Clip-OGD: An Experimental Design for Adaptive Neyman Allocation *to appear, NeurIPS 2023*

Jessica Dai
U.C. Berkeley

Chris Harshaw
MIT

Paula Gradu
U.C. Berkeley

Algorithm Design for Causal Inference, INFORMS 2023

Randomized Experiments: Standard

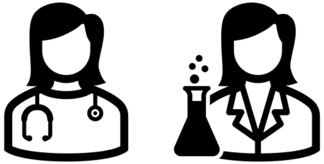

Experimenters

Randomized Experiments: Standard

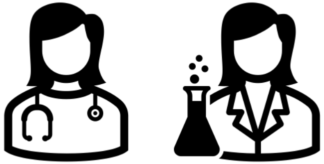

Treatment & control

Randomized Experiments: Standard

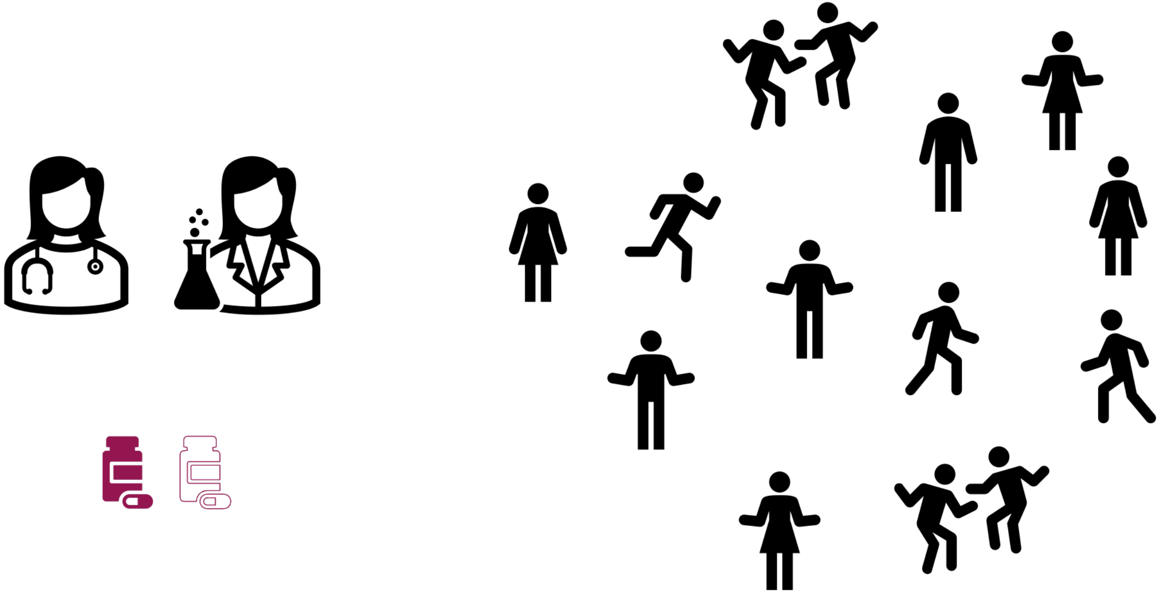

Experimental units

Randomized Experiments: Standard

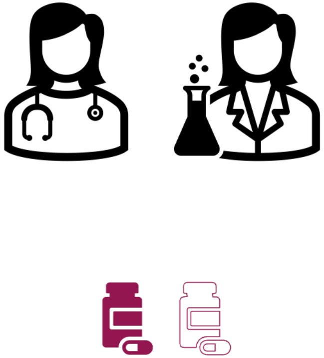

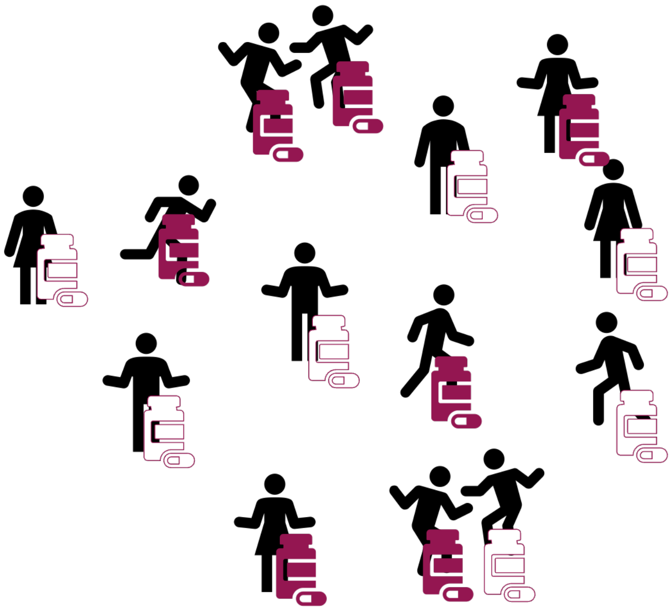

Treatment assignment

Randomized Experiments: Standard

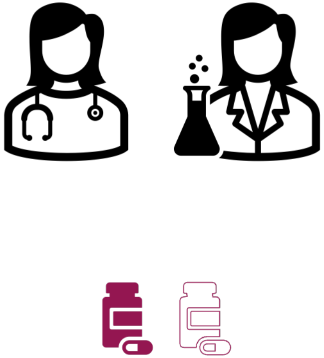

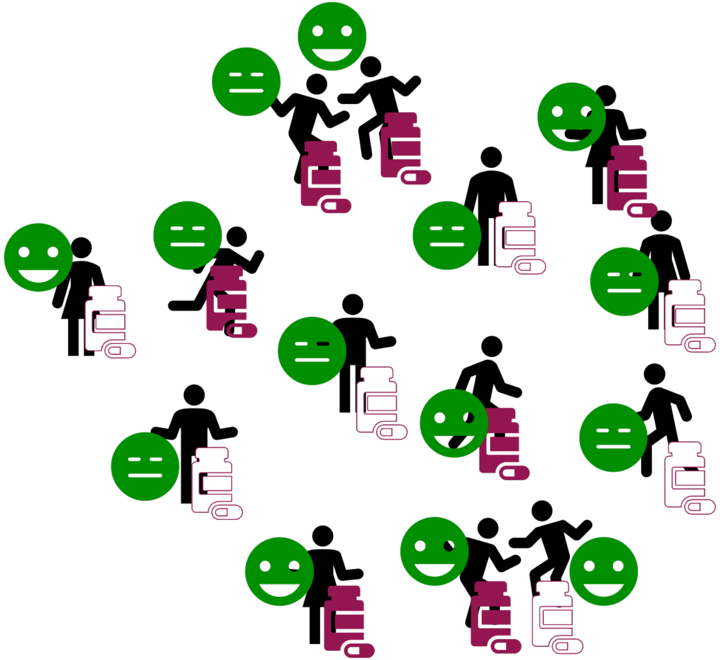

Observe outcomes

Randomized Experiments: Standard

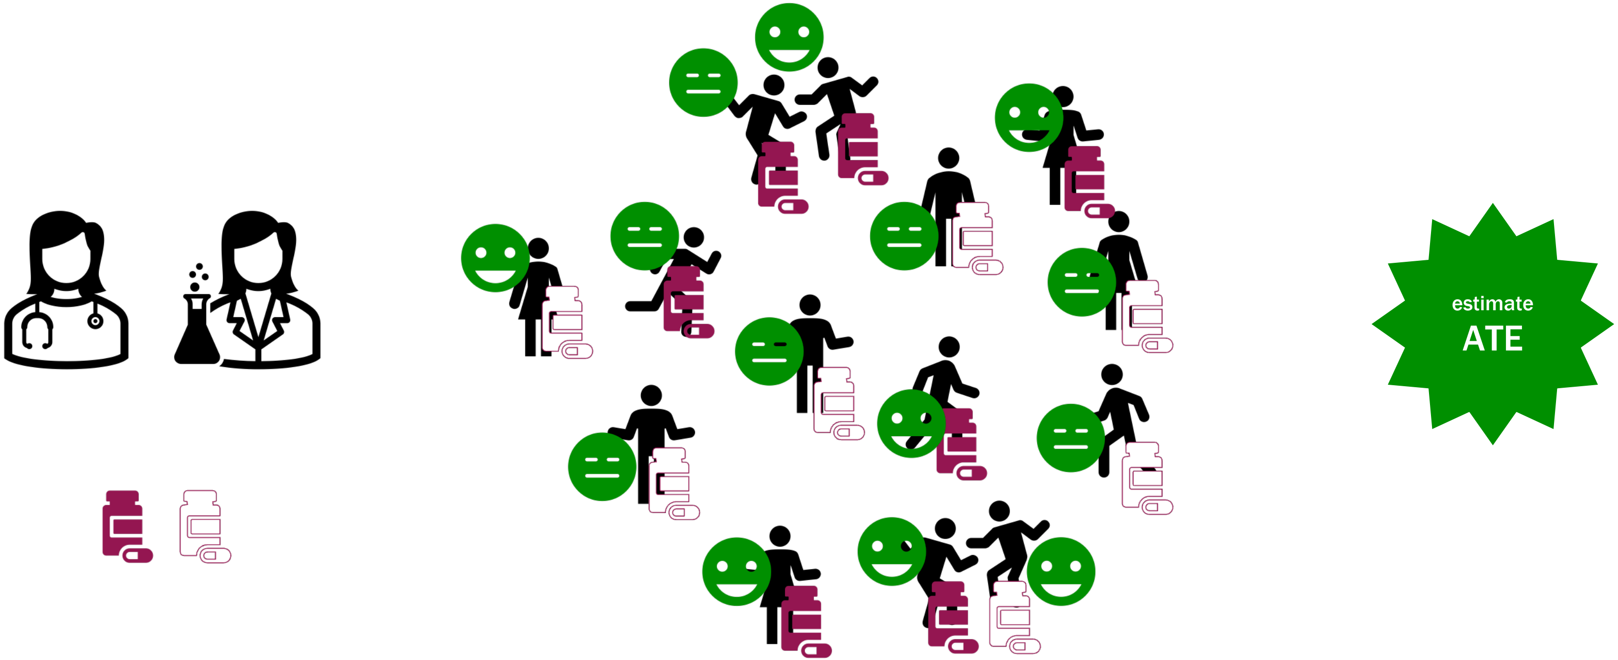

Randomized Experiments: Adaptive

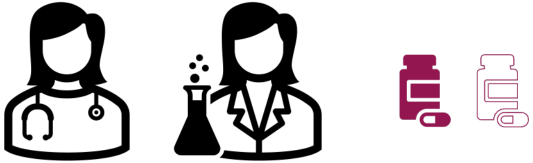

Experimenters and treatments, as before

Randomized Experiments: Adaptive

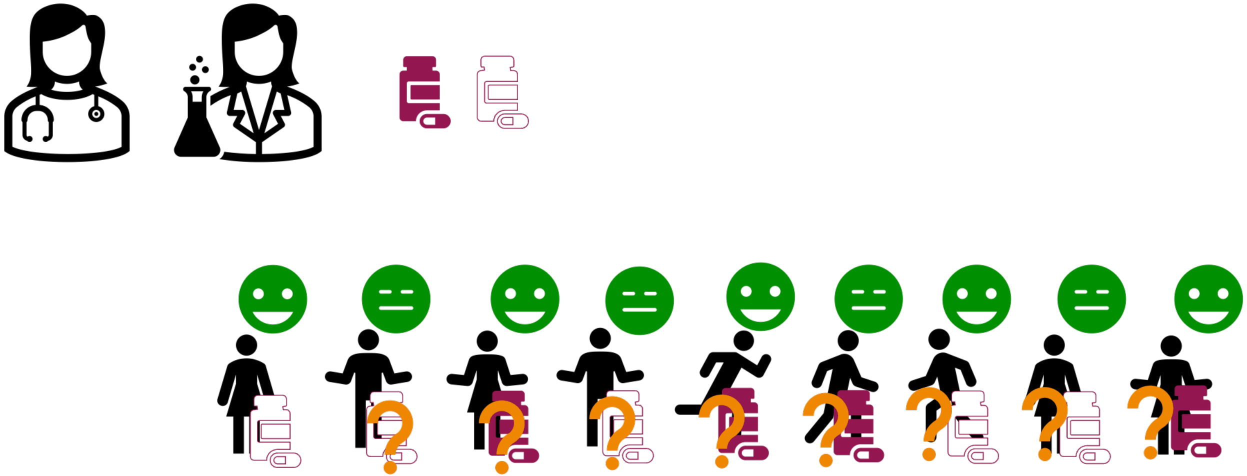

Sequential treatment assignments and outcome observations

Randomized Experiments: Adaptive

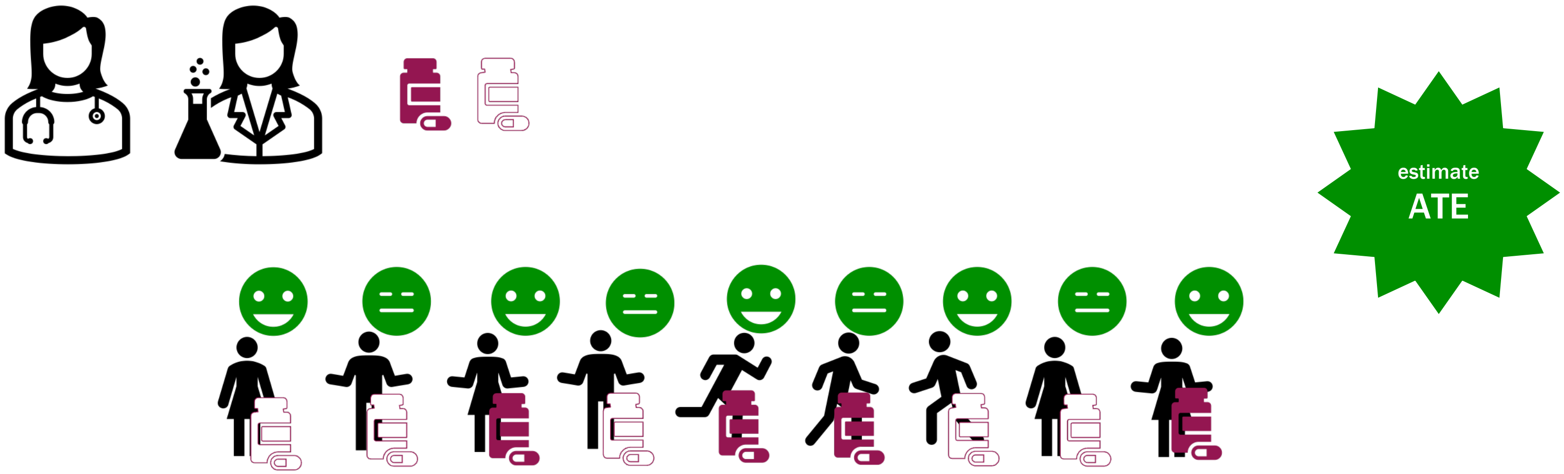

Randomized Experiments: Adaptive

Background: design-based causal inference

Background: design-based causal inference

We have T experimental units (indexed by t).

Background: design-based causal inference

We have T experimental units (indexed by t).

The experimenter has a **treatment** they would like to compare to **control.**

Background: design-based causal inference

We have T experimental units (indexed by t).

The experimenter has a **treatment** they would like to compare to **control.**

Each t has **deterministic potential outcomes** under treatment and control:

$$\big(y_t(1), y_t(0)\big)$$

Background: design-based causal inference

We have T experimental units (indexed by t).

The experimenter has a **treatment** they would like to compare to **control.**

Each t has **deterministic** **potential outcomes** under treatment and control:

$$\big(y_t(1), y_t(0)\big)$$

We want to know the **average treatment effect:**

$$\tau = \frac{1}{T} \sum_{t \in [T]} y_t(1) - y_t(0)$$

Background: standard (non-adaptive) experiments

Background: standard (non-adaptive) experiments

1. Decide treatment probability p

Background: standard (non-adaptive) experiments

1. Decide treatment probability p

2. Begin the experiment.

   For units t = 1, 2, ... T:

Background: standard (non-adaptive) experiments

1. Decide treatment probability p

2. Begin the experiment.

    For units t = 1, 2, … T:

    Sample treatment assignment $Z_t \sim \mathrm{Bern}(p)$

Background: standard (non-adaptive) experiments

1. Decide treatment probability p

2. Begin the experiment.

    For units t = 1, 2, ... T:

    Sample treatment assignment $Z_t \sim \mathrm{Bern}(p)$

    Observe one potential outcome $Y_t = Z_t \cdot y_t(1) + (1 - Z_t) \cdot y_t(0)$

Background: standard (non-adaptive) experiments

1. Decide treatment probability p

2. Begin the experiment.

   For units t = 1, 2, ... T:

   Sample treatment assignment $Z_t \sim \text{Bern}(p)$

   Observe one potential outcome $Y_t = Z_t \cdot y_t(1) + (1 - Z_t) \cdot y_t(0)$

3. Estimate τ with the Horvitz-Thompson estimator:

$$\hat{\tau} = \frac{1}{T} \sum_{t \in [T]} Y_t \left(\frac{Z_t}{p} + \frac{1 - Z_t}{1 - p} \right)$$

Background: standard (non-adaptive) experiments

1. Decide treatment probability p

2. Begin the experiment.

   For units t = 1, 2, ... T:

   Sample treatment assignment $Z_t \sim \mathrm{Bern}(p)$

   Observe one potential outcome $Y_t = Z_t \cdot y_t(1) + (1 - Z_t) \cdot y_t(0)$

3. Estimate τ with the Horvitz-Thompson estimator:

$$\hat{\tau} = \frac{1}{T} \sum_{t \in [T]} Y_t \left(\frac{Z_t}{p} + \frac{1 - Z_t}{1 - p} \right)$$

unbiased estimator

Background: standard (non-adaptive) experiments

1. Decide treatment probability p

2. Begin the

   Fo

   $- Z_t) \cdot y_t(0)$

How do we set *p*?

3. Estimate τ with the Horvitz-Thompson estimator:

$$\hat{\tau} = \frac{1}{T} \sum_{t \in [T]} Y_t \left(\frac{Z_t}{p} + \frac{1 - Z_t}{1 - p} \right)$$

unbiased estimator

How do we set p?

This work: minimize variance.

Neyman Variance: *the optimal allocation*

For any fixed p :

Neyman Variance: *the optimal allocation*

For any fixed p:

$$T \cdot \mathrm{Var}(\hat{\tau}) = S(1)^2\left(\frac{1}{p} - 1\right) + S(0)^2\left(\frac{1}{1-p} - 1\right) + 2\rho S(1)S(0)$$

Neyman Variance: *the optimal allocation*

For any fixed p:

$$T \cdot \mathrm{Var}(\hat{\tau}) = S(1)^2\left(\frac{1}{p} - 1\right) + S(0)^2\left(\frac{1}{1-p} - 1\right) + 2\rho S(1)S(0)$$

Second moment:

$$S(i) = \sqrt{\frac{1}{T}\sum_{t \in [T]} y_t(i)^2}$$

Cosine similarity:

$$\rho = \frac{\frac{1}{T}\sum_{t\in[T]} y_t(0)y_t(1)}{S(1)S(0)}$$

Neyman Variance: *the optimal allocation*

For any fixed p :

$$T \cdot \mathrm{Var}(\hat{\tau}) = S(1)^2\left(\frac{1}{p} - 1\right) + S(0)^2\left(\frac{1}{1-p} - 1\right) + 2\rho S(1)S(0)$$

Can't get rid of dependence on correlation

Second moment:

$$S(i) = \sqrt{\frac{1}{T}\sum_{t\in[T]} y_t(i)^2}$$

Cosine similarity:

$$\rho = \frac{\frac{1}{T}\sum_{t\in[T]} y_t(0)y_t(1)}{S(1)S(0)}$$

Neyman Variance: *the optimal allocation*

For any fixed p :

$$T \cdot \mathrm{Var}(\hat{\tau}) = S(1)^2\left(\frac{1}{p} - 1\right) + S(0)^2\left(\frac{1}{1-p} - 1\right) + 2\rho S(1)S(0)$$

Can't get rid of dependence on correlation

Then, set

$$p^* = \left(1 + \sqrt{\frac{S(0)}{S(1)}}\right)^{-1}$$

Second moment:

$$S(i) = \sqrt{\frac{1}{T}\sum_{t\in[T]} y_t(i)^2}$$

Cosine similarity:

$$\rho = \frac{\frac{1}{T}\sum_{t\in[T]} y_t(0)y_t(1)}{S(1)S(0)}$$

Neyman Variance: *the optimal allocation*

For any fixed p :

$$T \cdot \mathrm{Var}(\hat{\tau}) = S(1)^2\left(\frac{1}{p} - 1\right) + S(0)^2\left(\frac{1}{1-p} - 1\right) + 2\rho S(1)S(0)$$

Can't get rid of dependence on correlation

Then, set

$$p^* = \left(1 + \sqrt{\frac{S(0)}{S(1)}}\right)^{-1}$$

and get variance*

**normalized*

$$T \cdot \mathrm{Var}_N = 2(1+\rho)S(1)S(0)$$

Second moment:

$$S(i) = \sqrt{\frac{1}{T}\sum_{t\in[T]} y_t(i)^2}$$

Cosine similarity:

$$\rho = \frac{\frac{1}{T}\sum_{t\in[T]} y_t(0)y_t(1)}{S(1)S(0)}$$

Neyman Variance: *the optimal allocation*

For any fixed p:

$$T \cdot \mathrm{Var}(\hat{\tau}) = S(1)^2\left(\frac{1}{p} - 1\right) + S(0)^2\left(\frac{1}{1-p} - 1\right) + 2\rho S(1)S(0)$$

Then, set

$$p^* = \left(1 + \sqrt{\frac{S(0)}{S(1)}}\right)^{-1}$$

and get variance*

**normalized*

$$T \cdot \mathrm{Var}_N = 2(1+\rho)S(1)S(0)$$

Substantially improving the ***constant***,
not the ***rate*** of O(1/T)

Second moment:

$$S(i) = \sqrt{\frac{1}{T}\sum_{t\in[T]} y_t(i)^2}$$

Cosine similarity:

$$\rho = \frac{\frac{1}{T}\sum_{t\in[T]} y_t(0)y_t(1)}{S(1)S(0)}$$

Neyman Variance: *the optimal allocation*

For any fixed p :

$$T \cdot \mathrm{Var}(\hat{\tau}) = S(1)^2\left(\frac{1}{p} - 1\right) + S(0)^2\left(\frac{1}{1-p} - 1\right) + 2\rho S(1)S(0)$$

Sum of squares

Then, set

$$p^* = \left(1 + \sqrt{\frac{S(0)}{S(1)}}\right)^{-1}$$

and get variance*

**normalized*

$$T \cdot \mathrm{Var}_N = 2(1+\rho)S(1)S(0)$$

Substantially improving the ***constant***,
not the ***rate*** of O(1/T)

Second moment:

$$S(i) = \sqrt{\frac{1}{T}\sum_{t\in[T]} y_t(i)^2}$$

Cosine similarity:

$$\rho = \frac{\frac{1}{T}\sum_{t\in[T]} y_t(0)y_t(1)}{S(1)S(0)}$$

Neyman Variance: *the optimal allocation*

For any fixed p :

$$T \cdot \mathrm{Var}(\hat{\tau}) = S(1)^2\left(\frac{1}{p} - 1\right) + S(0)^2\left(\frac{1}{1-p} - 1\right) + 2\rho S(1)S(0)$$

Sum of squares

Then, set

$$p^* = \left(1 + \sqrt{\frac{S(0)}{S(1)}}\right)^{-1}$$

and get variance*

$$T \cdot \mathrm{Var}_N = 2(1+\rho)S(1)S(0)$$

Product

**normalized*

Substantially improving the ***constant***, not the ***rate*** of O(1/T)

Second moment:

$$S(i) = \sqrt{\frac{1}{T}\sum_{t\in[T]} y_t(i)^2}$$

Cosine similarity:

$$\rho = \frac{\frac{1}{T}\sum_{t\in[T]} y_t(0)y_t(1)}{S(1)S(0)}$$

Neyman Variance: *the optimal allocation*

For any fixed p:

$$T \cdot \mathrm{Var}(\hat{\tau}) = S(1)^2\left(\frac{1}{p} - 1\right) + S(0)^2\left(\frac{1}{1-p} - 1\right) + 2\rho S(1)S(0)$$

Then, set

$$p^* = \left(1 + \sqrt{\frac{S(0)}{S(1)}}\right)^{-1}$$

and get variance*

*normalized

$$T \cdot \mathrm{Var}_N = 2(1+\rho)S(1)S(0)$$

Second moment:

$$S(i) = \sqrt{\frac{1}{T}\sum_{t\in[T]} y_t(i)^2}$$

Cosine similarity:

$$\rho = \frac{\frac{1}{T}\sum_{t\in[T]} y_t(0)y_t(1)}{S(1)S(0)}$$

Neyman Variance: *the optimal allocation*

For any fixed p :

$$T \cdot \mathrm{Var}(\hat{\tau}) = S(1)^2\left(\frac{1}{p} - 1\right) + S(0)^2\left(\frac{1}{1-p} - 1\right) + 2\rho S(1)S(0)$$

Then, set

$$p^* = \left(1 + \sqrt{\frac{S(0)}{S(1)}}\right)^{-}$$

We don't know these outcomes!

and get variance*

**normalized*

$$T \cdot \mathrm{Var}_N = 2(1+\rho)S(1)S(0)$$

Cosine similarity:

$$\rho = \frac{\frac{1}{T}\sum_{t\in[T]} y_t(0)y_t(1)}{S(1)S(0)}$$

Neyman Variance: *the optimal allocation*

For any fixed p:

$$T \cdot \mathrm{Var}(\hat{\tau}) = S(1)^2\left(\frac{1}{\ } - 1\right) + S(0)^2\left(\frac{1}{\ } - 1\right) + 2\rho S(1)S(0)$$

Can we get close with adaptivity?

moment:

$$\sqrt{\frac{1}{T}\sum_{t\in[T]} y_t(i)^2}$$

Then, set

similarity:

$$\rho = \frac{\frac{1}{T}\sum_{t\in[T]} y_t(0)y_t(1)}{S(1)S(0)}$$

and get variance

$$T \cdot \mathrm{Var}_N = 2(1+\rho)S(1)S(0)$$

Background: Prior work

Background: Prior work

SOME ASPECTS OF THE SEQUENTIAL DESIGN OF EXPERIMENTS

HERBERT ROBBINS

Background: Prior work

Suppose we are dealing with two normally distributed populations with unknown means μ_1, μ_2 and variances σ_1^2, σ_2^2, and that we wish to estimate the value of the difference $\mu_1 - \mu_2$. In order to concentrate on the point at issue we shall suppose that the total sample size, n, is fixed. There remains the question of how the n observations are to be divided between the two populations. If $\bar{x}_1, \bar{x}_2$ denote the means of samples of sizes n_1, n_2 from the two populations, then $\bar{x}_1 - \bar{x}_2$ is an unbiased estimator of $\mu_1 - \mu_2$, with variance $\sigma^2 = (\sigma_1^2/n_1) + (\sigma_2^2/n_2)$. For fixed $n = n_1 + n_2$, σ^2 is a minimum when $n_1/n_2 = \sigma_1/\sigma_2$.

... gestures at Neyman allocation

Some Aspects of the Sequential Designs of Experiments [Robbins 1952]

Background: Prior work

Suppose we are dealing with two normally distributed populations with unknown means μ_1, μ_2 and variances σ_1^2, σ_2^2, and that we wish to estimate the value of the difference $\mu_1 - \mu_2$. In order to concentrate on the point at issue we shall suppose that the total sample size, n, is fixed. There remains the question of how the n observations are to be divided between the two populations. If $\bar{x}_1, \bar{x}_2$ denote the means of samples of sizes n_1, n_2 from the two populations, then $\bar{x}_1 - \bar{x}_2$ is an unbiased estimator of $\mu_1 - \mu_2$, with variance $\sigma^2 = (\sigma_1^2/n_1) + (\sigma_2^2/n_2)$. For fixed $n = n_1 + n_2$, σ^2 is a minimum when $n_1/n_2 = \sigma_1/\sigma_2$. If the latter ratio is known in advance, all is well. If this ratio is not known, but if the sampling can be done in two stages, then it would be reasonable to draw preliminary samples of some size m from each of the two populations and to use the values so obtained to estimate σ_1/σ_2; the remainder of the $n - 2m$ observations could then be allocated to the two populations in accordance with the sample estimate of σ_1/σ_2.

Two-stage design: recently studied in Hahn et.al. 2011, Blackwell et.al. 2022

Some Aspects of the Sequential Designs of Experiments [Robbins 1952]

Background: Prior work

Suppose we are dealing with two normally distributed populations with unknown means μ_1, μ_2 and variances σ_1^2, σ_2^2, and that we wish to estimate the value of the difference $\mu_1 - \mu_2$. In order to concentrate on the point at issue we shall suppose that the total sample size, n, is fixed. There remains the question of how the n observations are to be divided between the two populations. If $\bar{x}_1$, $\bar{x}_2$ denote the means of samples of sizes n_1, n_2 from the two populations, then $\bar{x}_1 - \bar{x}_2$ is an unbiased estimator of $\mu_1 - \mu_2$, with variance $\sigma^2 = (\sigma_1^2/n_1) + (\sigma_2^2/n_2)$. For fixed $n = n_1 + n_2$, σ^2 is a minimum when $n_1/n_2 = \sigma_1/\sigma_2$. If the latter ratio is known in advance, all is well. If this ratio is not known, but if the sampling can be done in two stages, then it would be reasonable to draw preliminary samples of some size m from each of the two populations and to use the values so obtained to estimate σ_1/σ_2; the remainder of the $n - 2m$ observations could then be allocated to the two populations in accordance with the sample estimate of σ_1/σ_2. The question then becomes, what is the best choice for m? If m is small, no accurate estimate of σ_1/σ_2 can be made. If m is large, then the remaining $n - 2m$ observations may be too few to permit full utilization of the approximate knowledge of σ_1/σ_2. (This kind of dilemma is characteristic of all sequential design problems.) More generally, we could consider schemes in which the observations are made one by one, with the decision as to which population each observation should come from being allowed to depend on all the previous observations; the total sample size n could be fixed or could be a random variable dependent on the observations.

Some Aspects of the Sequential Designs of Experiments [Robbins 1952]

Proposes a fully adaptive experiment…

Adaptive experimental setup

Adaptive experimental setup

1. Begin the experiment.

Adaptive experimental setup

1. Begin the experiment.

    For units t = 1, 2, … T:

Adaptive experimental setup

1. Begin the experiment.

   For units t = 1, 2, ... T:

   Determine treatment probability P_t

Adaptive experimental setup

1. Begin the experiment.

   For units t = 1, 2, … T:

   Determine treatment probability P_t

   Sample treatment assignment $Z_t \sim \mathrm{Bern}(P_t)$

Adaptive experimental setup

1. Begin the experiment.

   For units t = 1, 2, ... T:

   Determine treatment probability P_t

   Sample treatment assignment $Z_t \sim \mathrm{Bern}(P_t)$

   Observe one potential outcome $Y_t = Z_t \cdot y_t(1) + (1 - Z_t) \cdot y_t(0)$

Adaptive experimental setup

1. Begin the experiment.

   For units t = 1, 2, ... T:

Determine treatment probability	P_t
Sample treatment assignment	$Z_t \sim \text{Bern}(P_t)$
Observe one potential outcome	$Y_t = Z_t \cdot y_t(1) + (1 - Z_t) \cdot y_t(0)$

2. Estimate τ with the **adaptive** Horvitz-Thompson estimator:

$$\hat{\tau} = \frac{1}{T} \sum_{t \in [T]} Y_t \left(\frac{Z_t}{P_t} + \frac{1 - Z_t}{1 - P_t} \right)$$

Adaptive experimental setup

1. Begin the experiment.

   For units t = 1, 2, ... T:

   Determine treatment probability P_t

   Sample treatment assignment $Z_t \sim \mathrm{Bern}(P_t)$

   Observe one potential outcome $Y_t = Z_t \cdot y_t(1) + (1 - Z_t) \cdot y_t(0)$

2. Estimate τ with the **adaptive** Horvitz-Thompson estimator:

$$\hat{\tau} = \frac{1}{T} \sum_{t \in [T]} Y_t \left(\frac{Z_t}{P_t} + \frac{1 - Z_t}{1 - P_t} \right)$$

unbiased estimator

Adaptive experimental setup

1. Begin the experiment.

   For units t = 1, 2, ... T:

   Determine treatment probability P_t

   $(1 - Z_t) \cdot y_t(0)$

How do we set P_t?

2. Estimate

$$\hat{\tau} = \frac{1}{T} \sum_{t \in [T]} Y_t \left(\frac{Z_t}{P_t} + \frac{1 - Z_t}{1 - P_t} \right)$$

unbiased estimator

Neyman Regret: *Neyman allocation as a benchmark*

Neyman Regret: *Neyman allocation as a benchmark*

With adaptivity,

Neyman Regret: *Neyman allocation as a benchmark*

With adaptivity,

$$T \cdot \mathrm{Var}(\hat{\tau}) = \frac{1}{T} \sum_{t \in [T]} \left(y_t(1)^2 \cdot \mathbb{E}\left[\frac{1}{P_t}\right] + y_t(0)^2 \cdot \mathbb{E}\left[\frac{1}{1 - P_t}\right] \right) - \frac{1}{T} \sum_{t \in [T]} \left(y_t(1) - y_t(0)\right)^2$$

Neyman Regret: *Neyman allocation as a benchmark*

With adaptivity,

$$T \cdot \mathrm{Var}(\hat{\tau}) = \frac{1}{T} \sum_{t \in [T]} \left(y_t(1)^2 \cdot \mathbb{E}\left[\frac{1}{P_t}\right] + y_t(0)^2 \cdot \mathbb{E}\left[\frac{1}{1-P_t}\right] \right) - \frac{1}{T} \sum_{t \in [T]} \left(y_t(1) - y_t(0)\right)^2$$

Let's notate

$$f_t(p) = \frac{y_t(1)^2}{p} + \frac{y_t(0)^2}{1-p}$$

Neyman Regret: *Neyman allocation as a benchmark*

With adaptivity,

$$T \cdot \mathrm{Var}(\hat{\tau}) = \frac{1}{T} \sum_{t \in [T]} \left(y_t(1)^2 \cdot \mathbb{E}\left[\frac{1}{P_t}\right] + y_t(0)^2 \cdot \mathbb{E}\left[\frac{1}{1 - P_t}\right] \right) - \frac{1}{T} \sum_{t \in [T]} \left(y_t(1) - y_t(0)\right)^2$$

Let's notate

$$f_t(p) = \frac{y_t(1)^2}{p} + \frac{y_t(0)^2}{1 - p}$$

"controllable contribution of unit t to the total variance"

Neyman Regret: *Neyman allocation as a benchmark*

With adaptivity,

$$T \cdot \mathrm{Var}(\hat{\tau}) = \frac{1}{T} \sum_{t \in [T]} \left(y_t(1)^2 \cdot \mathbb{E}\left[\frac{1}{P_t}\right] + y_t(0)^2 \cdot \mathbb{E}\left[\frac{1}{1-P_t}\right] \right) - \frac{1}{T} \sum_{t \in [T]} (y_t(1) - y_t(0))^2$$

Let's notate

$$f_t(p) = \frac{y_t(1)^2}{p} + \frac{y_t(0)^2}{1-p}$$

"controllable contribution of unit t to the total variance"

Then, **Neyman Regret** is defined as

$$\mathcal{R}_T = \sum_{t \in [T]} f_t(P_t) - \min_{p \in [0,1]} \sum_{t \in [T]} f_t(p)$$

Neyman Regret: *Neyman allocation as a benchmark*

$$\mathcal{R}_T = \sum_{t \in [T]} f_t(P_t) - \min_{p \in [0,1]} \sum_{t \in [T]} f_t(p)$$

Neyman Regret: *Neyman allocation as a benchmark*

$$\mathcal{R}_T = \sum_{t \in [T]} f_t(P_t) - \min_{p \in [0,1]} \sum_{t \in [T]} f_t(p)$$

"incurred variance
from actual allocation"

Neyman Regret: *Neyman allocation as a benchmark*

$$\mathcal{R}_T = \sum_{t \in [T]} f_t(P_t) - \min_{p \in [0,1]} \sum_{t \in [T]} f_t(p)$$

"incurred variance from actual allocation"

"minimum variance from best possible allocation"

Neyman Regret: *Neyman allocation as a benchmark*

$$\mathcal{R}_T = \sum_{t \in [T]} f_t(P_t) - \min_{p \in [0,1]} \sum_{t \in [T]} f_t(p)$$

With this definition, minimizing Neyman regret also minimizes variance:

$$T \cdot V - T \cdot V_N = \frac{1}{T} \mathbb{E}[\mathcal{R}_T]$$

Neyman Regret: *Neyman allocation as a benchmark*

$$\mathcal{R}_T = \sum_{t \in [T]} f_t(P_t) - \min_{p \in [0,1]} \sum_{t \in [T]} f_t(p)$$

With this definition, minimizing Neyman regret also minimizes variance:

$$T \cdot V - T \cdot V_N = \frac{1}{T} \mathbb{E}[\mathcal{R}_T]$$

This turns a statistical problem into an algorithmic problem:

$$T \cdot V \to T \cdot V_N \iff \mathbb{E}\left[\mathcal{R}_T\right] = o(T)$$

Minimizing Neyman Regret

$$\mathcal{R}_T = \sum_{t \in [T]} f_t(P_t) - \min_{p \in [0,1]} \sum_{t \in [T]} f_t(p)$$

Minimizing Neyman Regret

$$\mathcal{R}_T = \sum_{t \in [T]} f_t(P_t) - \min_{p \in [0,1]} \sum_{t \in [T]} f_t(p)$$

Good: objective $f_t(p) = \frac{y_t(1)^2}{p} + \frac{y_t(0)^2}{1-p}$ is convex

Minimizing Neyman Regret

$$\mathcal{R}_T = \sum_{t \in [T]} f_t(P_t) - \min_{p \in [0,1]} \sum_{t \in [T]} f_t(p)$$

can borrow techniques from online convex opt

Good: objective $f_t(p) = \frac{y_t(1)^2}{p} + \frac{y_t(0)^2}{1-p}$ is convex

Minimizing Neyman Regret

$$\mathcal{R}_T = \sum_{t \in [T]} f_t(P_t) - \min_{p \in [0,1]} \sum_{t \in [T]} f_t(p)$$

can borrow techniques from online convex opt

Good: objective $f_t(p) = \frac{y_t(1)^2}{p} + \frac{y_t(0)^2}{1-p}$ is convex

Can we use online gradient descent?

Minimizing Neyman Regret

$$\mathcal{R}_T = \sum_{t \in [T]} f_t(P_t) - \min_{p \in [0,1]} \sum_{t \in [T]} f_t(p)$$

can borrow techniques from online convex opt

Good: objective $f_t(p) = \frac{y_t(1)^2}{p} + \frac{y_t(0)^2}{1-p}$ is convex

Can we use online gradient descent?

Less good: fundamental problem of causal inference

only observe $Y_t = y_t(Z_t)$, not $\big(y_t(1), y_t(0)\big)$

Minimizing Neyman Regret

$$\mathcal{R}_T = \sum_{t \in [T]} f_t(P_t) - \min_{p \in [0,1]} \sum_{t \in [T]} f_t(p)$$

can borrow techniques from online convex opt

Good: objective $f_t(p) = \frac{y_t(1)^2}{p} + \frac{y_t(0)^2}{1-p}$ is convex

Can we use online gradient descent?

Less good: fundamental problem of causal inference

only observe $Y_t = y_t(Z_t)$, not $\big(y_t(1), y_t(0)\big)$

can construct an unbiased estimator for the gradient

Minimizing Neyman Regret

$$\mathcal{R}_T = \sum_{t \in [T]} f_t(P_t) - \min_{p \in [0,1]} \sum_{t \in [T]} f_t(p)$$

can borrow techniques from online convex opt

Good: objective $f_t(p) = \frac{y_t(1)^2}{p} + \frac{y_t(0)^2}{1-p}$ is convex

Can we use online gradient descent?

Less good: fundamental problem of causal inference

only observe $Y_t = y_t(Z_t)$, not $\big(y_t(1), y_t(0)\big)$

can construct an unbiased estimator for the gradient

Challenging: gradients $\nabla f_t(p)$ blow up at the boundary

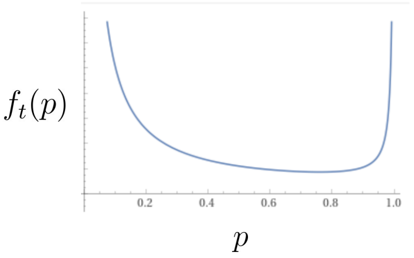

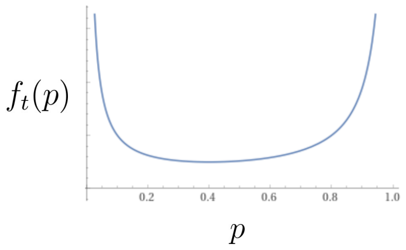

Minimizing Neyman Regret

$$\mathcal{R}_T = \sum_{t \in [T]} f_t(P_t) - \min_{p \in [0,1]} \sum_{t \in [T]} f_t(p)$$

Good: objective $f_t(p) = \frac{y_t(1)^2}{p} + \frac{y_t(0)^2}{1-p}$ is convex

Can we use online gradient descent?

can borrow techniques from online convex opt

Less good: fundamental problem of causal inference

only observe $Y_t = y_t(Z_t)$, not $\big(y_t(1), y_t(0)\big)$

can construct an unbiased estimator for the gradient

Challenging: gradients $\nabla f_t(p)$ blow up at the boundary

need to clip gradient updates

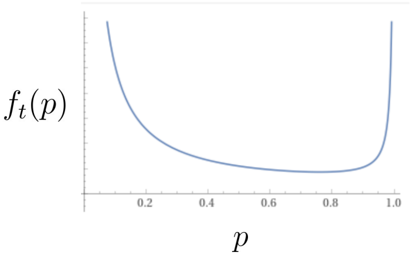

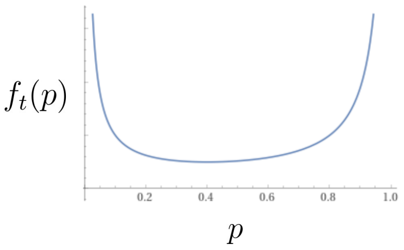

Our Algorithm: Clip-OGD (intuition)

Our Algorithm: Clip-OGD (intuition)

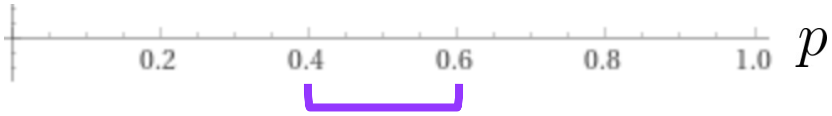

Our Algorithm: Clip-OGD (intuition)

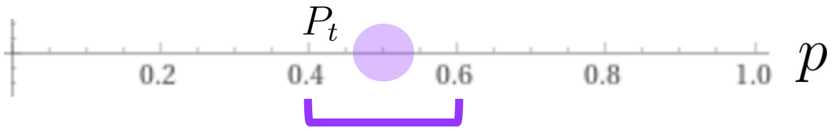

Our Algorithm: Clip-OGD (intuition)

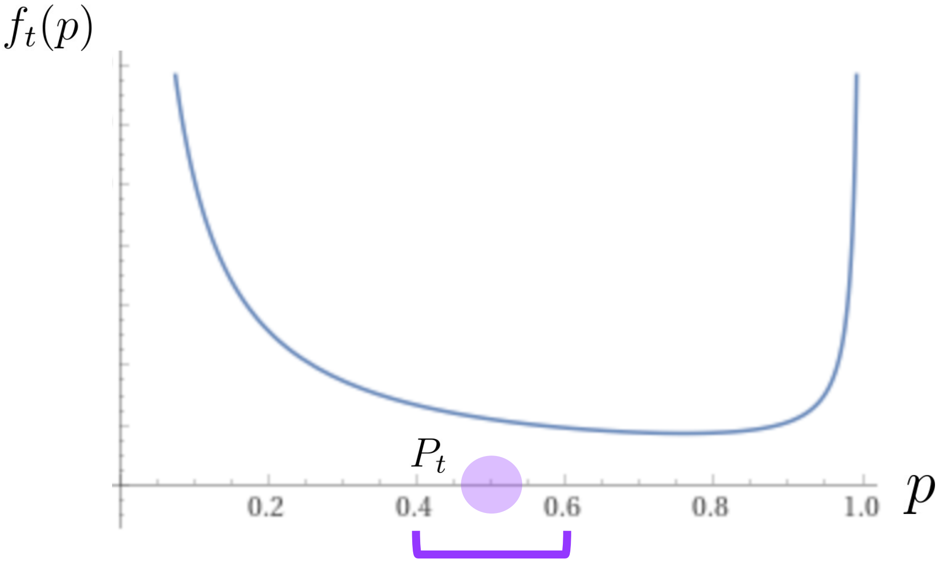

Our Algorithm: Clip-OGD (intuition)

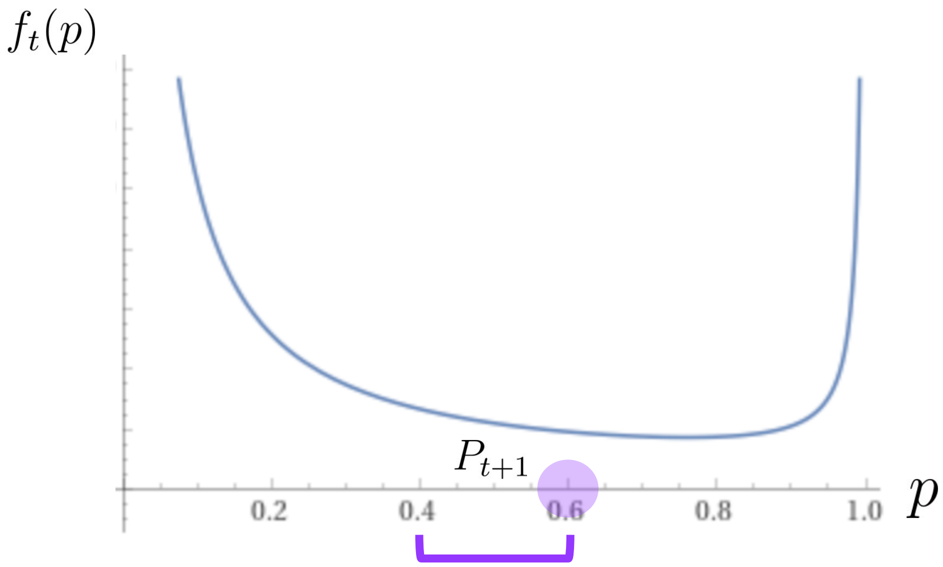

Our Algorithm: Clip-OGD (intuition)

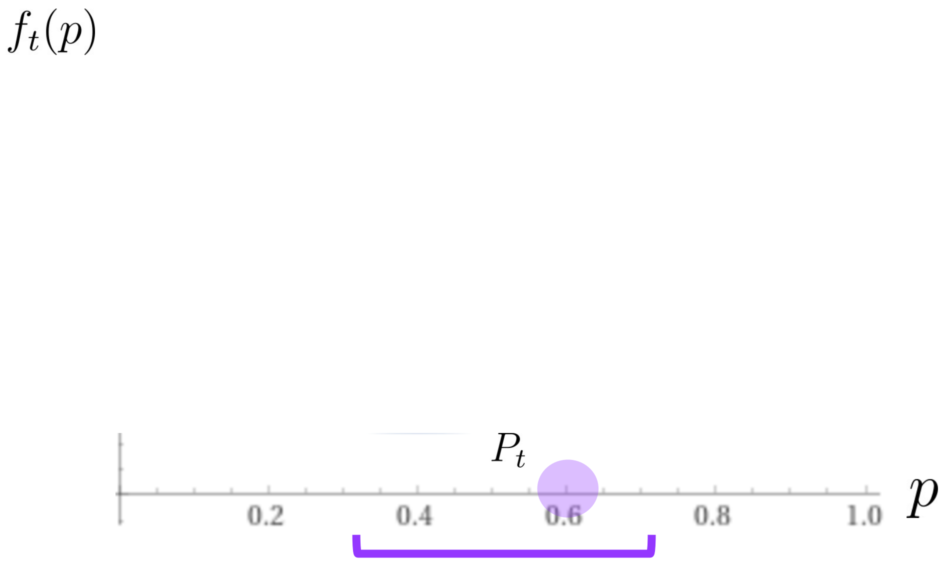

Our Algorithm: Clip-OGD (intuition)

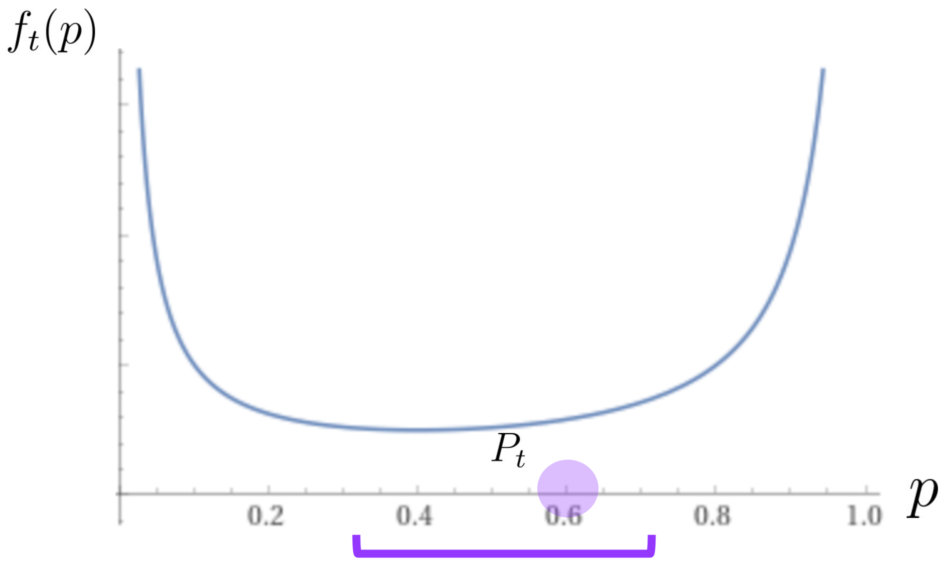

Our Algorithm: Clip-OGD (intuition)

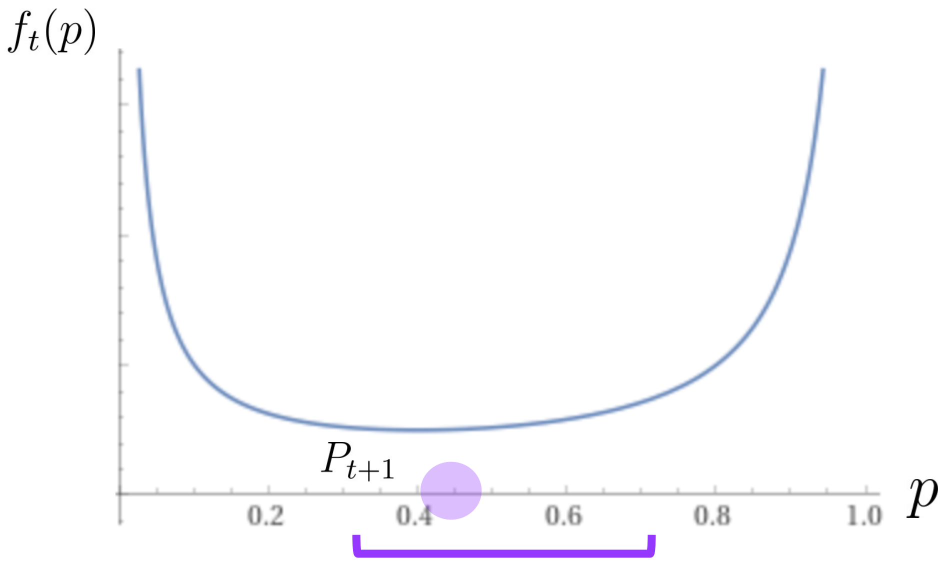

Our Algorithm: Clip-OGD

Algorithm 1: CLIP-OGD

Input: Step size η and decay parameter α

Initialize $P_0 \leftarrow 1/2$ and $G_0 \leftarrow 0$

for $t = 1 \dots T$ **do**

Set projection parameter $\delta_t = (1/2) \cdot t^{-1/\alpha}$

Compute new treatment probability $P_t \leftarrow \mathcal{P}_{\delta_t}(P_{t-1} - \eta \cdot G_{t-1})$

Sample treatment assignment Z_t as 1 with probability P_t and 0 with probability $1 - P_t$

Observe outcome $Y_t = \mathbf{1}[Z_t = 1]y_t(1) + \mathbf{1}[Z_t = 0]y_t(0)$

Construct gradient estimator $G_t = Y_t^2 \left(-\frac{\mathbf{1}[Z_t=1]}{P_t^3} + \frac{\mathbf{1}[Z_t=0]}{(1-P_t)^3} \right)$

end

Our Algorithm: Clip-OGD

Algorithm 1: CLIP-OGD

Input: Step size η and decay parameter α

Initialize $P_0 \leftarrow 1/2$ and $G_0 \leftarrow 0$

for $t = 1 \dots T$ **do**

clipped interval for current unit; grows with t

Set projection parameter $\delta_t = (1/2) \cdot t^{-1/\alpha}$

Compute new treatment probability $P_t \leftarrow \mathcal{P}_{\delta_t}(P_{t-1} - \eta \cdot G_{t-1})$

Sample treatment assignment Z_t as 1 with probability P_t and 0 with probability $1 - P_t$

Observe outcome $Y_t = \mathbf{1}[Z_t = 1]y_t(1) + \mathbf{1}[Z_t = 0]y_t(0)$

Construct gradient estimator $G_t = Y_t^2 \left(-\frac{\mathbf{1}[Z_t=1]}{P_t^3} + \frac{\mathbf{1}[Z_t=0]}{(1-P_t)^3} \right)$

end

Our Algorithm: Clip-OGD

Algorithm 1: CLIP-OGD

Input: Step size η and decay parameter α

Initialize $P_0 \leftarrow 1/2$ and $G_0 \leftarrow 0$

for $t = 1 \dots T$ **do**

  Set projection parameter $\delta_t = (1/2) \cdot t^{-1/\alpha}$

  Compute new treatment probability $P_t \leftarrow \mathcal{P}_{\delta_t}(P_{t-1} - \eta \cdot G_{t-1})$

  Sample treatment assignment Z_t as 1 with probability P_t and 0 with probability $1 - P_t$

  Observe outcome $Y_t = \mathbf{1}[Z_t = 1] y_t(1) + \mathbf{1}[Z_t = 0] y_t(0)$

  Construct gradient estimator $G_t = Y_t^2 \left(-\frac{\mathbf{1}[Z_t=1]}{P_t^3} + \frac{\mathbf{1}[Z_t=0]}{(1-P_t)^3} \right)$

end

Our Algorithm: Clip-OGD

Algorithm 1: CLIP-OGD

Input: Step size η and decay parameter α
Initialize $P_0 \leftarrow 1/2$ and $G_0 \leftarrow 0$
for $t = 1 \ldots T$ **do**
  Set projection parameter $\delta_t = (1/2) \cdot t^{-1/\alpha}$
  Compute new treatment probability $P_t \leftarrow \mathcal{P}_{\delta_t}(P_{t-1} - \eta \cdot G_{t-1})$
  Sample treatment assignment Z_t as 1 with probability P_t and 0 with probability $1 - P_t$
  Observe outcome $Y_t = \mathbf{1}[Z_t = 1]y_t(1) + \mathbf{1}[Z_t = 0]y_t(0)$
  Construct gradient estimator $G_t = Y_t^2 \left(-\frac{\mathbf{1}[Z_t=1]}{P_t^3} + \frac{\mathbf{1}[Z_t=0]}{(1-P_t)^3} \right)$
end

Our Algorithm: Clip-OGD

Algorithm 1: CLIP-OGD

Input: Step size η and decay parameter α
Initialize $P_0 \leftarrow 1/2$ and $G_0 \leftarrow 0$
for $t = 1 \ldots T$ **do**
  Set projection parameter $\delta_t = (1/2) \cdot t^{-1/\alpha}$
  Compute new treatment probability $P_t \leftarrow \mathcal{P}_{\delta_t}(P_{t-1} - \eta \cdot G_{t-1})$
  Sample treatment assignment Z_t as 1 with probability P_t and 0 with probability $1 - P_t$
  Observe outcome $Y_t = \mathbf{1}[Z_t = 1]y_t(1) + \mathbf{1}[Z_t = 0]y_t(0)$
  Construct gradient estimator $G_t = Y_t^2\left(-\frac{\mathbf{1}[Z_t=1]}{P_t^3} + \frac{\mathbf{1}[Z_t=0]}{(1-P_t)^3}\right)$
end

G_t estimates the gradient of $f_t(P_t)$

Main Result: Neyman Regret of Clip-OGD

Main Result: Neyman Regret of Clip-OGD

Assume: $\exists c \leq C$ *such that*

1. *For* $i \in \{0, 1\}$, $c \leq \left(\frac{1}{T} \sum_{t \in [T]} y_t(i)^2 \right)^{1/2} \leq \left(\frac{1}{T} \sum_{t \in [T]} y_t(i)^4 \right)^{1/4} \leq C$
2. *Cosine similarity bounded below:* $\rho \geq -(1 - c)$

Main Result: Neyman Regret of Clip-OGD

Assume: $\exists c \leq C$ *such that*

1. *For* $i \in \{0, 1\}$, $c \leq \left(\frac{1}{T} \sum_{t \in [T]} y_t(i)^2 \right)^{1/2} \leq \left(\frac{1}{T} \sum_{t \in [T]} y_t(i)^4 \right)^{1/4} \leq C$

2. *Cosine similarity bounded below:* $\rho \geq -(1 - c)$

Set: $\eta = \sqrt{1/T}$, $\alpha = \sqrt{5 \log(T)}$

Main Result: Neyman Regret of Clip-OGD

Assume: $\exists c \leq C$ such that

1. *For* $i \in \{0, 1\}$, $c \leq \left(\frac{1}{T} \sum_{t \in [T]} y_t(i)^2 \right)^{1/2} \leq \left(\frac{1}{T} \sum_{t \in [T]} y_t(i)^4 \right)^{1/4} \leq C$
2. *Cosine similarity bounded below:* $\rho \geq -(1 - c)$

Set: $\eta = \sqrt{1/T}$, $\alpha = \sqrt{5 \log(T)}$

Clip-OGD approaches the Neyman (optimal) variance at a rate of $\sqrt{T}$:

$$\mathbb{E}[\mathcal{R}_T] = \tilde{O}(\sqrt{T})$$

Main Result: Neyman Regret of Clip-OGD

Assume: $\exists c \leq C$ such that

1. *For* $i \in \{0, 1\}$, $c \leq \left(\frac{1}{T} \sum_{t \in [T]} y_t(i)^2 \right)^{1/2} \leq \left(\frac{1}{T} \sum_{t \in [T]} y_t(i)^4 \right)^{1/4} \leq C$

2. *Cosine similarity bounded below:* $\rho \geq -(1 - c)$

Set: $\eta = \sqrt{1/T}$, $\alpha = \sqrt{5 \log(T)}$

Clip-OGD approaches the Neyman (optimal) variance at a rate of $\sqrt{T}$:

(hides subpolynomial factors)

$$\mathbb{E}[\mathcal{R}_T] = \tilde{O}(\sqrt{T})$$

Variance Estimation: asymptotically valid confidence intervals

Variance Estimation: asymptotically valid confidence intervals

(Nonstandard) variance estimation approach:

Variance Estimation: asymptotically valid confidence intervals

(Nonstandard) variance estimation approach:

$$T \cdot V_{\text{Clip-OGD}} \to T \cdot V_N$$

by construction of Clip-OGD

Variance Estimation: asymptotically valid confidence intervals

(Nonstandard) variance estimation approach:

$$T \cdot V_{\text{Clip-OGD}} \to T \cdot V_N$$

by construction of Clip-OGD

$$= 2(1+\rho) S_1 S_0$$

by definition of Neyman variance

Variance Estimation: asymptotically valid confidence intervals

(Nonstandard) variance estimation approach:

$$
\begin{aligned}
T \cdot V_{\text{Clip-OGD}} &\to T \cdot V_N && \text{by construction of Clip-OGD} \\
&= 2(1+\rho) S_1 S_0 && \text{by definition of Neyman variance} \\
&\leq 4 S_1 S_0 && \text{upper bound on } \rho
\end{aligned}
$$

Variance Estimation: asymptotically valid confidence intervals

(Nonstandard) variance estimation approach:

$$
\begin{aligned}
T \cdot V_{\text{Clip-OGD}} &\to T \cdot V_N && \text{by construction of Clip-OGD} \\
&= 2(1+\rho) S_1 S_0 && \text{by definition of Neyman variance} \\
&\le \boxed{4 S_1 S_0} && \text{upper bound on } \rho \\
&\triangleq \boxed{T \cdot \mathrm{VB}} && \text{quantity to estimate}
\end{aligned}
$$

Variance Estimation: asymptotically valid confidence intervals

(Nonstandard) variance estimation approach:

$$
\begin{aligned}
T \cdot V_{\text{Clip-OGD}} &\to T \cdot V_N && \text{by construction of Clip-OGD} \\
&= 2(1+\rho) S_1 S_0 && \text{by definition of Neyman variance} \\
&\leq \boxed{4 S_1 S_0} && \text{upper bound on } \rho \\
&\triangleq \boxed{T \cdot \mathrm{VB}} && \text{quantity to estimate}
\end{aligned}
$$

Chebyshev-style confidence intervals:

$$
\lim_{T \to \infty} \inf \Pr \left[\tau \in \hat{\tau} \pm \alpha^{-1/2} \sqrt{\widehat{\mathrm{VB}}} \right] \geq 1 - \alpha
$$

Variance Estimation: asymptotically valid confidence intervals

(Nonstandard) variance estimation approach:

$$
\begin{aligned}
T \cdot V_{\text{Clip-OGD}} \to T \cdot V_N & \quad \text{by construction of Clip-OGD} \\
= 2(1+\rho) S_1 S_0 & \quad \text{by definition of Neyman variance} \\
\leq 4 S_1 S_0 & \quad \text{upper bound on } \rho \\
\triangleq T \cdot \mathrm{VB} & \quad \text{quantity to estimate}
\end{aligned}
$$

Chebyshev-style confidence intervals:

$$\lim_{T \to \infty} \inf \Pr \left[\tau \in \hat{\tau} \pm \alpha^{-1/2} \sqrt{\widehat{\mathrm{VB}}} \right] \geq 1 - \alpha$$

Numerical Results: Microfinance simulation (Groh & McKenzie 2016)

Numerical Results: Microfinance simulation (Groh & McKenzie 2016)

Treatment: provide a new insurance product

Outcome: amount of money invested in equipment

We impute missing potential outcomes & extend size from T=2961 to 14445

Numerical Results: Microfinance simulation (Groh & McKenzie 2016)

Treatment: provide a new insurance product

Outcome: amount of money invested in equipment

We impute missing potential outcomes & extend size from T=2961 to 17445

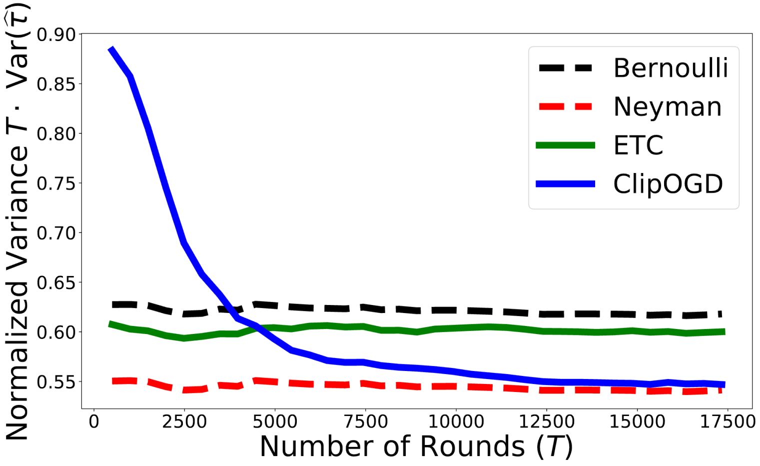

Numerical Results: Microfinance simulation (Groh & McKenzie 2016)

Treatment: provide a new insurance product

Outcome: amount of money invested in equipment

We impute missing potential outcomes & extend size from T=2961 to 17445

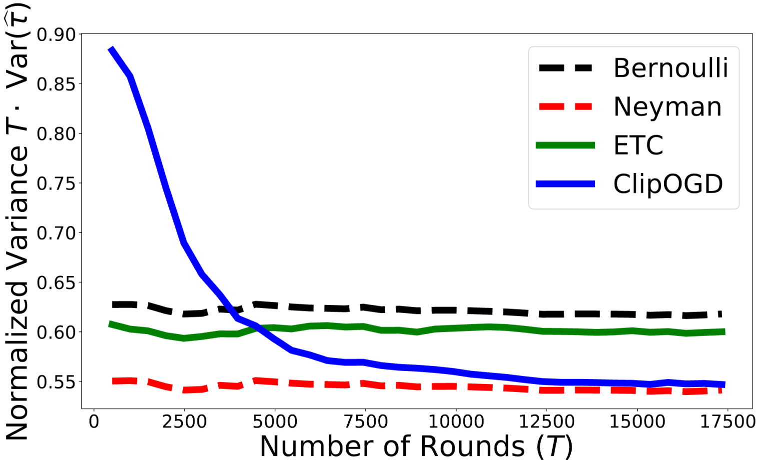

What if units came in a slightly different order?

Numerical Results: Microfinance simulation (Groh & McKenzie 2016)

Treatment: provide a new insurance product

Outcome: amount of money invested in equipment

We impute missing potential outcomes & extend size from T=2961 to 17445

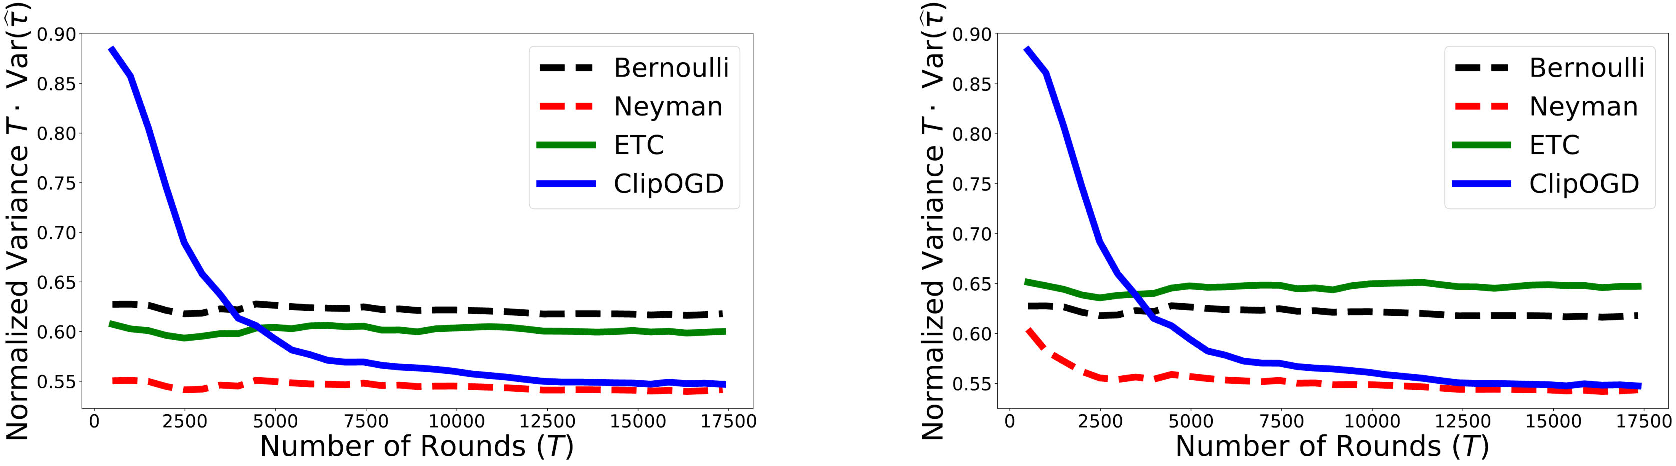

With first 100 outcomes corrupted

Additional Results: Other Designs?

***O(T)* lower bound** on explore-then-commit design

Impossibility result/ tradeoff for outcome and Neyman regret

Thank you!

Jessica Dai, UC Berkeley
jessicadai@berkeley.edu

Clip-OGD: An Experimental Design for Adaptive Neyman Allocation
Jessica Dai, Paula Gradu, Chris Harshaw
NeurIPS 2023
https://arxiv.org/abs/2305.17187

Chris Harshaw
MIT

Paula Gradu
U.C. Berkeley

Impossibility result: Neyman Regret and Outcome Regret

Let $\mathcal{R}_T^{\text{outcome}} = \frac{1}{T} \sum_{t \in [T]} Y_t - \min_{i \in \{0,1\}} \sum_{t \in [T]} Y_t(i)$.

Suppose algorithm ***A*** for sequential allocation is no-outcome-regret:

$$\mathbb{E}[\mathcal{R}_T^{\text{outcome}}] = \mathcal{O}(T^q) \quad \text{for } q \in (0,1)$$

Then ***A*** must suffer supralinear Neyman regret:

$$\mathbb{E}[\mathcal{R}_T^{\text{Neyman}}] \geq \Omega(T)$$

No math here

Any allocation which prioritizes outcomes ***within the experiment*** *will worsen the* ***information gained from the experiment.***

Proof sketch:

definition of outcome regret implies

$$\sum_{t\in[T]} \mathbb{E}[P_t] \leq O(T^q)$$

which means that

$$\sum_{t\in[T]} \frac{1}{\mathbb{E}[P_t]} \geq O(T^{2-q})$$

which lower bounds

$$\mathbb{E}\left[\sum_{t\in[T]} f_t(P_t)\right] \geq O(T^{2-q})$$

and therefore the Neyman regret.