
Staggered Rollout Designs Enable Causal Inference Under Interference Without Network Knowledge

Mayleen Cortez-Rodriguez
Center for Applied Math
Cornell University
Ithaca, NY 14850
mec383@cornell.edu

Matthew Eichhorn
Center for Applied Math
Cornell University
Ithaca, NY 14850
meichhorn@cornell.edu

Christina Lee Yu
Operations Research and Information Engineering
Cornell University
Ithaca, NY 14850
cleeyu@cornell.edu

Note from authors: This version contains the appendices.

1 Introduction

A cornerstone of much of the classic causal inference literature is the stable unit treatment value assumption (SUTVA), which posits that an individual’s potential outcome is a function only of their assigned treatment; there are no spillover effects due to the treatment of others. Such an assumption fails to account for the ways in which individuals interact in many real-world experimental settings. For instance, new features rolled out on social networking sites such as LinkedIn may alter these users’ behaviors, which in turn affect how their connections (who do not have access to the feature) interact with the platform. Individuals receiving a vaccine against an infectious disease may reduce the transmission probability of the disease to others they interact with. Implementing a different pricing policy for a subset of individuals in an online marketplace such as Airbnb or a platform such as Uber could impact the experience of other users, as they compete for the same resources or same customers. Public health measures instituted in one city can limit travel to nearby communities, indirectly affecting their health outcomes or transit related outcomes. These examples illustrate how network interference may arise naturally from the connectedness of our society. Unfortunately, the standard causal inference techniques which do not account for network interference may result in arbitrarily biased estimates.

As these issues come into greater focus, there is a growing research area in developing new tools for causal inference under *network interference*, when the outcome of an individual can be affected by the treatment of another. Many approaches either propose complex graph-based cluster randomized designs, or require strong parametric assumptions on the network interference effects. A limitation is that all these approaches require at least partial knowledge of the underlying network in order to implement the randomization or to compute the estimator. While structural knowledge is available to online social networks, other applications such as public health must reason about an unknown or potentially transient network. The additional effort required to collect or model network structure is both difficult and costly.

In this work, we explore the value of additional measurements which arise from a staggered rollout randomized design, in which the treatments are administered over a span of a few timepoints. For example, the experimentation team at LinkedIn may roll out an experiment over 5 days, increasing the fraction of treated individuals according to a schedule of 1%, 2%, 5%, 10%, 20%, where it continuously collects data and measurements before and during each day of the experiment. Not only

is this type of experiment easy to implement in such applications, it is often desirable to implement treatments according to such a staggered rollout design as it allows the system to first ensure safety of the proposed treatment on a smaller test group before implementing it on larger groups. This type of experimental design is also common for trials involving healthcare and medicine due to the requirement of certain safety considerations before testing for efficacy. A key contribution in our work is that we show the additional measurements from a *staggered rollout* design enable graph agnostic causal inference, lifting all requirements on knowledge of the network.

We focus on estimating the *total treatment effect* (TTE), informally defined as the difference in average outcomes across the population between two scenarios: when all individuals are treated and when no individuals are treated. It has also been referred to as the *global average treatment effect* (GATE). The TTE is particularly pertinent to applications where the decision maker must choose between entirely adopting the new treatment or remaining with the status quo. For example, LinkedIn would like to choose a single news feed recommendation algorithm, and Airbnb and Uber would like to choose a single dynamic pricing algorithm. We assume *neighborhood interference*, where each individual is only affected by the treatments of its direct neighbors; this is only mildly restrictive as the neighborhood can be defined with respect to an unknown network, which is neither used for the estimator nor the randomized design.

Related Work. In addressing the challenges that arise from network interference, a key tension arises between the model assumptions and the simplicity and efficiency of the proposed estimator. Previously proposed model assumptions can be generally classified into assumptions on exposure functions [1, 2, 12, 14, 24], interference neighborhoods [3, 5, 18, 22], parametric structure [4, 6, 8, 9, 20], or a combination of these. Each of these assumptions lead to different solution concepts. All of these approaches rely on knowledge of the network mediating the interference effects.

One class of approaches relies on assumptions about the network structure. They assume *partial interference*, meaning that the population can be partitioned into disjoint groups, such that all network interference effects can only occur within but not across the pre-specified groups [2, 5, 10, 13, 16, 17, 19, 23]. This assumption is motivated by scenarios where the network is naturally strongly clustered. A natural solution is to randomize treatments over the groups jointly, such that each group is assigned to be either fully treated or fully control. A drawback of this approach is that many networks are well-connected such that there is no clear clustering of the network which does not cut a significant fraction of the edges. The bias of standard estimators will scale with the number of edges that cross between groups, leading to proposed cluster randomized designs that randomize over clusters that are constructed to minimize the number of edges between clusters [8, 9]. Constructing good clusters itself can be computationally intensive. Additionally some applications may prohibit such nonuniform treatment assignment probabilities due to fairness considerations. Under *neighborhood interference* assumptions, [21, 22] analyze the Horvitz-Thompson estimator alongside a cluster randomized design, which involves both clustering the graph and computing probabilities of entire neighborhoods being assigned to treatment or control over the distribution of clusterings, which is computationally intensive. When one is willing to impose a distributional model on the network itself, [11] provides central limit theorem convergence results for a related but weaker estimand measuring the change in outcomes under small perturbations of the fraction of treated individuals.

An alternate approach is to impose structure on the form of the network interference effects. The most common assumption is that the network effects are linear with respect to a specified statistic of the local neighborhood [4, 6, 7, 9, 15, 20]. The assumptions reduce the number of unknown parameters in the model to a fixed dimension that does not grow with the population size, reducing the inference task to linear regression. As a result, the natural solution is to use a least squares estimate, shifting the focus to constructing randomized designs that minimize the variance of the estimate. A limitation of this approach is that it requires the correct choice of the the statistic governing the linearity, and it requires precise knowledge of the network structure to compute these neighborhood statistics. Furthermore, it assumes knowledge of the relevant covariate types that differentiate individual responses, or otherwise assumes homogeneity in the network effects.

The most similar work to our paper is the solution proposed in [25], which provides an estimator for the TTE under a heterogeneous linear interference model [8], also referred to as the joint assumptions of additivity of main effects and interference effects in [18]. Their estimator does not require knowledge of the network, but requires measurements over two time steps. Our work generalizes their results beyond linear to polynomial models, and we show that the staggered rollout experimental

design enables graph agnostic causal inference. The extension from linear to polynomial models introduces the possibility of non-trivial interactions within the neighborhood set, adding complexity to the model that necessitates a new analysis and algorithm.

Contributions. We show that under a staggered rollout experimental design, the task of estimating the total treatment effect reduces to polynomial extrapolation, where the degree of the polynomial is governed by the cardinality of interactions in the neighborhood interference model, bounded above by the degree of the graph. Our approach is the first in the literature to propose an estimator and randomized design that does not require any knowledge of the network structure, and yet is unbiased and consistent. We provide variance bounds on the estimator, showing that the variance only grows polynomially in the degree as opposed to the exponential growth that is exhibited in the Horvitz-Thompson estimator under simple Bernoulli randomized designs. We provide experiments that also illustrate that naively using regression models without allowing for heterogeneity could lead to significant bias, whereas our estimator is unbiased with significantly lower variance than the bias incurred due to a misspecified model. We are also the first to study the value of a staggered rollout experimental design in the presence of network interference, and we believe the overall framework could extend beyond polynomial models to other function classes, opening a new approach for handling network interference while allowing for flexible heterogeneity in the network effects.

2 Setup

Consider a population of n individuals, and assume that the network interference can be represented via an unknown directed graph with edge set $E \subset [n] \times [n]$. An edge $(j, i) \in E$ represents that individual i is affected by the treatment assignment of individual j ; as such, self-loops are expected. The in-neighborhood of individual i is denoted by $\mathcal{N}_i = \{j \in [n] : (j, i) \in E\}$, and we let d_{in} denote the maximum in-degree, d_{out} the maximum out-degree, and $d = \max\{d_{\text{in}}, d_{\text{out}}\}$. We posit that the outcome of individual i as a function of the entire population's exposure to treatment can be expressed by the potential outcomes function $Y_i : \{0, 1\}^n \rightarrow \mathbb{R}$.

Our task is to estimate the total treatment effect (TTE), which represents the difference in average outcomes when the entire population is fully under treatment as opposed to fully under control, denoted as

$$\text{TTE} := \frac{1}{n} \sum_{i=1}^n (Y_i(\mathbf{1}) - Y_i(\mathbf{0})). \quad (2.1)$$

We use $\mathbf{z} \in \{0, 1\}^n$ to denote the treatment assignment vector, where $z_i = 1$ if individual i is assigned to treatment, and $z_i = 0$ if i is assigned to control. By definition of E , it follows that the potential outcomes functions satisfy neighborhood interference with respect to the graph defined by E .

Assumption 1 (Neighborhood Interference). *$Y_i(\mathbf{z})$ only depends on the treatment of individuals in \mathcal{N}_i (including i). Equivalently, $Y_i(\mathbf{z}) = Y_i(\mathbf{z}')$ for any \mathbf{z} and \mathbf{z}' such that $z_j = z'_j$ for all $j \in \mathcal{N}_i$.*

Additionally, as the treatment variables z_i are binary, any potential outcomes function satisfying neighborhood interference can be written as a polynomial in the neighborhood treatment variables:

$$Y_i(\mathbf{z}) = \sum_{S \subseteq \mathcal{N}_i} a_S \prod_{j \in S} z_j \prod_{j' \in \mathcal{N}_i \setminus S} (1 - z_{j'}),$$

for some coefficients $\{a_S\}$. We use the degree of the polynomial to quantify the complexity of the model. In full generality, any model satisfying the neighborhood interference assumption will have polynomial degree bounded by $\max_i |\mathcal{N}_i|$, the maximum in-degree of the graph. In this work we consider the scenario where the polynomial degree may be significantly smaller than $\max_i |\mathcal{N}_i|$.

Assumption 2 (Low Polynomial Degree). *The potential outcomes model has polynomial degree at most β , i.e. there exist coefficients $\{c_{i,S}\}_{i \in [n], S \subseteq [n]}$ such that for all i and \mathbf{z} ,*

$$Y_i(\mathbf{z}) = \sum_{S \subseteq \mathcal{N}_i, |S| \leq \beta} c_{i,S} \cdot \mathbb{I}(S \text{ treated}) = \sum_{S \subseteq \mathcal{N}_i, |S| \leq \beta} c_{i,S} \prod_{j \in S} z_j. \quad (2.2)$$

The low-degree polynomial structure is perhaps better conceptualized as a constraint on the order of interactions among neighbors of an individual, as the potential outcomes function is polynomial with respect to the *binary* treatment vector. For general β , this assumption is not restrictive at all;

rather, a restriction is imposed by assuming a specific value for β , or more generally assuming that β is much smaller than the graph degree. We interpret the parameter $c_{i,S}$ as the effect that treating all individuals in S has on the outcome of individual i . The coefficient $c_{i,\emptyset}$ represents individual i 's outcome when everyone is assigned to control (i.e. their *baseline outcome*); this is unaffected by the treatment assignment. In the case of a singleton set $S = \{j\}$, we use the shorthand notation $c_{ij} := c_{i,\{j\}}$. It follows that the total treatment effect is the sum of all $c_{i,S}$ for nonempty subsets S , i.e. $\text{TTE} = \frac{1}{n} \sum_{i=1}^n \sum_{\substack{S \subseteq \mathcal{N}_i \\ 1 \leq |S| \leq \beta}} c_{i,S}$.

The number of unknown parameters in this model are $\sum_{i \in [n]} \sum_{k=0}^{\beta} \binom{|\mathcal{N}_i|}{k}$, which scales as nd^{β} . When $\beta = 1$, the network effects resulting from treated neighbors is additive, and is also equivalent to the heterogeneous linear outcomes model in [25]. This low degree assumption will not generally admit threshold models or saturation models, both of which would require the degree of Y_i to be $|\mathcal{N}_i|$.

An example in which the polynomial degree may be smaller than the neighborhood size would be a setting in which an individual's neighborhood can be further partitioned into smaller subcommunities: colleagues, university friends, high school friends, family, etc. Each subcommunity could have an additive affect on the individuals' outcome, but there may be nontrivial interactions among the treatments of individuals in the subcommunities. The polynomial degree would be bounded by the size of the largest subcommunity, which could be significantly smaller than the full neighborhood. As another example, suppose that a social platform is testing a "hangout room" feature that provides groups of up to 5 people a new environment to engage on the platform. One could posit a model with $\beta = 5$, as the change in any individual's usage on the platform can be attributed to experience with the new feature, which takes place in groups of up to 5 users.

We let Y_{\max} denote an upper bound on the absolute treatment effects for each individual, i.e.

$$Y_{\max} := \max_{i \in [n]} \sum_{S \subseteq \mathcal{N}_i, |S| \leq \beta} |c_{i,S}|.$$

It follows that the magnitude of the outcomes $Y_i(\mathbf{z})$ are bounded by Y_{\max} for any treatment vector \mathbf{z} . We let L_j denote the absolute effect or influence that individual j has on the population outcomes,

$$L_j := \sum_{i: j \in \mathcal{N}_i} \sum_{S \subseteq \mathcal{N}_i, |S| \leq \beta, j \in S} |c_{i,S}|.$$

Our boundedness assumption and the finiteness of our network imply the boundedness of the L_j . We denote the upper bound on the absolute effect or influence of any individual by $L_{\max} := \max_j \{L_j\}$.

Randomized Experiment Design. As it may be costly and/or detrimental to expose the entire population to treatment, we wish to estimate the total treatment effect after treating only a small random subset of individuals. In particular we assume that there may be an experimental budget that limits the proportion of individuals who may be treated. We will focus on two standard randomized designs. In Bernoulli design, a treatment vector \mathbf{z} is obtained by independently sampling each coordinate from a Bernoulli(p) distribution, so that the probability that a subset of individuals S are all treated is $p^{|S|}$. We assume that $p > \frac{1}{n}$ so that at least one individual is treated in expectation. In completely randomized design, a treatment vector \mathbf{z} is obtained by uniformly sampling a subset of k individuals to treat for some fixed k . Here, the probability that a subset of individuals S are all treated is

$$\prod_{i=0}^{|S|-1} \frac{k-i}{n-i} =: \left[\frac{k}{n} \right]^{|S|}. \quad (2.3)$$

Throughout the paper, we utilize a *staggered rollout* experimental design. Treatment is assigned to individuals in T stages throughout the experiment. Overall, the individuals' outcomes are measured $T + 1$ times: a baseline measurement before treatment, as well as a measurement after each treatment round. We'll use \mathbf{z}^t to denote the vector of treatment assignment in round t , and assume that each entry z_i^t is monotone increasing with t (individuals cannot be un-treated). This monotonicity requirement introduces significant correlation between the treatment vectors. Monotonicity is a constraint in many real-world scenarios where the experimental designer only has the option to introduce treatment to new individuals. For example, treatments in medication trials can have life-altering, irreversible effects, and the exposure of individuals to an advertising campaign cannot be "taken back." In other domains, such as the rollout of new interfaces on social media platforms, treatments are temporary or reversible and it may make sense to remove the monotonicity requirement.

Another assumption we make is that observations of outcomes are perturbed by iid Gaussian noise.

Assumption 3. We observe $Y_{i,t}^{\text{obs}} = Y_i(\mathbf{z}^t) + \varepsilon_{i,t}$ for $\varepsilon_{i,t} \stackrel{iid}{\sim} N(0, \sigma^2)$.

3 Graph Agnostic Estimators under Staggered Rollout Design

To motivate the design of our estimators, we begin with a high-level view of estimating the total treatment effect. We limit our attention to static networks. When we have no information about the underlying causal network, we do not know how much of each individual's neighborhood is treated, so have no systematic way to predict what their potential outcome would be if the entire population were treated. However, we can aggregate the average of the individuals' outcomes to obtain a meaningful statistic. In the following discussion we omit observation noise to make the intuition for our estimator clear. Consider the expected population average outcomes where the expectation is taken over the distribution of treatment vectors \mathbf{z} sampled from a parameterized class of distributions \mathcal{D}_x , where \mathcal{D}_0 refers to the distribution that deterministically assigns all individuals to control, and \mathcal{D}_1 refers to the distribution that deterministically assigns all individuals to treatment. Consider the underlying expected outcome function $F_{\mathcal{D}} : [0, 1] \rightarrow \mathbb{R}$ given by

$$F_{\mathcal{D}}(x) = \mathbb{E} \left[\frac{1}{n} \sum_{i=1}^n Y_i(\mathbf{z}) \right]$$

where the expectation is taken over the distribution of treatment vectors $\mathbf{z} \sim \mathcal{D}_x$. By construction, the TTE is exactly $F_{\mathcal{D}}(1) - F_{\mathcal{D}}(0)$.

If we can implement a staggered rollout design where at stage t of the experiment, the marginal distribution of the treatment vector is \mathcal{D}_{x_t} , the observed average outcomes collected in the experiment at stage t would give noisy estimates of $F_{\mathcal{D}}(x_t)$. Under this framing, our goal is to use these measurements to extrapolate the value of $F_{\mathcal{D}}(1)$. This provides a general framework for utilizing staggered rollout design to simplify estimation of the total treatment effect.

The simplest class of distributions we can consider is the Bernoulli(p) randomized design, in which each individual is independently assigned to treatment or control with probability p . For a degree- β polynomial potential outcomes model, the expected outcome function under this design is polynomial in the treatment probability p :

$$F_B(p) = \mathbb{E} \left[\frac{1}{n} \sum_{i=1}^n Y_i(\mathbf{z}) \right] = \frac{1}{n} \sum_{i=1}^n \sum_{\substack{\mathcal{S} \subseteq \mathcal{N}_i \\ |\mathcal{S}| \leq \beta}} c_{i,\mathcal{S}} \cdot \mathbb{E} \left[\prod_{j \in \mathcal{S}} z_j \right] = \frac{1}{n} \sum_{i=1}^n \sum_{\substack{\mathcal{S} \subseteq \mathcal{N}_i \\ |\mathcal{S}| \leq \beta}} c_{i,\mathcal{S}} \cdot p^{|\mathcal{S}|}.$$

To implement a staggered rollout Bernoulli design with treatment probabilities $p_1 < p_2 < \dots < p_T$ we independently sample $u_i \sim U[0, 1]$ for each individual i . Then, for each $t \in [T]$, we define treatment vector \mathbf{z}^t with $\mathbf{z}_i^t = \mathbb{I}(u_i \leq p_t)$. This both ensures that the marginal distribution of the treatment vector at stage t is equivalent to the Bernoulli(p_t) randomized design, and that the treatment assignments are monotone over the rounds.

Alternatively, we can consider a completely randomized design (CRD) in which we fix a number of treated individuals k , and sample a subset of k individuals uniformly at random among all size k subsets in the population. For a degree- β polynomial potential outcomes model, the expected outcome function under this design is polynomial in the treated fraction k/n :

$$F_C\left(\frac{k}{n}\right) = \mathbb{E} \left[\frac{1}{n} \sum_{i=1}^n Y_i(\mathbf{z}) \right] = \frac{1}{n} \sum_{i=1}^n \sum_{\substack{\mathcal{S} \subseteq \mathcal{N}_i \\ |\mathcal{S}| \leq \beta}} c_{i,\mathcal{S}} \cdot \mathbb{E} \left[\prod_{j \in \mathcal{S}} z_j \right] = \frac{1}{n} \sum_{i=1}^n \sum_{\substack{\mathcal{S} \subseteq \mathcal{N}_i \\ |\mathcal{S}| \leq \beta}} c_{i,\mathcal{S}} \cdot \left[\frac{k}{n} \right]^{|\mathcal{S}|}.$$

To implement a complete staggered rollout design, we sample a treatment vector from $\text{CRD}(k_1)$ at stage 1, and at stage $t > 1$, we sample a treatment vector from $\text{CRD}(k_t - k_{t-1})$ out of the remaining untreated individuals. The marginal distribution of the treatment vector at state t will be equivalent to the completely randomized design with parameter k_t .

To construct our estimators, we will make use of the Lagrange interpolation formula.

Definition 1 (Lagrange Interpolation). Given a dataset $\{(x_t, y_t)\}_{t=0}^T$ with distinct x -coordinates, the unique polynomial F of degree at most T with $F(x_t) = y_t$ for each t is given by

$$F(x) = \sum_{t=0}^T \ell_{t,\mathbf{x}}(x) \cdot y_t, \quad \ell_{t,\mathbf{x}}(x) = \prod_{\substack{s=0 \\ s \neq t}}^T \frac{x - x_s}{x_t - x_s}.$$

To estimate TTE, we require estimates of $F_B(x)$ or $F_C(x)$ at $x \in [0, 1]$. As both F_B and F_C have degree at most β , they can be recovered from $\beta + 1$ observations, requiring $T = \beta$ rounds of treatment. Given treatment targets $\mathbf{x} = (x_0, x_1, \dots, x_T)$ with realized treatment schedule $\{\mathbf{z}_t \sim \mathcal{D}_{x_t}\}$, we can utilize Lagrange interpolation to derive the following polynomial interpolation (PI) estimator:

$$\widehat{\text{TTE}}_{\text{PI}}(\mathbf{x}) := \begin{cases} \sum_{t=0}^T \left(\ell_{t,\mathbf{x}}(1) - \ell_{t,\mathbf{x}}(0) \right) \left(\frac{1}{n} \sum_{i=1}^n Y_{i,t}^{\text{obs}} \right) & x_0 < x_1 < \dots < x_T, \\ 0 & x_t = x_{t-1} \text{ for some } t \in [T]. \end{cases} \quad (3.1)$$

The separation into cases ensures that the Lagrange coefficients are well-defined. We assume that the degree β is known such that the experimenter can select $T = \beta$. We also assume that \mathbf{x} is monotone, and define $\Delta_{\mathbf{x}} = \min_{t=1 \dots m} \{x_t - x_{t-1}\}$. We can apply this estimator in both the Bernoulli and completely randomized design settings.

3.1 Theoretical Results and Discussion

For a potential outcomes model with degree β , we let the notation $\text{BRD}(\mathbf{p})$ refer to a staggered rollout Bernoulli design with distinct treatment probabilities $\mathbf{p} = (p_0, p_1, \dots, p_\beta)$. We let $\text{CRD}(\mathbf{k})$ refer to a staggered rollout completely randomized design with distinct treatment counts $\mathbf{k} = (k_0, k_1, \dots, k_\beta)$.

Theorem 2. *Consider a potential outcomes model with degree β . Under a $\text{BRD}(\mathbf{p})$ with $p_0 = 0$, the estimator $\widehat{\text{TTE}}_{\text{PI}}(\mathbf{p})$ is unbiased with variance*

$$O\left(\frac{d^2 \beta^2}{n} Y_{\max}^2 \Delta_{\mathbf{p}}^{-2\beta} + \frac{\sigma^2 \beta}{n} \Delta_{\mathbf{p}}^{-2\beta}\right).$$

Theorem 3. *Consider a potential outcomes model with degree β . Under a $\text{CRD}(\mathbf{k})$ with $k_0 = 0$, the TTE estimator $\widehat{\text{TTE}}_{\text{PI}}(\mathbf{k}/n)$ is unbiased with variance*

$$O\left(\beta^2 Y_{\max}^2 \left(\frac{d^2}{n} + \frac{\beta^2}{k_1}\right) \cdot \left(\frac{n}{\Delta_{\mathbf{k}}}\right)^{2\beta} + \frac{\sigma^2 \beta}{n} \left(\frac{n}{\Delta_{\mathbf{k}}}\right)^{2\beta}\right).$$

Proofs of both of these theorems are given in Appendix A. Notably, families of networks with $d = o(\log n)$ have variance asymptotically approaching 0. As such, our results can generally handle sparse networks. A key technical piece of the analysis is handling the strong correlation in the observations across measurements due to the monotonicity enforced by the staggered rollout design. In the case of a linear potential outcomes model, we can strengthen both of these variance bounds, which match the results from [25], differing only in an additive term coming from observation noise.

Corollary 4. *For a linear potential outcomes model:*

- The estimator $\widehat{\text{TTE}}_{\text{PI}}(\mathbf{p})$ under $\text{BRD}(0, p)$ has variance at most $\frac{1-p}{np} \cdot L_{\max}^2 + \frac{2\sigma^2}{np^2}$.
- The estimator $\widehat{\text{TTE}}_{\text{PI}}(\mathbf{k}/n)$ under $\text{CRD}(0, k)$ has variance at most $\frac{n-k}{(n-1)k} \cdot L_{\max}^2 + \frac{2\sigma^2 n}{k^2}$.

Observe that the Bernoulli estimator does not incorporate any information about the realized treatments. Notably, it does not account for the number of treated individuals. While this binomial random variable concentrates around its mean (especially for large values of n), it fails to account for significant deviations from this mean. Since this information is available at the time of estimation, it can be incorporated into an estimator. We let $\hat{\mathbf{k}} = (\hat{k}_0 = 0, \hat{k}_1, \dots, \hat{k}_\beta)$ be the realized number of treated individuals at each time step, and consider the estimator $\widehat{\text{TTE}}_{\text{PI}}(\hat{\mathbf{k}}/n)$.

Theorem 5. *Consider a potential outcomes model with degree β . Under a staggered rollout Bernoulli design with treatment probabilities $\mathbf{p} = (p_0 = 0, \dots, p_\beta)$, $\widehat{\text{TTE}}_{\text{PI}}(\hat{\mathbf{k}}/n)$ has bias decaying exponentially in n and variance $O\left(\beta^2 Y_{\max}^2 \left(\frac{d^2}{n} + \frac{\beta^2}{p_1 n}\right) \cdot \Delta_{\mathbf{p}}^{-2\beta} + \frac{\beta \sigma^2}{n} \Delta_{\mathbf{p}}^{-2\beta}\right)$.*

A proof is given in Appendix A. For large n , the performance of these three estimators will converge to each other. While our theoretical variance bound in Theorem 5 does not show improvement upon that from Theorem 2, our experimental results illustrate empirical improvements of this estimator.

Discussion. Our results illustrate a natural relationship between the complexity of the model (i.e. its degree β) and the complexity of the randomized design and corresponding estimator; we require

$T \geq \beta$, i.e. $\beta + 1$ outcome measurements, in order to construct an unbiased estimator. Intuitively, each of these measurements allows us to quantify one “degree” of the network effects. Given an overall treatment budget $p = p_T$ with a uniform treatment schedule where $p_t = tp/T$, the difference between treatment fractions is $\Delta_p = p/T$. As a result, for our setting in which $T = \beta$, the variance scales as $(\beta/p)^{2\beta}$, where β is always bounded above by the size of the neighborhood, i.e. graph degree. In comparison, under a fully general neighborhood model, the Horvitz-Thompson estimator has a variance that scales as $O(1/np^d)$, where d denotes the size of the largest neighborhood.

A practical question, critical to real-world experimental settings, is how one should determine the degree β if it is not known in advance. Even if we have many measurements, it may not always be wise to increase the degree of the interpolant, as this increases the magnitude of its slope outside of the interpolating region $[0, p]$. When the expected number of treated individuals np is small relative to the population size n (so $p \ll 1$), the value of the interpolant at 1 will be highly sensitive to any deviation of the later measurements from their expectation. On the other hand, choosing to fit a low degree polynomial may lead to bias if the underlying network effects exhibit higher order interactions. The study of the sensitivity of polynomial interpolation estimators under model misspecification in our randomized experiment setting is a captivating direction for future work. In heuristic settings we recommend choosing a conservative β erring on lower values, as is also common practice when using polynomial regression in supervised learning settings.

The low polynomial structure is primarily used to show that the expected total outcomes function F_D is a low degree polynomial of the treatment fraction. As F_D represents an expectation taken over the population, where treatments are assigned uniformly at random, it is plausible that this function varies in a smooth and simple way when the treatment level is changed. While the polynomial class is not the only hierarchy of function classes to capture complexity, it is a fairly natural one. However, continued study of this overall approach of interpolation for staggered rollout designs for other function classes beyond polynomial would also be incredibly interesting and relevant.

In this work, we also limit our attention to static network effects, but extensions to incorporate time-varying network effects or even time-varying network structures is an interesting direction for future work. When the total treatment budget p_T is small, such that a significant number of individuals are observed under the baseline outcomes across all stages of the experiment, then we could handle time-fixed effects via a simple modification of our estimator by using these baseline individuals to estimate the time-fixed effects and subtracting them from the current estimator.

4 Experiments

We provide simulations on synthetic data to illustrate the performance of our estimators relative to existing estimators. For a population of n individuals, we generate random directed networks of n nodes using a configuration model with in-degrees distributed as a power law with exponent 2.5, and out-degrees evenly shared among individuals. For degree β , we construct the following potential outcomes model:

$$Y_i(\mathbf{z}) = c_{i,\emptyset} + \sum_{j \in \mathcal{N}_i} \tilde{c}_{ij} z_j + \sum_{\ell=2}^{\beta} \left(\frac{\sum_{j \in \mathcal{N}_i} \tilde{c}_{ij} z_j}{\sum_{j \in \mathcal{N}_i} \tilde{c}_{ij}} \right)^{\ell}, \quad (4.1)$$

where $c_{i,\emptyset} \sim U[0, 1]$, $\tilde{c}_{ii} \sim U[0, 1]$, and for $i \neq j$, $\tilde{c}_{ij} = v_j |\mathcal{N}_i| / \sum_{k: (k,j) \in E} |\mathcal{N}_k|$ for $v_j \sim U[0, r]$, where r denotes a hyperparameter that governs the relative magnitude of the network effects relative to the direct effects. Essentially v_j represents the magnitude of individual j ’s influence, which is then shared among its out-neighbors proportional to their in-degrees. For simplicity we assume no observation noise in the experiments, i.e. $\sigma = 0$.

Other Algorithms. We benchmark our proposed estimators against least squares regression and difference-in-mean estimators. As these estimators don’t utilize the staggered rollout design, we evaluate them on the measurements taken at the last stage, T , of the experiment. We will use \mathbf{z} to denote the treatment vector at time T (suppressing the superscript). As a network sampled from a configuration model does not exhibit clustering, the solutions that propose cluster randomized designs perform poorly, and thus we omit them from the experiments.

The standard difference in means estimator is the difference between the average outcome of individuals assigned to treatment and the average outcome of individuals assigned to control, given by

$$\widehat{\text{TTE}}_{\text{DM}} = \frac{\sum_{i \in [n]} z_i Y_i(\mathbf{z})}{\sum_{i \in [n]} z_i} - \frac{\sum_{i \in [n]} (1 - z_i) Y_i(\mathbf{z})}{\sum_{i \in [n]} (1 - z_i)}. \quad (4.2)$$

This estimator is biased under the presence of network interference. Note that $\widehat{\text{TTE}}_{\text{DM}}$ does not take into account any information about each individual's neighborhood.

A modification of the difference in means estimator incorporates knowledge of the number of treated neighbors of each individual. Let U_i denote the number of individuals in $\mathcal{N}_i \setminus \{i\}$ assigned to treatment, and let \tilde{U}_i denote the number of neighbors individuals in $\mathcal{N}_i \setminus \{i\}$ assigned to control. This estimator is given by

$$\widehat{\text{TTE}}_{\text{DM}(\lambda)} = \frac{\sum_{i \in [n]} z_i \mathbb{I}(U_i \geq \lambda) Y_i(\mathbf{z})}{\sum_{i \in [n]} z_i \mathbb{I}(U_i \geq \lambda)} - \frac{\sum_{i \in [n]} (1 - z_i) \mathbb{I}(\tilde{U}_i \geq \lambda) Y_i(\mathbf{z})}{\sum_{i \in [n]} (1 - z_i) \mathbb{I}(\tilde{U}_i \geq \lambda)}, \quad (4.3)$$

for some user-defined tolerance $\lambda \in [0, 1]$. This estimator only counts an individual's outcome if at least λ of the individual's neighborhood is assigned to the same treatment as the individual itself. In our experiments, we set $\lambda = 0.75$.

Finally we compare against least squares regression models of degree β , which posit that the potential outcomes model can be described as

$$Y_i(\mathbf{z}) = g(z_i, \bar{z}_i) = \left(\rho + \sum_{k=1}^{\beta} \gamma_k X_i^k \right) + z_i \left(\tilde{\rho} + \sum_{k=1}^{\beta-1} \tilde{\gamma}_k X_i^k \right), \quad (4.4)$$

for some covariate X_i . In the two variations we consider, we set X_i equal to either the number of treated neighbors or the proportion of treated neighbors, where we do not include i itself. The two sets of coefficients $(\rho, \gamma_1, \dots, \gamma_{\beta})$ and $(\rho, \tilde{\gamma}_1, \dots, \tilde{\gamma}_{\beta})$ allow for the model to be different when i is treated vs not treated, and the second summation only goes until $\beta - 1$ since we want to only allow degree β interactions. The total number of coefficients in the model is $2\beta + 1$. Least squares regression finds the set of coefficients that minimizes the least squares predictive error on the dataset, which consists of $\{z_i, X_i, Y_i(\mathbf{z})\}_{i \in [n]}$. The estimated coefficients define an estimate for the function \hat{g} . For the variation which uses the number of treated neighbors as the covariates, setting $X_i = U_i$, the estimate is given by

$$\widehat{\text{TTE}}_{\text{LS-Num}} = \frac{1}{n} \sum_{i=1}^n (\hat{g}(1, |\mathcal{N}_i| - 1) - \hat{g}(0, 0)). \quad (4.5)$$

For the variation which uses the proportion of treated neighbors as the covariates, setting $X_i = U_i / (|\mathcal{N}_i| - 1)$, the estimate is given by

$$\widehat{\text{TTE}}_{\text{LS-Prop}} = \frac{1}{n} \sum_{i=1}^n (\hat{g}(1, 1) - \hat{g}(0, 0)). \quad (4.6)$$

As completely randomized design is more balanced than Bernoulli randomized design, we evaluate all the benchmark algorithms under a completely randomized design.

Results and Discussion. For each population size n , we sample G networks from the distribution described above. For each configuration of parameters in the experiment, we sample N treatment schedules $\{\mathbf{z}^0, \dots, \mathbf{z}^{\beta}\}$ from our parameterized distribution class (Bernoulli or CRD) compute the TTE using each estimator. For each estimator, we plot the relative bias of the TTE estimates averaged over the results from these GN samples and normalized by the magnitude of the TTE. The width of the shading in the figures depicts the standard deviation across the GN estimates. We ran all experiments on a Linux-based machine with 20 CPU(s) and 10 cores. The experiments for the linear setting took 8.3 minutes and the experiments with varying polynomial degree took 4.6 minutes.

In Figure 1, we visualize the effect of four network/estimator parameters on the quality of each of the five TTE estimators (the four described above, and our CRD estimator with treatment targets $k_t = \frac{tk}{\beta}$). Specifically, we consider the effects of the population size (n), the maximum proportion of treated individuals (k/n), the degree of the potential outcomes model (β), and the ratio between the network and direct effects (r). Each of the plots fixes three of these parameters and varies the fourth. Specific settings of the parameters are listed on each plot.

In Figure 1, our estimator (in blue) is unbiased as expected and the variance decreases as n and k/n increases. However, the other estimators remain significantly biased, with higher variances than ours,

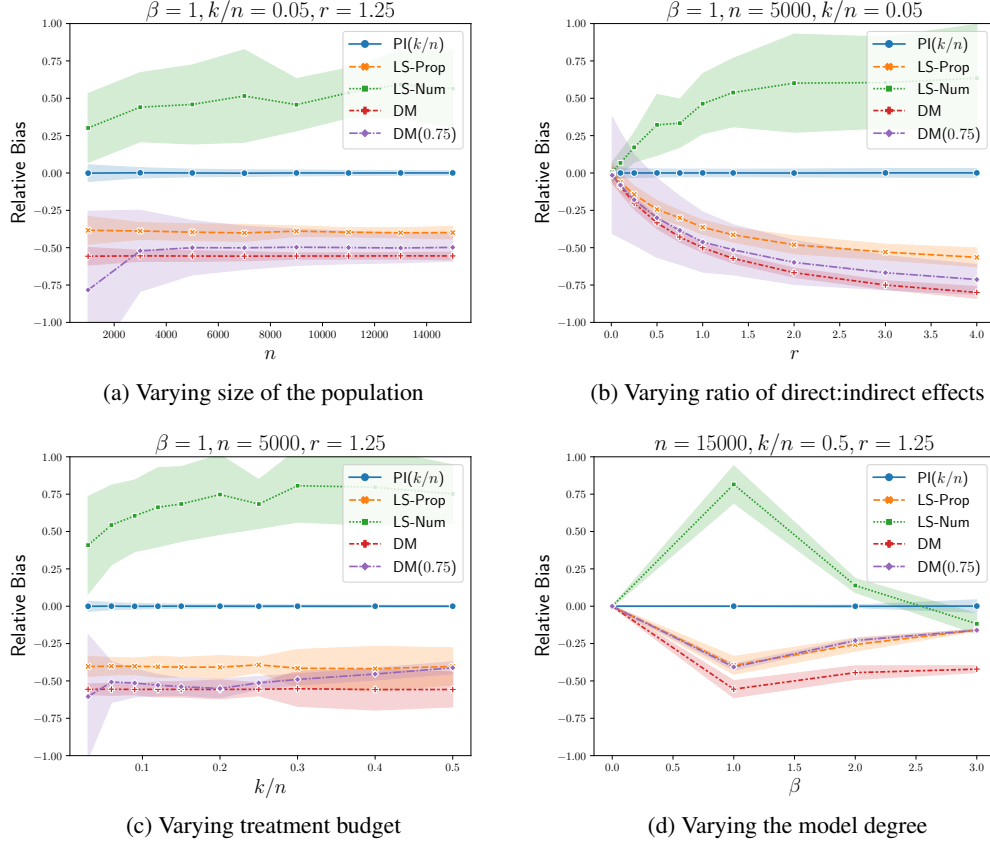


Figure 1: Four graphs visualizing the performance of various TTE estimators as various parameters are adjusted. The height of each graph depicts the experimental relative bias of the estimator and the shaded width depicts the experimental standard deviation.

regardless of treatment budget or population size. As the ratio r increases the network effects become more significant relative to the direct effect, and thus the bias of other estimators also increases. As a sanity check, when the ratio is close to 0, all estimators are unbiased as there are no network effects.

In Figure 2, we compare the variants of our estimator, evaluating $\widehat{\text{TTE}}_{\text{PI}}(\mathbf{k}/n)$ under CRD and evaluating $\widehat{\text{TTE}}_{\text{PI}}(\mathbf{p})$ and $\widehat{\text{TTE}}_{\text{PI}}(\hat{\mathbf{k}}/n)$ under Bernoulli(\mathbf{p}) randomized design, where $p_t = tp/\beta$ and $\hat{\mathbf{k}}$ is the vector of realized treatment counts. The estimators $\widehat{\text{TTE}}_{\text{PI}}(\mathbf{k}/n)$ and $\widehat{\text{TTE}}_{\text{PI}}(\hat{\mathbf{k}}/n)$ perform nearly identically. $\widehat{\text{TTE}}_{\text{PI}}(\hat{\mathbf{k}}/n)$ has lower variance than $\widehat{\text{TTE}}_{\text{PI}}(\mathbf{p})$, which is intuitive as it performs polynomial interpolation on the realized treatment fraction rather than the expected treatment fraction. We include additional experiments for higher degree models in Appendix C.

An additional point of comparison for these estimators is their computational complexity. Here, the most natural comparison is between our estimators and least squares, as these are the only approaches that make use of the various rounds of outcome measurements. Since our estimators require only an aggregated measurement of the individual’s outcomes, the $O(\beta^2)$ runtime of the interpolation is asymptotically dominated by the $O(n\beta)$ time to read in the outcome measurements. The least squares methods fit $O(\beta)$ parameters and have time complexity $O(\beta^2 n)$.

5 Conclusion

We propose a new approach for causal inference under network interference which performs significantly better than existing approaches without requiring knowledge of the graph. In particular, the additional measurements from a staggered rollout design enable us to reduce the task of estimating

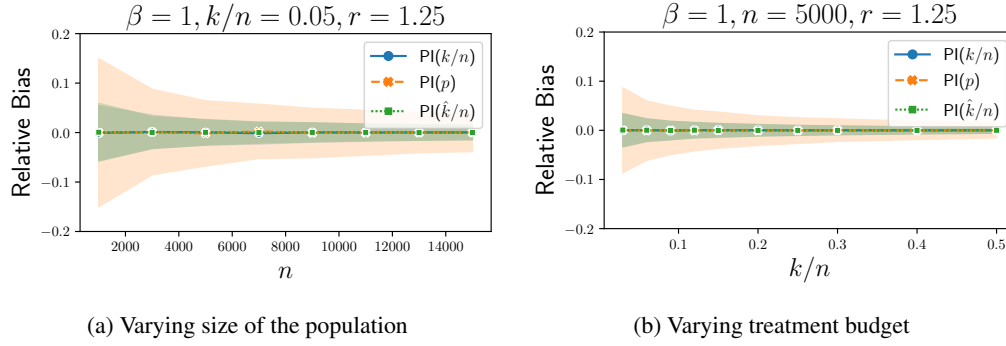


Figure 2: Two graphs visualizing the performance of our proposed TTE estimators as the size of the population (n) or treatment budget (k/n) is varied. The height of each graph depicts the experimental relative bias of the estimator and the shaded width depicts the experimental standard deviation. The blue and the green plots essentially overlap.

total treatment effect to that of polynomial interpolation. We show that under a flexible class of low degree polynomial potential outcomes our estimator is unbiased with variance scaling as $O(1/n)$. Future directions include how to optimally perform model selection when β is unknown, and generalizing to a dynamic setting by incorporating time-dependent noise to the model, considering time-varying effects, or allowing for time-varying networks. The staggered rollout design framework has implications towards estimation under other model classes beyond polynomial, such as sublinear or monotone functions, under which one may be able to construct bounds on TTE.

Acknowledgments and Disclosure of Funding

We gratefully acknowledge financial support from the National Science Foundation grants CCF-1948256 and CNS-1955997 and the National Science Foundation Graduate Research Fellowship grant DGE-1650441. Dr. Yu is also supported by an Intel Rising Stars award and a JPMorgan Faculty Research award. Mayleen is also supported in part by the Sloan Foundation grant 90855.

References

- [1] Peter M Aronow, Cyrus Samii, et al. Estimating average causal effects under general interference, with application to a social network experiment. *The Annals of Applied Statistics*, 11(4):1912–1947, 2017.
- [2] Eric Auerbach and Max Tabord-Meehan. The local approach to causal inference under network interference. *arXiv preprint arXiv:2105.03810*, 2021.
- [3] Falco J Bargagli-Stoffi, Costanza Tortù, and Laura Forastiere. Heterogeneous treatment and spillover effects under clustered network interference. *arXiv preprint arXiv:2008.00707*, 2020.
- [4] Guillaume W Basse and Edoardo M Airoldi. Model-assisted design of experiments in the presence of network-correlated outcomes. *Biometrika*, 105(4):849–858, 2018.
- [5] Rohit Bhattacharya, Daniel Malinsky, and Ilya Shpitser. Causal inference under interference and network uncertainty. In Ryan P. Adams and Vibhav Gogate, editors, *Proceedings of The 35th Uncertainty in Artificial Intelligence Conference*, volume 115 of *Proceedings of Machine Learning Research*, pages 1028–1038. PMLR, 22–25 Jul 2020. URL <https://proceedings.mlr.press/v115/bhattacharya20a.html>.
- [6] Jing Cai, Alain De Janvry, and Elisabeth Sadoulet. Social networks and the decision to insure. *American Economic Journal: Applied Economics*, 7(2):81–108, 2015.
- [7] Alex Chin. Regression adjustments for estimating the global treatment effect in experiments with interference. *Journal of Causal Inference*, 7(2), 2019.

- [8] Dean Eckles, Brian Karrer, and Johan Ugander. Design and analysis of experiments in networks: Reducing bias from interference. *Journal of Causal Inference*, 5(1), 2017.
- [9] Huan Gui, Ya Xu, Anmol Bhasin, and Jiawei Han. Network a/b testing: From sampling to estimation. In *Proceedings of the 24th International Conference on World Wide Web*, pages 399–409. International World Wide Web Conferences Steering Committee, 2015.
- [10] Michael G. Hudgens and M. Elizabeth Halloran. Toward causal inference with interference. *Journal of the American Statistical Association*, 103:832–842, 2008. URL <https://EconPapers.repec.org/RePEc:bes:jnlasa:v:103:y:2008:m:june:p:832-842>.
- [11] Shuangning Li and Stefan Wager. Random graph asymptotics for treatment effect estimation under network interference. *arXiv preprint arXiv:2007.13302*, 2020.
- [12] Wenrui Li, Daniel L Sussman, and Eric D Kolaczyk. Causal inference under network interference with noise. *arXiv preprint arXiv:2105.04518*, 2021.
- [13] Lan Liu and Michael G. Hudgens. Large sample randomization inference of causal effects in the presence of interference. *Journal of the American Statistical Association*, 109(505):288–301, 2014. doi: 10.1080/01621459.2013.844698. URL <https://doi.org/10.1080/01621459.2013.844698>. PMID: 24659836.
- [14] Charles F Manski. Identification of treatment response with social interactions. *The Econometrics Journal*, 16(1), 2013.
- [15] Ben M. Parker, Steven G. Gilmour, and John Schormans. Optimal design of experiments on connected units with application to social networks. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 66(3):455–480, 2017. doi: <https://doi.org/10.1111/rssc.12170>. URL <https://rss.onlinelibrary.wiley.com/doi/abs/10.1111/rssc.12170>.
- [16] Paul R Rosenbaum. Interference between units in randomized experiments. *Journal of the American Statistical Association*, 102(477):191–200, 2007. doi: 10.1198/016214506000001112. URL <https://doi.org/10.1198/016214506000001112>.
- [17] Michael E Sobel. What do randomized studies of housing mobility demonstrate? *Journal of the American Statistical Association*, 101(476):1398–1407, 2006. doi: 10.1198/016214506000000636. URL <https://doi.org/10.1198/016214506000000636>.
- [18] Daniel L Sussman and Edoardo M Airolidi. Elements of estimation theory for causal effects in the presence of network interference. *arXiv preprint arXiv:1702.03578*, 2017.
- [19] Eric J Tchetgen Tchetgen and Tyler J VanderWeele. On causal inference in the presence of interference. *Statistical Methods in Medical Research*, 21(1):55–75, 2012. doi: 10.1177/0962280210386779. URL <https://doi.org/10.1177/0962280210386779>. PMID: 21068053.
- [20] Panos Toulis and Edward Kao. Estimation of causal peer influence effects. In *International Conference on Machine Learning*, pages 1489–1497, 2013.
- [21] Johan Ugander and Hao Yin. Randomized graph cluster randomization. *arXiv preprint arXiv:2009.02297*, 2020.
- [22] Johan Ugander, Brian Karrer, Lars Backstrom, and Jon Kleinberg. Graph cluster randomization: Network exposure to multiple universes. In *Proceedings of the 19th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 329–337. ACM, 2013.
- [23] Tyler J. VanderWeele, Eric J. Tchetgen Tchetgen, and M. Elizabeth Halloran. Interference and sensitivity analysis. *Statist. Sci.*, 29(4):687–706, 11 2014. doi: 10.1214/14-STS479. URL <https://doi.org/10.1214/14-STS479>.
- [24] Davide Viviano. Experimental design under network interference. *arXiv preprint arXiv:2003.08421*, 2020.

- [25] Christina Lee Yu, Edoardo M Airolidi, Christian Borgs, and Jennifer T Chayes. Estimating total treatment effect in randomized experiments with unknown network structure. *arXiv preprint arXiv:2205.12803*, 2022. doi: 10.48550/ARXIV.2205.12803. URL <https://arxiv.org/abs/2205.12803>.

A Variance Calculations

In this section, we establish the unbiasedness and variance bounds of the estimators introduced throughout the paper. The following lemma will be useful for some of these calculations.

Lemma 6. *Suppose we have $Y_{i,t}^{\text{obs}} = Y_i(\mathbf{z}^t) + \varepsilon_{i,t}$ for iid noise $\varepsilon_{i,t} \sim N(0, \sigma^2)$ and our estimator has the form*

$$\widehat{\text{TTE}} = \frac{1}{n} \sum_{t=0}^{\beta} \sum_{i=1}^n \alpha_{i,t} \cdot Y_{i,t}^{\text{obs}},$$

with each $|\alpha_{i,t}| = O(\alpha)$. Further suppose that for any $t, t' \in 0, \dots, \beta$ and two subsets $\mathcal{S}, \mathcal{S}'$ of cardinality at most β ,

$$\left| \text{Cov} \left[\prod_{j \in \mathcal{S}} z_j^t, \prod_{j' \in \mathcal{S}'} z_{j'}^{t'} \right] \right| \leq \begin{cases} B_1 & \mathcal{S} \cap \mathcal{S}' \neq \emptyset, \\ B_2 & \mathcal{S} \cap \mathcal{S}' = \emptyset. \end{cases}$$

Then,

$$\text{Var}[\widehat{\text{TTE}}] = O\left(\alpha^2 \beta^2 Y_{\max}^2 \left(\frac{d^2}{n} \max\{B_1, B_2\} + B_2\right) + \frac{\sigma^2 \beta}{n} \alpha^2\right).$$

Proof. By the law of total variance, we have

$$\begin{aligned} \text{Var}[\widehat{\text{TTE}}] &= \text{Var} \left[\mathbb{E}[\widehat{\text{TTE}} \mid \mathbf{z}^t] \right] + \mathbb{E} \left[\text{Var}[\widehat{\text{TTE}} \mid \mathbf{z}^t] \right] \\ &= \text{Var} \left[\frac{1}{n} \sum_{t=0}^{\beta} \sum_{i=1}^n \alpha_{i,t} \cdot Y_i(\mathbf{z}^t) \right] \\ &\quad + \mathbb{E} \left[\text{Var} \left(\frac{1}{n} \sum_{t=0}^{\beta} \sum_{i=1}^n \alpha_{i,t} \cdot Y_i(\mathbf{z}^t) + \frac{1}{n} \sum_{t=0}^{\beta} \sum_{i=1}^n \alpha_{i,t} \cdot \varepsilon_{i,t} \mid \mathbf{z}^t \right) \right] \\ &= \text{Var} \left[\frac{1}{n} \sum_{t=0}^{\beta} \sum_{i=1}^n \alpha_{i,t} \cdot Y_i(\mathbf{z}^t) \right] + \mathbb{E} \left[\text{Var} \left(\frac{1}{n} \sum_{t=0}^{\beta} \sum_{i=1}^n \alpha_{i,t} \cdot \varepsilon_{i,t} \right) \right] \\ &= \text{Var} \left[\frac{1}{n} \sum_{t=0}^{\beta} \sum_{i=1}^n \alpha_{i,t} \cdot Y_i(\mathbf{z}^t) \right] + \mathbb{E} \left[\frac{1}{n^2} \sum_{t=0}^{\beta} \sum_{i=1}^n \text{Var}(\alpha_{i,t} \cdot \varepsilon_{i,t}) \right] \\ &= \text{Var} \left[\frac{1}{n} \sum_{t=0}^{\beta} \sum_{i=1}^n \alpha_{i,t} \cdot Y_i(\mathbf{z}^t) \right] + O\left(\frac{\sigma^2 \beta}{n} \alpha^2\right) \end{aligned}$$

Turning our attention to the first variance term, we introduce the notation $\mathcal{M}_i = \{i' : |\mathcal{N}_i \cap \mathcal{N}_{i'}| \geq 1\}$. Note that $|\mathcal{M}_i| \leq d^2$. In addition, for all $i' \notin \mathcal{M}_i$, all $\mathcal{S} \subseteq \mathcal{N}_i$, and all $\mathcal{S}' \subseteq \mathcal{N}_{i'}$, we have $\mathcal{S} \cap \mathcal{S}' = \emptyset$.

We may expand the variance,

$$\begin{aligned}
\text{Var}\left[\frac{1}{n} \sum_{t=0}^{\beta} \sum_{i=1}^n \alpha_{i,t} \cdot Y_i(\mathbf{z}^t)\right] &= \frac{1}{n^2} \sum_{i=1}^n \sum_{i'=1}^n \sum_{t=0}^{\beta} \sum_{t'=0}^{\beta} \alpha_{i,t} \cdot \alpha_{i',t'} \cdot \text{Cov}\left[Y_i(\mathbf{z}^t), Y_{i'}(\mathbf{z}^{t'})\right] \\
&\leq \frac{O(\alpha^2)}{n^2} \sum_{i=1}^n \sum_{i'=1}^n \sum_{t=0}^{\beta} \sum_{t'=0}^{\beta} \left| \text{Cov}\left[Y_i(\mathbf{z}^t), Y_{i'}(\mathbf{z}^{t'})\right] \right| \\
&\leq \frac{O(\alpha^2)}{n^2} \sum_{i=1}^n \sum_{i'=1}^n \sum_{t=0}^{\beta} \sum_{t'=0}^{\beta} \sum_{\substack{S \subseteq \mathcal{N}_i \\ |S| \leq \beta}} |c_{i,S}| \sum_{\substack{S' \subseteq \mathcal{N}_{i'} \\ |S'| \leq \beta}} |c_{i',S'}| \cdot \left| \text{Cov}\left[\prod_{j \in S} z_j^t, \prod_{j \in S'} z_{j'}^{t'}\right] \right| \\
&\leq \frac{O(\alpha^2)}{n^2} \left(\sum_{i=1}^n \sum_{i' \in \mathcal{M}_i} \sum_{t=0}^{\beta} \sum_{t'=0}^{\beta} \sum_{\substack{S \subseteq \mathcal{N}_i \\ |S| \leq \beta}} |c_{i,S}| \sum_{\substack{S' \subseteq \mathcal{N}_{i'} \\ |S'| \leq \beta}} |c_{i',S'}| \cdot \max\{B_1, B_2\} \right. \\
&\quad \left. + \sum_{i=1}^n \sum_{i' \notin \mathcal{M}_i} \sum_{t=0}^{\beta} \sum_{t'=0}^{\beta} \sum_{\substack{S \subseteq \mathcal{N}_i \\ |S| \leq \beta}} |c_{i,S}| \sum_{\substack{S' \subseteq \mathcal{N}_{i'} \\ |S'| \leq \beta}} |c_{i',S'}| \cdot B_2 \right) \\
&\leq \frac{O(\alpha^2)}{n^2} \left(\sum_{i=1}^n \sum_{i' \in \mathcal{M}_i} \beta^2 Y_{\max}^2 \cdot \max\{B_1, B_2\} + \sum_{i=1}^n \sum_{i' \notin \mathcal{M}_i} \beta^2 Y_{\max}^2 \cdot B_2 \right) \\
&\leq \frac{O(\alpha^2)}{n^2} \left(d^2 n \cdot \beta^2 Y_{\max}^2 \cdot \max\{B_1, B_2\} + n^2 \cdot \beta^2 Y_{\max}^2 \cdot B_2 \right) \\
&= O\left(\alpha^2 \beta^2 Y_{\max}^2 \left(\frac{d^2}{n} \max\{B_1, B_2\} + B_2\right)\right).
\end{aligned}$$

Therefore,

$$\text{Var}\left[\widehat{\text{TTE}}\right] = O\left(\alpha^2 \beta^2 Y_{\max}^2 \left(\frac{d^2}{n} \max\{B_1, B_2\} + B_2\right) + \frac{\sigma^2 \beta}{n} \alpha^2\right).$$

□

A.1 Graph Agnostic with Bernoulli Treatment

By plugging in the Bernoulli treatment probabilities into (3.1), we obtain the estimator:

$$\widehat{\text{TTE}}(\mathbf{p}) := \frac{1}{n} \sum_{i=1}^n \sum_{t=0}^{\beta} \left(\ell_{t,\mathbf{p}}(1) - \ell_{t,\mathbf{p}}(0) \right) \cdot Y_{i,t}^{\text{obs}}, \quad \ell_{t,\mathbf{p}}(x) = \prod_{\substack{s=0 \\ s \neq t}}^{\beta} \frac{x - p_s}{p_t - p_s}.$$

The following lemma will be useful in establishing a bound on the variance of this estimator.

Lemma 7. $\max_{t \in \{0 \dots \beta\}} \{|\ell_{t,\mathbf{p}}(1) - \ell_{t,\mathbf{p}}(0)|\} = O\left(\Delta_{\mathbf{p}}^{-\beta}\right).$

Proof. For each $t \in 0, \dots, \beta$, we have,

$$|\ell_{t,\mathbf{p}}(1) - \ell_{t,\mathbf{p}}(0)| \leq \left| \prod_{\substack{s=0 \\ s \neq t}}^{\beta} \frac{1 - p_s}{p_t - p_s} \right| + \left| \prod_{\substack{s=0 \\ s \neq t}}^{\beta} \frac{-p_s}{p_t - p_s} \right| \leq \prod_{\substack{s=0 \\ s \neq t}}^{\beta} \frac{|1 - p_s|}{\Delta_{\mathbf{p}}} + \prod_{\substack{s=0 \\ s \neq t}}^{\beta} \frac{|p_s|}{\Delta_{\mathbf{p}}} = O\left(\Delta_{\mathbf{p}}^{-\beta}\right).$$

Here, the first inequality is an application of the triangle inequality, the second uses the definition of $\Delta_{\mathbf{p}}$, and the third uses the fact that each $p_t \in [0, 1]$. □

Proof of Theorem 2. To establish the unbiasedness of the estimator, note that,

$$\begin{aligned}
\mathbb{E}[\widehat{\text{TTE}}(\mathbf{p})] &= \sum_{t=0}^{\beta} \left(\ell_{t,\mathbf{p}}(1) - \ell_{t,\mathbf{p}}(0) \right) \cdot \mathbb{E} \left[\frac{1}{n} \sum_{i=1}^n Y_i(\mathbf{z}^t) \right] \\
&= \sum_{t=0}^{\beta} \left(\ell_{t,\mathbf{p}}(1) - \ell_{t,\mathbf{p}}(0) \right) \cdot F_B(p_t) \\
&= \left(\sum_{t=0}^{\beta} \ell_{t,\mathbf{p}}(1) \cdot F_B(p_t) \right) - \left(\sum_{t=0}^{\beta} \ell_{t,\mathbf{p}}(0) \cdot F_B(p_t) \right) \\
&= F_B(1) - F_B(0) \\
&= \text{TTE}.
\end{aligned}$$

Now, we compute a bound on the variance. Since the entries of each \mathbf{z}^t are independent, $\text{Cov} \left[\prod_{j \in \mathcal{S}} z_j^t, \prod_{j' \in \mathcal{S}'} z_{j'}^{t'} \right] = 0$ for any disjoint $\mathcal{S}, \mathcal{S}'$. In addition, since both arguments of this covariance are indicator variables, we can upper bound the absolute value of each covariance by 1. We appeal to Lemma 6, with $B_1 = 1$, $B_2 = 0$, and $\alpha = \Delta_{\mathbf{p}}^{-\beta}$ (by Lemma 7), giving,

$$\text{Var}[\widehat{\text{TTE}}(\mathbf{p})] = O \left(\frac{d^2 \beta^2}{n} Y_{\max}^2 \Delta_{\mathbf{p}}^{-2\beta} + \frac{\sigma^2 \beta}{n} \Delta_{\mathbf{p}}^{-2\beta} \right).$$

□

A.2 Graph Agnostic with Completely Randomized Treatment

We'll make use of the following algebraic lemma to bound the variance; recall the bracket notation introduced in equation (2.3) in Section 2.

Lemma 8. *For any constants $a, b \in \mathbb{N}$ and any $p \in (0, 1]$,*

$$\left| \frac{\left[\frac{pn-a}{n-a} \right]^b}{\left[\frac{pn}{n} \right]^b} - 1 \right| = O \left(\frac{ab}{pn} \right).$$

Proof. Expanding the bracket notation, we have,

$$\begin{aligned}
\left| \frac{\left[\frac{pn-a}{n-a} \right]^b}{\left[\frac{pn}{n} \right]^b} - 1 \right| &= \left| \prod_{i=0}^{b-1} \left(\frac{pn-a-i}{pn-i} \right) \left(\frac{n-i}{n-a-i} \right) - 1 \right| \\
&= \left| \prod_{i=0}^{b-1} \left(1 - \frac{a}{pn-i} \right) \left(1 + \frac{a}{n-a-i} \right) - 1 \right| \\
&= \left| \prod_{i=0}^{b-1} \left(1 + O \left(\frac{a(p-1)}{pn} \right) \right) - 1 \right| \\
&\leq \sum_{j=1}^{b-1} \binom{b}{j} \cdot O \left(\frac{a}{pn} \right)^j \\
&\leq \sum_{j=1}^{b-1} O \left(\frac{ab}{pn} \right)^j \\
&= O \left(\frac{ab}{pn} \right).
\end{aligned}$$

□

Proof of Theorem 3. To establish the unbiasedness of the estimator, note that,

$$\begin{aligned}
\mathbb{E}[\widehat{\text{TTE}}(\mathbf{k})] &= \sum_{t=0}^{\beta} \left(\ell_{t,\mathbf{k}/n}(1) - \ell_{t,\mathbf{k}/n}(0) \right) \cdot \mathbb{E} \left[\frac{1}{n} \sum_{i=1}^n Y_i(\mathbf{z}^t) \right] \\
&= \sum_{t=0}^{\beta} \left(\ell_{t,\mathbf{k}/n}(1) - \ell_{t,\mathbf{k}/n}(0) \right) \cdot F_C\left(\frac{k}{n}\right) \\
&= \left(\sum_{t=0}^{\beta} \ell_{t,\mathbf{k}/n}(1) \cdot F_C\left(\frac{k}{n}\right) \right) - \left(\sum_{t=0}^{\beta} \ell_{t,\mathbf{k}/n}(0) \cdot F_C\left(\frac{k}{n}\right) \right) \\
&= F_C(1) - F_C(0) \\
&= \text{TTE}.
\end{aligned}$$

Next, we establish a bound on the variance of this estimator. We consider the covariance term $\left| \text{Cov} \left[\prod_{j \in \mathcal{S}} z_j^t, \prod_{j' \in \mathcal{S}'} z_{j'}^{t'} \right] \right|$ for various values of t, t', \mathcal{S} , and \mathcal{S}' . First, note that when t or $t' = 0$, an argument of this covariance is deterministically 0, so the covariance is 0 as well. Otherwise, when $\mathcal{S} \cap \mathcal{S}' \neq \emptyset$, we can bound $\left| \text{Cov} \left[\prod_{j \in \mathcal{S}} z_j^t, \prod_{j' \in \mathcal{S}'} z_{j'}^{t'} \right] \right| \leq 1$ by noting that both arguments are indicator variables. In the case that $\mathcal{S} \cap \mathcal{S}' = \emptyset$, we establish a stronger bound using Lemma 8. We have,

$$\begin{aligned}
\text{Cov} \left[\prod_{j \in \mathcal{S}} z_j^t, \prod_{j' \in \mathcal{S}'} z_{j'}^{t'} \right] &= \mathbb{E} \left[\prod_{j \in \mathcal{S}} z_j^t \prod_{j' \in \mathcal{S}'} z_{j'}^{t'} \right] - \mathbb{E} \left[\prod_{j \in \mathcal{S}} z_j^t \right] \mathbb{E} \left[\prod_{j' \in \mathcal{S}'} z_{j'}^{t'} \right] \\
&\leq \left[\frac{k_t}{n} \right]^{|S|} \left[\frac{k_{t'}}{n} \right]^{|S'|} \cdot \left| \frac{\left[\frac{k_{t'} - |S|}{n - |S|} \right]^{|S'|}}{\left[\frac{k_{t'}}{n} \right]^{|S'|}} - 1 \right| \\
&= O \left(\frac{|S||S'|}{k_{t'}} \right) \\
&= O \left(\frac{\beta^2}{k_1} \right).
\end{aligned}$$

In the second last line, we bound the first two factors by 1, and use Lemma 8 (with $p = \frac{k_{t'}}{n}$) to bound the third factor. Applying Lemma 6, with $B_1 = 1$, $B_2 = O\left(\frac{\beta^2}{k_1}\right)$, and $\alpha = \left(\frac{n}{\Delta_{\mathbf{k}}}\right)^{\beta}$ (by Lemma 7 using the substitution $\mathbf{p} = \mathbf{k}/n$), giving,

$$\text{Var}[\widehat{\text{TTE}}(\mathbf{k})] = O \left(\beta^2 Y_{\max}^2 \left(\frac{d^2}{n} + \frac{\beta^2}{k_1} \right) \cdot \left(\frac{n}{\Delta_{\mathbf{k}}} \right)^{2\beta} + \frac{\sigma^2 \beta}{n} \left(\frac{n}{\Delta_{\mathbf{k}}} \right)^{2\beta} \right).$$

□

A.3 Improved Variance Bounds in the Linear Setting

Proof of Corollary 4. In the linear setting ($\beta = 1$) for $\mathbf{x} = (0, x)$, the Lagrange polynomial coefficients evaluate to $\ell_{0,\mathbf{x}}(1) - \ell_{0,\mathbf{x}}(0) = -\alpha$ and $\ell_{1,\mathbf{x}}(1) - \ell_{1,\mathbf{x}}(0) = \alpha$ for $\alpha = \frac{1}{x}$, so that the estimator $\widehat{\text{TTE}}(\mathbf{x})$ is equal to

$$\begin{aligned}
\widehat{\text{TTE}}(\mathbf{x}) &= \frac{\alpha}{n} \left(\sum_{i=1}^n Y_{i,1}^{\text{obs}} - \sum_{i=1}^n Y_{i,0}^{\text{obs}} \right) \\
&= \frac{\alpha}{n} \sum_{i=1}^n \left(Y_i(\mathbf{z}^1) + \varepsilon_{i,1} - \varepsilon_{i,0} - c_{i,\emptyset} \right).
\end{aligned}$$

Using the Law of Total Variance, we get

$$\text{Var}[\widehat{\text{TTE}}(\mathbf{x})] = \text{Var} \left[\mathbb{E} \left[\widehat{\text{TTE}} \mid \mathbf{z}^1 \right] \right] + \mathbb{E} \left[\text{Var} \left[\widehat{\text{TTE}} \mid \mathbf{z}^1 \right] \right] = \text{Var} \left[\frac{\alpha}{n} \sum_{i=1}^n Y_i(\mathbf{z}^1) \right] + \frac{2\sigma^2 \alpha^2}{n}.$$

Rewriting the first term, we get

$$\begin{aligned}\text{Var}\left[\frac{\alpha}{n} \sum_{i=1}^n Y_i(\mathbf{z}^1)\right] &= \frac{\alpha^2}{n^2} \sum_{i=1}^n \sum_{i'=1}^n \text{Cov}[Y_i(\mathbf{z}^1), Y_{i'}(\mathbf{z}^1)] \\ &= \frac{\alpha^2}{n^2} \sum_{i=1}^n \sum_{i'=1}^n \sum_{j \in \mathcal{N}_i} c_{ij} \sum_{j' \in \mathcal{N}_{i'}} c_{i'j'} \text{Cov}[z_j, z_{j'}]\end{aligned}\quad (\text{A.1})$$

$$= \frac{\alpha^2}{n^2} \sum_{j=1}^n \sum_{j'=1}^n \left(\sum_{i: j \in \mathcal{N}_i} c_{ij} \right) \left(\sum_{i': j' \in \mathcal{N}_{i'}} c_{i'j'} \right) \text{Cov}[z_j, z_{j'}]. \quad (\text{A.2})$$

Here, we used the fact that $\mathbf{z}^0 = \mathbf{0}$ deterministically to remove covariance terms, as it has covariance 0 with any other random variable. Under BRD(0, p) we have $\alpha = \frac{1}{p}$. Additionally, $\text{Var}(z_j^1) = p(1-p)$ for each $j \in [n]$, and $\text{Cov}[z_j, z_{j'}] = 0$ for $j \neq j'$ so we may simplify the variance bound to

$$\begin{aligned}&= \frac{\alpha^2}{n^2} \sum_{j=1}^n \left(\sum_{i: j \in \mathcal{N}_i} c_{ij} \right)^2 \cdot \text{Var}(z_j^1) + \frac{2\sigma^2\alpha^2}{n} \\ &\leq \frac{\alpha^2}{n^2} \cdot L_{\max}^2 \cdot \sum_{j=1}^n \text{Var}(z_j^1) + \frac{2\sigma^2\alpha^2}{n} \\ &\leq \frac{1-p}{np} \cdot L_{\max}^2 + \frac{2\sigma^2}{np^2}.\end{aligned}$$

The analysis for the completely randomized design setting is presented in pg 32 of [25], and we include it here for convenience. Under CRD(0, k), we have $\alpha = \frac{n}{k}$. Additionally, $\text{Var}(z_j^1) = \frac{k(n-k)}{n^2}$ for each $j \in [n]$, and

$$\text{Cov}[z_j^1, z_{j'}^1] = \frac{k(k-1)n}{n^2(n-1)} - \frac{k^2(n-1)}{n^2(n-1)} = \frac{-k(n-k)}{n^2(n-1)} \leq 0.$$

Plugging into (A.2), we find that

$$\begin{aligned}\text{Var}[\widehat{\text{TTE}}(\mathbf{k})] &= \frac{1}{k^2} \sum_{j=1}^n \left(\sum_{i: j \in \mathcal{N}_i} c_{ij} \right)^2 \cdot \text{Var}(z_j^1) + \frac{1}{k^2} \sum_{j \neq j'} \left(\sum_{i: j \in \mathcal{N}_i} c_{ij} \right) \left(\sum_{i': j' \in \mathcal{N}_{i'}} c_{i'j'} \right) \cdot \text{Cov}(z_j^1, z_{j'}^1) + \frac{2\sigma^2 n}{k^2} \\ &= \frac{1}{k^2} \sum_{j=1}^n \left(\sum_{i: j \in \mathcal{N}_i} c_{ij} \right)^2 \left(\frac{k(n-k)}{n^2} + \frac{k(n-k)}{n^2(n-1)} \right) + \left(\frac{1}{k} \sum_{j=1}^n \sum_{i: j \in \mathcal{N}_i} c_{ij} \right)^2 \frac{-k(n-k)}{n^2(n-1)} + \frac{2\sigma^2 n}{k^2} \\ &\leq \frac{nL_{\max}^2}{k^2} \left(\frac{k(n-k)}{n^2} + \frac{k(n-k)}{n^2(n-1)} \right) + \frac{2\sigma^2 n}{k^2} \\ &\leq \frac{(n-k)}{(n-1)k} L_{\max}^2 + \frac{2\sigma^2 n}{k^2}.\end{aligned}$$

□

A.4 Bernoulli Estimator Utilizing Realized Treatment Counts

We will make use of the following lemma to bound the variance of this estimator.

Lemma 9. Suppose $X \sim \text{Binom}(n, p)$, and define

$$Y = \begin{cases} 0 & X = 0, \\ \frac{1}{X^\beta} & X > 0. \end{cases}$$

Then, $\mathbb{E}[Y] < (1 + o(1))(np)^{-\beta}$.

Proof. Using the law of total expectation, we can upper bound this expectation,

$$\begin{aligned}\mathbb{E}[Y] &\leq \Pr(X \leq (1-\delta)np) + \left(\frac{1}{(1-\delta)np}\right)^\beta \cdot \Pr(X > (1-\delta)np) \\ &\leq \Pr(X \leq (1-\delta)np) + \left(\frac{1}{(1-\delta)np}\right)^\beta.\end{aligned}\tag{A.3}$$

We apply Bernstein's inequality to compute this probability. Note that we can express

$$X = X_1 + \dots + X_n,$$

with each $X_i \sim \text{Bernoulli}(p)$. Now, define $Z = Z_1 + \dots + Z_n$ where each $Z_i = p - X_i$. Note that each $\mathbb{E}[Z_i] = 0$ and $|Z_i| \leq 1$. Thus,

$$\begin{aligned}\Pr(X \leq (1-\delta)np) &= \Pr(Z \geq \delta np) \\ &\leq \exp\left(\frac{-\frac{1}{2}(\delta np)^2}{\sum_{i=1}^n \mathbb{E}[Z_i^2] + \frac{1}{3}(\delta np)}\right) \\ &= \exp\left(\frac{-3\delta^2 n^2 p^2}{6np(1-p) + 2\delta np}\right) \\ &\leq \exp\left(\frac{-3\delta^2 np}{6+2\delta}\right).\end{aligned}$$

For $\delta = \log^{-1} n$ and large enough n , $\exp\left(\frac{-3\delta^2 np}{6+2\delta}\right) < (np)^{-2\beta}$, such that plugging into (A.3), we find $\mathbb{E}[Y] \leq ((1-\delta)np)^{-\beta} + (np)^{-2\beta} = (1+o(1))np^{-\beta}$. \square

Proof of Theorem 5. First, we reason about the bias of the estimator. We define the event \mathcal{E}_1 be the event $\{k_0 < k_1 < \dots < k_\beta\}$. By the argument from the proof of Theorem 3, $\widehat{\text{TTE}}(\hat{\mathbf{k}}/n)$ is unbiased on \mathcal{E}_1 . Thus, we can express the bias as

$$\mathbb{E}\left[\widehat{\text{TTE}}(\hat{\mathbf{k}}/n) - \text{TTE}\right] = -\Pr(\mathcal{E}_1^c) \cdot \text{TTE}.$$

However,

$$\begin{aligned}\Pr(\mathcal{E}_1^c) &= \Pr\left(\bigcup_{t=1}^{\beta} \{\hat{k}_t = \hat{k}_{t-1}\}\right) \\ &\leq \sum_{t=1}^{\beta} \Pr(\hat{k}_t = \hat{k}_{t-1}) \tag{Union Bound} \\ &= \sum_{t=1}^{\beta} \Pr(\hat{k}_t - \hat{k}_{t-1} \leq 0) \\ &\leq \sum_{t=1}^{\beta} \exp\left(\frac{-(p_t - p_{t-1})n}{2}\right) \tag{Chernoff Bound} \\ &\leq \beta \cdot \exp\left(\frac{-\Delta_{\mathbf{p}} n}{2}\right),\end{aligned}$$

so the bias decays exponentially with n .

To bound the variance, we apply the law of total variance:

$$\text{Var}\left[\widehat{\text{TTE}}\right] = \text{Var}\left[\mathbb{E}\left[\widehat{\text{TTE}} \middle| \sum_{j=1}^n z_j^t = \hat{k}_t \forall t\right]\right] + \mathbb{E}\left[\text{Var}\left[\widehat{\text{TTE}} \middle| \sum_{j=1}^n z_j^t = \hat{k}_t \forall t\right]\right]. \tag{A.4}$$

We bound these terms individually. For the first term, note that

$$\mathbb{E}\left[\widehat{\text{TTE}}(\hat{\mathbf{k}}/n) \middle| \sum_{j=1}^n z_j^t = \hat{k}_t \forall t\right] = \text{TTE} \cdot \mathbb{I}(\mathcal{E}_1).$$

which implies that,

$$\text{Var} \left[\mathbb{E} \left[\widehat{\text{TTE}}(\hat{\mathbf{k}}/n) \mid \sum_{j=1}^n z_j^t = \hat{k}_t \forall t \right] \right] = \text{TTE}^2 \cdot \text{Var}(\mathbb{I}(\mathcal{E}_1)) = \text{TTE}^2 \cdot \Pr(\mathcal{E}_1) \cdot \Pr(\mathcal{E}_1^c).$$

This term decays exponentially as n grows large, so (A.4) will be dominated by the second term.

Next, we define the event

$$\mathcal{E}_2 := \mathcal{E}_1 \cap \bigcap_{t=1}^{\beta} \left\{ |\hat{k}_t - p_t n| \leq \delta p_t n \right\}.$$

Then,

$$\begin{aligned} \Pr(\mathcal{E}_2^c) &= \Pr \left(\mathcal{E}_1^c \cup \bigcup_{t=1}^{\beta} \left\{ |\hat{k}_t - p_t n| > \delta p_t n \right\} \right) \\ &\leq \Pr(\mathcal{E}_1^c) + \sum_{t=1}^{\beta} \Pr \left(|\hat{k}_t - p_t n| \geq \delta p_t n \right) \quad (\text{Union Bound}) \\ &\leq \Pr(\mathcal{E}_1^c) + \sum_{t=1}^{\beta} \exp \left(-\frac{\delta^2 p_t n}{3} \right) \quad (\text{Chernoff Bound}) \\ &\leq \Pr(\mathcal{E}_1^c) + \beta \cdot \exp \left(-\frac{\delta^2 p_1 n}{3} \right). \end{aligned}$$

Notice that by a different application of the law of total variance, we get

$$\begin{aligned} \text{Var} \left[\widehat{\text{TTE}}(\hat{\mathbf{k}}/n) \right] &= \text{Var} \left[\mathbb{E} \left[\widehat{\text{TTE}}(\hat{\mathbf{k}}/n) \mid \mathbf{z}^t \right] \right] + \mathbb{E} \left[\text{Var} \left[\widehat{\text{TTE}}(\hat{\mathbf{k}}/n) \mid \mathbf{z}^t \right] \right] \\ &= \text{Var} \left[\frac{1}{n} \sum_{t=0}^{\beta} \sum_{i=1}^n \left(\ell_{t, \hat{\mathbf{k}}/n}(1) - \ell_{t, \hat{\mathbf{k}}/n}(0) \right) Y_i(\mathbf{z}^t) \right] + \mathbb{E} \left[\frac{\sigma^2}{n} \sum_{t=0}^{\beta} \left(\ell_{t, \hat{\mathbf{k}}/n}(1) - \ell_{t, \hat{\mathbf{k}}/n}(0) \right)^2 \right]. \end{aligned}$$

We can bound $\ell_{t, \hat{\mathbf{k}}/n}(1) - \ell_{t, \hat{\mathbf{k}}/n}(0) \leq n^{\beta}$ independently of the realized treatment counts to get

$$\mathbb{E} \left[\frac{\sigma^2}{n} \sum_{t=0}^{\beta} \left(\ell_{t, \hat{\mathbf{k}}/n}(1) - \ell_{t, \hat{\mathbf{k}}/n}(0) \right)^2 \right] \leq \frac{\beta \sigma^2}{n} \cdot n^{2\beta}.$$

$$\text{Let } \widehat{\text{TTE}}_{-\varepsilon} := \frac{1}{n} \sum_{t=0}^{\beta} \sum_{i=1}^n \left(\ell_{t, \hat{\mathbf{k}}/n}(1) - \ell_{t, \hat{\mathbf{k}}/n}(0) \right) Y_i(\mathbf{z}^t).$$

Using the fact that the variance of Bernoulli random variables is always bounded above by 1, and again using the bound $\ell_{t, \hat{\mathbf{k}}/n}(1) - \ell_{t, \hat{\mathbf{k}}/n}(0) \leq n^{\beta}$, we get

$$\begin{aligned} \text{Var} \left[\widehat{\text{TTE}}_{-\varepsilon} \right] &\leq \frac{1}{n^2} \sum_{i=1}^n \sum_{i'=1}^n \sum_{t=0}^{\beta} \sum_{t'=0}^{\beta} \sum_{\substack{S \subseteq \mathcal{N}_i \\ |S| \leq \beta}} \sum_{\substack{S' \subseteq \mathcal{N}_{i'} \\ |S'| \leq \beta}} |c_{i,S}| \cdot |c_{i',S'}| \cdot |\ell_{t, \frac{\hat{\mathbf{k}}}{n}}(1) - \ell_{t, \frac{\hat{\mathbf{k}}}{n}}(0)| \cdot |\ell_{t', \frac{\hat{\mathbf{k}}}{n}}(1) - \ell_{t', \frac{\hat{\mathbf{k}}}{n}}(0)| \\ &\leq \beta^2 \cdot Y_{\max}^2 \cdot n^{2\beta}. \end{aligned}$$

Then, to bound the second term of (A.4), we use the unconditional bound

$$\text{Var} \left[\widehat{\text{TTE}}(\hat{\mathbf{k}}/n) \right] \leq \beta^2 \cdot Y_{\max}^2 \cdot n^{2\beta} + \frac{\beta \sigma^2}{n} \cdot n^{2\beta}. \quad (\text{A.5})$$

Applying the definition of expectation, we have

$$\begin{aligned} &\mathbb{E} \left[\text{Var} \left[\widehat{\text{TTE}} \mid \sum_{j=1}^n z_j^t = \hat{k}_t \right] \right] \\ &\leq \sum_{\mathbf{k} \in \mathcal{E}_2} \Pr \left(\sum_{j=1}^n z_j^t = \hat{k}_t \forall t \right) \cdot \text{Var} \left[\widehat{\text{TTE}} \mid \sum_{j=1}^n z_j^t = \hat{k}_t \right] + \Pr(\mathcal{E}_2^c) \cdot (\beta^2 Y_{\max}^2 n^{2\beta} + \frac{\beta \sigma^2}{n} n^{2\beta}) \\ &\leq O \left(\beta^2 Y_{\max}^2 \left(\frac{d^2}{n} + \frac{\beta^2}{(1-\delta)p_1 n} \right) \cdot \left(\frac{n}{(\Delta_{\mathbf{p}} - \delta p)n} \right)^{2\beta} + \frac{\sigma^2 \beta}{n} \left(\frac{n}{(\Delta_{\mathbf{p}} - \delta p)n} \right)^{2\beta} \right) + \Pr(\mathcal{E}_2^c) \cdot (\beta^2 Y_{\max}^2 n^{2\beta} + \frac{\beta \sigma^2}{n} n^{2\beta}). \end{aligned}$$

Here, the first equality makes use of our unconditional bound on the variance, given in inequality A.5. The second inequality plugs the variance bound from Theorem 3 for the most pessimistically perturbed treatment count vector in \mathcal{E}_2 . The probability $\Pr(\mathcal{E}_2^c)$ decays exponentially in n . Therefore, choosing $\delta = \Theta(\frac{1}{\log(n)})$ and letting n get sufficiently large, the upper bound for this estimator is

$$O\left(\beta^2 Y_{\max}^2 \left(\frac{d^2}{n} + \frac{\beta^2}{p_1 n}\right) \cdot \Delta_{\mathbf{p}}^{-2\beta} + \frac{\beta \sigma^2}{n} \Delta_{\mathbf{p}}^{-2\beta}\right).$$

□

B Unbiased Estimation with Additional Observations

A natural question is whether we continue to see improvements in the estimator when we increase the number of estimates beyond $\beta + 1$. Note that we restrict our attention to unbiased estimators, as we desire the asymptotic reduction in mean-squared error as the population grows large. We may thus assess the quality of an estimator by its variance. While in general, with noisy data, more measurements may result in improved estimates, we show that in the linear setting, under perfect observations (i.e. no observation noise), these extra measurements do not help to reduce variance. In fact, we'll argue that the unbiased estimator with minimum variance is the one that ignores all but its first and last observations and then performs polynomial interpolation on these endpoints. We record this result in Theorem 10.

Theorem 10. *Suppose that the potential outcomes model is linear, and a staggered rollout Bernoulli design is implemented with a set of $T + 1$ distinct treatment probabilities $p_0 < p_1 < \dots < p_T$. Then, the unbiased estimator for TTE of the form*

$$\widehat{\text{TTE}} = \frac{1}{n} \sum_{i=1}^n \sum_{t=0}^T \alpha_t Y_i(\mathbf{z}^t)$$

that minimizes variance has $\alpha_0 = \frac{-1}{p_T - p_0}$, $\alpha_T = \frac{1}{p_T - p_0}$ and $\alpha_1, \dots, \alpha_{T-1} = 0$.

On one hand, such a result seems surprising: having more observations seems like it would only lead to a stronger estimator. However, what is overlooked is that there is strong correlation in the different measurements due to the monotonicity of treatments enforced in the staggered rollout design, such that the information in the first and last measurements contain all the useful information one could construct from the intermediate measurements. When random noise is added, the trade-off between the noise-canceling effects of additional measurements and the increased sensitivity of higher-degree interpolating polynomials adds an additional level of complexity.

Proof. To begin, we derive the constraints on $(\alpha_0, \dots, \alpha_T)$ needed to ensure unbiasedness. We have,

$$\mathbb{E}[\widehat{\text{TTE}}] = \frac{1}{n} \sum_{i=1}^n \sum_{t=0}^T \alpha_t \left(c_{i,\emptyset} + p_t \sum_{j \in \mathcal{N}_i} c_{ij} \right) = \frac{1}{n} \sum_{i=1}^n \left[c_{i,\emptyset} \left(\sum_{t=0}^T \alpha_t \right) + \sum_{j \in \mathcal{N}_i} c_{ij} \left(\sum_{t=0}^T \alpha_t p_t \right) \right].$$

Comparing to our expression for TTE in terms of the $c_{i,S}$ coefficients:

$$\text{TTE} = \frac{1}{n} \sum_{i=1}^n \sum_{\substack{S \subseteq \mathcal{N}_i \\ 1 \leq |S| \leq \beta}} c_{i,S},$$

we see that we must have,

$$\sum_{t=0}^T \alpha_t = 0, \quad \sum_{t=0}^T \alpha_t p_t = 1. \quad (\text{B.1})$$

Now, we consider the variance of this family of estimators. We have,

$$\begin{aligned}
\text{Var}[\widehat{TTE}] &= \frac{1}{n^2} \sum_{i=1}^n \sum_{i'=1}^n \sum_{t=0}^T \sum_{t'=0}^T \alpha_t \alpha_{t'} \cdot \text{Cov}[Y_i(\mathbf{z}^t), Y_{i'}(\mathbf{z}^{t'})] \\
&= \frac{1}{n^2} \sum_{i=1}^n \sum_{i'=1}^n \sum_{t=0}^T \sum_{t'=0}^T \sum_{j \in \mathcal{N}_i \cap \mathcal{N}_{i'}} \alpha_t \alpha_{t'} \cdot c_{ij} c_{i'j} \cdot (p_{\min(t,t')} - p_t p_{t'}) \\
&= \left(\frac{1}{n^2} \sum_{i=1}^n \sum_{i'=1}^n \sum_{j \in \mathcal{N}_i \cap \mathcal{N}_{i'}} c_{ij} c_{i'j} \right) \left(\sum_{t=0}^T \sum_{t'=0}^T \alpha_t \alpha_{t'} \cdot (p_{\min(t,t')} - p_t p_{t'}) \right). \quad (\text{B.2})
\end{aligned}$$

Note that the first factor is a constant depending only on the network (i.e. not on the α and p parameters of the estimator). Thus, to minimize the variance, it suffices to locate critical values of this second factor, subject to our unbiasedness constraints. We can rewrite this factor

$$\sum_{t=0}^T \alpha_t^2 \cdot p_t(1-p_t) + 2 \sum_{t=0}^T \sum_{t'=t+1}^T \alpha_t \alpha_{t'} \cdot p_t(1-p_{t'}) = \sum_{t=0}^T \alpha_t p_t \left(\alpha_t(1-p_t) + 2 \sum_{t'=t+1}^T \alpha_{t'}(1-p_{t'}) \right).$$

Then, we consider the Lagrangian,

$$\mathcal{L} := \sum_{t=0}^T \alpha_t p_t \left(\alpha_t(1-p_t) + 2 \sum_{t'=t+1}^T \alpha_{t'}(1-p_{t'}) \right) + \lambda \sum_{t=0}^T \alpha_t + \mu \left(1 - \sum_{t=0}^T \alpha_t p_t \right). \quad (\text{B.3})$$

We compute the partial derivatives of this Lagrangian with respect to each α_t as,

$$\frac{\partial \mathcal{L}}{\partial \alpha_t} = 2(1-p_t) \sum_{t'=0}^{t-1} \alpha_{t'} p_{t'} + 2p_t \sum_{t''=t}^T \alpha_{t''}(1-p_{t''}) + \lambda - p_t \mu.$$

We will set each of these partial derivatives equal to 0 sequentially to fix each of the variables at the critical point. First, we consider the partial derivative with respect to α_0 . We have,

$$\frac{\partial \mathcal{L}}{\partial \alpha_0} = 2p_0 \sum_{t''=0}^T \alpha_{t''}(1-p_{t''}) + \lambda - p_0 \mu = -p_0(2+\mu) + \lambda.$$

Here, the second inequality uses the unbiasedness constraints. Setting this partial derivative equal to 0, we must have $\lambda = p_0(2+\mu)$. Next, we consider the partial derivative with respect to α_1 :

$$\begin{aligned}
\frac{\partial \mathcal{L}}{\partial \alpha_1} &= 2\alpha_0 p_0(1-p_1) - 2p_1 \sum_{t''=1}^T \alpha_{t''}(1-p_{t''}) + \lambda - p_1 \mu \\
&= 2\alpha_0 p_0(1-p_1) + 2p_1 \left(-1 - \alpha_0(1-p_0) \right) + \lambda - p_1 \mu \quad (\text{unbiasedness}) \\
&= 2\alpha_0 p_0(1-p_1) - 2p_1 - 2p_1 \alpha_0(1-p_0) + p_0(2+\mu) - p_1 \mu \\
&= (p_0 - p_1)(2\alpha_0 + 2 + \mu).
\end{aligned}$$

Note that $p_0 - p_1 \neq 0$ by our distinct probabilities assumption. Thus, setting this partial derivative equal to 0, we must have $2 + \mu = -2\alpha_0$. In addition, combining with the previous constraint, we can re-express $\lambda = -2\alpha_0 p_0$. Next, we consider the partial derivative with respect to α_2 :

$$\begin{aligned}
\frac{\partial \mathcal{L}}{\partial \alpha_2} &= 2\alpha_0 p_0(1-p_2) + 2\alpha_1 p_1(1-p_2) - 2p_2 - 2p_2 \alpha_0(1-p_0) - 2p_2 \alpha_1(1-p_1) + \lambda - p_2 \mu \\
&= 2\alpha_0(p_0 - p_2) + 2\alpha_1(p_1 - p_2) - 2\alpha_0 p_0 - p_2(2 + \mu) \\
&= 2\alpha_0(p_0 - p_2) + 2\alpha_1(p_1 - p_2) - 2\alpha_0(p_0 - p_2) \\
&= 2\alpha_1(p_1 - p_2).
\end{aligned}$$

Setting this partial derivative equal to 0, we must have $\alpha_1 = 0$, since $p_1 - p_2 \neq 0$. We can iterate this process on the partial derivatives with respect to $\alpha_3, \dots, \alpha_T$, concluding that $\alpha_2, \dots, \alpha_{T-1} = 0$.

We are left with the system of two linear equations given by the unbiasedness constraints:

$$\alpha_0 + \alpha_T = 0, \quad \alpha_0 p_0 + \alpha_T p_T = 1.$$

The unique solution to this system is $\alpha_0 = \frac{-1}{p_T - p_0}$, $\alpha_T = \frac{1}{p_T - p_0}$. \square

C Experimental Results under a Quadratic Outcomes Model

In this section, we discuss the results of our experiments¹ under a quadratic potential outcomes model ($\beta = 2$). As in the linear setting (see Section 4), for each population size n , we sample G networks from the distribution described in Section 4. For each configuration of parameters in the experiment, we sample N treatment schedules $\{\mathbf{z}^0, \dots, \mathbf{z}^\beta\}$ from our parameterized distribution class (Bernoulli or CRD) and compute the TTE using each estimator. In the experiments for both this setting and the linear setting, we set $G = 30$ and $N = 100$.

For each estimator, we plot the relative bias of the TTE estimates averaged over the results from these GN samples and normalized by the magnitude of the TTE. The width of the shading in the figures depicts the standard deviation across the GN estimates. The experiments in the quadratic setting ran for 29.4 minutes on the same Linux machine.

In Figure 3, we visualize the effect of three network or estimator parameters on the quality of each of the five TTE estimators (the four described in the Other Algorithms paragraph of Section 4, and our CRD estimator with treatment targets $k_t = \frac{tk}{\beta}$). Specifically, we consider the effects of the population size (n), the maximum proportion of treated individuals (k/n) and the degree of the potential outcomes model (β). Each of the plots fixes two of these parameters and varies the third. Specific settings of the parameters are listed on each plot.




Figure 3: Three graphs visualizing the performance of various TTE estimators as different parameters are adjusted. The height of each graph depicts the experimental relative bias of the estimator and the shaded width depicts the experimental standard deviation.

Our estimator is the blue line with blue shading on each of the plots. As expected, the estimator is unbiased and the variance decreases as n or k/n increases. On the other hand, regardless of population size or treatment budget, the rest of the estimators remain biased. In general, the variances of these other estimators remains higher than ours, although it is worth noting that when the treatment budget k/n is lower, the variance of our estimator is higher. As the ratio r increases, the network (aka indirect) effects become greater relative to the direct effect. This is exhibited by the increase in the bias of all the estimators, besides ours, as shown in Figure 3b. As expected, when the ratio is near 0, all estimators are unbiased as this corresponds to the case where there is no network interference.

In Figure 4, we compare the variants of our estimator when $\beta = 2$, evaluating $\widehat{\text{TTE}}_{\text{PI}}(\mathbf{k}/n)$ under CRD and evaluating $\widehat{\text{TTE}}_{\text{PI}}(\mathbf{p})$ and $\widehat{\text{TTE}}_{\text{PI}}(\hat{\mathbf{k}}/n)$ under Bernoulli(\mathbf{p}) randomized design, where $p_t = tp/\beta$ and $\hat{\mathbf{k}}$ is the vector of realized treatment counts.

The estimators $\widehat{\text{TTE}}_{\text{PI}}(\mathbf{k}/n)$ and $\widehat{\text{TTE}}_{\text{PI}}(\hat{\mathbf{k}}/n)$ perform nearly identically as we vary the size of the population. They differ for lower treatment budgets, with $\widehat{\text{TTE}}_{\text{PI}}(\hat{\mathbf{k}}/n)$ having lower bias than $\widehat{\text{TTE}}_{\text{PI}}(\mathbf{k}/n)$ but about the same variance. As the treatment budget increases, they perform almost identically. $\widehat{\text{TTE}}_{\text{PI}}(\hat{\mathbf{k}}/n)$ has lower variance than $\widehat{\text{TTE}}_{\text{PI}}(\mathbf{p})$, which is intuitive as it performs polynomial interpolation on the realized treatment fraction rather than the expected treatment fraction.

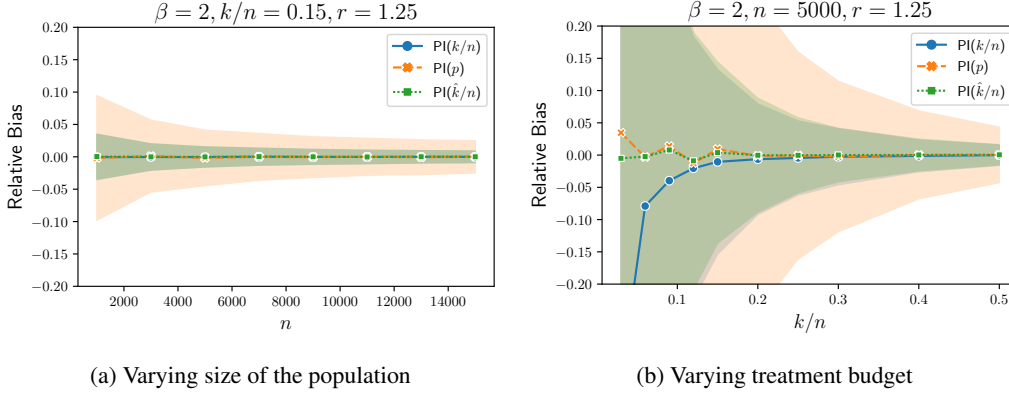


Figure 4: Two graphs visualizing the performance of our proposed TTE estimators as the size of the population (n) or treatment budget (k/n) is varied. The height of each graph depicts the experimental relative bias of the estimator and the shaded width depicts the experimental standard deviation. The blue and the green plots essentially overlap.

D Experimental Results under Bernoulli Design

We performed similar experiments to Section 4 and Appendix C for the Bernoulli randomized design setting. The main difference is that our parameterization on the budget in the realized fraction of treated individuals, k/n , has been replaced by an upper threshold on the treatment probability, p . The results we find in this Bernoulli design setting exhibit the same trends as those under completely randomized design. We include these plots for completeness and refer the reader to earlier sections for discussion and analysis.

¹Code can be found at <https://tinyurl.com/kee88h6d>

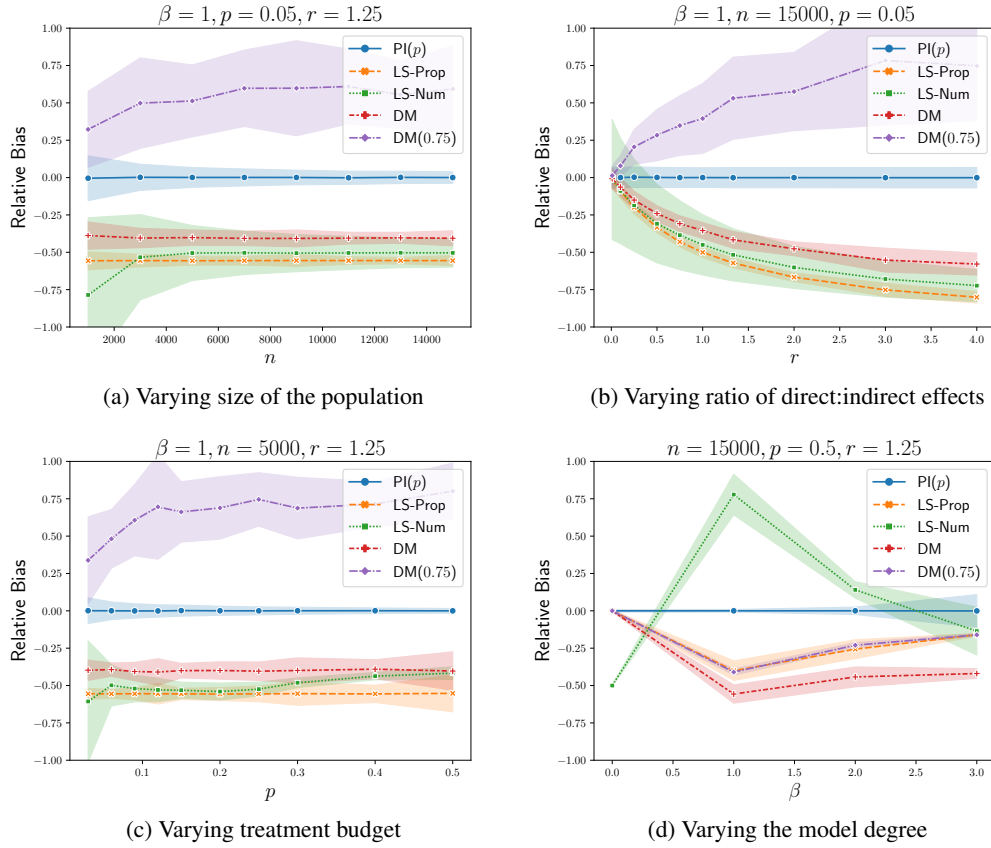


Figure 5: Four graphs visualizing the performance of various TTE estimators, under Bernoulli randomized design, as various parameters are adjusted. The height of each graph depicts the experimental relative bias of the estimator and the shaded width depicts the experimental standard deviation.

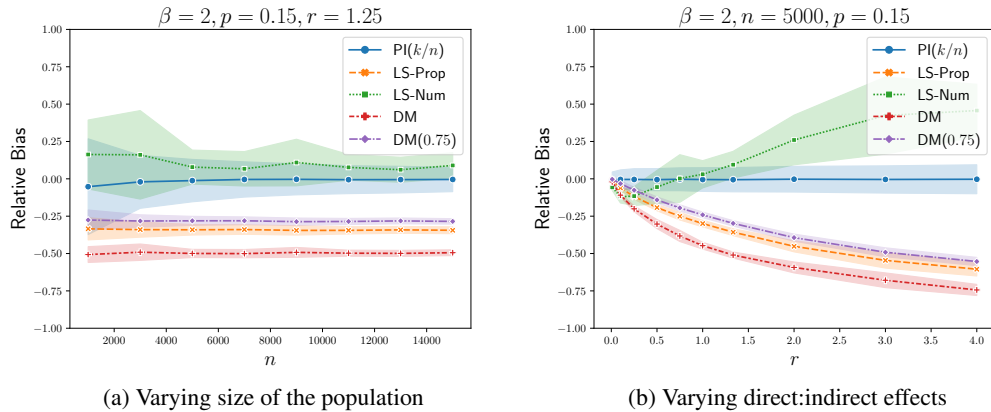


Figure 6: Two graphs visualizing the performance of our proposed TTE estimators under Bernoulli randomized design as the size of the population (n) or ratio between direct and indirect effects (r) is varied. The height of each graph depicts the experimental relative bias of the estimator and the shaded width depicts the experimental standard deviation.