1-p)}{t}) \leq q \leq \widehat{Q}_t(p + \frac{u_t(p)}{t})) \geq 1 - \alpha$.

for any $q \in [Q^-(p), Q(p)]$.

(Also, a numerical “mixture” boundary that has better constants.)

Confidence sequence for all quantiles simultaneously

Confidence sequence for all quantiles simultaneously

Define $u_t := \sqrt{\frac{\log \log(et) + 0.75 \log(34/\alpha)}{t}}$



$$\Pr(\forall t \in \mathbb{N}, p \in (0,1) : \widehat{Q}_t(p - u_t) \leq Q(p) \leq \widehat{Q}_t(p + u_t)) \geq 1 - \alpha.$$

Confidence sequence for all quantiles simultaneously

Define $u_t := \sqrt{\frac{\log \log(et) + 0.75 \log(34/\alpha)}{t}}$

$$\Pr(\forall t \in \mathbb{N}, p \in (0,1) : \widehat{Q}_t(p - u_t) \leq Q(p) \leq \widehat{Q}_t(p + u_t)) \geq 1 - \alpha.$$

(also, a “practical” uniform DKW inequality, and numerical methods)



Confidence sequence for fixed quantiles



Define $\ell_t := 1.4 \log \log(2.04t) + \log(9.97/\alpha)$ and $u_t := \sqrt{0.52t\ell(t)}$


Then $\Pr(\forall t \in \mathbb{N} : \widehat{Q}_t(1/2 - \frac{u_t}{t}) \leq Q(1/2) \leq \widehat{Q}_t(1/2 + \frac{u_t}{t})) \geq 1 - \alpha$.

Confidence sequence for all quantiles simultaneously

Define $u_t := \sqrt{\frac{\log \log(et) + 0.75 \log(34/\alpha)}{t}}$

$$\Pr(\forall t \in \mathbb{N}, p \in (0,1) : \widehat{Q}_t(p - u_t) \leq Q(p) \leq \widehat{Q}_t(p + u_t)) \geq 1 - \alpha.$$





UPPER confidence bound radius (scaled)

Quantile ϵ -best arm identification

For a fixed $\pi \in (0,1)$ and $\epsilon \in (0,1 - \pi)$, we call an arm ϵ -optimal

$$Q_{k^*}(\pi + \epsilon) \geq Q_j(\pi) \text{ for all } j \neq k^* .$$

Task: identify any ϵ -optimal arm w.p. $1 - \delta$, for given $\delta \in (0,1)$.

Input target quantile $\pi \in (0,1)$, approximation error $\epsilon \in (0,1 - \pi)$, and error probability $\delta \in (0,1)$.

Sample each arm once, set $N_{k,1} = 1$ for all $k \in [K]$ and set $t = 1$.

while $L_{k,t}^{\pi+\epsilon} < \max_{j \neq k} U_{j,t}^{\pi}$ for all $k \in [K]$ **do**,

 Set $h_t \in \arg \max_{k \in [K]} L_{k,t}^{\pi+\epsilon}$ and $l_t \in \arg \max_{k \in [K] \setminus h_t} U_{k,t}^{\pi}$.

 Sample arms h_t and l_t .

 Set $N_{k,t+1} = N_{k,t} + 1$ if $k = h_t$ or $k = l_t$, and $N_{k,t+1} = N_{k,t}$ otherwise.

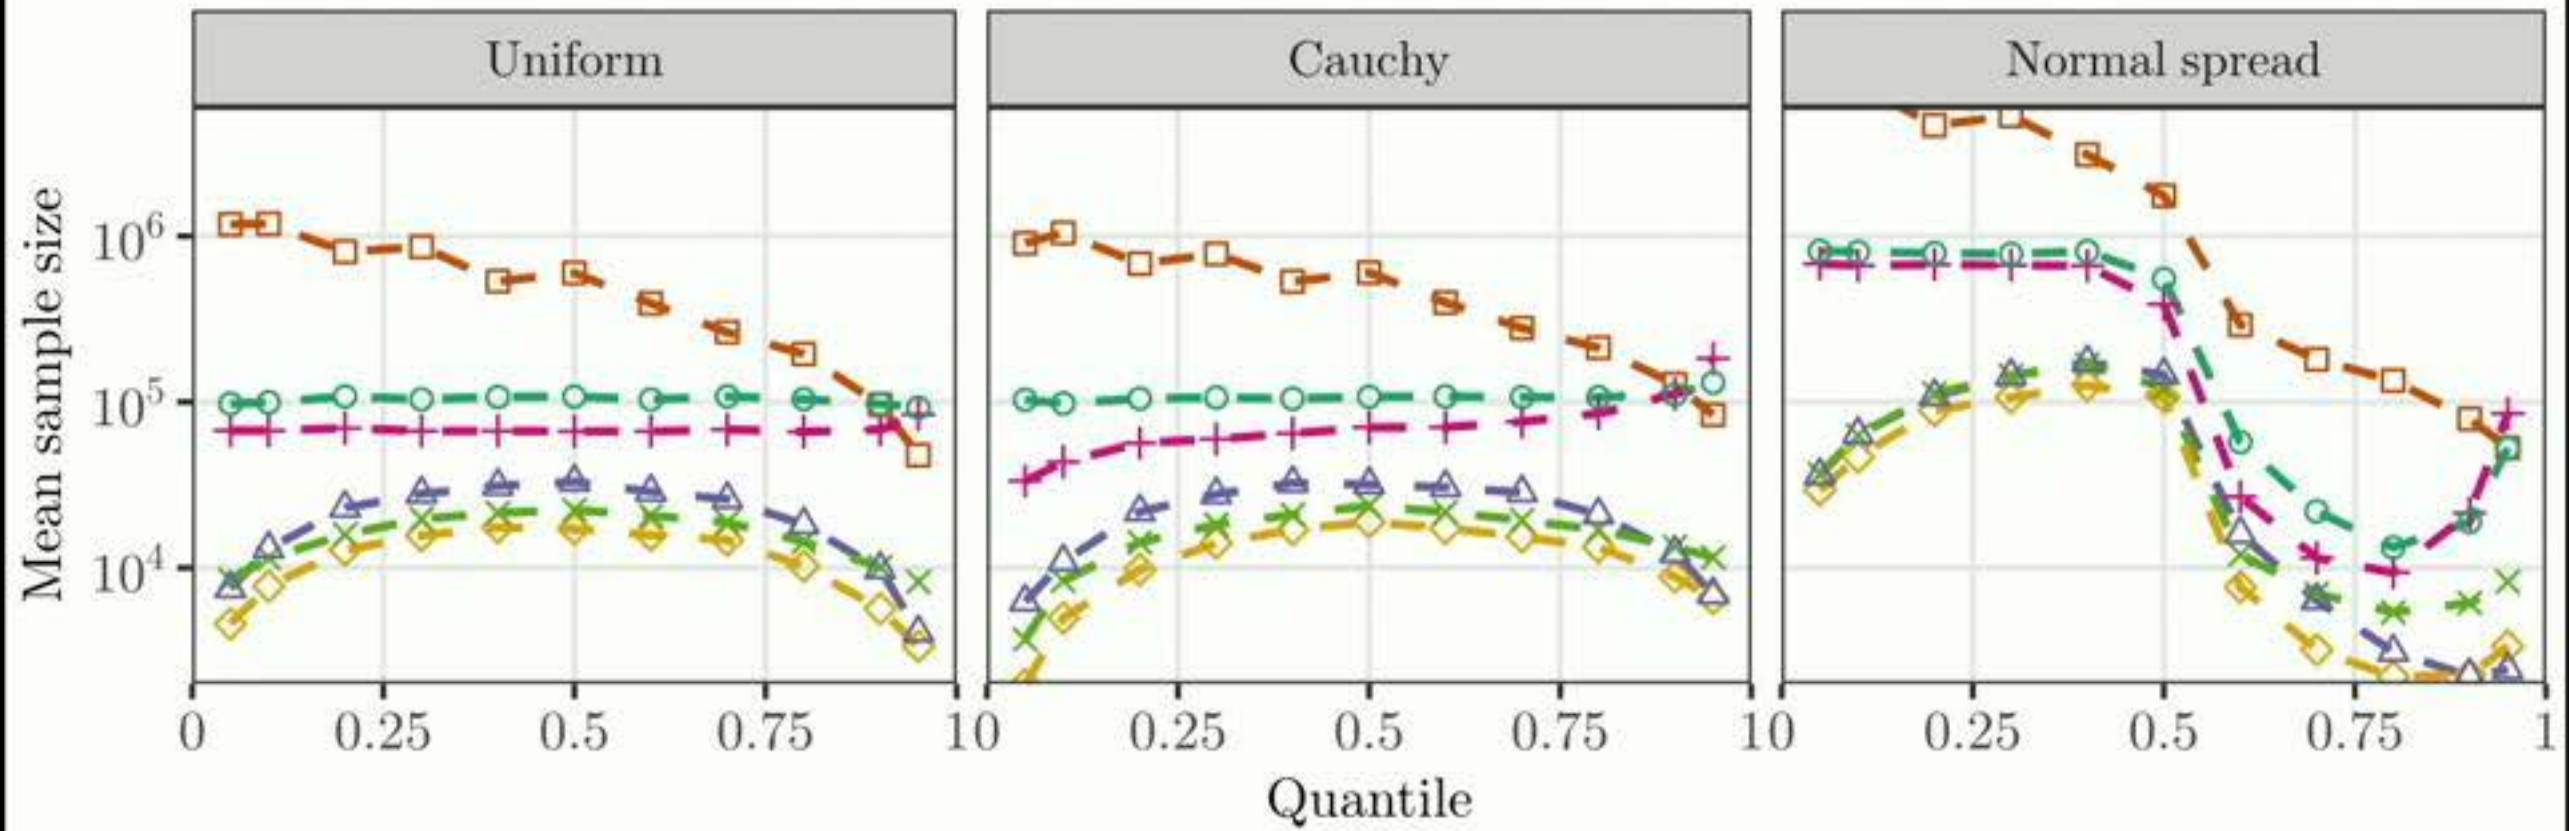
 Increment $t \leftarrow t + 1$.

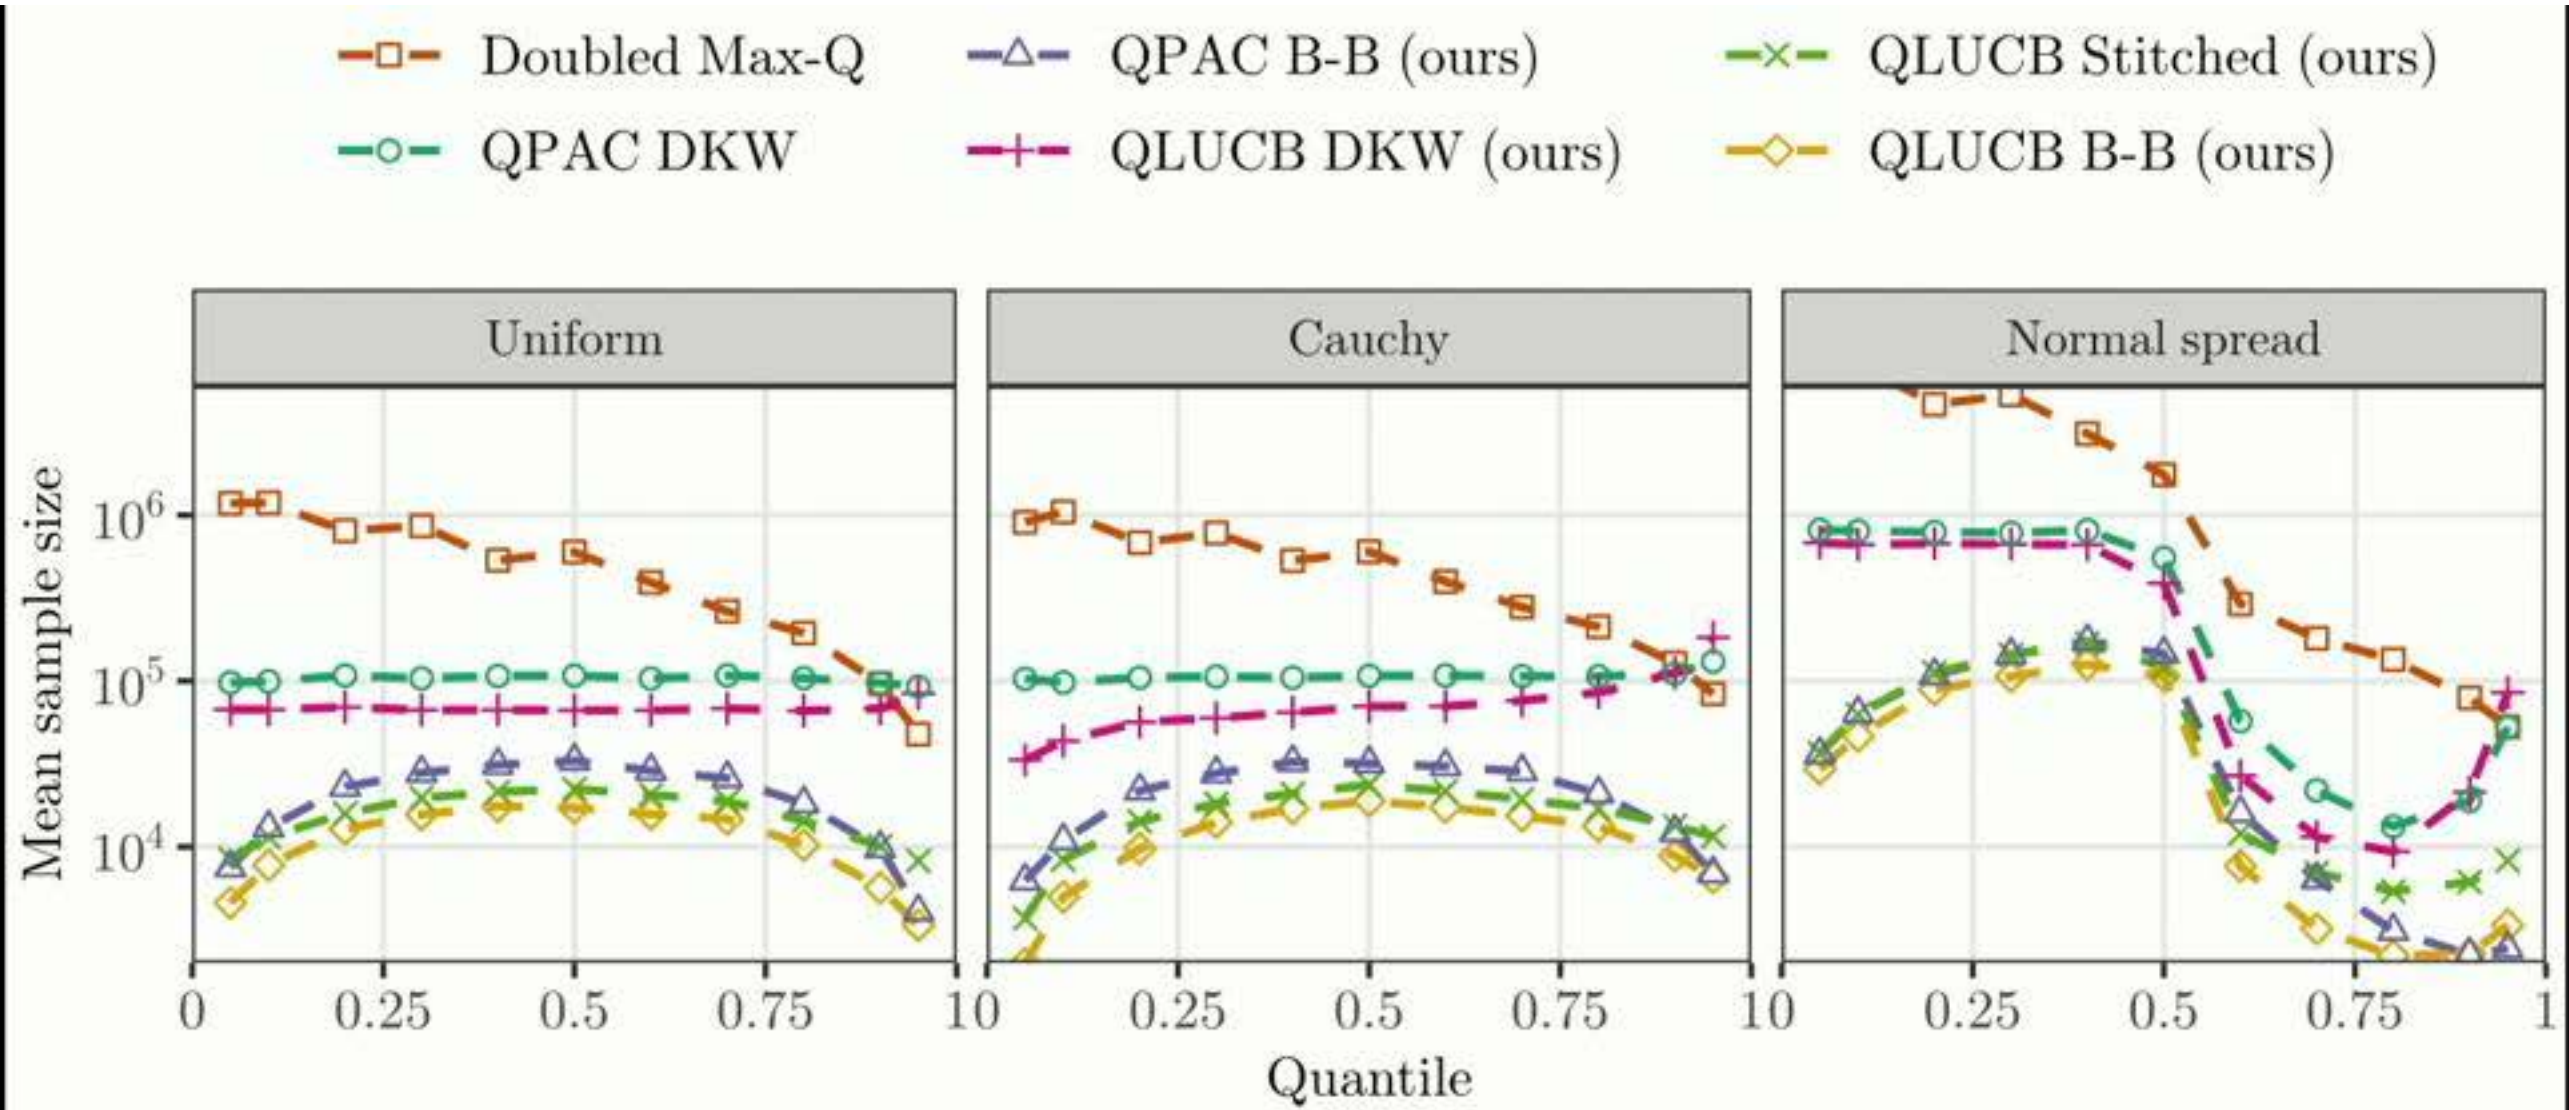
end while

Output any element of $\arg \max_{k \in [K]} L_{k,t}^{\pi+\epsilon}$.

QLUCB (our variant of standard LUCB)

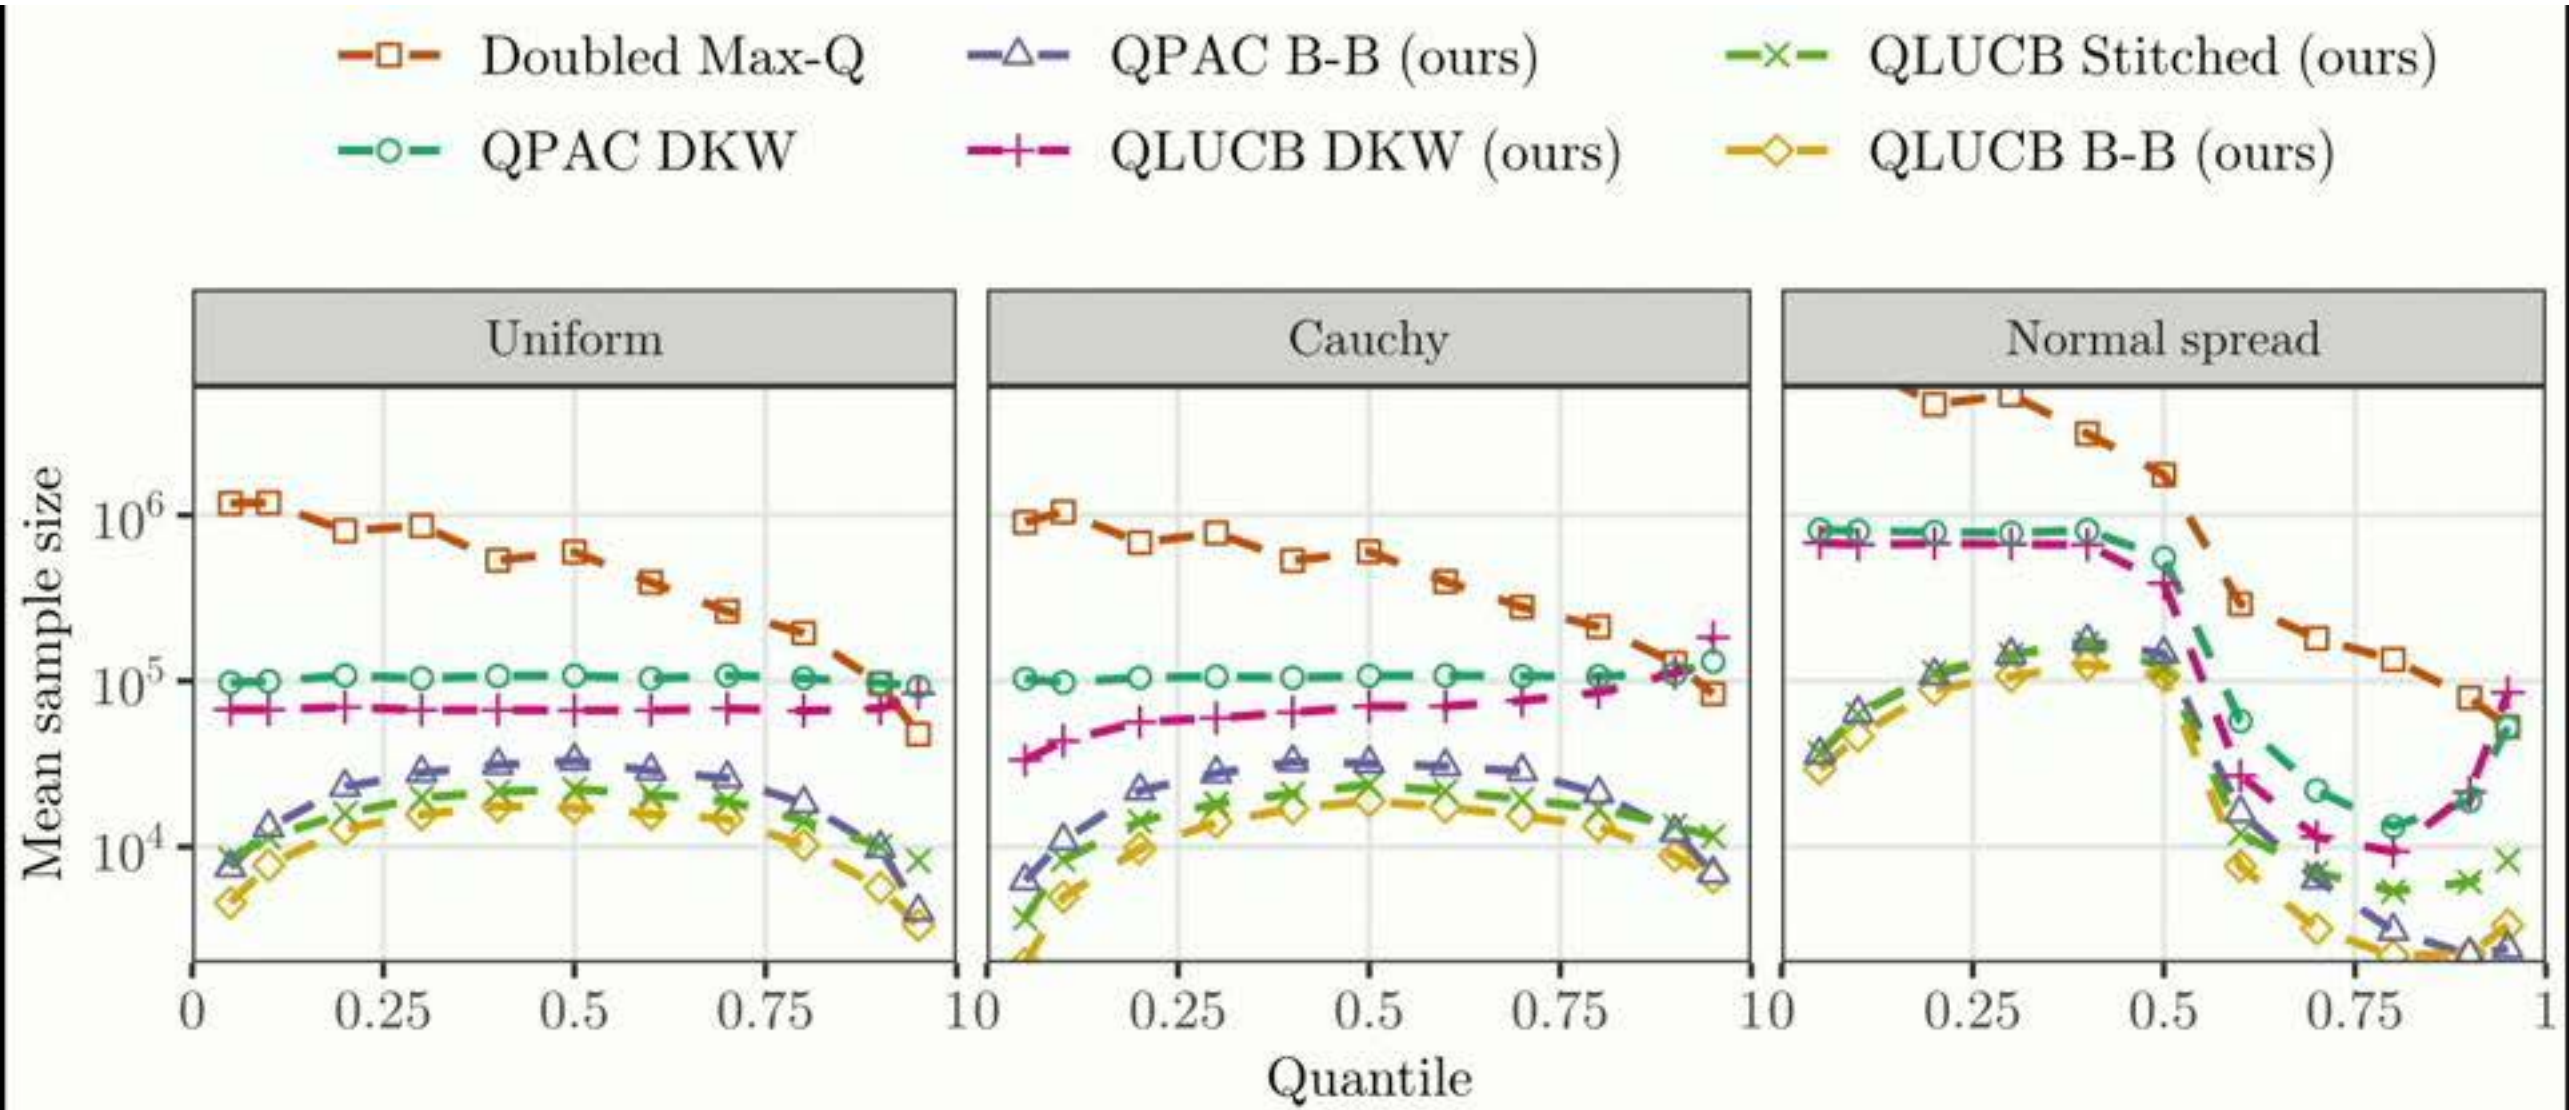
- Doubled Max-Q
- △— QPAC B-B (ours)
- ×— QLUCB Stitched (ours)
- QPAC DKW
- +— QLUCB DKW (ours)
- ◇— QLUCB B-B (ours)





Theorem:
$$\Delta_k := \sup \left\{ \Delta \geq 0 : Q_k(\pi + \Delta) < \max_{j \in [K]} Q_j(\pi) \right\}$$

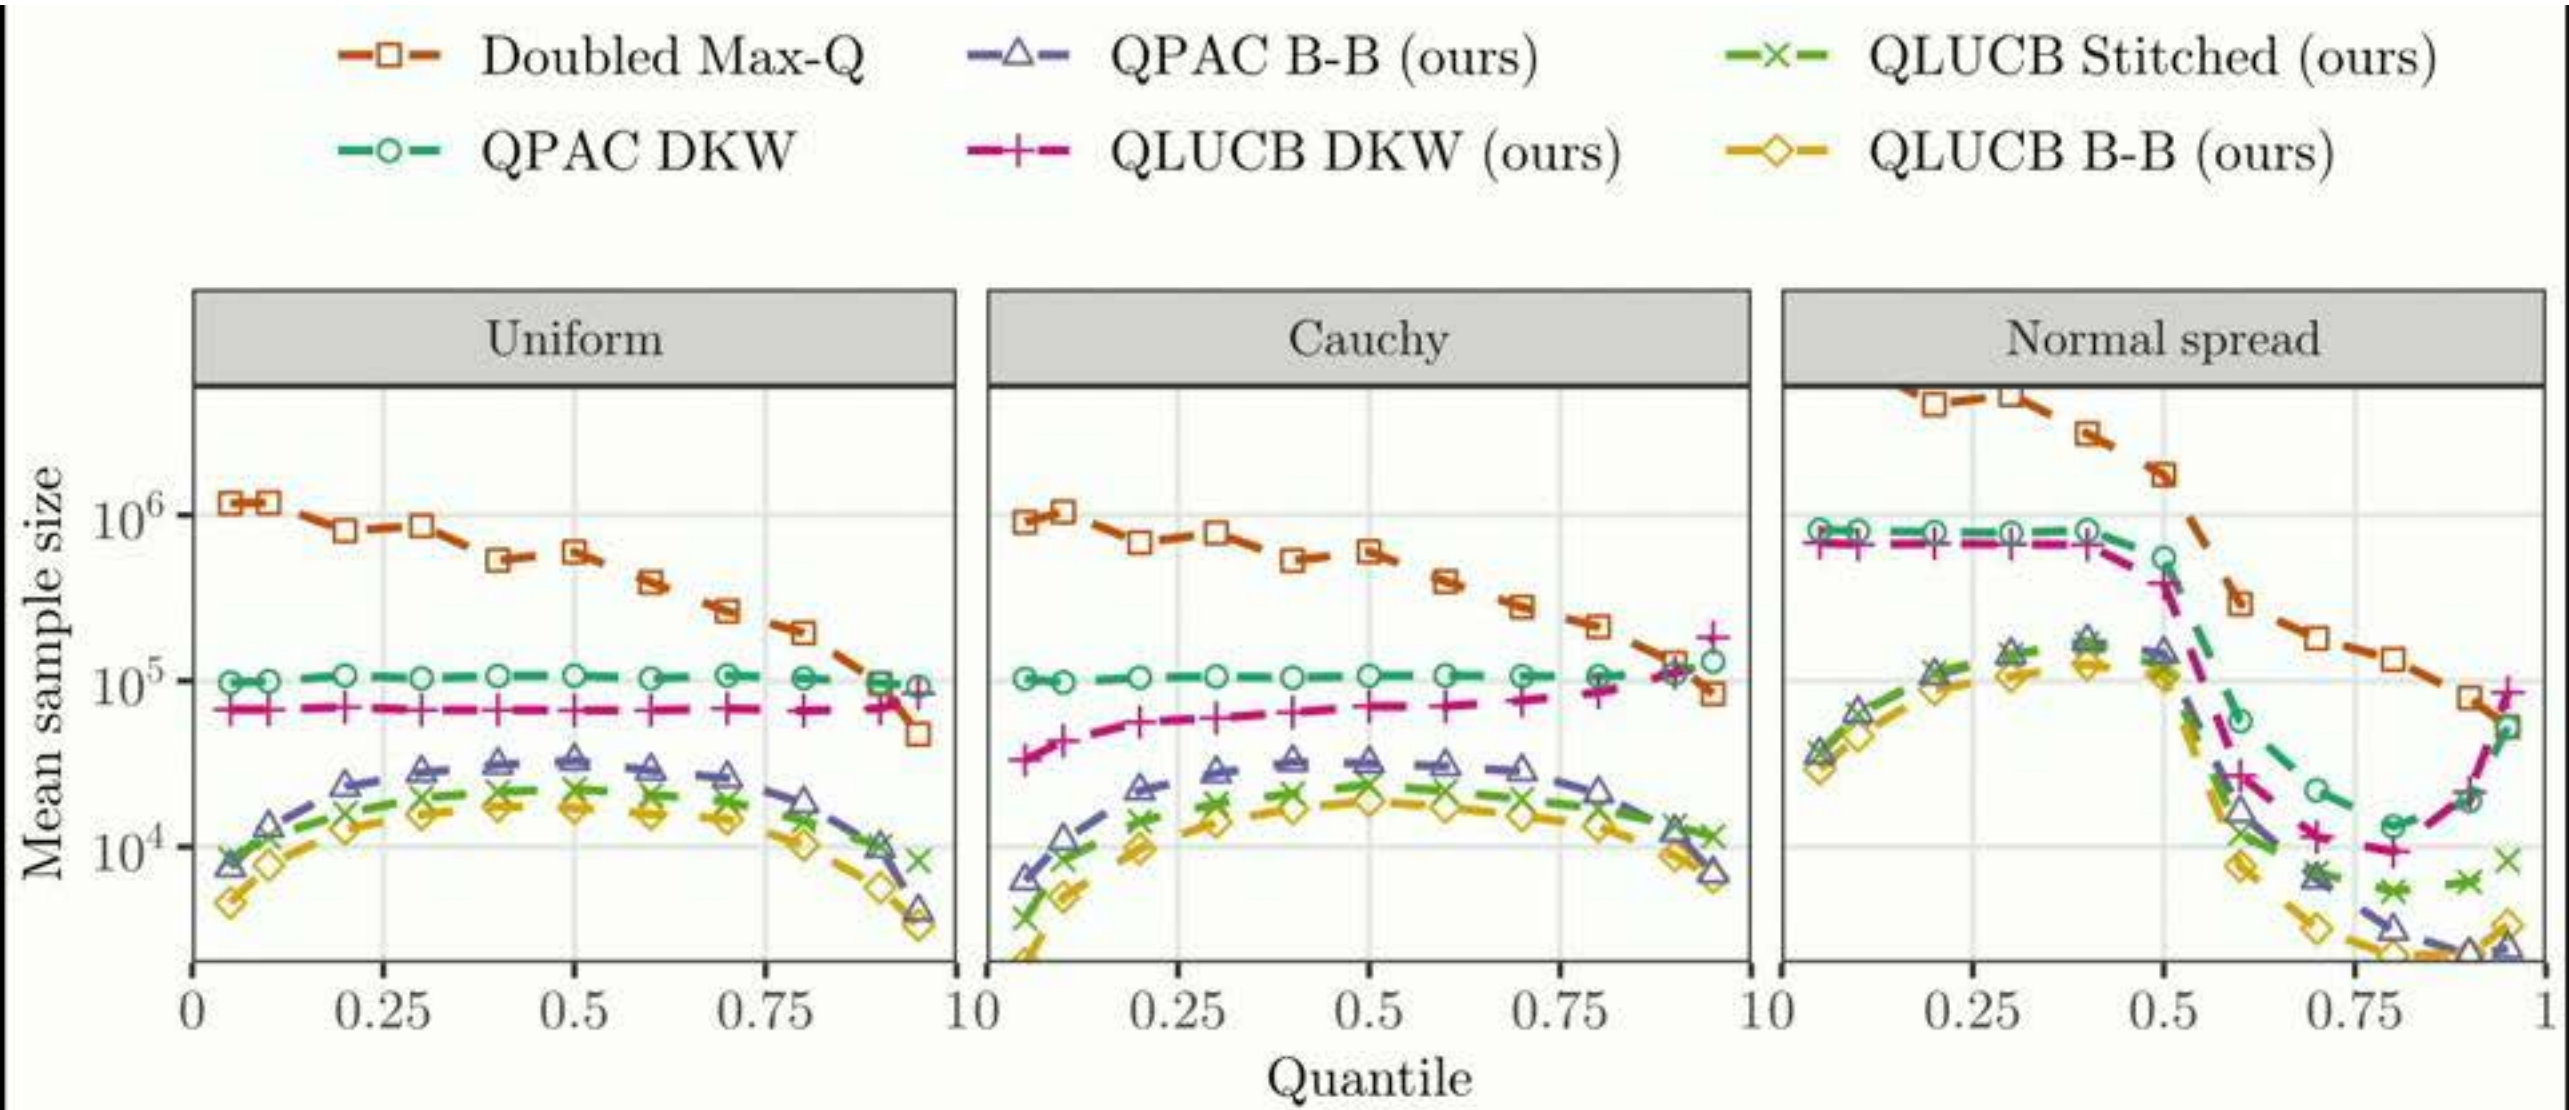
QLUCB stops w.p. one, and w.p. $1 - \delta$, it picks an ϵ -optimal arm.



Theorem:
$$\Delta_k := \sup \left\{ \Delta \geq 0 : Q_k(\pi + \Delta) < \max_{j \in [K]} Q_j(\pi) \right\}$$

QLUCB stops w.p. one, and w.p. $1 - \delta$, it picks an ϵ -optimal arm.

W.p. $1 - \delta$, the number of samples T satisfies



Theorem: $\Delta_k := \sup \left\{ \Delta \geq 0 : Q_k(\pi + \Delta) < \max_{j \in [K]} Q_j(\pi) \right\}$

QLUCB stops w.p. one, and w.p. $1 - \delta$, it picks an ϵ -optimal arm.

W.p. $1 - \delta$, the number of samples T satisfies

$$T = \mathcal{O} \left(\sum_{k=1}^K (\epsilon \vee \Delta_k)^{-2} \log \left(\frac{K |\log(\epsilon \vee \Delta_k)|}{\delta} \right) \right).$$

Other applications in the paper

- A/B testing: how to derive an always valid p-value by inverting and combining confidence sequences
- Sequential Kolmogorov-Smirnov test (classical KS test)
- Sequential test for stochastic dominance

(all have controlled false positive rate, and power tending to one at the optimal rate)

SOFTWARE

Python package called “**confseq**”

Maintained by [Steven R. Howard](#)

Frequent updates + wrappers for months to come

Takeaway messages

Takeaway messages

Reasons to use quantiles include:

- Quantiles always exist for any distribution, while means do not always exist (eg: Cauchy).

Takeaway messages

Reasons to use quantiles include:

- Quantiles always exist for any distribution, while means do not always exist (eg: Cauchy).
- Quantiles can be defined for any totally ordered space, eg: ratings A-F, where “distance between ratings” undefined.

Takeaway messages

Reasons to use quantiles include:

- Quantiles always exist for any distribution, while means do not always exist (eg: Cauchy).
- Quantiles can be defined for any totally ordered space, eg: ratings A-F, where “distance between ratings” undefined.
- Estimating quantiles can be done sequentially.

Takeaway messages

Reasons to use quantiles include:

- Quantiles always exist for any distribution, while means do not always exist (eg: Cauchy).
- Quantiles can be defined for any totally ordered space, eg: ratings A-F, where “distance between ratings” undefined.
- Estimating quantiles can be done sequentially.
- Can get p-values for testing difference in quantiles.

Takeaway messages

Reasons to use quantiles include:

- Quantiles always exist for any distribution, while means do not always exist (eg: Cauchy).
- Quantiles can be defined for any totally ordered space, eg: ratings A-F, where “distance between ratings” undefined.
- Estimating quantiles can be done sequentially.
- Can get p-values for testing difference in quantiles.
- Can run bandit experiments, including best-arm identification.

Takeaway messages

Reasons to use quantiles include:

- Quantiles always exist for any distribution, while means do not always exist (eg: Cauchy).
- Quantiles can be defined for any totally ordered space, eg: ratings A-F, where “distance between ratings” undefined.
- Estimating quantiles can be done sequentially.
- Can get p-values for testing difference in quantiles.
- Can run bandit experiments, including best-arm identification.

Can also estimate all quantiles simultaneously.

Takeaway messages

Reasons to use quantiles include:

- Quantiles always exist for any distribution, while means do not always exist (eg: Cauchy).
- Quantiles can be defined for any totally ordered space, eg: ratings A-F, where “distance between ratings” undefined.
- Estimating quantiles can be done sequentially.
- Can get p-values for testing difference in quantiles.
- Can run bandit experiments, including best-arm identification.

Can also estimate all quantiles simultaneously.

Open problems

Takeaway messages

Reasons to use quantiles include:

- Quantiles always exist for any distribution, while means do not always exist (eg: Cauchy).
- Quantiles can be defined for any totally ordered space, eg: ratings A-F, where “distance between ratings” undefined.
- Estimating quantiles can be done sequentially.
- Can get p-values for testing difference in quantiles.
- Can run bandit experiments, including best-arm identification.

Can also estimate all quantiles simultaneously.

Open problems

- Extend to RL, connected to notions of “safety”.

Takeaway messages

Reasons to use quantiles include:

- Quantiles always exist for any distribution, while means do not always exist (eg: Cauchy).
- Quantiles can be defined for any totally ordered space, eg: ratings A-F, where “distance between ratings” undefined.
- Estimating quantiles can be done sequentially.
- Can get p-values for testing difference in quantiles.
- Can run bandit experiments, including best-arm identification.

Can also estimate all quantiles simultaneously.

Open problems

- Extend to RL, connected to notions of “safety”.
- Perhaps start with contextual bandits or MDPs.

Sequential estimation of quantiles with applications to *A/B* testing and bandits



Aaditya Ramdas

Assistant Professor
Dept. of Statistics and Data Science
Machine Learning Dept.
Carnegie Mellon University



Steven R. Howard

PhD student
Dept. of Statistics
UC Berkeley

Sequential estimation of quantiles with applications to A/B testing and bandits



Aaditya Ramdas

Funding welcomed!

Assistant Professor
Dept. of Statistics and Data Science
Machine Learning Dept.
Carnegie Mellon University



Steven R. Howard

PhD student
Dept. of Statistics
UC Berkeley

Sequential estimation of quantiles with applications to A/B testing and bandits



Aaditya Ramdas

Funding welcomed!

Assistant Professor
Dept. of Statistics and Data Science
Machine Learning Dept.
Carnegie Mellon University



Steven R. Howard

PhD student
Dept. of Statistics
UC Berkeley

Thank you! Questions?