# Optimal Experimental Design for Staggered Rollouts

Ruoxuan Xiong, Susan Athey, Mohsen Bayati, Guido Imbens

Emory University and Stanford University

# Introduction

# Designing experiments with staggered rollouts

- Estimating treatment effects in panel data with **staggered rollouts**
  - Units $i \in \{1, \cdots, N\}$ observed in time periods $s \in \{1, \cdots, T\}$
  - Design: Treatment assignment $Z_{is} \in \{0, 1\}$
  - Potential outcomes: $Y_{is}(z_{i,s-\ell}, \cdots, z_{is})$ may depend on the history of treatment to date, with known $\ell$ periods of history that matter
  - Observed outcomes: $Y_{is} = Y_{is}(Z_{i,s-\ell}, \cdots, Z_{is})$
- Staggered rollout designs commonly encountered in observational data:
  - Products/promotions released in different regions at different times
  - State regulations adopted over time
- Question: How should analyst **design** a staggered rollout experiment?
  - How fast should rollout occur?
  - How does rollout depend on hypothesized maximum duration of carryover effects?
  - How can historical data be used to optimize design?
  - Can an **adaptive design**, where analyst updates speed of rollout and termination based on data collected during experiment, improve performance?

# Panel experiments with staggered rollouts

**Formal objective**: Propose experimental designs that optimize the precision of post-experiment estimates of treatment effects

**Focus on environment with**: Irreversible treatment adoption pattern ($Z_{is} \leq Z_{i,s+1}$)

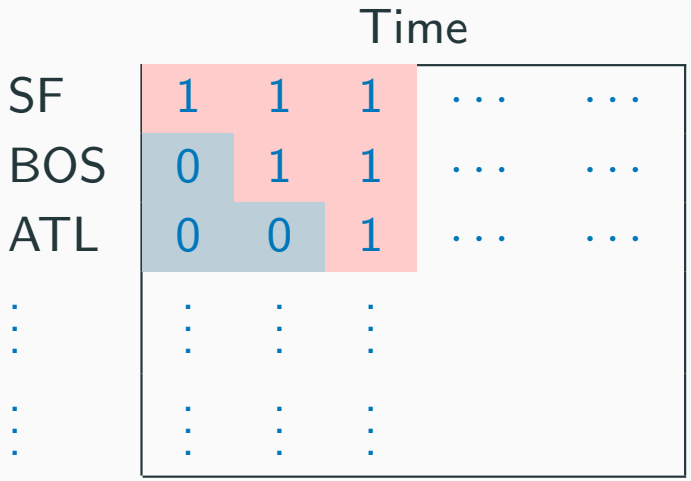


0 denotes control and 1 denotes treated

# Contribution: Non-adaptive experiments

**Non-adaptive experiments**: $N$ and $T$ are set, and treatment decisions are made, pre-experiment

- Assume after experiment will use GLS to estimate instantaneous and lagged treatment effects from nonstationary observed outcomes
- Analytical optimality conditions for the designs that maximize linearly combined precisions of estimated instantaneous and lagged effects
- Propose an algorithm to choose a treatment design based on the optimality conditions. The design has two features
  - $\Rightarrow$ Fraction of treated units per period takes an $S$-shaped curve: Treatment rollouts slowly at the beginning and end, and quickly in the middle
    - Bigger $\ell$ leads to more pronounced $S$
  - $\Rightarrow$ This rollout pattern is imposed for each stratum of units with the same observed and estimated latent covariate values

# Illustration of optimal assignment

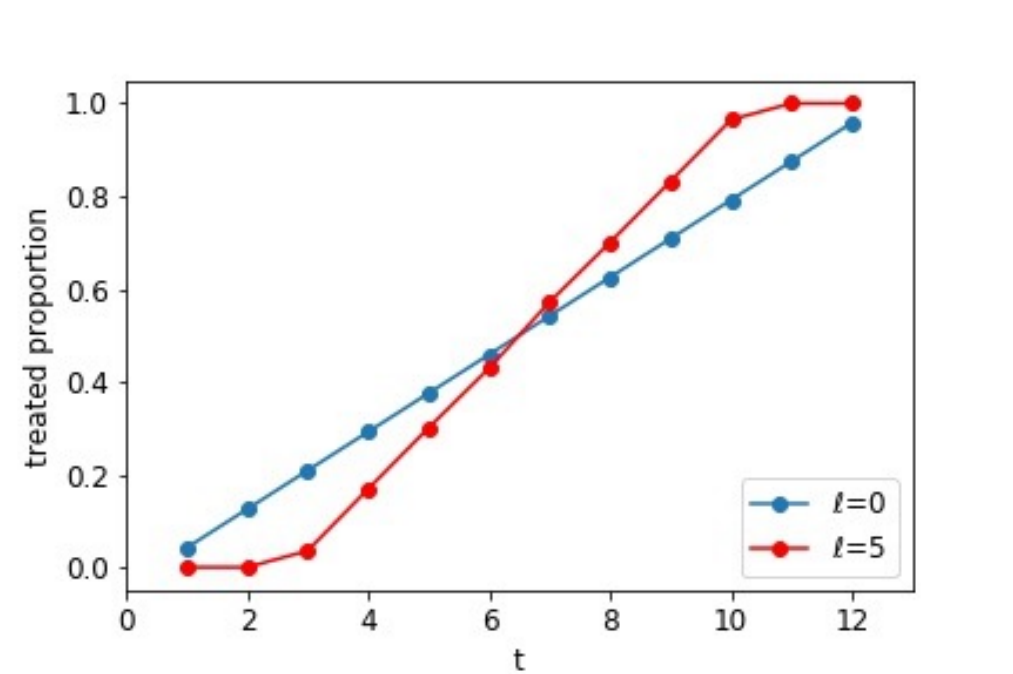

# Contribution: Adaptive experiments

**Adaptive experiments**: $N$ is fixed, but the experiment can be terminated early. Treatment decisions are updated after each period's data is collected

- Propose the Precision-Guided Adaptive Experiment (PGAE) algorithm
  - adaptively terminates the experiment based on the estimated precision
  - adaptively optimizes speed of rollout using dynamic programming
  - an estimation scheme of treatment effects based on sample splitting
- Derive the asymptotic normal distribution of final treatment effect and variance estimates from PGAE
  - Optimal convergence rate and no efficiency loss of final treatment effect estimate, as compared to an oracle with access to the same design a-priori

## Related literature (partial list)

- Most closely related to stepped wedge designs in clinical trials (Hussey and Hughes 2007, Hemming et al. 2015, Li, Turner, and Preisser 2018)
  - ⇒ We study the design under a more general outcome specification, where cumulative effects can vary with treatment duration
- Recently proposed alternative designs for estimation of carryover effects
  - Minimax temporal experimental design (Basse, Ding, and Toulis 2019)
  - Switchback design (Bojinov, Simchi-Levi, and Zhao 2020)
  - Synthetic control design that selects units for (simultaneous) treatment, anticipating synthetic control estimation (Doudchenko et al. 2021a,b, Abadie and Zhao 2021)
  - ⇒ Our design leverages variation of treatment times across units and maximizes the precision of treatment effect estimates
- Recently proposed designs in settings with interference
  - Multiple randomization designs (Bajari et al. 2021, Johari et al. 2022)
  - Equilibrium designs (Wager and Xu 2021)
  - ⇒ Our experiment is run at the aggregate level and leverages the time dimension to increase power

⇒ We also consider **adaptive** designs; above papers pre-specify design

# Two examples for staggered rollout experiments

**Example 1** (marketplace experiments): A ride-hailing platform plans to test the impact of a new app feature that improves driver experience

**Example 2** (public health intervention): A country aims to measure the effect of a new public health intervention (e.g., encouraging the use of masks or social distancing policies) on the spread of an infectious disease

Staggered rollout experiments run at the city level for multiple time periods can

- avoid bias from interference
- facilitate the estimation of cumulative effects
- better design can improve the estimation precision of cumulative effects

# Setup

# Potential outcomes and treatment effects

- The potential outcomes for unit $i$ at time $s$ can be written as

$$Y_{is}(z_{i,s-\ell}, \cdots, z_{i,s-1}, z_{is})$$

  for a nonnegative, known integer $\ell$ ($\ell$: duration of treatment effects)

- Let the average instantaneous effect $\tau_0$ and $j$-th period lagged effect $\tau_j$ be

$$\tau_j := \frac{1}{NT} \sum_{i,s} \left[ Y_{is}(0, \cdots, 0, \underbrace{1}_{z_{i,s-j}}, 1, \cdots, 1) - Y_{is}(0, \cdots, 0, \underbrace{0}_{z_{i,s-j}}, 1, \cdots, 1) \right],$$

  for all $j \in \{0, 1, \cdots, \ell\}$

- Let the average cumulative effect of treatment for $j$ periods be

$$\tau_0 + \cdots + \tau_{\min(\ell, j-1)}$$

  that is constant for $j > \ell$

# Illustrative examples of cumulative effects

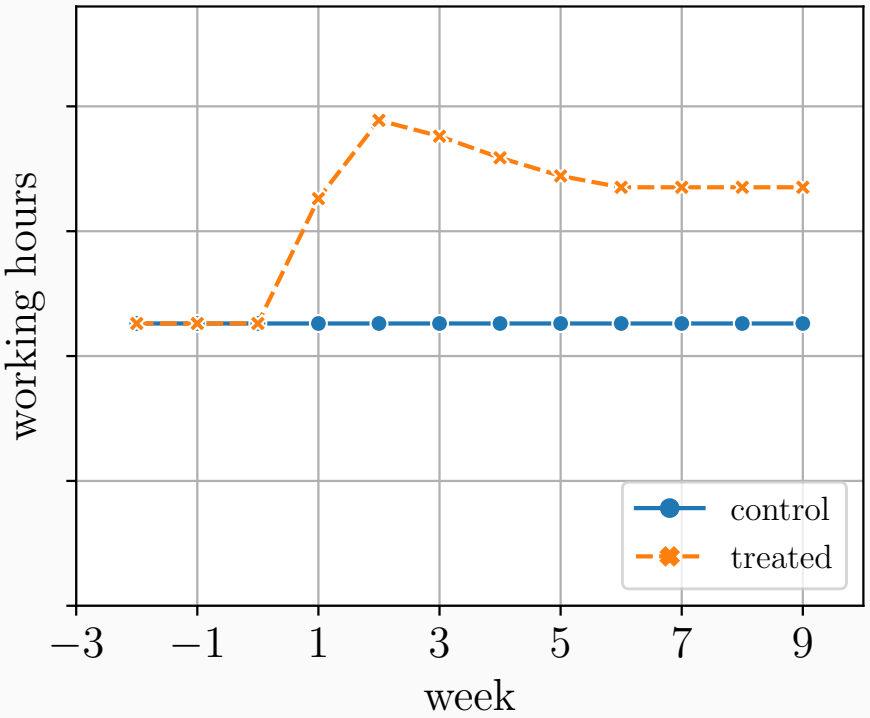


Cumulative effect of treatment for $j$ periods with $\ell = 5$, $\tau_0, \tau_1 > 0$ and $\tau_2, \tau_3, \tau_4, \tau_5 < 0$

# Illustrative examples of cumulative effects

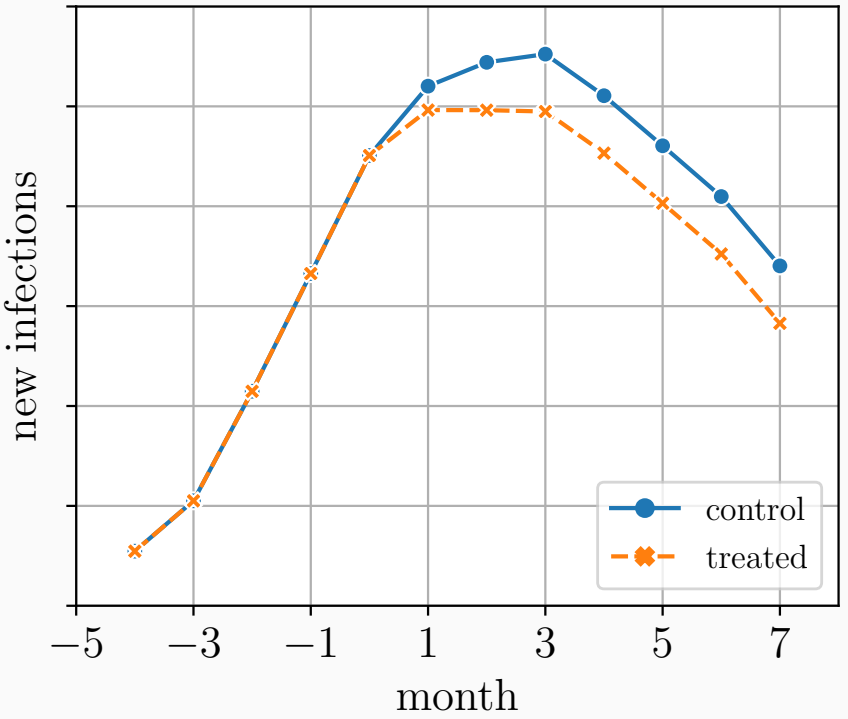


Cumulative effect of treatment for $j$ periods with $\ell = 2$ and $\tau_0, \tau_1, \tau_2 < 0$

# Outcome specification

A general outcome specification for treatment effect estimation post-experiment

$$Y_{is} = \alpha_i + \beta_s + \mathbf{X}_i^\top \boldsymbol{\theta}_s + \tau_0 z_{is} + \tau_1 z_{i,s-1} + \cdots + \tau_\ell z_{i,s-\ell} + \underbrace{\mathbf{u}_i^\top \mathbf{v}_s + \varepsilon_{is}}_{e_{is}}$$

- $\alpha_i$: unknown unit fixed effect
- $\beta_s$: unknown time fixed effect
- $\mathbf{X}_i$: observed covariates; $\boldsymbol{\theta}_s$: unknown time-varying coefficients
- $\mathbf{u}_i$: latent covariates; $\mathbf{v}_s$: latent coefficients
- $\varepsilon_{is}$: iid residual with mean 0 and variance $\sigma^2$

# Decision making problem

**Decision**: Optimally choose the treatment times for each unit

**Goal**: Most precisely estimate average instantaneous and lagged effects

**Implication**: Reduce sample size requirement and lower the experimental cost!

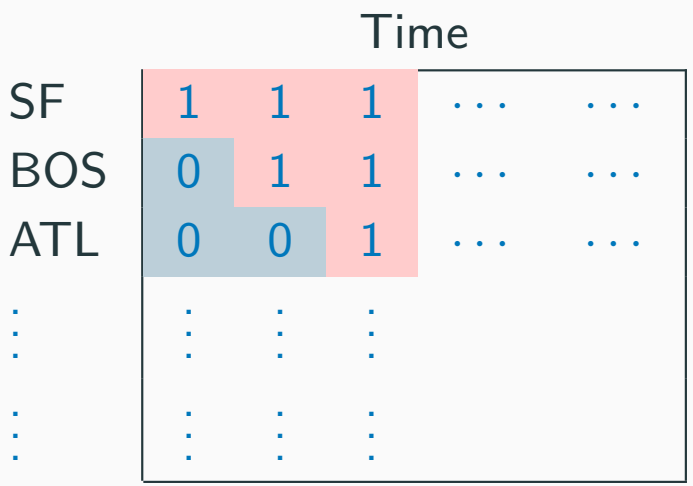


0 denotes control and 1 denotes treated

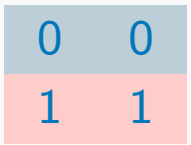


$Z_{ff}$?

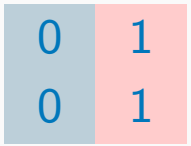


$Z_{ba}$?

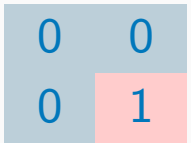


$Z_{ffba}$?

# Non-adaptive experiments

# Estimation of average treatment effects

GLS estimator $\hat{\tau}_0, \cdots, \hat{\tau}_\ell$ from the specification

$$Y_{is} = \alpha_i + \beta_s + \mathbf{X}_i^\top \boldsymbol{\theta}_s + \tau_0 z_{is} + \tau_1 z_{i,s-1} + \cdots + \tau_\ell z_{i,s-\ell} + e_{is} \,,$$

- GLS is the best linear unbiased estimator (BLUE)
- Precision matrix (inverse of variance-covariance matrix) of $\hat{\tau}_0, \cdots, \hat{\tau}_\ell$, denoted by $\mathrm{Prec}(\hat{\tau}_0, \cdots, \hat{\tau}_\ell; Z)$, is a quadratic function of $Z = [z_{is}]_{(i,s)\in[N]\times[T]}$, where $[N]$ stands for $\{1, 2, \cdots, N\}$

# Designs of non-adaptive experiments

Trace(T)-optimal design: Choose $Z = [z_{is}]_{(i,s)\in[N]\times[T]}$ pre-experiment to maximize the trace of the precision matrix (Pukelsheim, 2016)

$$\begin{aligned} \max_{Z} \quad & \operatorname{trace}(\operatorname{Prec}(\hat{\tau}_0, \cdots, \hat{\tau}_\ell; Z)) \\ \text{s.t.} \quad & z_{is} \leq z_{i,s+1} \\ & z_{is} \in \{0, 1\} \end{aligned}$$

Other objective functions, for example, determinant(D)-optimal design and A-optimal design

- No analytical solutions in general
- Numerical solutions for D-optimal design in the paper

# Optimal solution (No covariates)

$$Y_{is} = \alpha_i + \beta_s + \tau_0 z_{is} + \tau_1 z_{i,s-1} + \cdots + \tau_\ell z_{i,s-\ell} + \varepsilon_{is} \tag{1}$$

**Theorem 1: Optimal solution (no covariates)**

Under the specification (1), $\varepsilon_{is} \overset{i.i.d.}{\sim} (0, \sigma^2)$ and $\tau_j$ is estimated from OLS. Then any treatment design is optimal if it satisfies

$$\omega_s = \frac{1}{N} \sum_i Z_{is} = \omega^*_{\ell,s} \,.$$

If $\ell = 0$, then $\omega^*_{\ell,s} = (2s-1)/(2T)$.

For general $\ell$, $\omega^*_{\ell,s}$ has five stages, and the expression of $\omega^*_{\ell,s}$ is provided in the paper.

# Visualization of $\omega^*_{\ell,s}$ in optimal solution

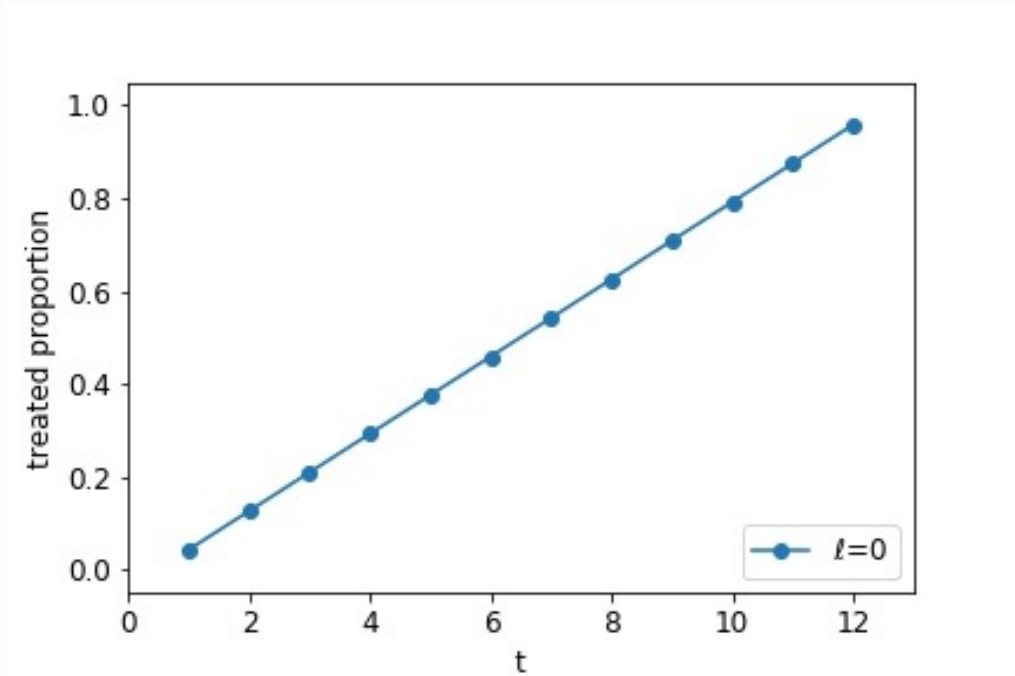

# Visualization of $\omega^*_{\ell,s}$ in optimal solution

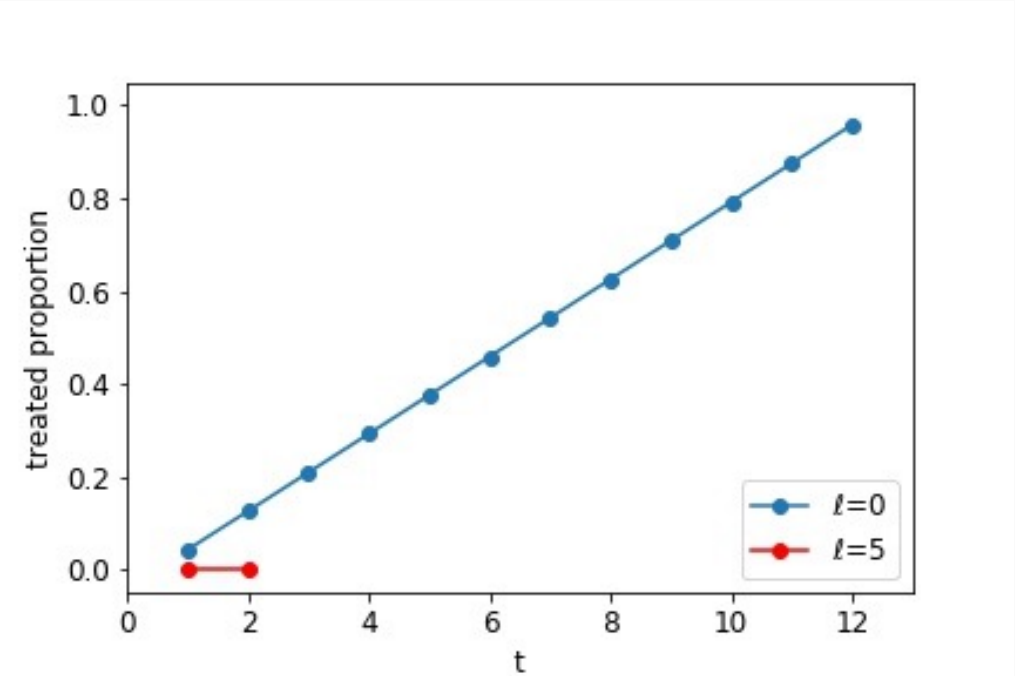

# Visualization of $\omega^*_{\ell,s}$ in optimal solution

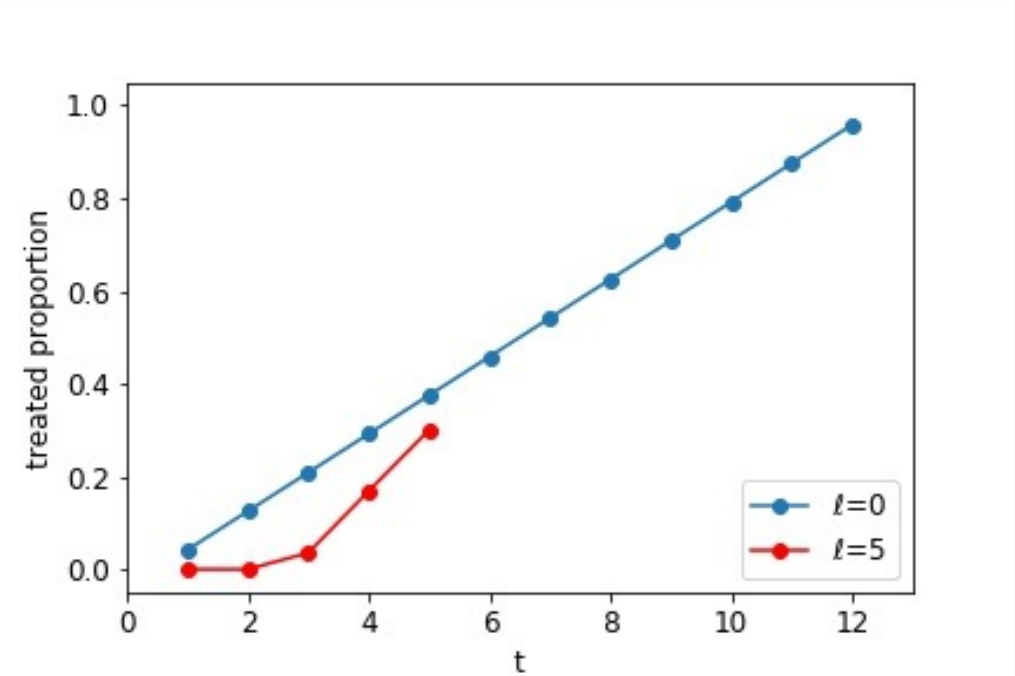

# Visualization of $\omega^*_{\ell,s}$ in optimal solution

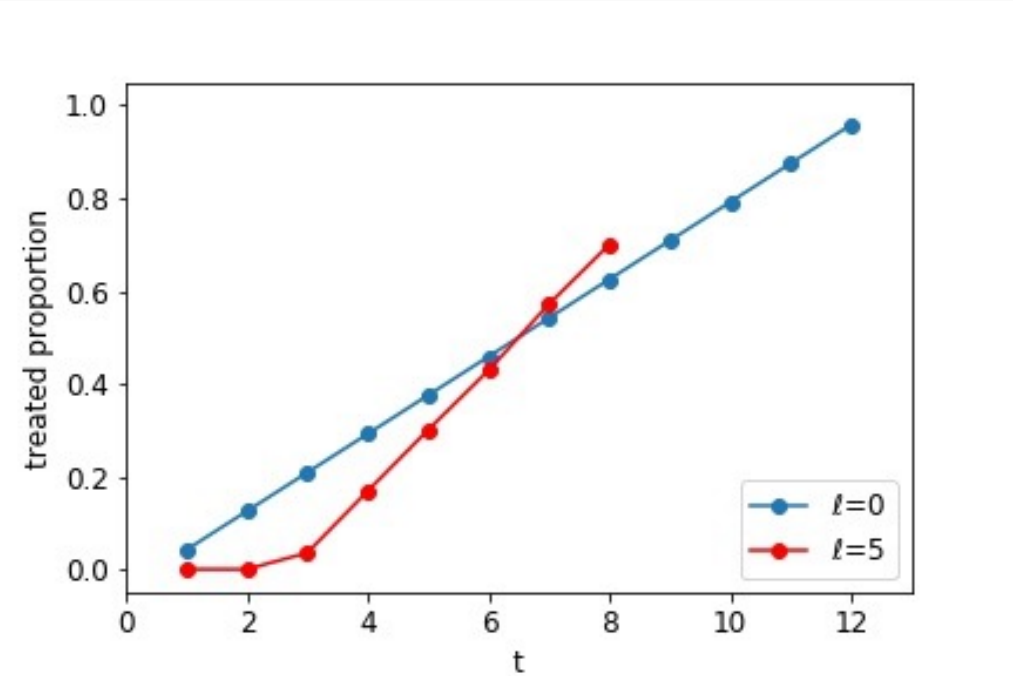

# Visualization of $\omega^*_{\ell,s}$ in optimal solution

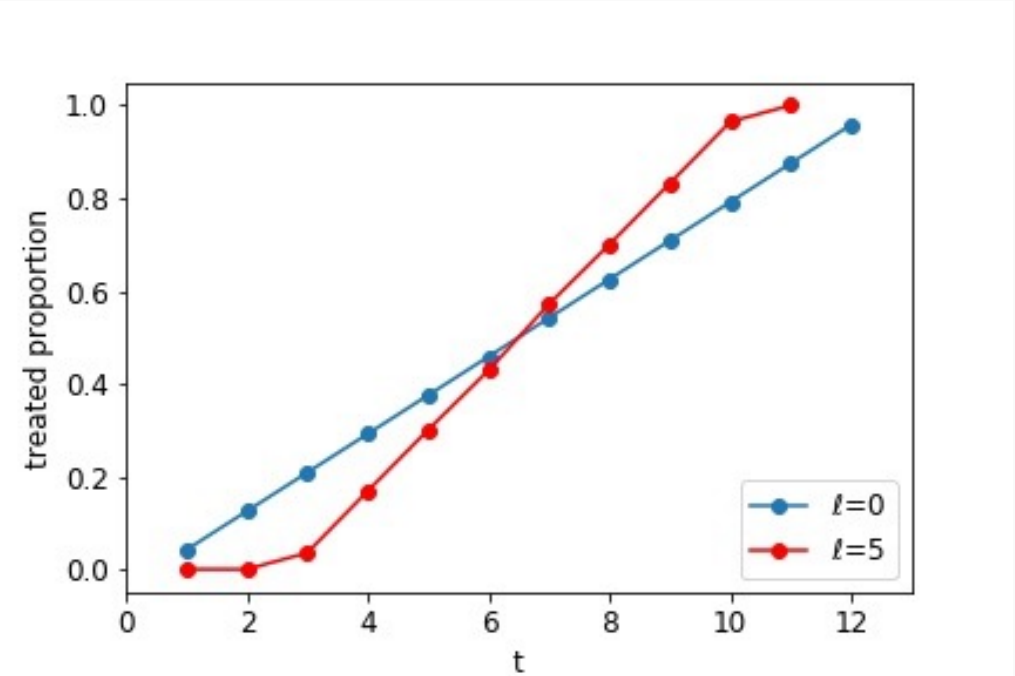

# Visualization of $\omega^*_{\ell,s}$ in optimal solution

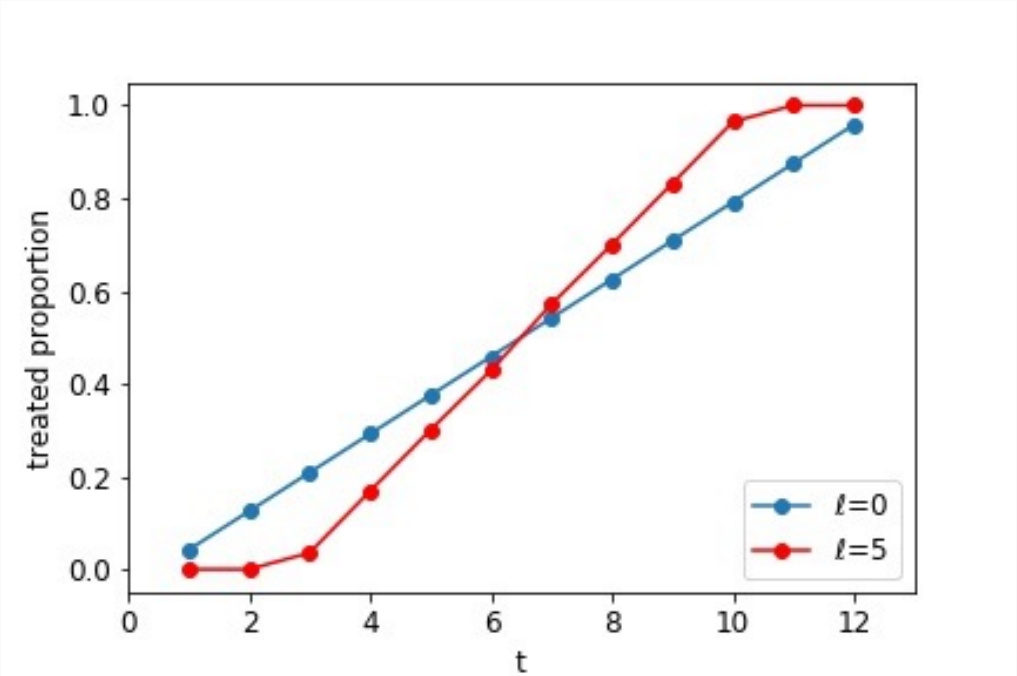

# Optimal solution (Adding observed and/or latent covariates)

$$Y_{is} = \alpha_i + \beta_s + \mathbf{X}_i^\top \boldsymbol{\theta}_s + \tau_0 z_{is} + \tau_1 z_{i,s-1} + \cdots + \tau_\ell z_{i,s-\ell} + \underbrace{\mathbf{u}_i^\top \mathbf{v}_s + \varepsilon_{is}}_{e_{is}} \quad (2)$$

**Theorem 1: Optimal solution (with covariates)**

Under the specification (2), $\varepsilon_{is} \overset{i.i.d.}{\sim} (0, \sigma^2)$, both $\mathbf{X}_i$ and $\mathbf{u}_i$ are demeaned, and $\tau_j$ is estimated from infeasible GLS. Then any treatment design is optimal if it satisfies

- $\omega_s = N^{-1} \sum_i Z_{is} = \omega^*_{\ell,s}$
- $N^{-1} \sum_i \mathbf{X}_i Z_{is}$ is fixed for all $s$
- $N^{-1} \sum_i \mathbf{u}_i Z_{is}$ is fixed for all $s$

# Interpretation of optimal solution (with covariates)

With $\mathbf{X}_i$ only: Stratification if $\mathbf{X}_i$ is discrete-valued

- Each stratum (group of units with the same $\mathbf{X}_i$) satisfies the treated fraction conditions $\omega^*_{\ell,s}$ (possibly with rounding)

With $\mathbf{u}_i$: $\mathbf{u}_i$ is unknown in practice

- Estimate $\mathbf{u}_i$ using historical data
- Partition units into strata based on $\hat{\mathbf{u}}_i$

An algorithm proposed in the paper to choose a treatment design

$$\mathbf{X}_i = \mathbf{x}_1 \begin{cases} 0 & 0 \\ 0 & 1 \\ 0 & 1 \\ 1 & 1 \end{cases}$$

$$\mathbf{X}_i = \mathbf{x}_2 \begin{cases} 0 & 0 \\ 0 & 1 \\ 0 & 1 \\ 1 & 1 \end{cases}$$

# Adaptive experiments

# Decisions for adaptive experiments

Goal: Most precisely estimate average treatment effects with valid inference, using the least sample size

Two adaptive decisions:

- Stop the experiment early if the desired precision is achieved (i.e., max duration is $T_{\max}$, and duration $\tilde{T} \in [T_{\max}]$ is a random variable)
- Speed of treatment rollout for the next time period is determined after each period's outcomes are collected

This talk: Focus on a simpler specification

$$Y_{is} = \alpha_i + \beta_s + \tau_0 z_{is} + \varepsilon_{is}$$

# Decision 1: Experiment termination rule

Terminate the experiment if the precision exceeds a target threshold $c$ at time $t$ (Glynn and White 1992)

$$\text{Prec}(\hat{\tau}_0; Z) \geq c$$

where $Z \in \{0,1\}^{N \times t}$ and

$$\text{Prec}(\hat{\tau}_0; Z) = \frac{Nt}{\sigma^2} \cdot \underbrace{(-2\boldsymbol{b}_t^\top \boldsymbol{\omega}_{1:t} - \boldsymbol{\omega}_{1:t}^\top \mathbf{P}_{\mathbf{1}_t} \boldsymbol{\omega}_{1:t})/t}_{g_\tau(\boldsymbol{\omega}, t)}$$

with

- $\boldsymbol{\omega}_{1:t} = [\omega_s]_{s \in [t]}$ and $\omega_s = N^{-1} \sum_i Z_{is}$
- $\mathbf{P}_{\mathbf{1}_t} = I_t - \mathbf{1}_t \mathbf{1}_t^\top / t$ and $\boldsymbol{b}_t$ is a vector of constants
- $\sigma^2 = \mathbb{E}[\varepsilon_{it}^2]$

$\Rightarrow$ Termination rule needs key unknown parameter $\sigma^2$

$\Rightarrow$ Implement termination rule in a way that allows for valid inference of $\tau_0$ (due to the peeking challenge in sequential testing (Johari et al. 2017))

# Decision 2: Treatment assignment

$\tilde{T}$ is unknown for adaptive experiments, therefore infeasible to optimally choose the speed of treatment rollout, pre-experiment

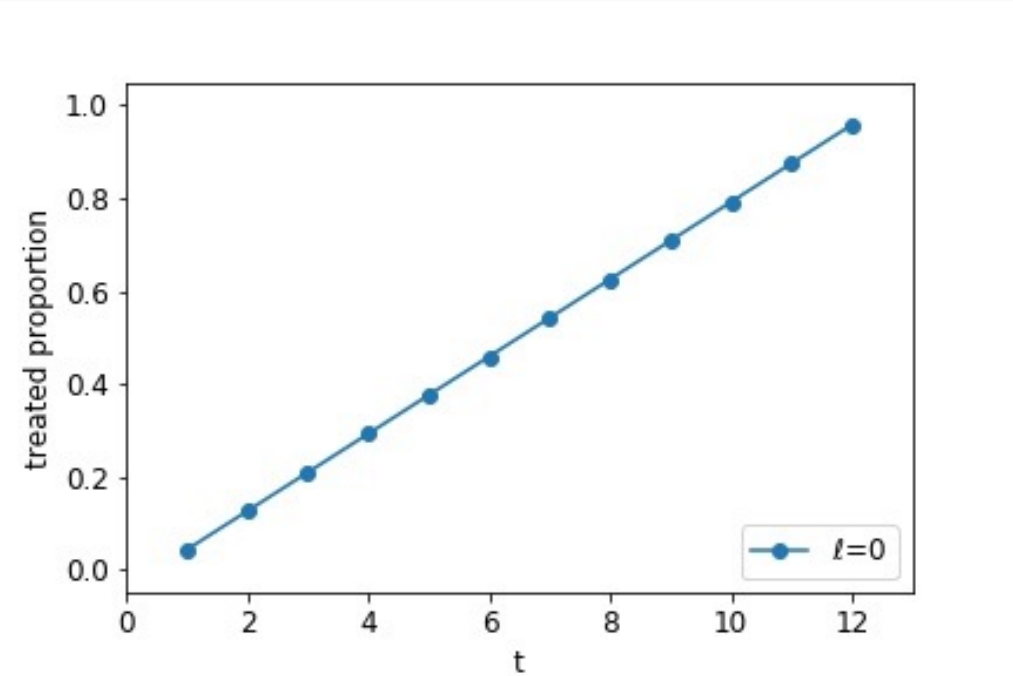


$$\omega^*_{0,s} = (2s-1)/(2T)$$

# Three competing goals in adaptive experiments

Goal 1: Choosing a treatment design

- Adaptively choose the speed of rollout, as we gather more information about $\sigma^2$ during the experiment

Goal 2: Implementing the termination rule

- Estimate $\sigma^2$ to make the next challenge manageable

Goal 3: Efficient estimation and valid inference for $\tau_0$

- Use as many observations as possible

Propose the Precision-Guided Adaptive Experiment (PGAE) algorithm

- simultaneously achieves the three goals
- uses sample splitting and dynamic programming

# PGAE: Sample splitting

Partition units into non-adaptive treatment units (NTU) and adaptive treatment units (ATU)

- NTU: Treatment design set pre-experiment (a small set)
  - Set as $\omega_{bm,s} = (2s-1)/(2T_{\max})$ (optimal solution for $T_{\max}$)
- ATU: Treatment design chosen adaptively

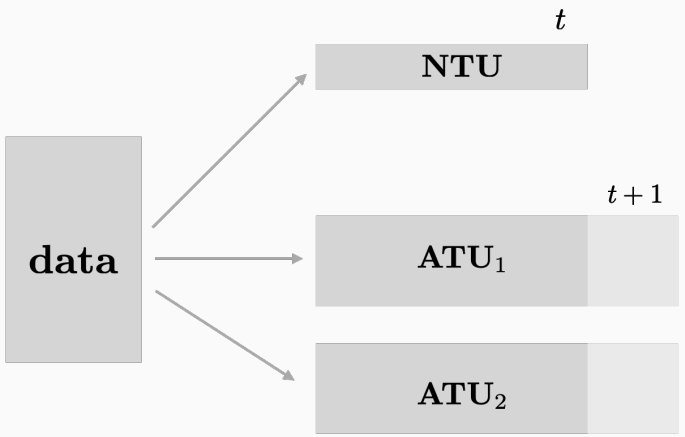

# Component 1 in PGAE: Choosing a treatment design

At time $t$, estimate distribution of $\sigma^2$ from NTU

- Estimate $\sigma^2 = \mathbb{E}[\varepsilon_{it}^2]$ and variance of $\varepsilon_{it}^2$, i.e., $\xi^2 = \mathbb{E}[(\varepsilon_{it}^2 - \sigma^2)^2]$
- Normal approximation of the distribution of $\sigma^2$ (based on the asymptotic normality of $\widehat{\sigma^2}$)

Update belief about $\tilde{T}$, denoted by $P_t(\tilde{T})$, using the estimated distribution of $\sigma^2$

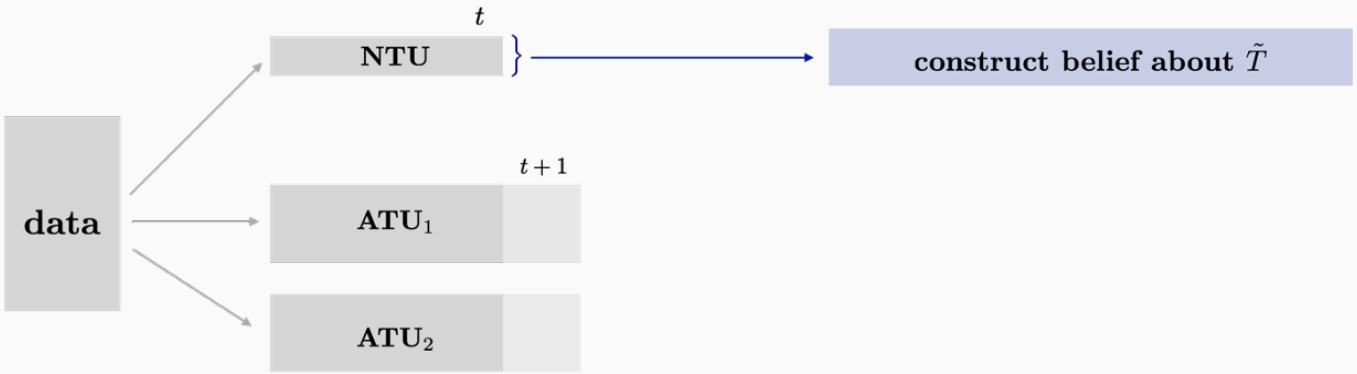

# Component 1 in PGAE: Choosing a treatment design

At time $t$, optimize $\omega_{t+1}$ for $\text{ATU}_1$ and $\text{ATU}_2$ through dynamic programming (DP)

- In the DP, no intermediate cost and terminal cost is the precision at termination, i.e., $\text{Prec}(\hat{\tau}_0; Z_{:,1:\tilde{T}}) = (N\tilde{T}/\sigma^2) \cdot g_\tau(\boldsymbol{\omega}, \tilde{T})$
- Solve $\omega_{t+1}$ from DP based on the belief about $\tilde{T}$

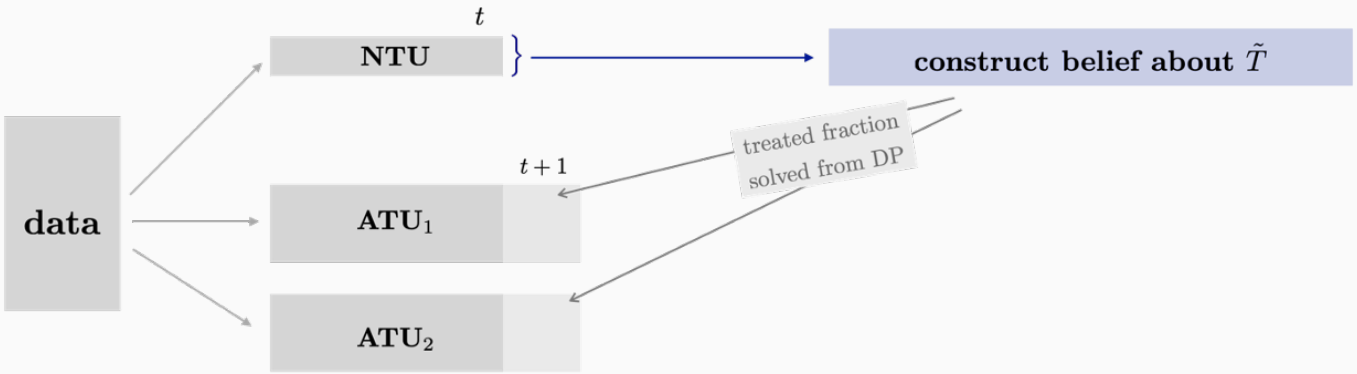

# Component 2 in PGAE: Implementing the termination rule

Estimate $\sigma^2$ from $\mathrm{ATU}_1$ and $\widehat{\mathrm{Prec}(\hat{\tau}_0; Z_{:,1:t})} = (Nt/\widehat{\sigma^2}) \cdot g_\tau(\omega, t)$

If $\widehat{\mathrm{Prec}(\hat{\tau}_0; Z_{:,1:t})} \geq c$, terminate the experiment; otherwise, keep running the experiment

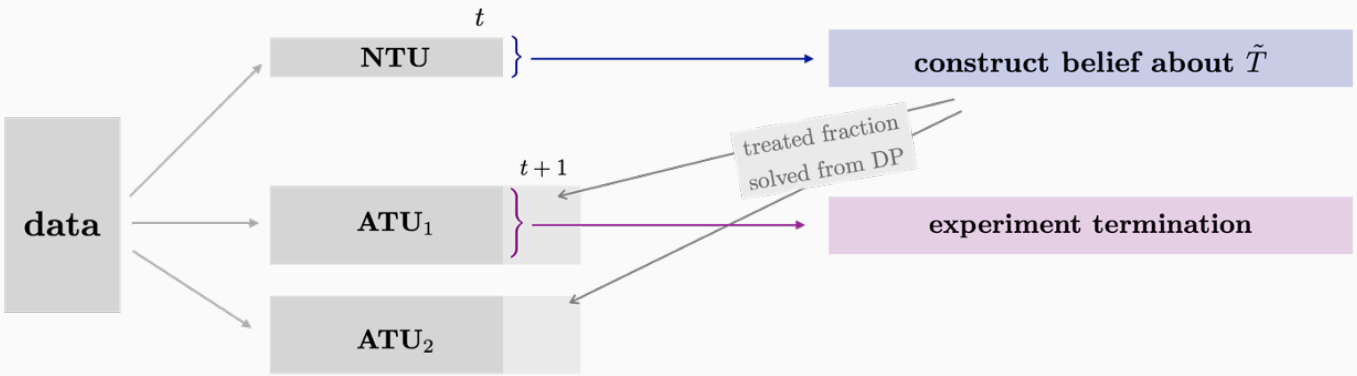

# Component 3 in PGAE: Efficient estimation and valid inference

Post-experiment,

- $\hat{\tau}_{\text{all},\tilde{T}}$: estimator of $\tau_0$ using all $N$ units and $\tilde{T}$ periods of data (no efficiency loss)
- $\widehat{\sigma^2}_{\text{atu},2,\tilde{T}}$: estimator of $\sigma^2$ using $\tilde{T}$ periods of data of $\text{ATU}_2$

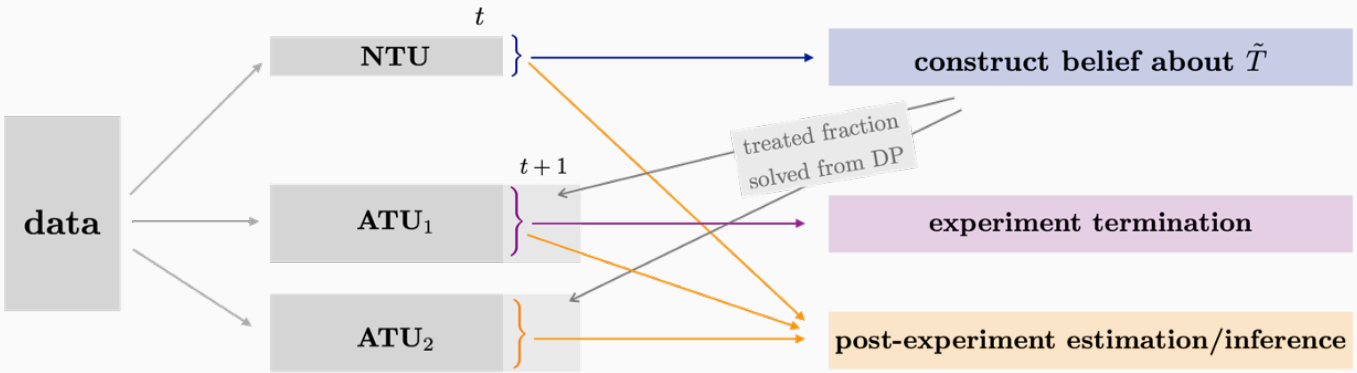

# Asymptotic properties of PGAE

**Theorem 2: Asymptotic distribution of estimators from PGAE**

Suppose $\varepsilon_{is}$ is bounded with a symmetric distribution around 0. As $N \to \infty$,

$$\sqrt{N} \cdot \begin{bmatrix} \left(\tilde{T} g_\tau(\boldsymbol{\omega}_{\text{all},1:\tilde{T}}, \tilde{T})/\sigma^2\right)^{1/2} \cdot \left(\hat{\tau}_{\text{all},\tilde{T}} - \tau_0\right) \\ \left(\tilde{T} p_{\text{atu},2}/\xi^{\dagger 2}_{\tilde{T}}\right)^{1/2} \cdot \left(\widehat{\sigma^2}_{\text{atu},2,\tilde{T}} - \sigma^2\right) \end{bmatrix} \xrightarrow{d} \mathcal{N}(\mathbf{0}, I_2) \,, \tag{3}$$

where $\xi^{\dagger}_{\tilde{T}} = \left[\xi^2 + \sigma^4/(\tilde{T} - 1)\right]^{1/2}$ and $\xi^2 = \mathbb{E}[(\varepsilon_{it}^2 - \sigma^2)^2]$.

- $\hat{\tau}_{\text{all},\tilde{T}}$ is consistent for $\tau$ with the optimal convergence rate $\sqrt{N}$
  - Intuition: Asymptotic conditional mean of $\varepsilon_{is}$ on estimated even moments of $\varepsilon_{is}$ is zero (due to the symmetric distribution of $\varepsilon_{is}$)
- $\widehat{\sigma^2}_{\text{atu},2,\tilde{T}}$ is consistent for $\sigma^2$
  - Intuition: A different sample is used to estimate $\widehat{\sigma^2}_{\text{atu},2,\tilde{T}}$

# Asymptotic properties of PGAE

**Theorem 2: Asymptotic distribution of estimators from PGAE**

Suppose $\varepsilon_{is}$ is bounded with a symmetric distribution around 0. As $N \to \infty$,

$$\sqrt{N} \cdot \begin{bmatrix} \left(\tilde{T} g_\tau(\boldsymbol{\omega}_{\text{all},1:\tilde{T}}, \tilde{T})/\sigma^2\right)^{1/2} \cdot \left(\hat{\tau}_{\text{all},\tilde{T}} - \tau_0\right) \\ \left(\tilde{T} p_{\text{atu},2}/\xi^{\dagger 2}_{\tilde{T}}\right)^{1/2} \cdot \left(\widehat{\sigma^2}_{\text{atu},2,\tilde{T}} - \sigma^2\right) \end{bmatrix} \xrightarrow{d} \mathcal{N}(\mathbf{0}, I_2) \,, \qquad (4)$$

where $\xi^{\dagger}_{\tilde{T}} = \left[\xi^2 + \sigma^4/(\tilde{T} - 1)\right]^{1/2}$ and $\xi^2 = \mathbb{E}[(\varepsilon_{it}^2 - \sigma^2)^2]$.

- The adaptivity of the design, with the termination time depending on early values of the outcomes, comes at no cost in the estimation of $\tau_0$
  - Compare with a series of experiments with the same distribution of termination times, the average variance of $\hat{\tau}_{\text{all},\tilde{T}}$ is the same
- Adaptive treatment decisions improve the estimation precision of $\tau_0$
  - $g_\tau(\boldsymbol{\omega}_{\text{all},1:\tilde{T}}, \tilde{T})$ is increased through adaptive treatment decisions

# Empirical application

# Empirical application

MarketScan medical claims databases

- Inpatient and outpatient claim records from early 2007 to mid 2017
- Primary diagnosis is influenza $21,277$ inpatient and $9,678,572$ outpatient admissions

Study effect of interventions (e.g., face cover, social distancing, and vaccine) on flu occurrence rate

- Aggregate at the Metropolitan Statistical Area (MSA) level and month
- Focus on the flu peak season (October to April)

Other applications (medical home visits, grocery expenditure, and Lending Club loans) are in the paper

# Comparison of non-adaptive designs

Benchmark designs

- $Z_{ff}$: 50% control and 50% treated for all time periods
- $Z_{ba}$: first half time periods all control, and second half all treated
- $Z_{ffba}$: first half time periods all control, and second half half treated

Non-adaptive staggered designs

- $Z_{opt}$: nonlinear staggered design with $\omega_s = \omega^*_{\ell,s}$
- $Z_{opt,linear}$: linear staggered design with $\omega_s = \omega^*_{0,s} = (2s-1)/(2T)$
- $Z_{opt,stratified}$: nonlinear staggered design with $\omega_s = \omega^*_{\ell,s}$ and historical data used for stratification

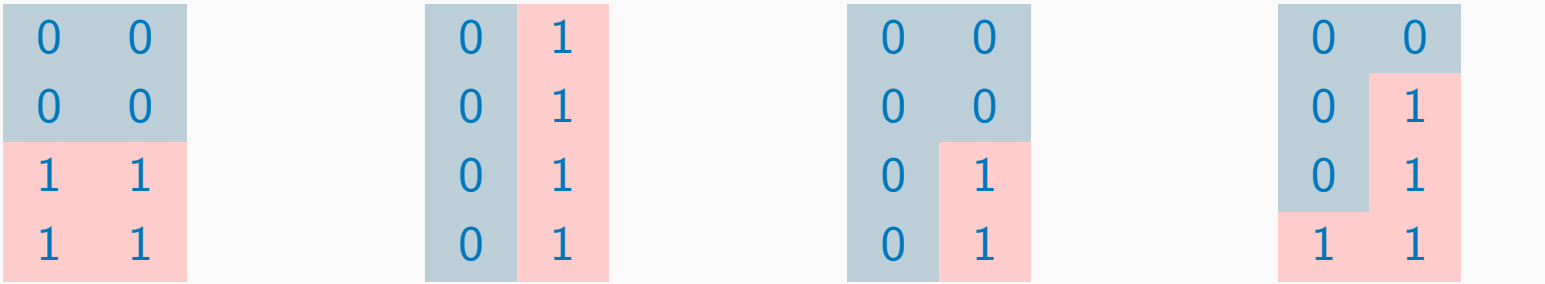


$Z_{ff}$ $Z_{ba}$ $Z_{ffba}$ staggered designs

# Synthetic non-adaptive experimental data and evaluation metrics

## Synthetic non-adaptive experimental data

- Assume the synthetic treatment was not applied to the original data, so $\text{observed value} = Y_{is}(\mathbf{0})$
- Apply a synthetic treatment using $Z$ and obtain synthetic experimental data

$$Y_{is} = Y_{is}(\mathbf{0}) + \tau_0 \cdot Z_{is} + \tau_1 \cdot Z_{i,s-1} + \tau_2 \cdot Z_{i,s-2}$$

## Evaluation metrics

- Estimate $\tau_0$, $\tau_1$ and $\tau_2$ from $Y_{it}$, and compare $\sum_j (\hat{\tau}_j - \tau_j)^2$ from the data generated by various $Z$
- Other evaluation metrics (estimation error of cumulative effects, recall and "precision") in the paper

# Results for synthetic non-adaptive experiments

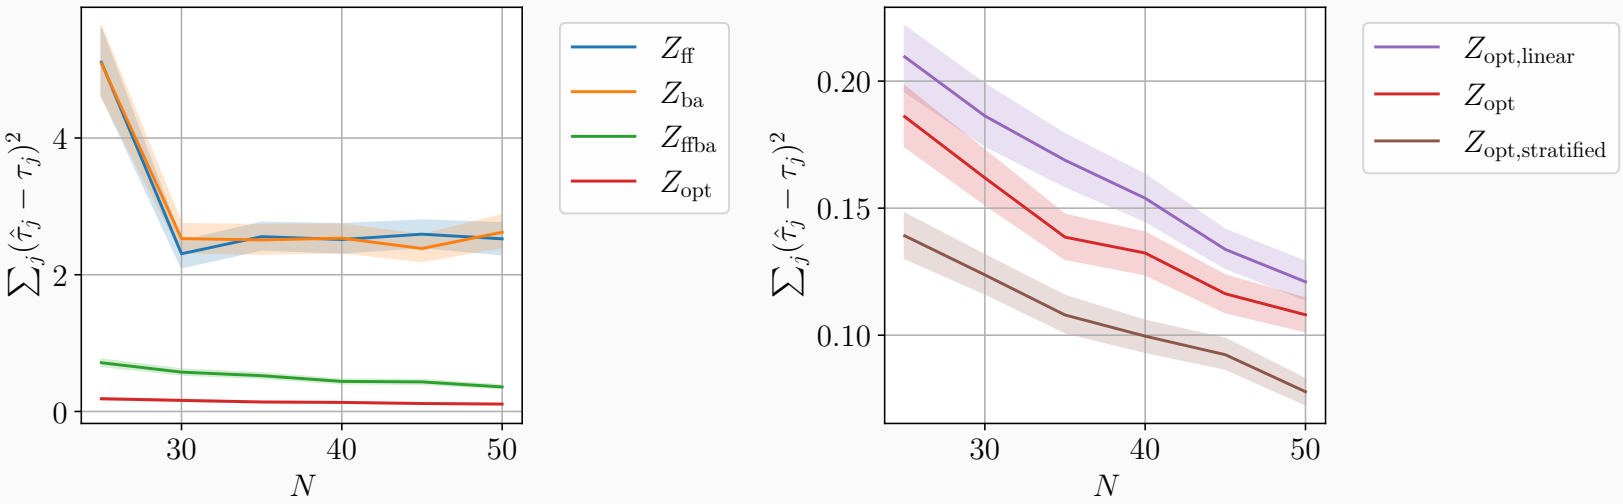


- $Z_{opt}$ requires fewer than 50% units to achieve the same estimation error as $Z_{ff}$, $Z_{ba}$, and $Z_{ffba}$
- $Z_{opt,stratified}$ further saves at least 20% units to achieve the same estimation error as $Z_{opt}$ and $Z_{opt,linear}$

$\Rightarrow$ Using our solution with historical data can substantially reduce the experimental cost

# Synthetic adaptive data

## Synthetic adaptive experimental data

- Run PGAE: The adaptive experiment is run for $\tilde{T}$ periods with precision threshold $c$
- Apply a synthetic treatment using $Z$ and obtain synthetic experimental data

$$Y_{is} = Y_{is}(\mathbf{0}) + \tau_0 \cdot Z_{is}$$

## Three designs

- $Z_{adaptive}$: design produced by PGAE with dimension $N \times \tilde{T}$
- $Z_{benchmark}$: design with $\omega_s = (2s-1)/(2T_{\max})$ with dimension $N \times \tilde{T}$ (optimal when $\tilde{T} = T_{\max}$)
- $Z_{oracle}$: design with $\omega_s = (2s-1)/(2\tilde{T})$ with dimension $N \times \tilde{T}$ (assuming $\tilde{T}$ is known ex-ante)

# Results for adaptive experiments

- Estimation error of the adaptive design always below variance threshold $1/c$
- Adaptive design $Z_{adaptive}$ reduces errors by 20% compared to benchmark design $Z_{benchmark}$

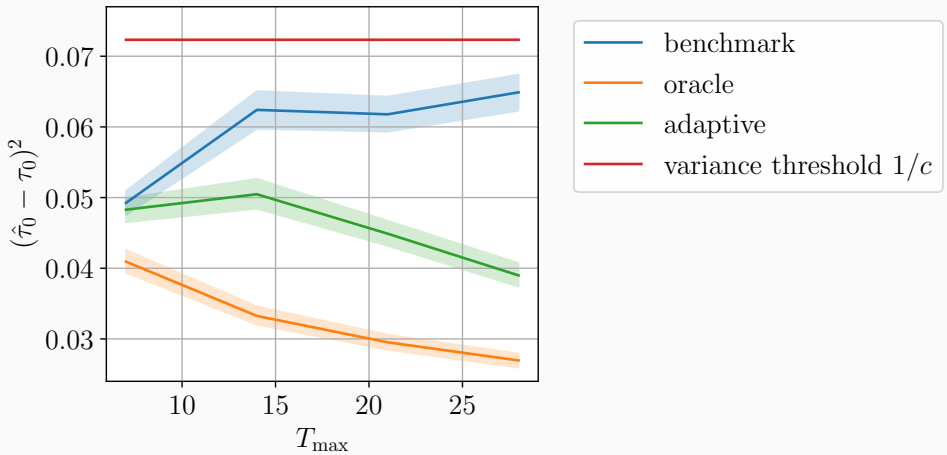

# Termination time in adaptive experiments

For $T_{\max} > 7$, the experiment is always terminated quite early

$\Rightarrow$ Desired precision threshold $c$ achieved with less than $T_{\max}/2$ duration

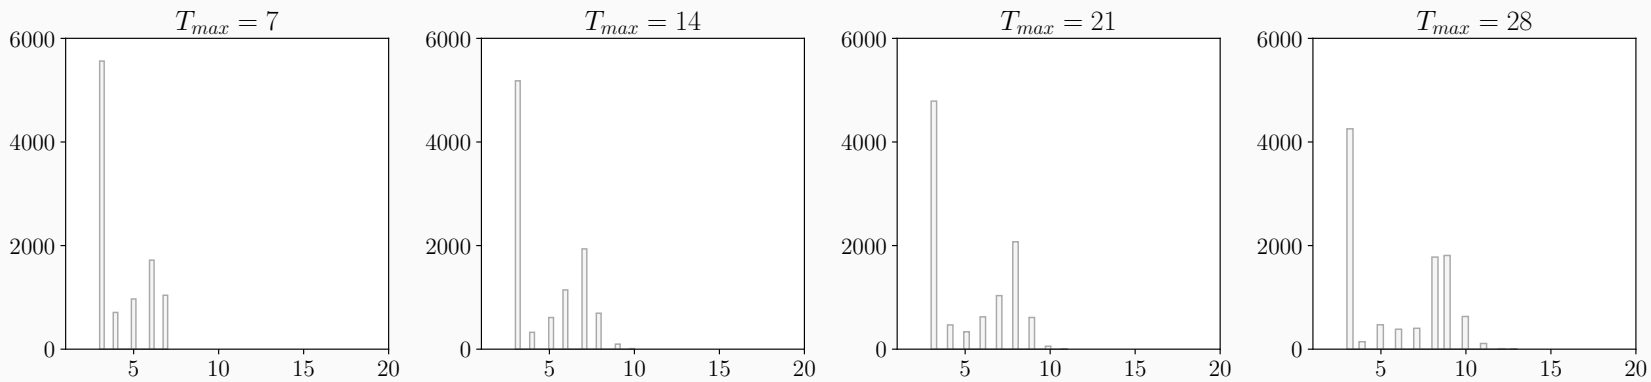

# Conclusion

# Conclusion

**Non-adaptive experiments**: $N$, $T$ and treatment decisions are determined, pre-experiment

- Analyze the statistical properties of GLS estimator of instantaneous and lagged effects from a general outcome specification
- Provide analytical optimality conditions that maximize a linear combination of precisions of estimated treatment effects
- Propose the treatment design that has two features: (1) treatment fraction takes an $S$-shaped curve in time; (2) stratification

**Adaptive experiments**: $N$ is fixed, and experiment duration and treatment decisions are determined during the experiment

- Propose the Precision-Guided Adaptive Experiment (PGAE) algorithm for adaptive treatment design and post-experiment inference
  - Combines ideas from dynamic programming and sample splitting
- Derive the asymptotic normal distribution of final treatment effect and variance estimates from PGAE
  - Final treatment effect estimate is efficient and achieves the optimal convergence rate

# Supplementary material

# D-optimal design

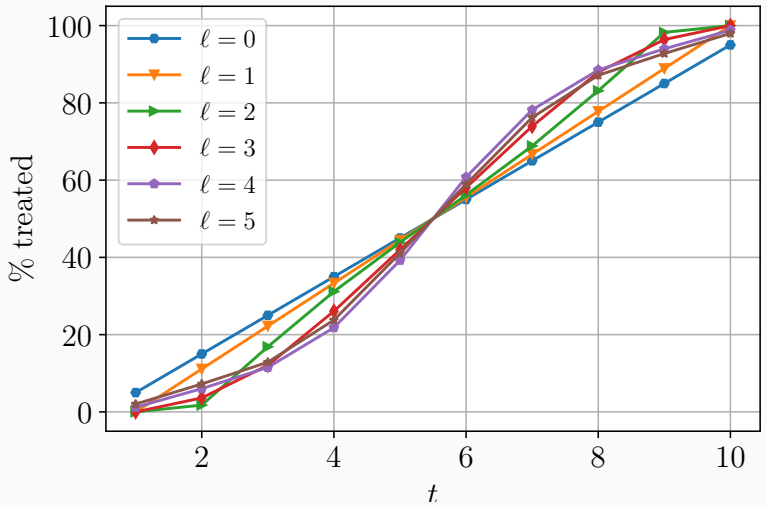


D-optimal treatment design: Optimal treated proportion $\omega_t$ at each period for a $T$-period treatment design and various $\ell$, where $T = 10$. Different colors represent different $\ell$.

# T-optimal design

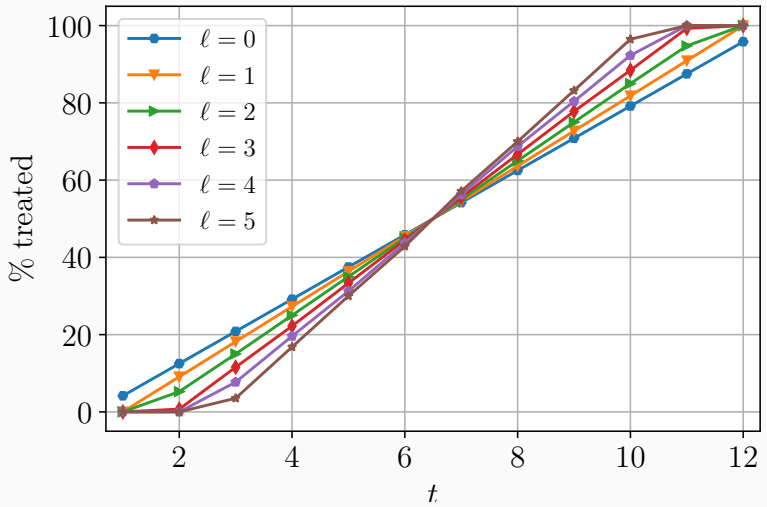


T-optimal treatment design: Optimal treated proportion $\omega_t$ at each period for a $T$-period treatment design and various $\ell$, where $T = 12$. Different colors represent different $\ell$.

## Expression of $\omega^*_{\ell,s}$

$$\omega^*_{\ell,s} = \begin{cases} 0 & s \le \lfloor \ell/2 \rfloor \\ a^{(\ell)}_{s-\lfloor \ell/2 \rfloor} & \lfloor \ell/2 \rfloor < s \le \ell \\ \big(2s - (\ell+1)\big)/(2(T-\ell)) & \ell < s \le T - \ell \\ 1 - \omega^*_{\ell,T+1-s} & T - \ell < s \le T - \lfloor \ell/2 \rfloor \\ 1 & T - \lfloor \ell/2 \rfloor < s \end{cases} \tag{5}$$

# Expression of $\omega^*_{\ell,s}$

$a^{(\ell)}$ is defined as

$$a^{(\ell)} = (1 + (M^{(\ell)})^{-1} b^{(\ell)})/2,$$

where $M^{(\ell)}$ and $b^{(\ell)}$ are defined as

$$M^{(\ell)} = \begin{bmatrix} \lfloor \ell/2 \rfloor + 1 & & & \\ & \lfloor \ell/2 \rfloor + 2 & & \\ & & \ddots & \\ & & & \ell \end{bmatrix} - \frac{1}{T-\ell} \begin{bmatrix} \ell - \lfloor \ell/2 \rfloor & \ell - 1 - \lfloor \ell/2 \rfloor & \ell - 2 - \lfloor \ell/2 \rfloor & \cdots & 1 \\ \ell - 1 - \lfloor \ell/2 \rfloor & \ell - 1 - \lfloor \ell/2 \rfloor & \ell - 2 - \lfloor \ell/2 \rfloor & \cdots & 1 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & 1 & 1 & \cdots & 1 \end{bmatrix}$$

$$b^{(\ell)} = -\begin{bmatrix} \lfloor \ell/2 \rfloor + 1 \\ \vdots \\ \ell - 1 \\ \ell \end{bmatrix} + \frac{1}{T-\ell} \begin{bmatrix} (\lfloor \ell/2 \rfloor + 1)^2 \\ \vdots \\ (\ell-1)^2 \\ \ell^2 \end{bmatrix} - \frac{1}{T-\ell} \begin{bmatrix} \sum_{l=1}^{\ell - \lfloor \ell/2 \rfloor} (\lfloor \ell/2 \rfloor + 1 - l) \\ \vdots \\ 2\lfloor \ell/2 \rfloor - 1 \\ \lfloor \ell/2 \rfloor \end{bmatrix}$$

## Examples of $\omega^*_{\ell,s}$

If $\ell = 1$, then

$$\omega^*_{\ell,s} = (s-1)/(T-1)$$

If $\ell = 2$, then

$$\omega^*_{\ell,1} = 0, \quad \omega^*_{\ell,2} = 1/(2T-5)$$
$$\omega^*_{\ell,s} = (2t-3)/2(T-2) \text{ for } t = 4, \cdots, T-3,$$
$$\omega^*_{\ell,T-1} = 1 - 1/(2T-5), \quad \omega^*_{\ell,T} = 1.$$

If $\ell = 3$, then

$$\omega^*_{\ell,1} = 0, \quad \omega^*_2 = \frac{3}{6T^2-44T+79}, \quad \omega^*_3 = \frac{6(T-4)}{6T^2-44T+79},$$
$$\omega^*_t = \frac{t-2}{T-3} \text{ for } t = 4, \cdots, T-3,$$
$$\omega^*_{T-2} = 1 - \frac{6(T-4)}{6T^2-44T+79}, \quad \omega^*_{T-1} = 1 - \frac{3}{6T^2-44T+79}, \quad \omega^*_T = 1.$$

# An algorithm to choose a treatment design

**Algorithm 1:** Choose a treatment design for each stratum $g$

**Inputs:** $|\mathcal{O}_g|$, $[\omega^*_{\ell,t}]_{t\in[T]}$
**for** $t = 1, \cdots, T$ **do**
  $N^{\text{int}}_{\text{treated},g,t} \leftarrow \lfloor |\mathcal{O}_g| \cdot \omega^*_{\ell,t} \rfloor$;
  $N^{\text{dec}}_{\text{treated},g,t} \leftarrow |\mathcal{O}_g| \cdot \omega^*_{\ell,t} - N^{\text{int}}_{\text{treated},g,t}$ ;
  **if** $N^{\text{dec}}_{\text{treated},g,t} < 0.5$ *or* $N^{\text{dec}}_{\text{treated},g,t} = 0.5$ *with* $t < T/2$ **then**
    $N_{g,t} \leftarrow N^{\text{int}}_{\text{treated},g,t}$ ;
  **else**
    $N_{g,t} \leftarrow N^{\text{int}}_{\text{treated},g,t} + 1$ ;
  **end**
**end**
$f(\cdot) \leftarrow$ a random function that shuffles $\{1, 2, \cdots, |\mathcal{O}_g|\}$;
$Z_g \leftarrow [0]^{|\mathcal{O}_g| \times T}$ ;
**for** $i = 1, \cdots, |\mathcal{O}_g|$ **do**
  **for** $t = 1, \cdots, T$ **do**
    **if** $f(i) \leq N_{g,t}$ **then**
      $z_{g,it} \leftarrow 1$ ;
    **else**
      $z_{g,it} = 0$ ;
    **end**
  **end**
**end**
**return** $Z_g$;

# Estimators in adaptive experiments

Three estimators are used in adaptive experiments

Suppose The estimators use the data of units in a set $\mathcal{S}$ over $t$ periods collected so far, where $t$ is small, but set size $|\mathcal{S}|$ can be large

1. within estimator for $\tau_0$

   - Regresses $\dot{Y}_{is}$ on $\dot{z}_{is}$ based on the specification $\dot{Y}_{is} = \tau_0 \dot{z}_{is} + \dot{\varepsilon}_{is}$, where for any variables $\{x_{is}\}_{(i,s)\in\mathcal{S}\times[t]}$ (*e.g.*, $Y_{is}$ and $z_{is}$), and $\dot{x}_{is}$ denotes the within transformed $x_{is}$

     $$\dot{x}_{is} = x_{is} - \bar{x}_{i\cdot} - \bar{x}_{\cdot s} + \bar{x}\,,$$

     in which $\bar{x}_{i\cdot}$, $\bar{x}_{\cdot s}$, and $\bar{x}$ are averages of $x_{is}$'s over $t$ time periods, units in $\mathcal{S}$, and both of them, respectively

## Estimators in adaptive experiments

2. Plug-in estimator for $\sigma^2$

$$\widehat{\sigma^2}_{\mathcal{S},t} = \frac{1}{|\mathcal{S}| \cdot (t-1)} \sum_{i \in \mathcal{S}} \sum_{s=1}^{t} \left(\dot{y}_{is} - \hat{\tau}_{\mathcal{S},t} \cdot \dot{z}_{is}\right)^2$$

- The factor $1/(t-1)$ is for finite $t$ correction
- $\widehat{\sigma^2}_{\mathcal{S},t}$ is consistent and asymptotically normal for any finite $t$

3. A new estimator for $\xi^2 = \mathbb{E}[(\varepsilon_{is}^2 - \sigma^2)^2]$

$$\widehat{\xi^2}_{\mathcal{S},t} = \underbrace{\frac{t^2}{(t-1)^2}}_{\substack{\text{correction} \\ \text{multiplier}}} \cdot \underbrace{\frac{1}{|\mathcal{S}| \cdot t} \sum_{i \in \mathcal{S}} \left( \sum_{s=1}^{t} \left[ (\dot{y}_{is} - \hat{\tau}_{\mathcal{S},t} \cdot \dot{z}_{is})^2 - \widehat{\sigma^2}_{\mathcal{S},t} \right] \right)^2}_{\text{plug-in estimator of } \xi^2} - \underbrace{\frac{3t-2}{(t-1)^2} \cdot (\widehat{\sigma^2}_{\mathcal{S},t})^2}_{\text{correction term}}$$

- $\widehat{\xi^2}_{\mathcal{S},t}$ is consistent for any finite $t$

# Asymptotic distribution for non-adaptive data

**Lemma: Asymptotic distribution of estimators from non-adaptive data**

Suppose $\varepsilon_{is}$ is i.i.d. for any $i$ and $s$ with $\mathbb{E}[\varepsilon_{is}] = 0$, $\mathbb{E}[\varepsilon_{is}^2] = \sigma_\varepsilon^2$, $\mathbb{E}[\varepsilon_{is}^3] = 0$, and $\mathbb{E}[(\varepsilon_{is}^2 - \sigma^2)^2] = \xi^2$. $\hat{\tau}_{\text{ntu},t}$ and $\widehat{\sigma^2}_{\text{ntu},t}$ are consistent. As $|\mathcal{S}_{\text{ntu}}| \to \infty$, for any finite $t$, conditional on $Z_{\text{ntu}}$, we have

$$\sqrt{|\mathcal{S}_{\text{ntu}}|}\left(\begin{bmatrix}\hat{\tau}_{\text{ntu},t}\\ \widehat{\sigma^2}_{\text{ntu},t}\end{bmatrix} - \begin{bmatrix}\tau\\ \sigma^2\end{bmatrix}\right) \xrightarrow{d} \mathcal{N}\left(\begin{bmatrix}0\\0\end{bmatrix}, \begin{bmatrix}\sigma^2/(t \cdot g_\tau(\boldsymbol{\omega}_{\text{ntu},1:t}, t)) & 0\\ 0 & \xi_t^{\dagger 2}/t\end{bmatrix}\right),$$

where $\xi_t^{\dagger 2} = \xi^2 + 2\left(\sigma^2\right)^2/(t-1)$.

Furthermore, $\sqrt{|\mathcal{S}_{\text{ntu}}|}\left(\widehat{\xi}_t^2 - \xi^2\right) = O_p(1)$.

$\Rightarrow$ This lemma is used to prove Theorem 2

# Finite sample properties of Theorem 2

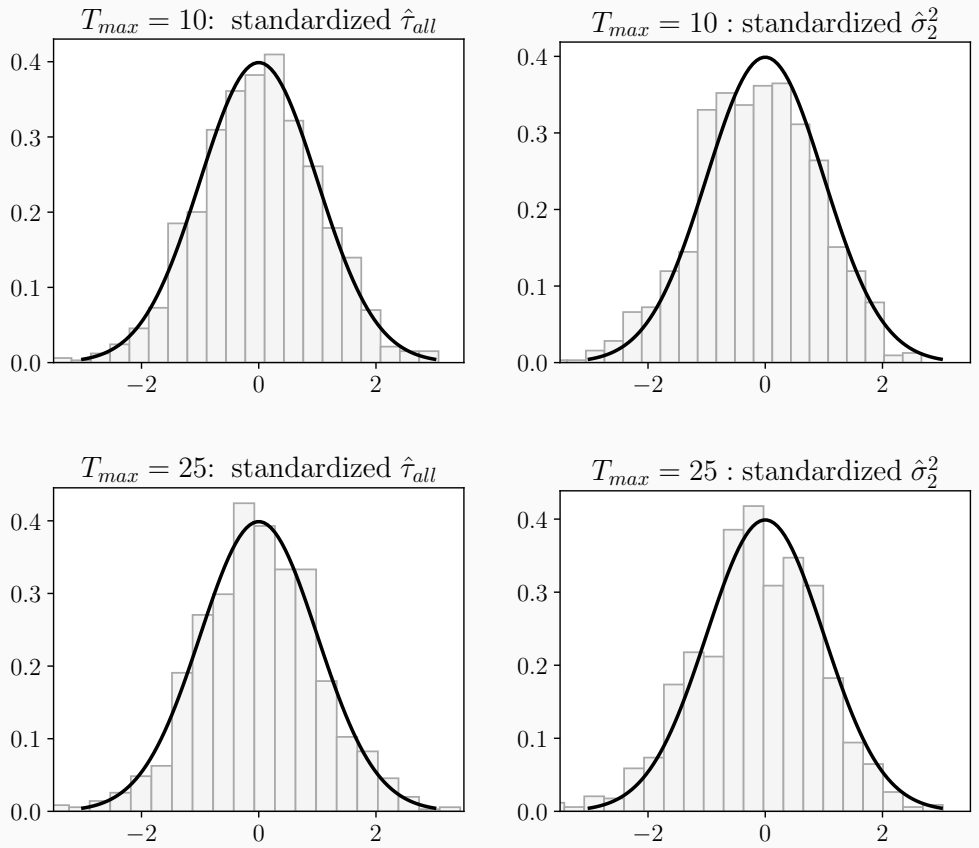


Finite sample properties of Theorem 2: Histograms of $\hat{\tau}_{\text{all,ss}}$ and $\widehat{\sigma^2}_{\text{atu,2,ss}}$. The standard normal density function is superimposed on the histograms. $N = 500$, $\tau_0 = 1$, and $\sigma_\varepsilon = 1$.

## Within estimator for $\tau_0$

Least-squares estimator of $\tau_0$ from the specification

$$Y_{is} = \alpha_i + \beta_s + \tau_0 z_{is} + \varepsilon_{is}$$

is equivalent to the within estimator that regresses $\dot{Y}_{is}$ on $\dot{z}_{is}$ based on the specification

$$\dot{Y}_{is} = \tau \dot{z}_{is} + \dot{\varepsilon}_{is},$$

where for any variables $\{x_{is}\}_{(i,s)\in\mathcal{S}\times[t]}$ (*e.g.*, $Y_{is}$ and $z_{is}$), and $\dot{x}_{is}$ denotes the within transformed $x_{is}$

$$\dot{x}_{is} = x_{is} - \bar{x}_{i\cdot} - \bar{x}_{\cdot s} + \bar{x}\,,$$

in which $\bar{x}_{i\cdot}$, $\bar{x}_{\cdot s}$, and $\bar{x}$ are averages of $x_{is}$'s over $t$ time periods, units in $\mathcal{S}$, and both of them, respectively

## Proof of Theorem 2: Key challenge

The estimation error of $\hat{\tau}_{\mathrm{all},t}(N)$ depends on $\varepsilon_{is}$ (using data of $N$ units and $t$ periods)

$$\hat{\tau}_{\mathrm{all},\tilde{T}}(N) - \tau = \left( \sum_{i\in[N],s\le\tilde{T}} \dot{z}_{is}^2 \right)^{-1} \sum_{i\in[N],s\le\tilde{T}} \dot{z}_{is}\varepsilon_{is}.$$

The estimation error of the plug-in estimator for $\sigma^2$ also depends on $\varepsilon_{is}$

$$\begin{aligned}
\widehat{\sigma^2}_{\mathcal{S},t}(N) &= \frac{1}{|\mathcal{S}|\cdot(t-1)} \sum_{i\in\mathcal{S}} \sum_{s=1}^{t} \left(\dot{y}_{is} - \hat{\tau}_{\mathcal{S},t}\cdot\dot{z}_{is}\right)^2 \\
&= \frac{1}{|\mathcal{S}|(t-1)} \sum_{i,s} \varepsilon_{is}^2 - \frac{t}{|\mathcal{S}|(t-1)} \sum_{i} \bar{\varepsilon}_{i,\cdot}^2 - \frac{1}{t-1} \sum_{s} \bar{\varepsilon}_{\cdot,s}^2 + \frac{t}{t-1}\bar{\varepsilon}^2 \\
&\quad - \left(\hat{\tau}_{\mathcal{S},t}(N) - \tau\right)^2 \cdot \frac{1}{|\mathcal{S}|(t-1)} \sum_{i,s} \dot{z}_{is}^2
\end{aligned}$$

- Key challenge: We need to show $\hat{\tau}_{\mathrm{all},\tilde{T}}(N)$ is "well-behaved" even if we condition on $\widehat{\sigma^2}_{\mathcal{S},t}(N)$ that is used to make adaptive treatment decisions ($\mathcal{S} = \mathrm{NTU}$) and experiment termination ($\mathcal{S} = \mathrm{ATU}_2$)

# Proof of Theorem 2: Two key properties

We leverage two critical properties

- First property: Given that $\varepsilon_{is}$ has a symmetric distribution, $\mathbb{E}[\varepsilon_{is} \mid \widehat{\sigma^2}_{\mathcal{S},t}(N)] = 0$
  - $\Rightarrow$ The asymptotic mean of $\hat{\tau}_{\text{all},\tilde{T}}(N) - \tau_0$ is zero
- Second property: Given that $\widehat{\sigma^2}_{\mathcal{S},t}(N)$ is consistent, $\mathbb{E}[\varepsilon_{is}^2 - \sigma^2 \mid \widehat{\sigma^2}_{\mathcal{S},t}(N)] = \widehat{\sigma^2}_{\mathcal{S},t}(N) - \sigma^2$ converges to zero in probability
  - $\Rightarrow$ The asymptotic variance of $\left(\tilde{T} g_\tau(\boldsymbol{\omega}_{\text{all},1:\tilde{T}}, \tilde{T})/\sigma^2\right)^{1/2} \cdot \left(\hat{\tau}_{\text{all},\tilde{T}} - \tau_0\right)$ is 1 (with probability approaching one, the variance is sufficiently close one)