

On the Performance of the Neyman Allocation with Small Pilots

Yong Cai ^{*}, Ahnaf Rafi ^{†‡}

June 5, 2024

Abstract

The Neyman Allocation is used in many papers on experimental design, which typically assume that researchers have access to large pilot studies. This may be unrealistic. To understand the properties of the Neyman Allocation with small pilots, we study its behavior in an asymptotic framework that takes pilot size to be fixed even as the size of the main wave tends to infinity. Our analysis shows that the Neyman Allocation can lead to estimates of the ATE with higher asymptotic variance than with (non-adaptive) balanced randomization. In particular, this happens when the outcome variable is relatively homoskedastic with respect to treatment status or when it exhibits high kurtosis. We provide a series of empirical examples showing that such situations can arise in practice. Our results suggest that researchers with small pilots should not use the Neyman Allocation if they believe that outcomes are homoskedastic or heavy-tailed. Finally, we examine some potential methods for improving the finite sample performance of the FNA via simulations.

JEL Classification: C21, C90

Keywords: Experiment design, treatment effect estimation, alternative asymptotics

^{*}Becker Friedman Institute, University of Chicago, Chicago, IL 60615, USA

[†]Department of Economics, Northwestern University, Evanston, IL 60208, USA

[‡]Corresponding author. Email address: ahnafrafi@u.northwestern.edu

1 Introduction

A growing literature on experiment design provides researchers with tools for reducing the asymptotic variance of their average treatment effect (ATE) estimates. Many do so in the context of two-wave experiments, where the researcher has access to a pilot study that can be used to improve the main study. Pilots are typically assumed to be large, allowing population parameters to be well-estimated. However, large pilots may not be realistic in practice. In this paper, we study the implications of small pilots for experiment design through the lens of the Neyman Allocation.

The Neyman Allocation (Neyman 1934) is a simple method for minimizing the variance of the difference-in-means estimator of ATE. In a setting without covariates, suppose that the standard deviations of the treated and control outcomes are known. The Neyman Allocation assigns more units to either treatment or control in proportion to the ratio of their standard deviations. Intuitively, the optimal experiment entails more measurements of the noisier quantity. Since the variances are not known in practice, the feasible Neyman Allocation (FNA) estimates the variances using the pilot study and then plugs the estimates into the assignment rule.

The FNA is an important part of many experiment design procedures. In the econometrics literature alone there are several notable works. Hahn et al. (2011) propose to estimate the variance of outcome and control groups conditional on covariate value, implementing the FNA conditional on covariates. In a similar vein, Tabord-Meehan (2021) employs tree-based techniques to stratify units based on their covariates. Units are then assigned to treatment and control based on the FNA conditional on strata. Cytrynbaum (2021) proposes local randomization to select representative units for participation and treatment in experiments. In what Cytrynbaum (2021) terms the “fully efficient” case, treatment proportion conditional on the randomization group is chosen by the FNA. Despite their differences, the above papers study their proposals in asymptotic frameworks that take the pilot size to infinity. Their analyses, appropriate for large pilots, essentially assume that population parameters are arbitrarily well-estimated from the pilots alone. In practice, pilots are often conducted for logistical reasons and may be small. In such settings, accurately estimating the relevant variances may be difficult. Indeed, all of the above authors caution against the use of their methods when pilot sizes are small.

To understand the implications of small pilots, we study the properties of the Neyman Allocation in an asymptotic framework that takes pilot size to be fixed even as the size of the main wave tends to infinity. In this setting, uncertainty in parameter estimation is non-negligible in the limit, and the FNA may not be the optimal allocation. In fact, we show

that the FNA can do worse than the naive, balanced allocation that assigns half the units to treatment and half to control. This occurs when outcomes have similar variances across treatment and control – that is, when outcomes are relatively homoskedastic with respect to treatment status. To assess how much homoskedasticity exists in practice, we examine the first 10 completed experiments in the AER RCT Registry. We ask the hypothetical question: if researchers conducted these studies as a two-wave experiment with a random sample from the same population, would they do better with the FNA or with balanced randomization? We find that the treatment and control groups are often highly homoskedastic across a range of outcomes and across experiments. This suggests that if faced with a small pilot, the authors of those studies would likely not have benefited from implementing the FNA. Finally, we show that as the pilot size increases, the amount of heteroskedasticity needed for the FNA to be preferable to the balanced allocation decreases, but at a rate that depends on the kurtosis of the outcome variables. Hence, even when researchers believe they are in a setting with high heteroskedasticity, they may want to avoid the FNA if they also believe that the outcomes are fat-tailed.

The above findings suggest that researchers should be cautious when designing experiments using the FNA with small pilots. However, even when pilots are large, methods which condition on many covariates may end up estimating the FNA using a small conditional sample. Furthermore, if researchers believe that units exhibit cluster-dependence – a common assumption in empirical work – the number of “effective” observations may be smaller still, impeding the estimation of the FNA.

However, our results do not imply that the FNA should never be used. Instead, researchers may consider alternative procedures with better performance under homoskedasticity. We explore three candidate solutions via simulations. One is based on testing for homoskedasticity, and the other two continuously regularize the FNA towards the balanced allocation. Our simulations suggest that the procedure we term “exponential regularization” offers the best trade-offs in that performance under heteroskedasticity is the least compromised in exchange for better performance under homoskedasticity. Alternatively, researchers might consider optimizing other aspects of the design, such as how strata are formed, or in the selection of covariates for stratification. [Bai \(2022\)](#) and [Cytrynbaum \(2023\)](#) discuss methods to robustify these approaches for small pilots.

The challenges that small pilots pose for the FNA were noted as early as [Sukhatme \(1935\)](#), which is concerned with both small pilots and small main wave experiments. Given the relative intractability of the set-up, the author concludes using simulations that the FNA performs well. Authors such as [Hahn et al. \(2011\)](#); [Tabord-Meehan \(2021\)](#); [Bai \(2022\)](#) have noted the limitations of their method in a small pilot setting. In particular, [Cytrynbaum](#)

(2021, 2023) highlights that the FNA leads to estimators with residual variation from variance estimation when pilots are small. [Cytrynbaum \(2023\)](#) further shows that estimating stratification variables using small pilots may be counterproductive. In a network setting, [Viviano \(2022\)](#) characterizes the regret of the designed experiment as a function of pilot size, providing a rule for choosing pilot size. Our paper builds on this body of work by setting up balanced randomization as the natural alternative to the FNA when pilots are small, focusing on the relative efficiency of the two methods. Our analysis specifically addresses the use of the FNA in reducing the asymptotic variance of the difference-in-means estimator in two-wave experiments. It does not speak to papers which take the treatment assignment probability as given, such as [Bai \(2022\)](#).

Our work stands in contrast to [Blackwell et al. \(2022\)](#), who argue that the FNA can perform well even when pilots are small. However, their theoretical results rely on pilot sizes going to infinity, which we consider less suitable for addressing the small pilot case. Their simulations, which consider only normally distributed outcomes, may understate the challenges arising from high kurtosis.

Our paper is closely related to papers pointing out a similar issue in the design of sequential experiments. [Melfi and Page \(1998\)](#) argue by simulation that treatment assignment rules based on estimated outcomes can do worse than non-adaptive rules due to estimation noise. Theoretical analysis is provided in [Hu and Rosenberger \(2003\)](#), in an asymptotic framework that does not nest ours. Our paper is also related to those studying small sample problems in experiments. [de Chaisemartin and Ramirez-Cuellar \(2019\)](#) are concerned with the problems small strata pose for inference. [Bruhn and McKenzie \(2009\)](#) consider the effectiveness of various randomization strategies in achieving balance when a single-wave experiment is small. Finally, we note that there is a large literature discussing the large-pilot properties of the Neyman Allocation for alternative criteria such as power (e.g. [Brittain and Schlesselman 1982](#), [Azriel et al. 2012](#)), minimax optimality (e.g. [Bai 2021](#)) or ethical considerations (e.g. Chapter 8 of [Hu and Rosenberger 2006](#)). These criteria fall outside the scope of this present paper.

The remainder of this paper is organized as follows. Section 2 presents the theoretical framework. Section 3 contains analytical and simulation results using a stylized toy example. Our main theoretical results can be found in Section 4. We assess the level of homoskedasticity in selected empirical applications in Section 5. Section 7 concludes the paper. All proofs as well as additional empirical examples are contained in the online appendix.

2 Framework

We use a standard binary treatment potential outcomes framework assuming an infinite superpopulation. The potential outcomes are $(Y(0), Y(1))$, where $Y(0)$ denotes the potential outcome under control or status quo and $Y(1)$ denotes the potential outcome under treatment or the innovation.

Assumption 2.1. Potential outcomes $(Y(0), Y(1)) \sim F$ have finite second moments. The vector of means is μ and the covariance matrix is Σ where

$$\begin{aligned}\mu &= \mathbb{E} \left[\begin{pmatrix} Y(0) \\ Y(1) \end{pmatrix} \right] = \begin{pmatrix} \mu(0) \\ \mu(1) \end{pmatrix} \\ \Sigma &= \text{Var} \left[\begin{pmatrix} Y(0) \\ Y(1) \end{pmatrix} \right] = \begin{pmatrix} \sigma^2(0) & \rho \cdot \sigma(0) \cdot \sigma(1) \\ \rho \cdot \sigma(0) \cdot \sigma(1) & \sigma^2(1) \end{pmatrix}.\end{aligned}$$

Additionally, assume potential outcomes have variances that are positive so that $\sigma^2(a) > 0$ for each $a \in \{0, 1\}$.

The estimand of interest is the Average Treatment Effect (ATE), $\theta = \mathbb{E}[Y(1) - Y(0)]$. To estimate the ATE, the experimenter conducts a two-wave experiment. The smaller first wave, also known as the pilot, is used to inform the experimenter about aspects of the design of the larger main wave (i.e. the second wave). The following assumptions about the two experimental waves will be maintained throughout the paper. For notational clarity, the tilde symbol (e.g. \tilde{X}) refers to quantities associated with the pilot.

Assumption 2.2. Potential outcomes in the pilot, denoted $\left\{ \tilde{Y}_i(0), \tilde{Y}_i(1) \right\}_{i=1}^m$, consist of m i.i.d. draws from the distribution of the random vector $(Y(0), Y(1))'$. Treatment is randomly assigned in the pilot so that denoting assignments by $\left\{ \tilde{A}_i \right\}_{i=1}^m$,

$$\left\{ \tilde{Y}_i(0), \tilde{Y}_i(1) \right\}_{i=1}^m \perp\!\!\!\perp \left\{ \tilde{A}_i \right\}_{i=1}^m.$$

Assumption 2.3. Potential outcomes in the main wave, denoted by $\{Y_i(0), Y_i(1)\}_{i=1}^n$, are n i.i.d. draws from the distribution of the random vector $(Y(0), Y(1))'$ and are independent to potential outcomes and treatment assignments in the pilot. That is,

$$\{Y_i(0), Y_i(1)\}_{i=1}^n \perp\!\!\!\perp \left\{ \tilde{Y}_i(0), \tilde{Y}_i(1), \tilde{A}_i \right\}_{i=1}^m.$$

There are various ways in which the experimenter could implement treatment assignment in the main wave. We describe a few common ones below. Note that in all of these allocation

schemes, treatment assignment is independent of potential outcomes.

- **Simple Random Assignment:** For a given $p \in (0, 1)$, let

$$A_{p,i}^{\text{rand}} \stackrel{\text{i.i.d.}}{\sim} \text{Bernoulli}(p) .$$

The associated observed outcomes in this case are denoted $Y_{p,i}^{\text{rand}} = Y_i(0) (1 - A_{p,i}^{\text{rand}}) + Y_i(1)A_{p,i}^{\text{rand}}$ and the estimator for the average treatment effect is the difference-in-means estimator:

$$\hat{\theta}_p^{\text{rand}} = \frac{\frac{1}{n} \sum_{i=1}^n Y_{p,i} A_{p,i}^{\text{rand}}}{\frac{1}{n} \sum_{i=1}^n A_{p,i}^{\text{rand}}} - \frac{\frac{1}{n} \sum_{i=1}^n Y_{p,i} (1 - A_{p,i}^{\text{rand}})}{\frac{1}{n} \sum_{i=1}^n (1 - A_{p,i}^{\text{rand}})} .$$

- **Block/Complete Randomization:** For a given $p \in (0, 1)$, let $n_1 = \lfloor np \rfloor$. Randomly assign exactly n_1 units to treatment so that all $\binom{n}{n_1}$ assignments are equally likely. In other words, let $A_p = (A_{p,1}, \dots, A_{p,n})'$ denote the random vector of treatment assignments. Then, under block/complete randomization, the distribution of A_p satisfies

$$P(A_p = \mathbf{a}) = 1 / \binom{n}{n_1} \text{ for all } \mathbf{a} \in \{0, 1\}^n \text{ such that } \sum_{i=1}^n a_i = n_1 .$$

Correspondingly, we observe $Y_{p,i} = Y_i(0) (1 - A_{p,i}) + Y_i(1)A_{p,i}$ and form the difference-in-means estimator:

$$\hat{\theta}_p = \frac{\frac{1}{n} \sum_{i=1}^n Y_{p,i} A_{p,i}}{\frac{1}{n} \sum_{i=1}^n A_{p,i}} - \frac{\frac{1}{n} \sum_{i=1}^n Y_{p,i} (1 - A_{p,i})}{\frac{1}{n} \sum_{i=1}^n (1 - A_{p,i})} .$$

Balanced randomization corresponds to choosing $p = \frac{1}{2}$. We will refer to the associated treatment assignment rule as the balanced allocation.

- **The Infeasible Neyman Allocation:** For any given $p \in (0, 1)$, it can be shown that

$$\sqrt{n} (\hat{\theta}_p - \theta) \xrightarrow{d} \mathcal{N} \left(0, \frac{\sigma^2(1)}{p} + \frac{\sigma^2(0)}{1-p} \right) .$$

The optimal choice of p to minimize the variance of the limiting distribution is the Neyman Allocation:

$$p_* = \frac{\sigma(1)}{\sigma(1) + \sigma(0)} .$$

The optimal treatment scheme is therefore A_{p_*} . The associated observed outcomes and

difference-in-means estimator are denoted by

$$Y_{p^*,i} = Y_i(0)(1 - A_{p^*,i}) + Y_i(1)A_{p^*,i}$$

$$\hat{\theta}_{p^*} = \frac{\frac{1}{n} \sum_{i=1}^n Y_{p^*,i} A_{p^*,i}}{\frac{1}{n} \sum_{i=1}^n A_{p^*,i}} - \frac{\frac{1}{n} \sum_{i=1}^n Y_{p^*,i} (1 - A_{p^*,i})}{\frac{1}{n} \sum_{i=1}^n (1 - A_{p^*,i})} .$$

Implementing the Neyman Allocation requires knowledge of the quantities $\sigma(0), \sigma(1)$ and as such is infeasible.

- **The Feasible Neyman Allocation (FNA):** One feasible implementation is to use the pilot data to form a plug-in estimator for p_* . We start by using pilot data to estimate potential outcome variances:

$$\tilde{\sigma}_m^2(a) = \frac{1}{m_a - 1} \sum_{i=1}^m \left(\tilde{Y}_i \mathbb{I} \{ \tilde{A}_i = a \} - \frac{1}{m_a} \sum_{i=1}^m \tilde{Y}_i \mathbb{I} \{ \tilde{A}_i = a \} \right)^2 ,$$

where $\tilde{Y}_i = \tilde{Y}_i(1)\tilde{A}_i + \tilde{Y}_i(0)(1 - \tilde{A}_i)$,

$$\text{and } m_a = \sum_{i=1}^m \mathbb{I} \{ \tilde{A}_i = a \} .$$

The Feasible Neyman Allocation is

$$\tilde{p} = \begin{cases} \frac{\tilde{\sigma}_m(1)}{\tilde{\sigma}_m(1) + \tilde{\sigma}_m(0)} & \text{if } \tilde{\sigma}_m(1), \tilde{\sigma}_m(0) > 0 , \\ \frac{1}{2} & \text{otherwise .} \end{cases} \quad (2.1)$$

The latter case in (2.1) (where at least one $\tilde{\sigma}_m(a) = 0$) avoids division by zero and additionally avoids the case where the entire main wave sample gets assigned to a single treatment arm. If the potential outcomes are continuously distributed, this latter case happens with zero probability. The distribution of \tilde{p} depends on m , but we omit the sample size subscript for notational convenience. The associated treatment allocation scheme is $A_{\tilde{p}}$ and the observed outcomes and difference-in-means estimator are denoted

$$Y_{\tilde{p},i} = Y_i(0)(1 - A_{\tilde{p},i}) + Y_i(1)A_{\tilde{p},i} ,$$

$$\hat{\theta}_{\tilde{p}} = \frac{\frac{1}{n} \sum_{i=1}^n Y_{\tilde{p},i} A_{\tilde{p},i}}{\frac{1}{n} \sum_{i=1}^n A_{\tilde{p},i}} - \frac{\frac{1}{n} \sum_{i=1}^n Y_{\tilde{p},i} (1 - A_{\tilde{p},i})}{\frac{1}{n} \sum_{i=1}^n (1 - A_{\tilde{p},i})} .$$

Remark 2.1. We can also define INA and FNA with simple random assignment instead of block/complete randomization. However, the latter is generally considered to be preferable to simple random assignment (see e.g. [Lachin \(1988\)](#)) and is the standard in fields such as

Development Economics (see examples in [Bugni et al. \(2018\)](#)).

3 Toy Example

In this section, we illustrate the main problems with the FNA in small samples with a toy model. In the context of this simple example, we ask the question: when does the FNA do worse than the balanced allocation? It turns out that this happens for a range of plausible values of population parameters. Section 4 describes the extension of our findings into more general settings.

We assume in this section that the potential outcomes are bivariate normal:

$$\begin{pmatrix} Y(1) \\ Y(0) \end{pmatrix} \sim N(\mu, \Sigma) .$$

Suppose we have a pilot of size m where m is even. Suppose treatment is assigned deterministically as follows:

$$\tilde{A}_i = \begin{cases} 1 & \text{if } 1 \leq i \leq \frac{m}{2} \\ 0 & \text{otherwise.} \end{cases}$$

Using standard arguments, it follows that the variance estimates are distributed as independent χ^2 random variables:

$$\tilde{\sigma}_m^2(1) \perp\!\!\!\perp \tilde{\sigma}_m^2(0) \quad , \quad \left(\frac{m}{2} - 1\right) \left(\frac{\tilde{\sigma}_m^2(a)}{\sigma^2(a)}\right) \sim \chi_{\frac{m}{2}-1}^2 .$$

Conditional on the pilot sample, the FNA assigns

$$\tilde{p} = \frac{\tilde{\sigma}_m(1)}{\tilde{\sigma}_m(0) + \tilde{\sigma}_m(1)}$$

proportion of the units in the main wave to treatment. Suppose for simplicity again these are the first $n\tilde{p}$ units (rounding $n\tilde{p}$ if necessary). Then,

$$\hat{\theta}_{\tilde{p}} = \left(\frac{1}{n\tilde{p}} \sum_{i=1}^{n\tilde{p}} Y_i(1) \right) - \left(\frac{1}{n - n\tilde{p}} \sum_{i=n\tilde{p}+1}^n Y_i(0) \right) .$$

Terms in the above expression have conditional distributions:

$$\begin{aligned} \left[\frac{1}{\sqrt{n\tilde{p}}} \sum_{i=1}^{n\tilde{p}} Y_i(1) \mid \tilde{p} \right] &\sim N(0, \sigma^2(1)) , \\ \left[\frac{1}{\sqrt{n - n\tilde{p}}} \sum_{i=n\tilde{p}+1}^n Y_i(0) \mid \tilde{p} \right] &\sim N(0, \sigma^2(0)) . \end{aligned}$$

Furthermore, they are independent. Hence,

$$\left[\sqrt{n} (\hat{\theta}_{\tilde{p}} - \theta) \mid \tilde{p} \right] \sim N \left(0, \frac{\sigma^2(1)}{\tilde{p}} + \frac{\sigma^2(0)}{1 - \tilde{p}} \right) .$$

Taking expectation over \tilde{p} , we have that:

$$\text{Var} \left[\sqrt{n} (\hat{\theta}_{\tilde{p}} - \theta) \right] = \mathbb{E} \left[\frac{\sigma^2(1)}{\tilde{p}} + \frac{\sigma^2(0)}{1 - \tilde{p}} \right] . \quad (3.1)$$

Letting \tilde{p} be constant at p , we obtain the variance of the difference-in-means estimator under simple random assignment as a special case:

$$\text{Var} \left[\sqrt{n} (\hat{\theta}_p - \theta) \right] = \frac{\sigma^2(1)}{p} + \frac{\sigma^2(0)}{1 - p} .$$

Our goal is then to compare the two expressions above when we set $p = \frac{1}{2}$. First, note that we can rewrite (3.1) as

$$\mathbb{E} \left[\frac{\sigma^2(1)}{\tilde{p}} + \frac{\sigma^2(0)}{1 - \tilde{p}} \right] = \mathbb{E} \left[\left(1 + \frac{1}{Z_m} \frac{\sigma(0)}{\sigma(1)} \right) \sigma^2(1) + \left(1 + Z_m \frac{\sigma(1)}{\sigma(0)} \right) \sigma^2(0) \right] ,$$

where $Z_m \sim \sqrt{F(\frac{m}{2} - 1, \frac{m}{2} - 1)}$.

The Neyman Allocation does worse than balanced randomization whenever the following obtains:

$$\begin{aligned} &\mathbb{E} \left[\left(1 + \frac{1}{Z_m} \frac{\sigma(0)}{\sigma(1)} \right) \sigma^2(1) + \left(1 + Z_m \frac{\sigma(1)}{\sigma(0)} \right) \sigma^2(0) \right] \geq 2\sigma^2(1) + 2\sigma^2(0) \\ \Leftrightarrow &\mathbb{E} \left[\frac{1}{Z_m} + Z_m \right] \sigma(1)\sigma(0) \geq \sigma^2(1) + \sigma^2(0) \\ \Leftrightarrow &\frac{\sigma^2(1)}{\sigma^2(0)} - \mathbb{E} \left[\frac{1}{Z_m} + Z_m \right] \frac{\sigma(1)}{\sigma(0)} + 1 \leq 0 \\ \Leftrightarrow &\frac{\sigma^2(1)}{\sigma^2(0)} - 2\mathbb{E}[Z_m] \frac{\sigma(1)}{\sigma(0)} + 1 \leq 0 . \end{aligned}$$

The final implication uses the fact that $Z_m \stackrel{d}{=} 1/Z_m$ under bivariate normality and balanced randomization. By the quadratic formula, the above inequality is satisfied if and only if

$$\frac{\sigma(1)}{\sigma(0)} \in C_m := \left[\mathbb{E}[Z_m] - \sqrt{\mathbb{E}[Z_m]^2 - 1}, \mathbb{E}[Z_m] + \sqrt{\mathbb{E}[Z_m]^2 - 1} \right]. \quad (3.2)$$

Note that reciprocal symmetry together with Jensen's inequality guarantees that the discriminant is strictly positive:

$$\mathbb{E}[Z_m]^2 = \mathbb{E}[Z_m] \mathbb{E} \left[\frac{1}{Z_m} \right] > 1.$$

Furthermore, we have that

$$\left(\mathbb{E}[Z_m] + \sqrt{\mathbb{E}[Z_m]^2 - 1} \right) \left(\mathbb{E}[Z_m] - \sqrt{\mathbb{E}[Z_m]^2 - 1} \right) = 1.$$

In other words, the interval has the form

$$C_m = \left[\frac{1}{c_m}, c_m \right],$$

where $c_m > 1$ is the upper bound in (3.2). Hence, there is a range of parameter values under which the FNA does strictly worse than balanced randomization. We first note that $1 \in C_m$ for all m . This is intuitive since $p = \frac{1}{2}$ is the infeasible Neyman Allocation when $\sigma(1)/\sigma(0) = 1$. Secondly, $x \in C_m \Leftrightarrow 1/x \in C_m$. That is, the relative performance of the FNA to the balanced allocation does not change when we relabel treatment and control.

Simulation Evidence

Given an underlying distribution, it is simple to compute C_m by Monte Carlo integration. In this subsection, we present the values of C_m for some simple models and argue that for plausible values of $\sigma(1)/\sigma(0)$, the FNA performs worse than the balanced allocation.

We start with the toy model, where $Y(1) \sim N(\mu(1), \sigma^2(1))$ and $Y(0) \sim N(\mu(0), \sigma^2(0))$. Results are shown in the left panel of Figure 1. The set of parameter values over which the FNA does worse is colored in blue. This area is larger when m is smaller. In fact, we show in Section 4 that if the data generating process (DGP) is sub-Gaussian, the length of C_m is $O(m^{-1/2})$. In the toy model, $C_{20} = [0.70, 1.43]$, while C_{50} is $[0.81, 1.23]$. Suppose instead that $Y(1) \sim \mu(1) + \sigma(1)\chi_1^2$, $Y(0) \sim \mu(0) + \sigma(0)\chi_1^2$. The right panel of Figure 1 shows that C_m is wider across the range of m . In particular, $C_{20} = [0.49, 2.22]$, while $C_{50} = [0.64, 1.61]$.




Figure 1: C_m when $Y(a) \sim \mu(a) + \sigma(a) \cdot N(0, 1)$ and $Y(a) \sim \mu(a) + \sigma(a) \cdot \chi_1^2$ for $a \in \{0, 1\}$. Monte Carlo integration using 10,000 draws.

While the intervals may appear rather narrow at first glance, we provide numerous examples in Section 5 in which the amount of heteroskedasticity falls within this range.

As our toy model shows, estimation of $\tilde{\sigma}(1)$ and $\tilde{\sigma}(0)$ causes problems when m is small. It is therefore intuitive that C_m will be wide when the distributions are fat tailed, that is, when kurtosis is high. As such, we consider the following parametrization: $Y(1) \sim \mu(1) + \sigma(1) \cdot \text{Pareto}(l, s)$ and $Y(0) \sim \mu(0) + \sigma(0) \cdot \text{Pareto}(l, s)$. Here, l and s are the location and scale parameters respectively. Figure 2 plots C_m for $l = 1$ and $s \in \{3, 4, 5\}$. Indeed, we see that the bands are much wider than in the previous models. For $s = 3$, even when $m = 50$, $C_m = [0.50, 2.15]$. Given the examples in Section 5, this appears to be fairly extreme amounts of heteroskedasticity. Finally, we note that, C_m decreases in width as we move from $s = 3$ to $s = 5$.

In sum, we see that the FNA does worse than the balanced allocation when the treatment and control groups are relatively homoskedastic. Furthermore, C_m can be quite large when m is small. The small pilot problem is exacerbated when observations exhibit cluster dependence, so that the “effective observations” are fewer in number. Small pilot issues can also arise when researchers perform stratified randomization with many strata, so that each stratum ends up with few observations. In Section 5, we argue that many empirical applications in fact have fairly homoskedastic outcomes. As such, unless a researcher has reason to believe that their outcomes are highly heteroskedastic, they should exercise caution in using the FNA with small pilots.

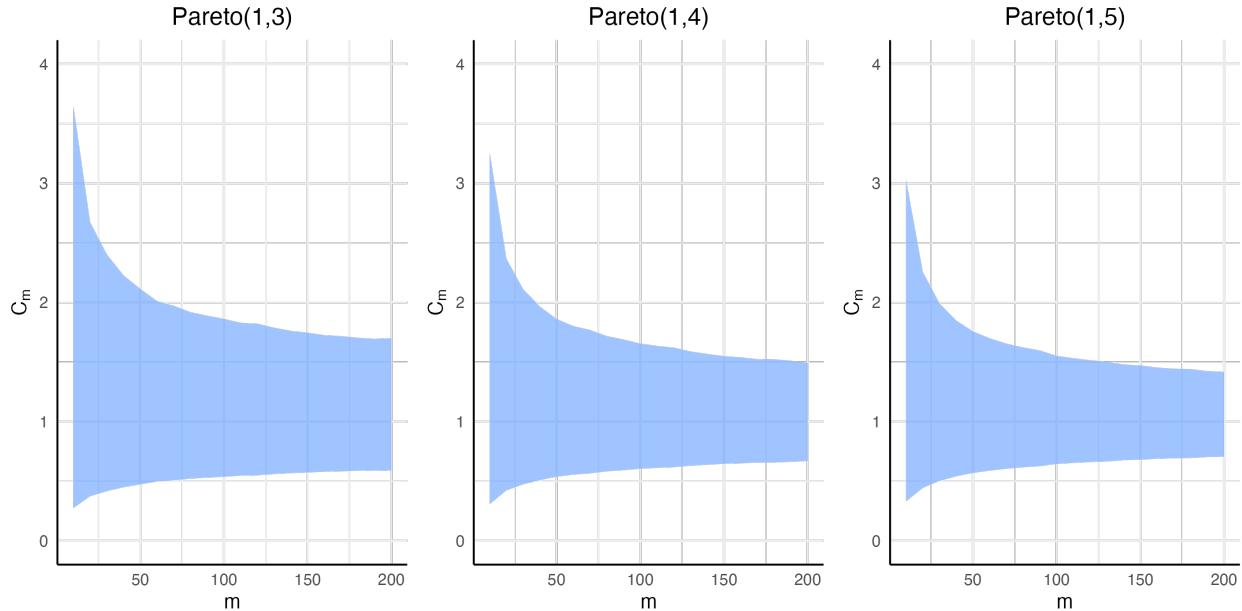


Figure 2: C_m^m when $Y(a) \sim \mu(a) + \sigma(a) \cdot \text{Pareto}(1, s)$, for $a \in \{0, 1\}$ and $s \in \{3, 4, 5\}$. Monte Carlo integration using 10,000 draws.

4 Theoretical Results

In this section, we study the theoretical properties of the FNA. We first review known results on properties of the FNA under large and small pilot asymptotics. We then present novel results in which we compare the efficiency of FNA to balanced randomization.

4.1 Review on Properties of the FNA

Large- m asymptotics

Consider an asymptotic regime where both $m, n \rightarrow \infty$, corresponding to situations where both the pilot and the main wave are large. It is straightforward to see that $\tilde{p} \xrightarrow{p} p_*$. Furthermore, we have that:

Proposition 4.1 (Hahn et al. (2011), Theorem 1). *Under Assumptions 2.1 to 2.3, as $m, n \rightarrow \infty$*

$$\sqrt{n} \left(\hat{\theta}_{\tilde{p}} - \theta \right) \xrightarrow{d} \mathcal{N} \left(0, \Sigma_* \right),$$

where

$$\Sigma_* = \frac{\sigma^2(1)}{p_*} + \frac{\sigma^2(0)}{1 - p_*} = (\sigma(1) + \sigma(0))^2.$$

The above result obtains by applying Theorem 1 in Hahn et al. (2011) to the case with only one stratum. In words, $\hat{\theta}_{\tilde{p}}$ and $\hat{\theta}_{p_*}$ have the same limiting distribution after suitable

centering and scaling. This occurs because in the limit, noise coming from estimating p_* with \tilde{p} is negligible in comparison to the sampling error of the difference-in-means estimator. Researchers employing the large m framework essentially assume that \tilde{p} is an arbitrarily good estimator for p_* so that its sampling error can be ignored. In practice, error in \tilde{p} can be large, particularly when m is small. Asymptotic approximations that do not take this into account will likely perform poorly in finite sample.

Remark 4.1. We note two minor differences in the set-up of [Hahn et al. \(2011\)](#): (1) Their difference-in-means estimator uses both main wave and pilot observations. The main wave allocation therefore needs to be adjusted since half of the pilot is treated. (2) They assume that m is of the same order as n . These differences do not affect the result.

Fixed- m Asymptotics

To better understand the behavior of the FNA under small pilots, we follow [Cytrynbaum \(2021\)](#) in studying its properties in an asymptotic framework that takes m to be fixed, even as $n \rightarrow \infty$.

A growing literature in econometrics uses fixed-sample asymptotics to study settings in which the “effective sample size” is small. For example, when data exhibit cluster-dependence, settings with few clusters pose unique challenges for estimation and inference. To tailor their analyses to these problems, papers such as [Ibragimov and Müller \(2010\)](#), [Canay et al. \(2017\)](#) and [Canay et al. \(2021\)](#) employ asymptotic frameworks in which the number of clusters is fixed in the limit. Similarly, to model difference-in-differences studies involving few treated units, [Conley and Taber \(2011\)](#) keep the number of treated units fixed even as the number of untreated units tend to infinity. In the same vein, inference and specification testing for regression discontinuity designs typically involves few observations around the discontinuity. To capture this, [Canay and Kamat \(2017\)](#) analyze a permutation test for continuity in the distribution of baseline covariates under an asymptotic regime with a fixed number of observations on either side of the discontinuity.

The small pilot approach is similar in spirit to these papers. In keeping m fixed, \tilde{p} is a noisy estimate of p_* even in the limit. Preserving this important feature of the statistical problem makes our chosen framework more appropriate for analyzing experiments with small pilots. In this setting, $\hat{\theta}_{\tilde{p}}$ converges in distribution to a mixture of Gaussians instead of $\mathcal{N}(0, \Sigma_*)$. Specifically, the form of the limiting mixture distribution depends on the distribution of \tilde{p} . This is the content of the following result due to [Cytrynbaum \(2021\)](#) (also see the discussion following equation (3.5) of [Cytrynbaum \(2023\)](#)):

Proposition 4.2 (Cytrynbaum (2021), Theorem 3.17). *Under Assumptions 2.1 to 2.3, if m remains fixed as $n \rightarrow \infty$,*

$$\sqrt{n} \left(\widehat{\theta}_{\tilde{p}} - \theta \right) \xrightarrow{d} \mathcal{L}_m ,$$

where \mathcal{L}_m is a random variable whose distribution takes the form

$$\mathbb{P}(\mathcal{L}_m \leq t) = \int_0^1 \Phi \left(\frac{t}{s(p)} \right) G_m(dp),$$

Φ is the CDF of $N(0, 1)$, G_m is the distribution of \tilde{p} and $s(\cdot)$ is defined by

$$s(p) = \sqrt{\frac{\sigma^2(1)}{p} + \frac{\sigma^2(0)}{1-p}} .$$

Remark 4.2. The above proposition can be obtained from Theorem 3.17 of Cytrynbaum (2021) by specializing it to the case with unstratified sampling and assignment ($q = 1, \psi = 1$). Conditional normality then implies a mixed normal limiting distribution by a dominated convergence argument. Also see the related intuition below.

Remark 4.3. Earlier versions of this paper present the same mixed normality result but derived under simple random assignment, which is not directly implied by Cytrynbaum (2021).

Remark 4.4. The weak limit \mathcal{L}_m in Proposition 4.2 has mean zero. To see this, note that conditional on a value of the assignment probability p , i.e. conditional on the event $\tilde{p} = p$, \mathcal{L}_m has a $N(0, s(p))$ distribution (where $s(p)$ is as defined in Proposition 4.2). The conclusion then holds by the Law of Iterated Expectations.

To see the intuition for this result, recall that for each $p \in (0, 1)$,

$$\sqrt{n} \left(\widehat{\theta}_p - \theta \right) \xrightarrow{d} N \left(0, \frac{\sigma^2(1)}{p} + \frac{\sigma^2(0)}{1-p} \right) = N(0, s^2(p)) .$$

When m is held fixed, in the limit as $n \rightarrow \infty$, \tilde{p} remains a non-degenerate random variable. In particular, the limiting distribution of \tilde{p} is its finite-sample distribution. Thus, the distribution of $\sqrt{n} \left(\widehat{\theta}_{\tilde{p}} - \theta \right)$ becomes a mixture of the marginal distributions of the process

$$\left\{ \sqrt{n} \left(\widehat{\theta}_p - \theta \right) : p \in (0, 1) \right\} ,$$

where the mixing distribution is the distribution of \tilde{p} , denoted here by G_m .

4.2 When is the FNA preferred?

Proposition 4.2 shows that when pilots are small, $\hat{\theta}_{\tilde{p}}$ has asymptotic distribution that is different from $\hat{\theta}_{p^*}$. It stands to reason that the FNA may not be optimal in this regime. In this sub-section, we compare the asymptotic variances of $\hat{\theta}_{\tilde{p}}$, $\hat{\theta}_{p^*}$ as well as $\hat{\theta}_p$ with $p = \frac{1}{2}$. We first show that $\hat{\theta}_{\tilde{p}}$ has larger asymptotic variance in the fixed- m regime than in the large- m regime. We next show that under reasonable ranges of parameter values, the asymptotic variance of $\hat{\theta}_{\tilde{p}}$ can exceed that of $\hat{\theta}_p$ with $p = \frac{1}{2}$.

Remark 4.5. It can be shown that $\hat{\theta}_{\tilde{p}}$, $\hat{\theta}_{p^*}$ and $\hat{\theta}_p$ are all unbiased estimators of θ in finite samples. Therefore, comparing asymptotic mean squared errors is equivalent to comparing the variances of the limit distributions.

We begin with the following observation:

Corollary 4.3. *Under Assumptions 2.1 to 2.3, suppose m remains fixed as $n \rightarrow \infty$. \mathcal{L}_m has variance:*

$$\mathbb{E} \left[\frac{\sigma^2(1)}{\tilde{p}} + \frac{\sigma^2(0)}{1 - \tilde{p}} \right] > \Sigma_* .$$

In words, the asymptotic variance of $\hat{\theta}_{\tilde{p}}$ is larger under the fixed- m regime than under the large- m regime. When pilots are small, uncertainty in \tilde{p} may be large and could affect the asymptotic variance of $\hat{\theta}_{\tilde{p}}$. In particular, $\hat{\theta}_{\tilde{p}}$ will not be able to attain the optimal asymptotic variance of the infeasible allocation p^* . Conventional large- m asymptotics may be too optimistic about the effectiveness of the Neyman Allocation with small pilots.

In addition to not attaining Σ_* , the $\hat{\theta}_{\tilde{p}}$ can do worse than $\hat{\theta}_p$ for certain values of $\sigma^2(1)$ and $\sigma^2(0)$, as our next two results asserts. For convenience, define the following:

Definition 4.1. Let C_m be the set such that for a pilot of size m , $\hat{\theta}_{\tilde{p}}$ has higher asymptotic variance than $\hat{\theta}_p$ if and only if $\frac{\sigma(1)}{\sigma(0)} \in C_m$. Let $|C_m|$ denote the length of C_m .

Definition 4.2. Let

$$Z_m = \frac{\tilde{\sigma}_m(1)}{\sigma(1)} \bigg/ \frac{\tilde{\sigma}_m(0)}{\sigma(0)} .$$

and

$$B_m = \frac{1}{2} \mathbb{E} \left[\frac{1}{Z_m} + Z_m \right] .$$

Here, Z_m is our estimator of the ratio of $\sigma(1)/\sigma(0)$, appropriately normalized. $B_m - 1$ is therefore the average bias of the estimators for $\sigma(1)/\sigma(0)$ and $\sigma(0)/\sigma(1)$, up to normalization. Our next proposition characterizes the region C_m in terms of B_m :

Proposition 4.4. Under [Assumptions 2.1 to 2.3](#), suppose m remains fixed as $n \rightarrow \infty$. Then

$$C_m = \left[\frac{1}{c_m}, c_m \right] .$$

Furthermore,

$$c_m = B_m + \sqrt{B_m^2 - 1} > 1 \quad \text{and} \quad |C_m| = 2\sqrt{B_m^2 - 1} > 0 .$$

The properties of C_m are intuitive. Firstly, $x \in C_m$ implies that $1/x \in C_m$. That is, the relative performance of the FNA to the balanced allocation does not change when we relabel treatment and control. Secondly, $1 \in C_m$. This is because when $\sigma(1)/\sigma(0) = 1$, the balanced allocation is optimal. Finally, note that $|C_m|$ depends on the bias of $\tilde{\sigma}_m(1)/\tilde{\sigma}_m(0)$ and $\tilde{\sigma}_m(0)/\tilde{\sigma}_m(1)$. In particular, if both terms are unbiased, $B_m = 1$ and $|C_m| = 0$. However, $|C_m|$ is strictly positive as long as \tilde{p} is not degenerate.

The exact properties of C_m depend on the underlying distributions of potential outcomes. To understand its behavior in a more general setting, we study its first-order approximation under a sub-Gaussian assumption. Recall that F is the distribution of $(Y(0), Y(1))$ in [Assumption 2.1](#).

Definition 4.3 (Sub-Gaussian). Let Y be a random variable. We say that Y is sub-Gaussian if there exists K such that

$$\mathbb{P}(|X| \geq t) \leq 2 \exp(-t^2/K^2) \quad \text{for all } t \geq 0 .$$

We say that F is sub-Gaussian if $Y(1)$ and $Y(0)$ are sub-Gaussian.

That is, a random variable is sub-Gaussian if its tails decay at least as quickly as a normal/Gaussian distribution. This yields the following:

Proposition 4.5. Under [Assumptions 2.1 to 2.3](#), suppose $n \rightarrow \infty$. Suppose additionally that F is sub-Gaussian. Then,

$$C_m = \left[1 - \sqrt{\frac{V}{m}} + \delta_m^-, 1 + \sqrt{\frac{V}{m}} + \delta_m^+ \right] ,$$

where $\delta_m^+, \delta_m^- = o\left(\frac{1}{\sqrt{m}}\right)$ and

$$V = \frac{1}{4} \left(\frac{\mathbb{E}[(Y_i(1) - \mu(1))^4]}{\sigma^4(1)} + \frac{\mathbb{E}[(Y_i(0) - \mu(0))^4]}{\sigma^4(0)} - 2 \right) .$$

Provided that the potential outcomes are sub-Gaussian, the relative efficiency of $\widehat{\theta}_{\widetilde{p}}$ and $\widehat{\theta}_p$ under $p = 1/2$ is, to a first order, determined by the kurtosis of $Y_i(1)$ and $Y_i(0)$. Intuitively, if the potential outcomes have fatter tails, \widetilde{p} will be poorly estimated, leading to larger variance in $\widehat{\theta}_{\widetilde{p}}$.

Furthermore, $|C_m|$ is shrinking to 0 at the rate $1/\sqrt{m}$. Letting $m \rightarrow \infty$, $\widehat{\theta}_{\widetilde{p}}$ has weakly lower asymptotic variance across all parameter values. Hence, we recover the classic result concerning the optimality of the Neyman Allocation. When m is small, however, C_m can be wide, as Proposition 4.4 suggests. As we will argue in Section 5, many empirical applications have $\sigma(1)/\sigma(0)$ close to 1, so that for small m , these ratios fall within the range in which balanced randomization is preferred.

While the sub-Gaussian assumption limits the applicability of Proposition 4.5, it covers binary and bounded outcomes, which are relevant in empirical work. Furthermore, we consider it to be a negative result: even when potential outcomes are well-behaved, the FNA is sensitive to the kurtosis of the potential outcomes. It can still perform poorly relative to the balanced allocation as a result.

4.3 Efficiency Loss from the FNA

The previous section showed that the FNA can perform worse than the balanced allocation when the pilots are small. This raises the question: what is the efficiency loss from using the FNA in such scenarios? To answer this question, consider the following criteria:

Definition 4.4 (Efficiency Loss of FNA under F). For $(Y(0), Y(1)) \sim F$, denote

$$\begin{aligned} \mathcal{L}^d(F) &= \mathbb{E} \left[\frac{\sigma^2(1)}{\widetilde{p}} + \frac{\sigma^2(0)}{1-\widetilde{p}} \right] - 2(\sigma^2(1) + \sigma^2(0)) \ , \\ \mathcal{L}^r(F) &= \mathbb{E} \left[\frac{\sigma^2(1)}{\widetilde{p}} + \frac{\sigma^2(0)}{1-\widetilde{p}} \right] \Big/ 2(\sigma^2(1) + \sigma^2(0)) \ . \end{aligned}$$

In words, we measure efficiency loss using either the difference or the ratio of the asymptotic variances arising from the FNA and the balanced allocation. Larger values of \mathcal{L}^d and \mathcal{L}^r indicate that the FNA is doing worse relative to balanced randomization.

Clearly, the efficiency losses depend on F . For a given F , the two criteria are equivalent in the sense that the difference criterion is positive if and only if the ratio criterion exceeds one. Depending on the set-up, one will be more convenient to work with than the other. From Definition 4.4, it is straightforward to see that:

Proposition 4.6. *Under Assumptions 2.1 to 2.3, suppose m remains fixed as $n \rightarrow \infty$. Then*

$$\begin{aligned}\mathcal{L}^d(F) &= 2(B_m - 1)\sigma(1)\sigma(0) - (\sigma(1) - \sigma(0))^2, \\ \mathcal{L}^r(F) &= \frac{1}{2} + B_m \cdot \frac{\sigma(1)\sigma(0)}{\sigma^2(1) + \sigma^2(0)}.\end{aligned}$$

Furthermore, if F is sub-Gaussian, then

$$\begin{aligned}\mathcal{L}^d(F) &= \frac{V}{m} \cdot \sigma(1)\sigma(0) - (\sigma(1) - \sigma(0))^2 + \delta_m^d, \\ \mathcal{L}^r(F) &= \frac{1}{2} + \left(1 + \frac{V}{2m}\right) \cdot \frac{\sigma(1)\sigma(0)}{\sigma^2(1) + \sigma^2(0)} + \delta_m^r,\end{aligned}$$

where $\delta_m^d, \delta_m^r = o(1/m)$.

The above proposition shows us how the efficiency loss from using the FNA depends on B_m as well as the amount of heteroskedasticity in F . To better understand the trade-off from the two terms, suppose we fix B_m and reparametrize $\sigma(0) = 1/\sqrt{1 + \rho^2}$, $\sigma(1) = \rho/\sqrt{1 + \rho^2}$. Now let us consider the effect of changing ρ . This is the thought experiment in which we change the relative scaling of $\varepsilon(1) = Y(1) - \mu(1)$ and $\varepsilon(0) = Y(0) - \mu(0)$ while keeping their functional forms fixed. Additionally, this entails keeping $\text{AVar}(\hat{\theta}_{\frac{1}{2}}) = \text{Var}(Y(1)) + \text{Var}(Y(0))$ fixed at 2. This scaling is useful for discussing the properties of \mathcal{L}^d since it scales linearly with $\sigma(0)$ and $\sigma(1)$. \mathcal{L}^r is invariant to the scale of $\sigma(0)$ and $\sigma(1)$. With this parametrization, we have that:

$$\frac{\partial \mathcal{L}^d}{\partial \rho} = \frac{2B_m(1 - \rho^2)}{(1 + \rho^2)^2} \quad \text{and} \quad \frac{\partial \mathcal{L}^r}{\partial \rho} = \frac{B_m(1 - \rho^2)}{(1 + \rho^2)^2}.$$

The above expressions are intuitive. When ρ exceeds 1, increasing ρ amounts to increasing heteroskedasticity. FNA is more preferable to balanced randomization and the ratio loss decreases.

Furthermore, it is straightforward to see that both \mathcal{L}^d and \mathcal{L}^r are maximized when $\rho = 1$, with the respective values of B_m and $\frac{1}{2}(1 + B_m)$. On the other hand, \mathcal{L}^d and \mathcal{L}^r have minimum values -1 and $\frac{1}{2}$ respectively, which are obtained when either $\rho = 0$ or $\rho \rightarrow \infty$. In summary, for a given functional form of $\varepsilon(0)$ and $\varepsilon(1)$, the loss from using FNA depends on B_m but the maximum possible gain does not. This asymmetry foreshadows the next result.

Since F is unknown, we proceed by comparing the supremum and infimum of $\mathcal{L}^d(F)$ and $\mathcal{L}^r(F)$ (i.e. maximum losses and gains) over appropriate classes of DGPs. To that end, consider the following.

Definition 4.5. For $K < \infty$, let

$$\mathcal{F}(K) = \{F \text{ s.t. } 0 < \sigma^2(0), \sigma^2(1) \leq K\},$$

where it should be understood that $\sigma^2(0), \sigma^2(1)$ both depend on F . Furthermore, let

$$\mathcal{F}(\infty) = \{F \text{ s.t. } 0 < \sigma^2(0), \sigma^2(1) < \infty\}.$$

Then we have that:

Proposition 4.7. Under *Assumptions 2.1 to 2.3*, suppose m remains fixed as $n \rightarrow \infty$. Then,

$$\begin{aligned} \sup_{F \in \mathcal{F}(K)} \mathcal{L}^d(F) = \infty \quad , \quad \inf_{F \in \mathcal{F}(K)} \mathcal{L}^d(F) = -K \quad , \\ \sup_{F \in \mathcal{F}(\infty)} \mathcal{L}^r(F) = \infty \quad , \quad \inf_{F \in \mathcal{F}(\infty)} \mathcal{L}^r(F) = \frac{1}{2} \quad . \end{aligned}$$

In words, the worst case loss from using the FNA – whether measured as a difference or as a ratio relative to the variance of the balanced allocation – is unbounded. On the other hand, the gains are bounded. Most strikingly, the fact that $\inf_{F \in \mathcal{F}(\infty)} \mathcal{L}^r(F) = \frac{1}{2}$ suggests that the best a researcher can do with the FNA is to obtain main wave confidence intervals for $\hat{\theta}$ that are $1/\sqrt{2} \approx 70\%$ the length of that from balanced randomization. This comes at the cost of potentially doing much worse as a result of noise from the pilot estimation.

We can also interpret the result in terms of regret. Suppose a researcher is constrained to either using the FNA or the balanced allocation. If their utility over the second stage estimator is given by $U(\hat{\theta}) = -(\hat{\theta} - \theta)^2$, then $\mathcal{L}^d(F)$ is the asymptotic regret of choosing the FNA, while $-\mathcal{L}^d(F)$ is the asymptotic regret of choosing the balanced allocation. A similar argument can be made for $\mathcal{L}^r(F)$ with $U(\hat{\theta}) = -\log(\hat{\theta} - \theta)$. Proposition 4.7 then says that the regret of using the FNA over the balanced allocation is unbounded, while the regret of using the balanced allocation over the FNA is not. The balanced allocation is therefore preferred in terms of maximum regret. This is admittedly a very conservative way of choosing an estimation procedure. However, it is coherent in the following sense. Researchers may prefer the FNA if they want to guard against large amounts of heteroskedasticity in the population. However, if researchers are being agnostic about $\sigma(1)/\sigma(0)$, they should also be agnostic about F , in which case the balanced allocation may ultimately be preferable.

Remark 4.6. Proposition 4.7 continues to hold if we further restrict F so that $Y(0)$ and $Y(1)$ have uniformly bounded a^{th} moments for any a .

Remark 4.7. In the proof of Proposition 4.7, we construct a sequence of DGPs over which the variance of the $\hat{\theta}_{\bar{p}}$ diverge. This implies that $\hat{\theta}_{\bar{p}}$ is not uniformly consistent with large pilots.

Remark 4.8. It is also known in the large pilot case that $\inf_{F \in \mathcal{F}(\infty)} \mathcal{L}^r(F) = \frac{1}{2}$. See for example Obradovic (2023) and Zhao (2023).

4.4 Simulation Evidence

The efficiency losses recorded in Proposition 4.7 are pessimistic and may not be relevant in practice. To better appreciate the losses that researchers might experience, we next present mean-squared errors for $\hat{\theta}_{\bar{p}}$ under various data-generating processes, computed via Monte Carlo simulations. Our choice of parameter values is inspired by Blackwell et al. (2022).

Let $Y(1) = 0.075 + \sigma(1)\varepsilon(1)$ and $Y(0) = \sigma(0)\varepsilon(0)$. We are interested in the MSE of $\hat{\theta}_{\bar{p}}$ as $\sigma(1)/\sigma(0)$ varies over $[0, 1]$. We do not need to consider values greater than 1 given the reciprocal symmetry. We set $\sigma^2(1) + \sigma^2(0) = 2$ so that the variance of $\hat{\theta}_{\frac{1}{2}}$ is fixed. As in Section 3, consider the following models:

Model 1: $\varepsilon(1), \varepsilon(0) \sim N(0, 1)$

Model 2: $\varepsilon(1), \varepsilon(0) \sim (\chi_1^2 - 1)/\sqrt{2}$

Model 3: $\varepsilon(1), \varepsilon(0) \sim (\text{Pareto}(1, 3) - \frac{3}{2})/\sqrt{3/4}$

Model 4: $\varepsilon(1), \varepsilon(0) \sim (\text{Pareto}(1, 4) - \frac{4}{3})/\sqrt{2/9}$

Model 5: $\varepsilon(1), \varepsilon(0) \sim (\text{Pareto}(1, 5) - \frac{5}{4})/\sqrt{5/48}$

where $\varepsilon(1)$ and $\varepsilon(0)$ are standardized to have mean 0 and variance 1. We set the pilot to have $m = 20$, while main wave has $n = 1000$. We run the simulations with 100,000 draws for each set of parameter values.

Results for Models 1 and 2 are presented in Figure 3, while those for Models 3, 4 and 5 are presented in Figure 4. First note that the MSE curves follow the theoretical description provided in Section 4.2. The MSE of the FNA is monotonically increasing in $\sigma(1)/\sigma(0)$. As $\sigma(1)/\sigma(0) \rightarrow 0$, the MSE of $\hat{\theta}_{\bar{p}}$ approaches 1/2 that of $\hat{\theta}_{\frac{1}{2}}$. As $\sigma(1)/\sigma(0) \rightarrow 1$, the MSE of $\hat{\theta}_{\bar{p}}$ exceeds that of $\hat{\theta}_{\frac{1}{2}}$ by B_m .

Next, consider Model 1. With a normal distribution, the MSE of the FNA is 2.9% larger than the balanced allocation when $\sigma(1)/\sigma(0) = 1$. Weighed against a potential reduction of 50%, the FNA is definitely preferred in this setting. However, the picture becomes less clear once we move to Model 2. Here, the MSE of the FNA is around 17% larger than that of the

balanced allocation. This is true up to $\sigma(1)/\sigma(0) = 0.75$, which is more heteroskedasticity in many of the outcomes considered in our empirical examples. Similar results obtain with Models 3, 4 and 5, in which the FNA has MSEs that are 26%, 20% and 17% more than of the balanced allocation when $\sigma(1)/\sigma(0) = 1$. The above results suggest that the FNA may not be advisable when $m = 20$. On the other hand, the worst-case efficiency loss shrinks quickly as m increases, so that the FNA performs favorably. Results for $m = 50$ can be found in Appendix C.1. At $m = 50$, the loss in efficiency over Models 1 to 5 are, respectively, 1.5%, 6.4%, 13%, 11% and 8.6%.

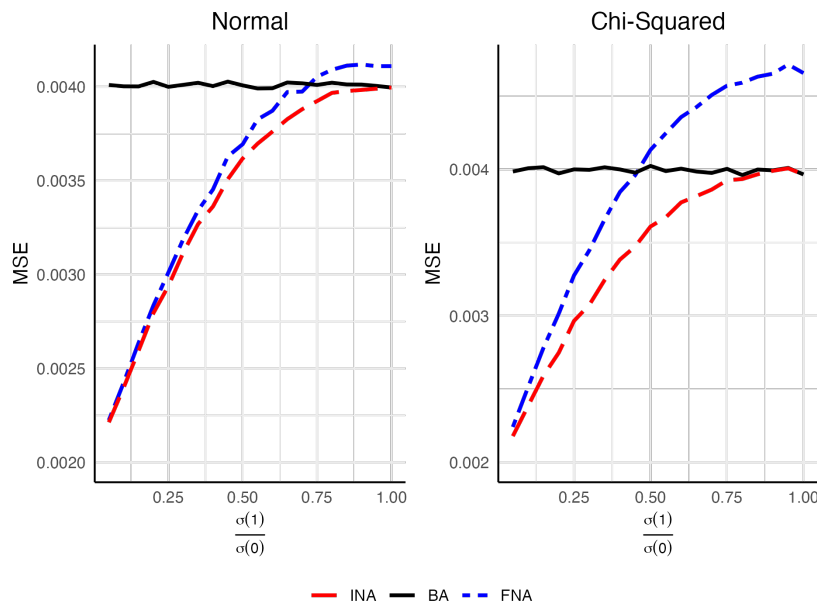


Figure 3: MSE under the Infeasible Neyman Allocation (INA), Balanced Allocation (BA) and Feasible Neyman Allocation (FNA) when $\varepsilon(1), \varepsilon(0)$ has the standard Normal distribution and the standardized Chi-Squared distribution. $m = 20$.

5 Empirical Evidence of Approximate Homoskedasticity

To assess the amount of heteroskedasticity that empirical researchers face, we revisit the first 10 completed experiments in the AER RCT Registry. In each experiment, we ask the following question: suppose the authors had access to a small pilot prior to the main study, would they have done better using the FNA instead of balanced randomization? In each experiment presented, we use the full experimental sample to estimate the standard deviations of each treatment arm and compute the corresponding ratios to see if these are close to one. In practice, researchers cannot do this given a small pilot since they do not have access to consistent variance estimators. Our findings suggest that these authors would likely




Figure 4: MSE under the Infeasible Neyman Allocation (INA), Balanced Allocation (BA) and Feasible Neyman Allocation (FNA) when $\varepsilon(1), \varepsilon(0)$ has the standardized Pareto(1, s) distribution, with $s \in \{3, 4, 5\}$. $m = 20$.

not have done better with the FNA. We present two examples in this section: [Avvisati et al. \(2014\)](#) conduct an experiment in which the outcomes are relatively homoskedastic. [Ashraf et al. \(2006\)](#) study outcome variables which are heteroskedastic. In this latter example, we also provide estimates of the interval C_m and show that it will be wide even when m is large. The remaining eight experiments are qualitatively similar to [Avvisati et al. \(2014\)](#) and can be found in Section B of the online appendix.

5.1 [Avvisati et al. \(2014\)](#)

There is a significant body of research examining the impact of school-level factors such as class size or teacher quality on educational performance of students. These are typically seen as the primary instruments for educational policy intervention. A large body of work also examines the impact of parental inputs on educational outcomes. [Avvisati et al. \(2014\)](#) study whether or not parental inputs can be effectively manipulated through simple participation programs at schools. They do so via a large-scale randomized control trial in middle schools in the Créteil educational district of Paris. The experiment targeted families of 6th graders and the program consisted of a sequence of three meetings with parents every 2–3 weeks. The sessions focused on how parents can help their children by participating at school and at home in their education and additionally included advice on how to adapt to results in end-of-term report cards. Participation in the program was randomized at the class level

– half of the classes at a given school were assigned to the participation program. Classes were groups of 20–30 students. The overall sample comprised 183 classes and a total of 4,308 students. The study tracked three types of outcomes: (1) parental involvement attitudes and behavior; (2) children’s behavior, namely truancy, disciplinary record and work effort; and (3) children’s academic results. Since randomization was done at the class-level, we examine heteroskedasticity with respect to treatment status at both the individual level and at the class level.

Table 1 reports student-level standard deviations in treatment and control groups, as well as their ratios, for the main outcomes of interest. These are the outcomes considered in Tables 2, 3 and 5 in [Avvisati et al. \(2014\)](#). The ratios are all close to one, so that by and large the treatment and control groups are relatively homoskedastic at the student level. This indicates that if the experimenters were to run a randomized control trial in the same population with treatment assigned at the individual-level, the FNA would likely yield no improvement relative to the balanced allocation.

Table 2 reports standard deviations and their ratios in class-level means of outcomes. This corresponds to a scenario in which classes are the units of interest, with class-level means as the relevant outcomes. We first calculate class-level means and then compute their standard deviations across classes for the treatment and control groups respectively. The standard deviation ratios for class-level means are by and large also close to one. Hence, if the hypothetical experiment were to be conducted at the class-level, the treatment and control groups would still be relatively homoskedastic. In this case, the FNA would again not improve upon the balanced allocation.

5.2 [Ashraf et al. \(2006\)](#)

A large body of economic models posit that individuals have time inconsistent preferences, exhibiting more impatience for near-term trade-offs than for future trade-offs. The implication of these models is that those who engage in commitment devices ex ante may improve their welfare. To test this hypothesis, [Ashraf et al. \(2006\)](#) conducted an RCT in the Philippines, in which individuals were randomly offered the chance to open a SEED (Save, Earn, Enjoy Deposits) account. Money deposited into the account cannot be withdrawn until the owner reached a goal, such as reaching a savings amount or until a pre-specified month in which they expected large expenditures.

Partnering with a rural bank in Mindanao, the authors surveyed 1,777 of their existing or former clients, of which 842 were placed into the treatment group, while 469 were placed in the control group. As treatment involved receiving a briefing on the importance of savings,

Table 1: Student-Level Heteroskedasticity in [Avvisati et al. \(2014\)](#).

	Outcome Variables	Treatment	Control	Treat./Cont.
Parental Involvement	Global parenting score	0.34	0.34	1.01
	School-based involvement score	0.66	0.63	1.05
	Home-based involvement score	0.59	0.57	1.04
	Understanding and perceptions score	0.53	0.55	0.97
	Parent-school interaction	0.40	0.40	1.01
	Parental monitoring of school work	0.43	0.41	1.05
Behavior	Absenteeism	6.29	8.63	0.73
	Pedagogical team: Behavioral score	0.73	0.74	0.98
	Pedagogical team: Discipl. sanctions	1.20	1.18	1.02
	Pedagogical team: Good conduct	0.49	0.47	1.04
	Pedagogical team: Honors	0.28	0.32	0.89
	Teacher assessment: Behavior in class	0.48	0.49	0.98
	Teacher assessment: School work	0.49	0.50	1.00
Test Scores	French (Class grade)	3.73	3.70	1.01
	Mathematics (Class grade)	4.26	4.25	1.00
	Average across subjects (Class grade)	2.87	2.88	1.00
	Progress over the school year	0.49	0.49	0.99
	French (Uniform test)	0.99	1.01	0.98
	Mathematics (Uniform test)	0.99	1.02	0.98

Table 2: Class-Level Heteroskedasticity in [Avvisati et al. \(2014\)](#).

	Outcome Variables	Treatment	Control	Treat./Cont.
Parental Involvement	Global parenting score	0.12	0.12	0.97
	School-based involvement score	0.24	0.35	0.69
	Home-based involvement score	0.20	0.15	1.33
	Understanding and perceptions score	0.20	0.22	0.94
	Parent-school interaction	0.13	0.12	1.09
	Parental monitoring of school work	0.13	0.14	0.96
Behavior	Absenteeism	2.21	3.45	0.64
	Pedagogical team: Behavioral score	0.24	0.27	0.88
	Pedagogical team: Discipl. sanctions	0.36	0.36	1.01
	Pedagogical team: Good conduct	0.22	0.23	0.95
	Pedagogical team: Honors	0.10	0.11	0.87
	Teacher assessment: Behavior in class	0.16	0.16	1.01
	Teacher assessment: School work	0.15	0.14	1.02
Test Scores	French (Class grade)	1.32	1.31	1.01
	Mathematics (Class grade)	1.76	1.84	0.96
	Average across subjects (Class grade)	0.83	0.93	0.89
	Progress over the school year	0.16	0.14	1.11
	French (Uniform test)	0.42	0.42	1.01
	Mathematics (Uniform test)	0.40	0.43	0.93

the remaining 466 individuals were placed in the marketing group, receiving the briefing but not access to SEED. We focus on the Table VI of [Ashraf et al. \(2006\)](#), containing results on saving behavior. The parameter of interest is the Intent-to-Treat effect, with approximately 25% of the treated taking up treatment. Here, the authors find that relative to the control group, the treatment group had a higher change in savings 6 months (6m) and 12 months (12m) after treatment. Comparing the treatment to the marketing group led to weaker but still positive results.

We present standard deviations of the outcomes as well as their ratios in Table 3. The outcomes of interest are

1. Change in Total Balance (6m) ($\Delta \text{ Tot. Bal. (6m)}$),
2. Change in Total Balance (12m) ($\Delta \text{ Tot. Bal. (12m)}$),
3. Change in Total Balance exceeds 0% (12m) ($\Delta \text{ Tot. Bal. } > 0\% \text{ (12m)}$),
4. Change in Total Balance exceeds 20% (12m) ($\Delta \text{ Tot. Bal. } > 20\% \text{ (12m)}$).

We first note that $\Delta \text{ Tot. Bal. } > 0\% \text{ (12m)}$ and $\Delta \text{ Tot. Bal. } > 20\% \text{ (12m)}$ are binary outcomes which are relatively homoskedastic. $\Delta \text{ Tot. Bal. (6m)}$ and $\Delta \text{ Tot. Bal. (12m)}$, measured in Philippine pesos, exhibit more heteroskedasticity. In particular, comparing the treatment group to the marketing group at the 12 month period, we observe a standard deviation ratio of 3.13. At first glance, this suggests that the FNA might outperform balanced randomization, at least with respect to this specific outcome. This turns out to be false once we investigate other features of $\Delta \text{ Tot. Bal. (6m)}$ and $\Delta \text{ Tot. Bal. (12m)}$. Quantiles of these variables are displayed in Table 4. Clearly, they have extremely fat right tails, which we confirm by computing the kurtosis, contained in Table 5. Fat tails worsen the performance of a variety of statistical techniques, including the FNA, as our analysis in Section 4 shows.

Table 3: Heteroskedasticity in [Ashraf et al. \(2006\)](#)

	Treat.	Cont.	Market.	Treat./Cont.	Treat./Market.
$\Delta \text{ Tot. Bal. (6m)}$	2347.60	2880.70	1335.98	1.76	0.81
$\Delta \text{ Tot. Bal. (12m)}$	6093.24	1945.00	2690.65	2.26	3.13
$\Delta \text{ Tot. Bal. } > 0\% \text{ (12m)}$	0.47	0.45	0.42	1.11	1.05
$\Delta \text{ Tot. Bal. } > 20\% \text{ (12m)}$	0.40	0.35	0.31	1.30	1.16

Researchers cannot estimate C_m if they only have access to a small pilot. However, this is feasible with data from the full experiment. To do so, we draw samples of size m with replacement and compute Z_m for each draw. We then take the average of $Z_m + 1/Z_m$

Table 4: Quantiles of Outcome Variables in [Ashraf et al. \(2006\)](#).

Variable	Group	1%	5%	10%	50%	90%	95%	99%	99.5%	99.9%
Δ Tot. Bal. (6m)	Treat.	-1100	-500	-300	0	500	1500	7200	13100	28900
	Market.	-1000	-600	-400	0	100	900	5600	12200	40800
	Cont.	-1600	-600	-500	0	0	600	2700	6100	18400
Δ Tot. Bal. (12m)	Treat.	-1300	-900	-500	0	500	1600	8500	18100	102300
	Market.	-1300	-900	-500	-100	300	1600	10500	15700	19900
	Cont.	-2000	-1200	-800	-100	100	900	6500	8100	34300

Table 5: Kurtosis of the Outcome Variables in [Ashraf et al. \(2006\)](#)

Variable	Treat.	Market.	Cont.
Δ Tot. Bal. (6m)	252.78	218.92	156.33
Δ Tot. Bal. (12m)	258.56	66.56	309.89

over 10,000 draws to evaluate B_m , essentially evaluating the expectation with respect to the empirical CDF. Proposition 4.4 then allows us to compute C_m . We evaluate C_m for $m \leq 2000$, using grids of size 20. Results are displayed in Figure 5. Given the high kurtosis, there is a relatively large range of ratios of standard deviations for which the balanced allocation is preferred to the FNA.

We can further compute the m necessary before the FNA outperforms the balanced allocation. Given the upper and lower bounds in Figure 5, we smooth over the grid points using cubic interpolation. We then find the minimum m at which the either bound attains the values in Table 3. These are the ‘‘Exact’’ intervals presented in Table 6. Here, we see that the necessary pilot sizes are between 25–50% of the full experiment. For the comparison of the treatment and marketing groups at 6 months, the necessary pilot size exceeds 2,000, falling outside the set of grid points we explored. To complete the analysis, we use Proposition 4.5 to obtain approximations of the necessary m . Specifically, we estimate V using their sample analogs computed on the full experiment:

$$\hat{V} = \frac{\frac{1}{n_1} \sum_{A_i=1} (Y_i - \bar{Y}(1))^4}{\left(\frac{1}{n_1-1} \sum_{A_i=1} (Y_i - \bar{Y}(1))^4\right)^2} + \frac{\frac{1}{n_0} \sum_{A_i=0} (Y_i - \bar{Y}(0))^4}{\left(\frac{1}{n_0-1} \sum_{A_i=0} (Y_i - \bar{Y}(0))^4\right)^2} - 2,$$

where n_a is the number of units with treatment status a in the full experiment and $\bar{Y}(a) =$

$\frac{1}{n_a} \sum_{A_i=a} Y_i$. We can then estimate m by

$$\hat{m} = \frac{\hat{V}}{\left(1 - \frac{\hat{\sigma}(1)}{\hat{\sigma}(0)}\right)^2}.$$

Comparing the asymptotic intervals to the exact ones, we see that the former is far too optimistic for the fat-tailed DGP in [Ashraf et al. \(2006\)](#). This is likely because the sub-Gaussian assumption is inappropriate for this data. Nonetheless, applying the asymptotic bounds to the case of the treatment vs marketing groups at 6 months, we find that a pilot size of 7,000 would be necessary for the FNA to outperform balanced randomization. This is nearly 4 times the size of the actual experiment.

Table 6: Necessary Pilot Sizes

	Treat./Cont.	Treat./Market.
Exact 6m	930.70	–
Exact 12m	1014.94	475.29
Asympt. 6m	355.02	6857.59
Asympt. 12m	177.10	35.52

In sum, our analysis suggests that the FNA would have performed poorly in the context of [Ashraf et al. \(2006\)](#). Even though the outcomes of interest exhibit stronger heteroskedasticity, the fat tails of the outcome distributions also impedes the estimation of \tilde{p} , so that ultimately, very large pilot sizes are needed for the FNA to outperform the balanced allocation in this example.

6 Potential Solutions

In this section, we use simulations to explore methods that can improve the performance of the FNA with small pilots. For definitions of the models, see [Section 4.4](#). We consider in turn (1) testing for homoskedasticity, (2) additive regularization and (3) exponential regularization. We conclude with a comparison of these results and speculate that exponential regularization is preferable to the other two methods. Additional simulation results can be found in [Appendix C.2](#).




Figure 5: C_m for the outcomes of interest estimated using the full experiment data. Interval computed under the assumption that the pilot assigns half of the units to treatment and half to control.

6.1 Testing for Homoskedasticity

Since balanced randomization is preferred when outcomes are relatively homoskedastic, consider the following two-step procedure. First test for homoskedasticity in the pilot. If the test rejects, use the FNA in the main wave. Otherwise, use the balanced allocation instead.

Specifically, we will use the (linear) Wald test for the equality of variance, which has the following test statistic:

$$\tilde{W} = \sqrt{\frac{m}{2}} \frac{(\tilde{\sigma}^2(1) - \tilde{\sigma}^2(0))}{\sqrt{\tilde{V}_W(1) + \tilde{V}_W(0)}} ,$$

where

$$\begin{aligned} \tilde{V}_W(a) &= \begin{pmatrix} 1 & -2\tilde{Y}(a) \end{pmatrix} \tilde{\Sigma}_W(a) \begin{pmatrix} 1 \\ -2\tilde{Y}(a) \end{pmatrix} , \\ \tilde{\Sigma}_W(a) &= \begin{pmatrix} \frac{1}{m/2} \sum_{\tilde{A}_i=a} (\tilde{Y}_i^2 - \tilde{Y}^2(a))^2 & \frac{1}{m/2} \sum_{\tilde{A}_i=a} (\tilde{Y}_i - \tilde{Y}(a)) (\tilde{Y}_i^2 - \tilde{Y}^2(a)) \\ * & \frac{1}{m/2} \sum_{A_i=a} (\tilde{Y}_i - \tilde{Y}(a))^2 \end{pmatrix} , \\ \tilde{Y}(a) &= \frac{1}{m/2} \sum_{A_i=a} \tilde{Y}_i \quad , \quad \tilde{Y}^2(a) = \frac{1}{m/2} \sum_{A_i=a} \tilde{Y}_i^2 . \end{aligned}$$

It is well known that $\tilde{W} \xrightarrow{d} N(0, 1)$ as $m \rightarrow \infty$. We can therefore form a two-sided, level α test for homoskedasticity with the rejection rule:

$$\phi = \mathbf{1} \left\{ \tilde{W} > \Phi^{-1} \left(1 - \frac{\alpha}{2} \right) \right\} + \mathbf{1} \left\{ \tilde{W} < \Phi^{-1} \left(\frac{\alpha}{2} \right) \right\} ,$$

where $\Phi(\cdot)$ is the CDF of the standard normal distribution. Under small pilot asymptotics, this test will not be consistent. However, the two-step procedure will be asymptotically equivalent to the FNA when pilots are large, since for any fixed level α , the test will reject with probability approaching 1 as $m \rightarrow \infty$. However, under such a ‘‘pre-testing procedure’’, the main-wave difference-in-means estimator will not be uniformly consistent.

Nonetheless, the procedure appears to perform well in practice. Figure 6 presents the MSEs that obtain using first stage tests with $\alpha \in \{1\%, 5\%, 10\%, 20\%, 50\%\}$. When $\sigma(1)/\sigma(0) \geq 0.5$, decreasing α decreases MSE across all distributions. Conversely, when $\sigma(1)/\sigma(0) \leq 0.5$, decreasing α increases MSE across all distributions. The appropriate level of α then depends on researcher’s preference for the performance of the method over the range of $\sigma(1)/\sigma(0)$.

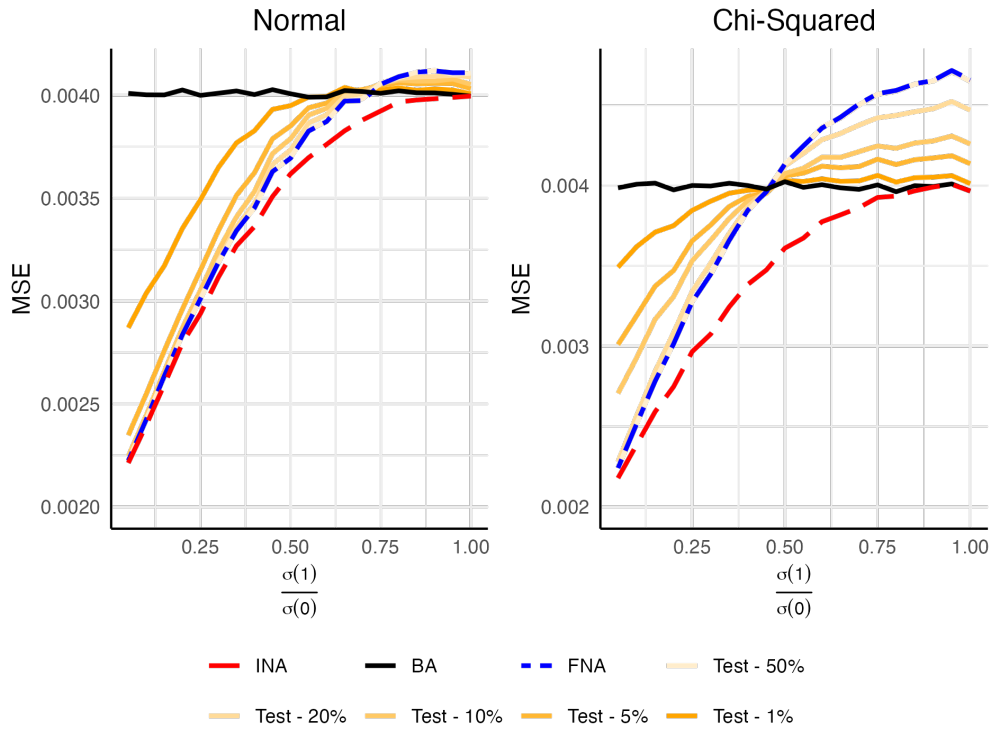


Figure 6: MSE under testing for homoskedasticity. We compare the Infeasible Neyman Allocation (INA), Balanced Allocation (BA), Feasible Neyman Allocation (FNA), as well as testing for homoskedasticity with various values of α (Test - $\alpha\%$).

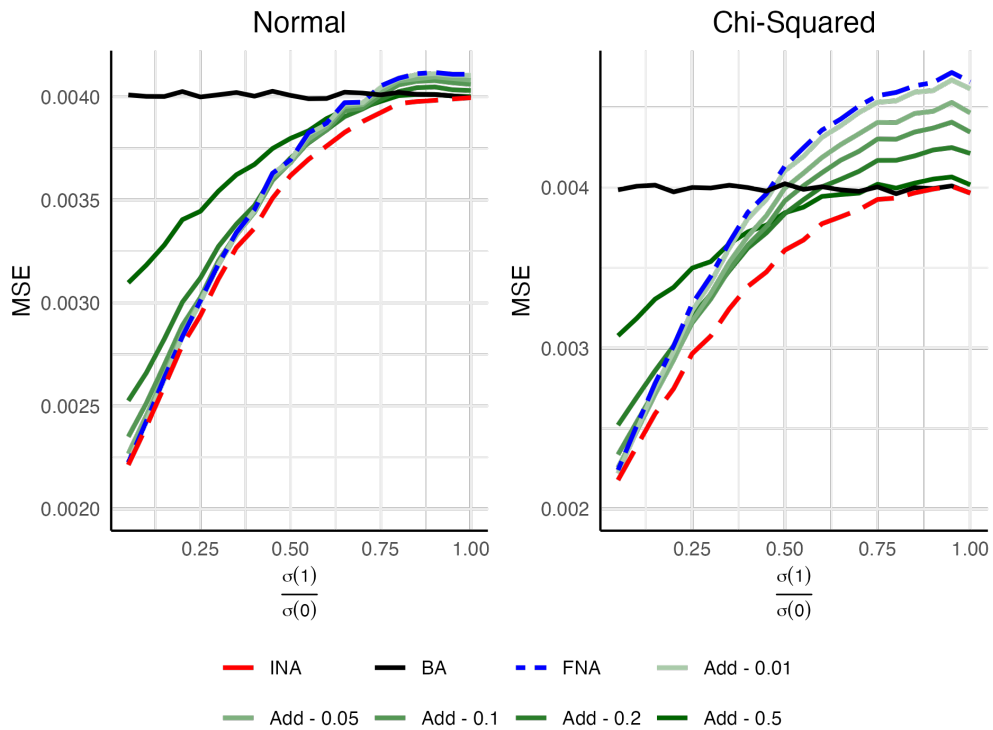


Figure 7: MSE under additive regularization. We compare the Infeasible Neyman Allocation (INA), Balanced Allocation (BA), Feasible Neyman Allocation (FNA), as well as additive regularization with various values of ν (Add - ν).

6.2 Additive Regularization

Rather than using a test to make a binary decision between the FNA and the balanced allocation, we can consider regularizing the FNA towards the balanced allocation. In particular, for a given tuning parameter ν , let

$$\tilde{p}^{\text{add}} = \frac{\tilde{\sigma}(1) + \tilde{\sigma}(0)\nu}{(\tilde{\sigma}(1) + \tilde{\sigma}(0))(1 + \nu)} .$$

In other words, we regularize $\tilde{\sigma}(1)$ by adding to it $\nu\tilde{\sigma}(0)$, and similarly for $\tilde{\sigma}(0)$. Pre-multiplying ν with $\tilde{\sigma}(1)$ and $\tilde{\sigma}(0)$ ensures that \tilde{p}^{add} is invariant to their scale. Furthermore,

$$\begin{aligned} \tilde{p}^{\text{add}} &\rightarrow 1/(1 + \nu) \text{ as } \sigma(1)/\sigma(0) \rightarrow \infty , \\ \tilde{p}^{\text{add}} &\rightarrow \nu/(1 + \nu) \text{ as } \sigma(1)/\sigma(0) \rightarrow 0 . \end{aligned}$$

Figure 7 presents results for $\nu \in \{0.01, 0.05, 0.1, 0.2, 0.5\}$. As we increase ν , the MSE decreases for the homoskedastic region but increases for the heteroskedastic region.

6.3 Exponential Regularization

Another method by which we can continuously regularize the FNA is as follows. For a given tuning parameter τ , let

$$\tilde{p}^{\text{exp}} = \frac{\left(\frac{\tilde{\sigma}(1)}{\tilde{\sigma}(0)}\right)^\tau}{1 + \left(\frac{\tilde{\sigma}(1)}{\tilde{\sigma}(0)}\right)^\tau} .$$

Observe that when $\tau = 1$, \tilde{p}^{add} is FNA. For $0 \leq \tau < 1$, \tilde{p}^{exp} regularizes the FNA towards the balanced allocation. In particular, for $\tau = 0$, $\tilde{p}^{\text{add}} = \frac{1}{2}$. Figure 7 presents results for $\tau \in \{0.99, 0.95, 0.9, 0.8, 0.5\}$. As before, when we decrease τ , the MSE decreases for the homoskedastic region but increases for the heteroskedastic region.

6.4 Comparison of Solutions

As the above sections demonstrate, each potential solution entails different trade-offs between MSE in the homoskedastic and heteroskedastic region. These trade-offs also vary depending on the tuning parameter chosen. In a crude attempt at comparing these methods, we consider the following. Choose tuning parameter values so that the MSEs of the methods are made as close as possible in the homoskedastic region. Then compare MSEs in the heteroskedastic region.

One such comparison is displayed in Figure 9. Here, we see that having fixed performance

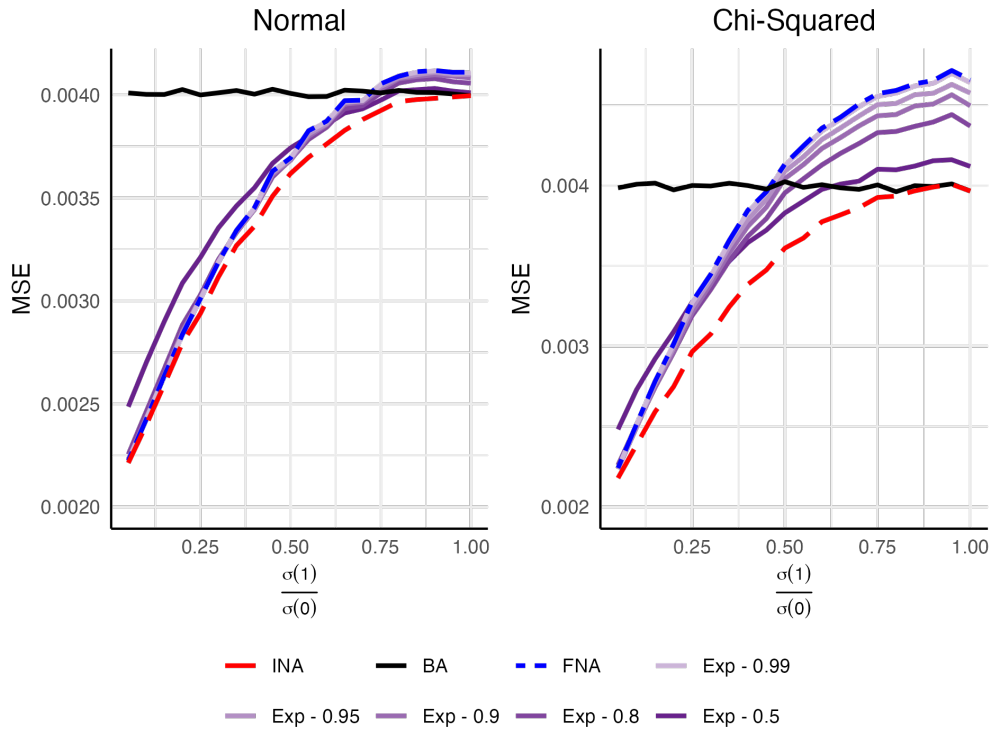


Figure 8: MSE under exponential regularization. We compare the Infeasible Neyman Allocation (INA), Balanced Allocation (BA), Feasible Neyman Allocation (FNA), as well as exponential regularization with various values of τ (Exp - τ).

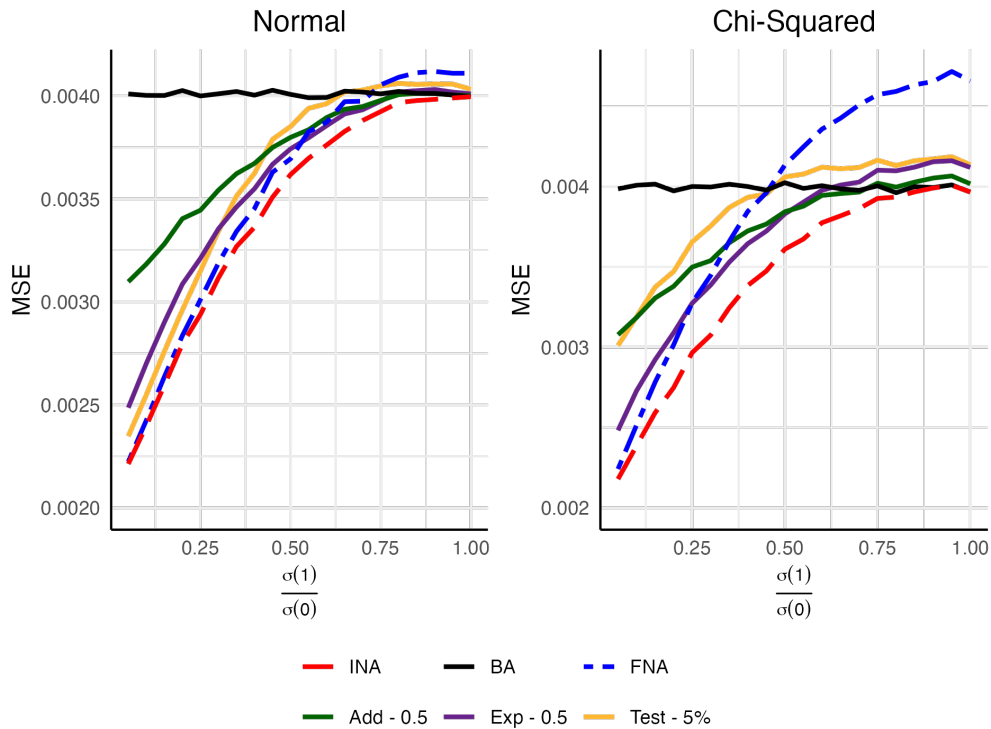


Figure 9: MSE under various candidate solutions. The tuning parameter for each solution is chosen so that their performance in the homoskedastic region is approximately comparable.

in the homoskedastic region, exponential regularization seems to do a better job in the heteroskedastic region. This conclusion ostensibly extends to the Pareto distributions, as well as other choices of tuning parameters that keep the solutions appropriately comparable. We therefore speculate that exponential regularization is an effective method for improving the small sample performance of the FNA.

7 Conclusion

We study the properties of the Feasible Neyman Allocation (FNA) in an asymptotic framework for two-wave experiments that takes pilot size to be fixed as the size of the main wave tends to infinity. In this setting, the estimated allocation has error that is not negligible even in the limit. Our asymptotic model therefore corresponds more closely to the finite sample statistical problem in which pilots are small and the optimal allocation may be poorly estimated. We characterize the conditions under which the difference-in-means estimator has larger asymptotic variance compared to balanced randomization, where half of the main wave is assigned to treatment and the remainder to control. This happens when the potential outcomes are relatively homoskedastic with respect to treatment status or exhibit high kurtosis – situations that may arise in practice. Our results suggest that the confluence of these factors – small pilots, homoskedasticity, heavy-tails – are disadvantageous to the FNA. Instead, researchers might find it useful to consider alternative procedures such as regularizing the FNA towards the balanced allocation.

Acknowledgements

We thank the editor, the associate editor and three anonymous referees for comments and suggestions that greatly improved the article. We are also grateful to Federico Bugni, Ivan Canay, Joel Horowitz, Dean Karlan, Max Tabord-Meehan and Chris Udry for helpful discussions. We report that there are no competing interests to declare.

References

Ashraf, N., D. Karlan, and W. Yin (2006). Tying odysseus to the mast: Evidence from a commitment savings product in the philippines. *The Quarterly Journal of Economics* 121(2), 635–672.

- Avvisati, F., M. Gurgand, N. Guyon, and E. Maurin (2014). Getting parents involved: A field experiment in deprived schools. *Review of Economic Studies* 81(1), 57–83.
- Azriel, D., M. Mandel, and Y. Rinott (2012). Optimal allocation to maximize the power of two-sample tests for binary response. *Biometrika* 99(1), 101–113.
- Bai, Y. (2021). Why randomize? minimax optimality under permutation invariance. *Journal of Econometrics*.
- Bai, Y. (2022). Optimality of matched-pair designs in randomized controlled trials.
- Blackwell, M., N. E. Pashley, and D. Valentino (2022). Batch adaptive designs to improve efficiency in social science experiments. Technical report, Working paper, Harvard University, 2022.
- Bloom, N., J. Liang, J. Roberts, and Z. J. Ying (2014, 11). Does Working from Home Work? Evidence from a Chinese Experiment. *The Quarterly Journal of Economics* 130(1), 165–218.
- Brittain, E. and J. J. Schlesselman (1982). Optimal allocation for the comparison of proportions. *Biometrics*, 1003–1009.
- Bruhn, M. and D. McKenzie (2009). In pursuit of balance: Randomization in practice in development field experiments. *American economic journal: applied economics* 1(4), 200–232.
- Bryan, G., D. Karlan, and J. Zinman (2015, 08). Referrals: Peer screening and enforcement in a consumer credit field experiment. *American Economic Journal: Microeconomics* 7(3), 174–204.
- Bugni, F. A., I. A. Canay, and A. M. Shaikh (2018). Inference under covariate-adaptive randomization. *Journal of the American Statistical Association* 113(524), 1784–1796.
- Canay, I. A. and V. Kamat (2017, 10). Approximate Permutation Tests and Induced Order Statistics in the Regression Discontinuity Design. *The Review of Economic Studies* 85(3), 1577–1608.
- Canay, I. A., J. P. Romano, and A. M. Shaikh (2017). Randomization tests under an approximate symmetry assumption. *Econometrica* 85(3), 1013–1030.
- Canay, I. A., A. Santos, and A. M. Shaikh (2021). The wild bootstrap with a “small” number of “large” clusters. *Review of Economics and Statistics* 103(2), 346–363.

- Chong, A., D. Karlan, J. Shapiro, and J. Zinman (2015). (ineffective) messages to encourage recycling: evidence from a randomized evaluation in peru. *The World Bank Economic Review* 29(1), 180–206.
- Conley, T. G. and C. R. Taber (2011). Inference with “difference in differences” with a small number of policy changes. *The Review of Economics and Statistics* 93(1), 113–125.
- Cytrynbaum, M. (2021). Designing representative and balanced experiments by local randomization. *arXiv preprint arXiv:2111.08157*.
- Cytrynbaum, M. (2023). Optimal stratification of survey experiments. *arXiv preprint arXiv:2111.08157*.
- de Chaisemartin, C. and J. Ramirez-Cuellar (2019). At what level should one cluster standard errors in paired and small-strata experiments? *arXiv preprint arXiv:1906.00288*.
- Deming, D. J., N. Yuchtman, A. Abulafi, C. Goldin, and L. F. Katz (2016, 03). The value of postsecondary credentials in the labor market: An experimental study. *American Economic Review* 106(3), 778–806.
- Dillon, M. R., H. Kannan, J. T. Dean, E. S. Spelke, and E. Duflo (2017). Cognitive science in the field: A preschool intervention durably enhances intuitive but not formal mathematics. *Science* 357(6346), 47–55.
- Finkelstein, A., S. Taubman, B. Wright, M. Bernstein, J. Gruber, J. P. Newhouse, H. Allen, K. Baicker, and O. H. S. Group (2012). The oregon health insurance experiment: evidence from the first year. *The Quarterly journal of economics* 127(3), 1057–1106.
- Galiani, S. and P. J. McEwan (2013). The heterogeneous impact of conditional cash transfers. *Journal of Public Economics* 103, 85–96.
- Hahn, J., K. Hirano, and D. Karlan (2011). Adaptive experimental design using the propensity score. *Journal of Business & Economic Statistics* 29(1), 96–108.
- Hu, F. and W. F. Rosenberger (2003). Optimality, variability, power: evaluating response-adaptive randomization procedures for treatment comparisons. *Journal of the American Statistical Association* 98(463), 671–678.
- Hu, F. and W. F. Rosenberger (2006). *The theory of response-adaptive randomization in clinical trials*. John Wiley & Sons.

- Ibragimov, R. and U. K. Müller (2010). t-statistic based correlation and heterogeneity robust inference. *Journal of Business & Economic Statistics* 28(4), 453–468.
- Lachin, J. M. (1988). Properties of simple randomization in clinical trials. *Controlled clinical trials* 9(4), 312–326.
- McKenzie, D. (2017). Identifying and spurring high-growth entrepreneurship: Experimental evidence from a business plan competition. *American Economic Review* 107(8), 2278–2307.
- Melfi, V. and C. Page (1998). Variability in adaptive designs for estimation of success probabilities. *IMS Lecture Notes - Monograph Series*, 106–114.
- Neyman, J. (1934). On the two different aspects of the representative method: the method of stratified sampling and the method of purposive selection. In *Journal of the Royal Statistical Society*, Volume 97, pp. 558–625.
- Obradovic, F. (2023). Using baseline data to guide experimental design.
- Sukhatme, P. (1935). Contribution to the theory of the representative method. *Supplement to the Journal of the Royal Statistical society* 2(2), 253–268.
- Tabord-Meehan, M. (2021). Stratification trees for adaptive randomization in randomized controlled trials. *arXiv preprint arXiv:1806.05127*.
- Viviano, D. (2022). Experimental design under network interference. *arXiv preprint arXiv:2003.08421*.
- Zhao, J. (2023). Adaptive neyman allocation. *arXiv preprint arXiv:2309.08808*.

Online Appendix for “On the Performance of the Neyman Allocation with Small Pilots”

A Proofs

A.1 Proof of Corollary 4.3

Let \mathcal{L}_m denote a random variable whose distribution is given by

$$\mathbb{P}(\mathcal{L}_m \leq t) = \int_0^1 \Phi\left(\frac{t}{s(p)}\right) G_m(dp) .$$

Using the Law of Iterated Expectations, it can be shown that \mathcal{L}_m has mean zero and variance given by

$$\mathbb{E}[\mathcal{L}_m^2] = \mathbb{E}[s^2(\tilde{p})] = \mathbb{E}\left[\frac{\sigma^2(1)}{\tilde{p}} + \frac{\sigma^2(0)}{1-\tilde{p}}\right] .$$

Next, note that for all $\tilde{p} \in [0, 1], \tilde{p} \neq p_*$,

$$\frac{\sigma^2(1)}{\tilde{p}} + \frac{\sigma^2(0)}{1-\tilde{p}} > \frac{\sigma^2(1)}{p_*} + \frac{\sigma^2(0)}{1-p_*} = \Sigma_* .$$

Since $\sigma^2(0) > 0$ and $\sigma^2(1) > 0$, $\mathbb{P}(\tilde{p} \neq p_*) > 0$. That is,

$$\mathbb{P}\left(\frac{\sigma^2(1)}{\tilde{p}} + \frac{\sigma^2(0)}{1-\tilde{p}} > \Sigma_*\right) > 0 .$$

By the strict monotonicity of expectation, we are done. □

A.2 Proof of Proposition 4.4

By our earlier derivation, $\hat{\theta}_p$ has lower asymptotic variance than $\hat{\theta}_{\tilde{p}}$ if and only if

$$\mathbb{E}\left[\frac{\sigma^2(1)}{\tilde{p}} + \frac{\sigma^2(0)}{1-\tilde{p}}\right] \geq \mathbb{E}\left[\frac{\sigma^2(1)}{1/2} + \frac{\sigma^2(0)}{1-1/2}\right] .$$

Define the variable Z_m :

$$Z_m = \frac{\hat{\sigma}(1)}{\sigma(1)} \bigg/ \frac{\hat{\sigma}(0)}{\sigma(0)} .$$

We can then rewrite the above condition as

$$\begin{aligned} & \mathbb{E} \left[\left(1 + \frac{1}{Z_m} \frac{\sigma(0)}{\sigma(1)} \right) \sigma^2(1) + \left(1 + Z_m \frac{\sigma(1)}{\sigma(0)} \right) \sigma^2(0) \right] \geq 2\sigma^2(1) + 2\sigma^2(0) \\ \Leftrightarrow & \mathbb{E} \left[\frac{1}{Z_m} + Z_m \right] \sigma(1)\sigma(0) \geq \sigma^2(1) + \sigma^2(0) \\ \Leftrightarrow & \frac{\sigma^2(1)}{\sigma^2(0)} - \mathbb{E} \left[\frac{1}{Z_m} + Z_m \right] \frac{\sigma(1)}{\sigma(0)} + 1 \leq 0 . \end{aligned}$$

By the quadratic formula, the above inequality is satisfied whenever

$$\frac{\sigma(1)}{\sigma(0)} \in \left[B_m - \sqrt{B_m^2 - 1}, B_m + \sqrt{B_m^2 - 1} \right] = C_m ,$$

where $B_m := \frac{1}{2} \mathbb{E} \left[\frac{1}{Z_m} + Z_m \right]$. Note that

$$B_m^2 - (B_m^2 - 1) = 1 \Rightarrow \left(B_m - \sqrt{B_m^2 - 1} \right) \left(B_m + \sqrt{B_m^2 - 1} \right) = 1 ,$$

so that

$$B_m - \sqrt{B_m^2 - 1} = \frac{1}{B_m + \sqrt{B_m^2 - 1}} = \frac{1}{x} .$$

Next note that on R_+ , $f(x) = \frac{1}{x} + x$ is strictly convex and attains its strict minimum at $x = 1$. Since Z_m is non-degenerate, Jensen's inequality yields

$$2B_m > \frac{1}{\mathbb{E}[Z_m]} + \mathbb{E}[Z_m] \geq 2 .$$

Hence, $|C_m| > 0$. We can then write:

$$|C_m| = 2\sqrt{B_m^2 - 1} = 2\sqrt{(B_m - 1)^2 + 2(B_m - 1)} ,$$

where

$$B_m - 1 = \frac{1}{2} \left[\frac{\sigma(1)}{\sigma(0)} \text{Bias}(\tilde{p}) + \frac{\sigma(0)}{\sigma(1)} \text{Bias}\left(\frac{1}{\tilde{p}}\right) \right] = W_m . \quad \square$$

A.3 Proof of Proposition 4.5

We start by evaluating the asymptotic distributions of Z_m . First,

$$\begin{aligned} \sqrt{m} (\hat{\sigma}^2(1) - \sigma^2(1)) &= \frac{1}{\sqrt{m}} \sum_{i=1}^m ((Y_i(1) - \mu(1))^2 - \sigma^2(1)) - \sqrt{n} (\bar{Y}(1) - \mu(1))^2 + o_p(1) \\ &\xrightarrow{d} \mathcal{N} \left(0, \mathbb{E} [(Y_i(1) - \mu(1))^4] - \sigma^4(1) \right) . \end{aligned}$$

By the Delta Method,

$$\sqrt{m} \left(\frac{\widehat{\sigma}^2(1)}{\sigma^2(1)} - 1 \right) \xrightarrow{d} \mathcal{N} \left(0, \frac{\mathbb{E} [(Y_i(1) - \mu(1))^4] - \sigma^4(1)}{4\sigma^4(1)} \right) .$$

Similarly,

$$\sqrt{m} \left(\frac{\widehat{\sigma}^2(0)}{\sigma^2(0)} - 1 \right) \xrightarrow{d} \mathcal{N} \left(0, \frac{\mathbb{E} [(Y_i(0) - \mu(0))^4] - \sigma^4(0)}{4\sigma^4(0)} \right) .$$

Since the above two displays contain independent random variables, another application of the Delta Method yields that

$$\sqrt{m} (Z_m - 1) \xrightarrow{d} \mathcal{N}(0, V) .$$

Next, let $f(x) = x + 1/x$. Note that

$$\begin{aligned} f'(x) &= 1 - \frac{1}{x^2} \quad , \quad f'(1) = 0 \quad , \\ f''(x) &= \frac{2}{x^3} \quad , \quad f''(1) = 2 \quad . \end{aligned}$$

By the second order Delta Method,

$$m \left(\frac{1}{2} \left(\frac{1}{Z_m} + Z_m \right) - 1 \right) \xrightarrow{d} \frac{1}{2} V \cdot \chi_1^2 .$$

Since the left hand side is an analytic function of sub-Gaussian random variables, all moments can be bounded uniformly in m . Conclude that:

$$m (B_m - 1) \rightarrow \frac{V}{2} . \tag{A.1}$$

Hence,

$$\sqrt{m} \left(\left(B_m + \sqrt{B_m^2 - 1} \right) - 1 \right) = \frac{m(B_m - 1)}{\sqrt{m}} + \sqrt{m(B_m - 1) \left(\frac{m(B_m - 1)}{m} + 2 \right)} \rightarrow \sqrt{V} .$$

□

A.4 Proof of Proposition 4.6

By definition, $\tilde{p} = \left(1 + \frac{1}{Z_m} \frac{\sigma(0)}{\sigma(1)} \right)^{-1}$. Substituting this term into Definition 4.4 yields the expressions in terms of B_m . Equation (A.1) the proof of Proposition 4.5 gives us the second set of expressions for when m is large.

A.5 Proof of Proposition 4.7

Recall from the proof of Proposition 4.4 that $B_m > 1$.

We begin by showing that the displayed infimums are indeed lower bounds. Observe that

$$\mathcal{L}^d(F) \geq (\sigma(1) - \sigma(0))^2 \geq -K .$$

In the above display, the first inequality follows because $B_m > 1$ (see proof of Proposition 4.4). The second inequality follows by minimizing the middle term over $\mathcal{F}(K)$. That $\frac{1}{2}$ is a lower bound $\mathcal{L}^d(F)$ is immediate since $B_m > 0$. To see that $\mathcal{L}^d(F)$ and $\mathcal{L}^r(F)$ are attained, first fix the distribution of $\varepsilon(1) := Y(1) - \mu(1)$ and $\varepsilon(0) := Y(0) - \mu(0)$. Now choose some sequence $\omega_k \rightarrow 0$. Define $Y_k(1) = \mu(1) + \omega_k \varepsilon(1)$ and $Y_k(0) = Y(0)$ and denote their joint distribution F_k . Notice that each F_k has the same $B_m(F_k) = \mathbb{E}_{F_k}[Z_m + \frac{1}{Z_m}]$. Then as $\omega_k \rightarrow 0$,

$$\mathcal{L}^d(F) \rightarrow -K \quad \text{and} \quad \mathcal{L}^r(F) \rightarrow \frac{1}{2} .$$

We next show that the variance of $\hat{\theta}_{\hat{p}}$ is unbounded over the class of distributions with uniformly bounded a^{th} moment for any a . We now construct a simple counterexample in which we essentially force division by 0. For a given a , let $Y(0) \sim N(0, 1)$. Let $Y(1) \sim X + \omega W$, where $W \sim N(0, 1)$, $W \perp\!\!\!\perp X$ and

$$X = \begin{cases} -1 & \text{w.p. } \frac{1}{2\kappa} \\ 0 & \text{w.p. } 1 - \frac{1}{\kappa} \\ 1 & \text{w.p. } \frac{1}{2\kappa} \end{cases} .$$

By construction, $E[Y(1)^p] \leq C_p(X^p + W^p) \leq K$ (we can scale $Y(1)$ if necessary). Now, let A be the event that $X_1 = X_2 = \dots = X_m = 0$. Then,

$$\begin{aligned} B_m &\geq E \left[Z_m + \frac{1}{Z_m} \mid A \right] P(A) \\ &= \left(\omega + \frac{1}{\omega} \right) \mathbb{E} \left[F_{\frac{m}{2}, \frac{m}{2}} \right] \left(1 - \frac{1}{\kappa} \right)^m \\ &\rightarrow \infty \text{ as } \omega \rightarrow 0 . \end{aligned}$$

To show non-uniformity in the large pilot regime, we can set $\kappa = m$ and $\omega = 1/m$. Then $B_m \approx me^{-1} \rightarrow \infty$.

B Additional Empirical Examples

We revisit the first 10 completed RCTs in the AER RCT Registry. This section contains the empirical examples omitted from the main text. They are:

- Section B.1: [Dillon et al. \(2017\)](#)
- Section B.2: [Finkelstein et al. \(2012\)](#)
- Section B.3: [McKenzie \(2017\)](#)
- Section B.4: [Chong et al. \(2015\)](#)
- Section B.5: [Deming et al. \(2016\)](#)
- Section B.6: [Bloom et al. \(2014\)](#)
- Section B.7: [Galiani and McEwan \(2013\)](#)
- Section B.8: [Bryan et al. \(2015\)](#)

B.1 [Dillon et al. \(2017\)](#)

[Dillon et al. \(2017\)](#) conduct an RCT to test the hypothesis that math game play in pre-school prepares poor children for formal math in primary school. Their study, conducted in Delhi, India, with the organization Pratham, involved 214 pre-schools with 1540 children, with treatment assigned at the school level. In the Math treatment arm, children were led by facilitators to play math games over the course of four months, while the control group received lessons according to Pratham’s usual curriculum. To distinguish the effect of the math games from the effect of engagement with adults, the experiment further involved a Social treatment arm, where social games were played. The outcomes of interest are Math Skills – subdivided into Symbolic Math Skills and Non-Symbolic Math Skills – as well as Social Skills, as measured by Pratham’s standardized tests. Since the authors are interested in persistence of treatment effects, they measure these outcomes at the following times after intervention: 0-3 months (Endline 1), 6-9 months (Endline 2) and 12-15 months (Endline 3). They find that the Math intervention has positive effects on Non-Symbolic Mathematical skills across all three Endlines, while Symbolic Mathematical skills only improves in Endline 1.

Table 7 displays the standard deviation of each outcome variable by treatment arm, computed at the individual level. In the absence of correlation among students in the same

pre-school, and assuming that treatment is assigned at the individual level, the numbers shown are the relevant empirical counterparts to $\sigma(1)$ and $\sigma(0)$. We see that the outcomes are relatively homoskedastic across the outcome variables and across time. The ratio $\sigma(1)\sigma(0)$ falls between 0.94 and 1.31, suggesting that naive experiment will do well when pilots are small.

Table 7: Individual Level Heteroskedasticity in [Dillon et al. \(2017\)](#)

Endline	Outcome	Math	Social	Control	Math/Control	Social/Control
1	Math	0.73	0.68	0.69	1.06	0.99
	Symbolic Math	0.74	0.78	0.77	0.96	1.01
	Non-Symbolic Math	0.94	0.81	0.77	1.21	1.04
	Social	1.18	1.41	1.07	1.10	1.31
2	Math	0.71	0.73	0.69	1.04	1.06
	Symbolic Math	0.72	0.74	0.71	1.02	1.05
	Non-Symbolic Math	0.98	0.99	0.92	1.06	1.07
	Social	0.92	0.96	0.99	0.94	0.97
3	Math	0.78	0.70	0.75	1.04	0.93
	Symbolic Math	0.72	0.68	0.74	0.98	0.92
	Non-Symbolic Math	1.17	1.09	1.06	1.11	1.03
	Social	1.01	1.05	1.07	0.94	0.98

Suppose we are concerned about correlation across students in the same pre-school. We can redefine our unit of observation to be the school by taking averages across students in the same school. The Neyman Allocation then tells us how many schools to allocate to treatment. In this case, the standard deviation in the mean across schools is the relevant counterpart to $\sigma(1)$ and $\sigma(0)$. They are presented in Table 8. As before, we see that outcomes are relatively homoskedastic across schools, though less so than in the individual level case. Nonetheless, the ratio falls between 0.84 and 1.55. Once we consider schools to be the unit of treatment, however, the effective pilot size also shrinks, such that the drawbacks of the estimated Neyman Allocation may be even more pronounced. All in all, the [Dillon et al. \(2017\)](#) example supports our case of relative homoskedasticity in empirical applications.

B.2 [Finkelstein et al. \(2012\)](#)

[Finkelstein et al. \(2012\)](#) study the Oregon Health Insurance Experiment, in which uninsured, low-income adults were randomly given the opportunity to apply for Medicaid. Over the course of a month in February 2008, Oregon conducted extensive public awareness campaign to encourage participation in the lottery. From a total of 89,824 sign-ups, 35,169 individ-

Table 8: School Level Heteroskedasticity in [Dillon et al. \(2017\)](#)

Endline	Outcome	Math	Social	Control	Math/Control	Social/Control
1	Math	0.41	0.38	0.30	1.34	1.24
	Symbolic Math	0.39	0.42	0.34	1.14	1.22
	Non-Symbolic Math	0.51	0.40	0.33	1.55	1.22
	All Social	0.51	0.71	0.46	1.10	1.55
2	All Math	0.39	0.35	0.36	1.10	0.97
	Symbolic Math	0.40	0.35	0.36	1.11	0.96
	Non-Symbolic Math	0.46	0.42	0.43	1.06	0.98
	Social	0.41	0.48	0.49	0.84	0.97
3	All Math	0.44	0.40	0.39	1.15	1.04
	Symbolic Math	0.40	0.35	0.37	1.06	0.94
	Non-Symbolic Math	0.65	0.63	0.53	1.23	1.20
	Social	0.55	0.59	0.50	1.10	1.20

uals (from 29,664 households) were selected. They, and any members of their households were then given the opportunity to apply for Medicaid. Hence, treatment occurred at the household level.

The authors used the data to study a variety of outcomes. In this section, we focus on their first set of results, which concern healthcare utilization. In particular, we revisit the outcome variables used in Tables V and VI of [Finkelstein et al. \(2012\)](#), which are obtained from survey data (as opposed to administrative data), and are hence publicly available. Inline with the authors’ results on the Intent-to-Treat effect, we define the treated group as those selected by the lottery. We note that the authors apply sampling weights to correct for differential response rates to the survey. We follow their weighting scheme in computing our results.

The standard deviation of individual-level outcome are presented in Table 9. Results taking household to be the unit of observation are presented in Table 10. Across both tables, we see that that the standard deviations in outcomes are remarkably similar across treatment and control groups. They are also very similar across individual and household level groups, since households with more than one person represented less than 5% of the survey sample.

B.3 [McKenzie \(2017\)](#)

Business plan competitions are growing in popularity as a way of fostering high growth entrepreneurship in developing countries. [McKenzie \(2017\)](#) studies the Youth Enterprise With Innovation in Nigeria (YouWiN!) program, which distributed up to US\$64,000 to winners. A portion of the awards were reserved for business plans that were clearly superior to the

Table 9: Individual Level Heteroskedasticity in [Finkelstein et al. \(2012\)](#)

	Outcome	Treatment	Control	Treat./Cont.
Extensive Margin	Prescription drugs currently	0.48	0.48	0.99
	Outpatient visits last six months	0.48	0.49	0.98
	ER visits last six months	0.44	0.44	1.00
	Inpatient hospital admissions last six months	0.26	0.26	1.00
Total Utilization	Prescription drugs currently	2.90	2.88	1.01
	Outpatient visits last six months	3.29	3.09	1.07
	ER visits last six months	1.01	1.04	0.97
	Inpatient hospital admissions last six months	0.42	0.40	1.04
Preventative Care	Blood cholesterol checked (ever)	0.48	0.48	0.98
	Blood tested for high blood sugar (ever)	0.48	0.49	0.99
	Mammogram within last 12 months (women ≥ 40)	0.48	0.46	1.04
	Pap test within last 12 months (women)	0.50	0.49	1.01

Table 10: Household Level Heteroskedasticity in [Finkelstein et al. \(2012\)](#)

	Outcome	Treatment	Control	Treat./Cont.
Extensive Margin	Prescription drugs currently	0.46	0.47	0.99
	Outpatient visits last six months	0.47	0.48	0.97
	ER visits last six months	0.43	0.44	0.99
	Inpatient hospital admissions last six months	0.26	0.26	0.99
Total Utilization	Prescription drugs currently	2.88	2.86	1.01
	Outpatient visits last six months	3.30	3.09	1.07
	ER visits last six months	1.01	1.04	0.98
	Inpatient hospital admissions last six months	0.41	0.40	1.02
Preventative Care	Blood cholesterol checked (ever)	0.46	0.48	0.97
	Blood tested for high blood sugar (ever)	0.47	0.48	0.98
	Mammogram within last 12 months (women ≥ 40)	0.48	0.46	1.04
	Pap test within last 12 months (women)	0.50	0.49	1.01

rest. 1,841 entrepreneurs, determined to be of medium quality, were entered into a lottery, from which 729 were selected for the award. Three rounds of surveys were then conducted at 1,2 and 3 years after the application respectively.

We focus on the first set of results in [McKenzie \(2017\)](#) – presented in Table 2 – concerning the effect of the award on start-up and survival. Here, they find that the grant persistently increased the probability that the entrepreneur was operating a firm, as well as the number of hours they spent in self-employment. We present the standard deviations of these outcome variables in Table 11. Here we see that hours in self employment is roughly homoskedastic across all three periods. However, the outcome on whether the entrepreneur is operating a firm is arguably highly heteroskedastic, with standard deviations that is as small as 0.44 that of the control group. As in our earlier example, we find high kurtosis in these outcome variables, displayed in Table 12.

We do not estimate C_m in this example because almost all entrepreneurs in the treated group operate their own firms in the sample. As such, a small random sub-sample (e.g. of size below 200) from this group has variance 0 with high probability, impeding the estimation of C_m . These pathological cases are revealing. In a pilot, if the treated group has variance 0 in the outcome, the estimated Neyman Allocation assigns 0 units to treatment in the full experiment. The high probability of such an “extreme” outcome with small pilots is precisely the danger which we are warning against. We conclude that the estimated Neyman from a small pilot will likely lead to adverse results given the DGP in [McKenzie \(2017\)](#).

Table 11: Heteroskedasticity in [McKenzie \(2017\)](#)

Outcome	New Firms			Existing Firms		
	Treat.	Cont.	Ratio	Treat.	Cont.	Ratio
Operates a Firm at Round 1	0.43	0.50	0.86	0.21	0.34	0.63
Operates a Firm at Round 2	0.27	0.50	0.55	0.16	0.36	0.44
Operates a Firm at Round 3	0.28	0.50	0.56	0.20	0.43	0.48
Weekly Hours of Self Emp. at Round 1	29.40	29.75	0.99	25.74	27.81	0.93
Weekly Hours of Self Emp. at Round 2	24.98	28.62	0.87	24.51	29.71	0.82
Weekly Hours of Self Emp. at Round 3	24.77	25.85	0.96	25.09	26.10	0.96

B.4 [Chong et al. \(2015\)](#)

Partnering with the Peruvian nongovernmental organization PRISMA, [Chong et al. \(2015\)](#) conducted two RCTs to investigate the efficacy of various interventions in encouraging recycling. First, the Participation Study considers the following 9 different messaging strategies and their relative success in enrolling members into recycling programs:

Table 12: Kurtosis in [McKenzie \(2017\)](#)

Outcome	New Firms		Existing Firms	
	Treat.	Cont.	Treat.	Cont.
Operates a Firm at Round 1	2.52	1.04	19.68	5.91
Operates a Firm at Round 2	10.69	1.08	35.46	4.58
Operates a Firm at Round 3	9.62	1.03	21.05	2.47
Weekly Hours of Self Emp. at Round 1	1.99	3.08	3.93	3.08
Weekly Hours of Self Emp. at Round 2	2.77	2.92	3.03	2.42
Weekly Hours of Self Emp. at Round 3	2.09	3.28	3.47	2.15

1. Norms: Rich and Poor. Norm messaging focus on communicating high recycling rate of either a rich or poor reference neighborhoods, encouraging conformity.
2. Signal: Rich, Poor and Local. Signal messaging informs the targets that their recycling behavior will be known to either a nearby neighborhood (Local), a distal neighborhood of varying wealth (Rich or Poor), affecting the targets reputation.
3. Authority: Religious or Municipal. Authority messaging communicates that a higher authority, either religious or local governmental, advocates recycling.
4. Information: Environmental or Social. Informational messaging communicated the benefits of recycling, either to the environment or to the local society (e.g. by creating jobs).

Out of a total of 6,718 households, approximately 600 were assigned to each treatment arm, with the exception of Signal: Local, which were assigned 932 participants. 1,157 households were assigned to the control group. Three measures of participation were considered:

1. "Participates any time" is an indicator that takes the value 1 if a household turned in residuals over the course of the study.
2. "Participation Ratio" is the number of times a household turns in residual over the total number of opportunities they had to turn in residuals.
3. "Participates during either of last two visits" is an indicator that takes value 1 if the household turned in residual during one of the last two canvassing weeks.

The results in Table 3 of [Chong et al. \(2015\)](#) shows that messaging had no effect in increasing participation in the program. Table 13 displays the standard deviation of the various outcomes by treatment type. Table 14, shows the ratio of the standard deviation in the outcome variable of each treatment group, with respect to that of the control group. Here,

we see that the outcomes are highly homoskedastic, suggesting little scope for improvement over the naive experiment.

Table 13: S.D. of outcome by treatment type in the Participation Study. See text for definitions of outcome and treatment.

Outcome	Control	Norms		Signal			Authority		Info.	
		Rich	Poor	Rich	Poor	Local	Reli.	Muni.	Env.	Social
1	0.500	0.500	0.500	0.500	0.500	0.500	0.500	0.500	0.500	0.500
2	0.389	0.389	0.392	0.403	0.385	0.393	0.397	0.388	0.396	0.392
3	0.490	0.486	0.494	0.493	0.486	0.489	0.491	0.491	0.494	0.493

Table 14: Ratio of the S.D w.r.t. the control group in the Participation Study.

Outcome	Norms		Signal			Authority		Info.	
	Rich	Poor	Rich	Poor	Local	Reli.	Muni.	Env.	Social
1	1.000	0.999	1.000	1.000	1.000	1.000	0.999	1.000	1.000
2	1.000	1.008	1.036	0.989	1.012	1.022	0.998	1.018	1.008
3	0.992	1.007	1.005	0.990	0.997	1.000	1.001	1.007	1.005

The second experiment is the Participation Intensity Study. The outcomes of interest are the following measures recycling intensity:

1. Percentage of visits in which household turned in residuals
2. Average number of bins turned in per week
3. Average weight (in kg) of recyclables turned in per week
4. Average market value of recyclables given per week
5. Average percentage of contamination (non-recyclables mixed into recycling) per week.

The treatments of interest are (1) providing recycling bins to households and (2) sending SMS reminders for recycling.¹ Of the 1,781 households in this study, 182 were received Bin and SMS. 417 received the Bin only treatment. 369 received the SMS only treatment, leaving 817 in the control group. The authors find, in Table 4A that bin provision was highly effective in increasing recycling, though SMS reminders had no effect. We compute the standard deviation in each of their outcome variables by treatment type in Table 15. Also displayed is the ratio of these standard deviation with respect to the control group. Again, we see strong evidence of homoskedasticity, suggesting that the naive experiment will perform well in this scenario.

¹The authors also consider providing bins with and without instructions as well as generic vs personalized SMSes. However, these finer definitions leads to treatment arms with fewer 50 households. Hence we focus on the coarser definition of the treatments, as employed in panel 4A.

Table 15: S.D. of outcome by treatment type in the Participation Intensity Study. See text for definitions of outcome and treatment.

Outcome	Standard Deviation				Ratio of S.D. w.r.t. Control		
	Control	SMS only	Bin only	SMS & Bin	SMS only	Bin only	SMS & Bin
1	0.262	0.286	0.233	0.227	1.092	0.891	0.867
2	0.404	0.371	0.441	0.375	0.919	1.092	0.927
3	0.744	0.646	0.756	0.727	0.869	1.017	0.978
4	0.418	0.371	0.416	0.399	0.889	0.996	0.955
5	0.156	0.145	0.136	0.128	0.928	0.870	0.822

B.5 Deming et al. (2016)

Deming et al. (2016) study employers' perceptions of the value of post-secondary degrees using a field experiment. The experimental units are fictitious resumes to be used in applications to vacancies posted on a large online job board. Their focus is on degrees and certificates awarded in the two largest occupational categories in the United States: business and health. Resumes are randomly assigned sector and selectivity of (degree-awarding) institutions. Fictitious resumes are created using a vast online database of actual resumes of job seekers, with applicant characteristics varying randomly (i.e. characteristics are randomly assigned). Outcomes are callback rates. There are three main comparisons in the paper:

- for-profit vs. public institutions,
- for-profits that are online vs. brick-and-mortar (with a local presence),
- more selective vs. less selective public institutions.

Deming et al. (2016) find that BA degrees in business from large online for-profit institutions are more 22% less likely to receive a callback than applicants with similar degrees from non-selective public schools when the job vacancy requires a BA. When a business job opening does not list a BA requirement, they find no significant overall advantage to having a post-secondary degree. For health jobs, resumes with certificates from for-profit institutions are 57% less likely to receive a callback than those with similar certificates from public institutions when the job listing does not require a post-secondary certificate. No significant difference in callback rates are found when the health job listing requires a certificate.

Table 16 reports standard deviations across treatment arms across the various sub-populations of interest in Deming et al. (2016). Since in most sub-populations, there are more than two treatment arms, we do not report ratios. However, pairwise comparisons between treatment arms within any chosen subpopulation shows strong evidence of relative homoskedasticity across the board.

Table 16: Performance impact outcomes: standard deviations

Experimental Population	Treatment Arm	S.D.
Business jobs without BA requirements	No degree	0.3054
	AA (for profit)	0.3026
	AA (public)	0.3053
	BA (for profit)	0.3071
Business jobs with BA requirements	BA (for profit, online)	0.2522
	BA (for profit, not online)	0.2209
	BA (public, selective)	0.2879
	BA (public, not selective)	0.2595
Health job without cert. requirement	No certificate	0.2900
	Certificate (for profit)	0.2922
	Certificate (public)	0.3014
Health job with cert. requirement	Certificate (for profit)	0.2400
	Certificate (public)	0.2681

B.6 Bloom et al. (2014)

Bloom et al. (2014) study the effect of working from home on employees' productivity via a randomized experiment at Ctrip, a NASDAQ-listed Chinese travel agency with 16000 employees. The main concern is whether or not working from leads to shirking. Ctrip decided to run a nine-month experiment on working from home. They asked the 996 employees in the airfare and hotel departments of the Shanghai call center if they were interested in working from home four days a week and one day in the office. 503 of these employees were interested and of these, 249 were qualified to take part on the basis of tenure, broadband access and access to private work space at home. Qualified employees were assigned to working from home if they had even-numbered birthdays so that those with odd-numbered birthdays formed the control group. The treatment and control groups were comprised of 131 and 118 employees respectively. The only difference between the two groups was location of work – both groups used the same equipment, faced the same workload and were compensated under the same pay system. The authors find a 13% increase in productivity of which the main source of improvement was a 9% increase in the number of minutes worked during a shift. The remaining 4% came from an increase in the number of calls per minute worked.

Table 17 reports standard deviations for treatment and control groups in the experiment as well as the standard deviation ratios for the main outcomes of interest in Bloom et al. (2014). Strong evidence of relative homoskedasticity with respect to treatment status presents in this study as well – suggesting that the balanced allocation would outperform

the FNA.

Table 17: Performance impact outcomes: standard deviations

Outcome Variables	Control S.D.	Treated S.D.	Ratio
Overall Performance	1.0049	1.0035	0.9986
Phone calls	0.9775	0.7502	0.7675
Log phone calls	0.2476	0.1764	0.7123
Log call per sec	0.0217	0.0299	1.3786
Log call length	0.2701	0.2729	1.0105

B.7 Galiani and McEwan (2013)

Galiani and McEwan (2013) use the Honduran PRAF experiment to study the impact of conditional cash transfers (CCT) on the likelihood of children to work versus enrolling in school. The PRAF experiment randomly allocated CCT's among 70 municipalities. These 70 were chosen out of a total of 298 on the basis of mean heights-for-age z-scores of first graders. The 70 municipalities were further assigned to four treatment arms termed G1, G2, G3, G4. G1 received CCT's in education and health. G2 received CCT's in addition to direct investment in education and health centers. G3 received only direct investments and finally, G4 served as the control group and received no interventions. The 70 municipalities were further divided into 5 strata each consisting of 14 municipalities on the basis of quintiles of mean height-for-age. Random assignment was performed within these strata (stratified randomization). The final sample consisted of 20 municipalities in G1, 20 in G2, 10 in G3, and 20 in G4. The authors match the experimental data with census data and use the latter to construct the outcomes of interest which are three dummy variables. The first is an indicator for whether a child is enrolled in and attending school during the time of the census. The second indicates whether the child worked during the week prior to the census or conditional on a negative response to the former, whether they reported non-wage employment in a family farm or business. The third indicates whether the child worked exclusively on household chores. The authors find that overall, children eligible for CCTs were 8% more likely to enroll in school and 3% less likely to work.

Table 18 reports standard deviations for the outcomes of interest across treatment arms. Again, for a given outcome, pairwise comparisons across arms yields evidence of relative homoskedasticity with respect to treatment status. We conclude that in this case, the balanced allocation would outperform the FNA.

Table 18: St. dev. across treatment groups G1, G2, G3, G4

Outcome	G1	G2	G3	G4
Enrolled in school	0.4393	0.4474	0.4812	0.4769
Works outside home	0.2637	0.2267	0.2893	0.2986
Works only in home	0.3015	0.2841	0.3478	0.3409

B.8 Bryan et al. (2015)

Bryan et al. (2015) conduct a field experiment to study efficacy of peer intermediation in mitigating adverse selection and moral hazard in credit markets. To identify the effects of peer screening and enforcement, they use a two-stage referral incentive field experiment. The experiment was conducted through Opportunity Finance South Africa (Opportunity), a for-profit lender in the consumer micro-loan market. Over the period of February 2008 through July 2009, Opportunity offered individuals approved for a loan the option to participate in its “Refer-A-Friend” program. Referred individuals earned R40 if they brought in a referral card and were approved for a loan. The referrer could earn R100 for referring someone who was subsequently approved for and/or repaid the loan, depending on the referrer’s incentive contract. Referrers were randomly assigned to one of two ex-ante incentive contracts:

- Approval incentives: the referrer would be paid only if the referred was approved for a loan.
- Repayment incentive: the referrer would be paid only if the referred successfully repaid the loan.

Among referrers whose referred friends were approved for a loan, Opportunity randomly selected half to be surprised with an ex-post incentive change:

- Half among the ex-ante approval group were phoned and told that in addition to the R100 approval bonus, they would receive an additional R100 if the loan was successfully repaid by the referrer.
- Half among the ex-ante repayment group were phoned and told that they would receive the R100 now, and that receipt of the bonus would no longer be conditional on repayment of the loan by the referrer.

The overall incentive structure is as follows

- Ex-ante and ex-post approval (EA = A): no enforcement or screening incentive.

- Ex-ante repayment and ex-post approval (EA = R): screening incentive.
- Ex-ante approval and ex-post repayment (EA = A, EP = R): Enforcement incentive.
- Ex-ante repayment and ex-post repayment (EA = R, EP = R): Enforcement and screening incentive.

The authors find no evidence of screening but do find large enforcement effects.

Table 19 reports standard deviations in the main outcomes of interest in Bryan et al. (2015) across the four treatment arms. For a given outcome, pairwise comparisons across arms yields evidence of relative homoskedasticity with respect to treatment status. We conclude that in this case, the balanced allocation would outperform the FNA.

Table 19: St. dev. across treatment groups G1, G2, G3, G4

Outcome	EA = A	EA = A, EP = R	EA = R	EA = R, EP = R
Penalty interest	0.4919	0.4407	0.4488	0.5043
Positive balance owing at maturity	0.4086	0.2959	0.3613	0.4225
Proportion of value owing at maturity	0.4088	0.3106	0.3153	0.5526
Loan charged off	0.3652	0.2147	0.2917	0.3950

C Additional Simulations

C.1 MSE for $m = 50$

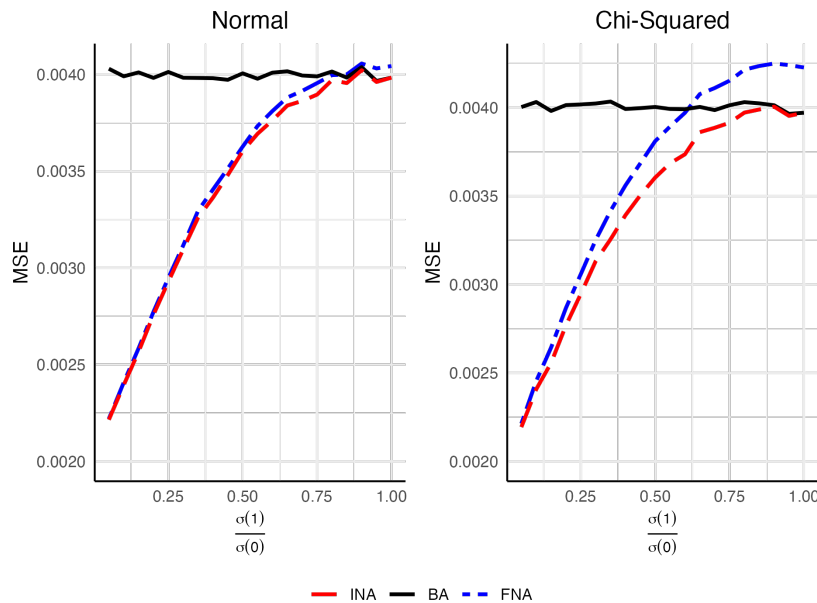


Figure 10: MSE under the Infeasible Neyman Allocation (INA), Balanced Allocation (BA) and Feasible Neyman Allocation (FNA) when $\varepsilon(1), \varepsilon(0)$ has the Normal distribution ($N(0, 1)$) and the standardized Chi-Squared distribution ($\chi_1^2 - 1$).




Figure 11: MSE under the Infeasible Neyman Allocation (INA), Balanced Allocation (BA) and Feasible Neyman Allocation (FNA) when $\varepsilon(1), \varepsilon(0)$ has the standardized Pareto(1, s) distribution, with $s \in \{3, 4, 5\}$.

C.2 MSE of Various Solutions with Pareto Distribution

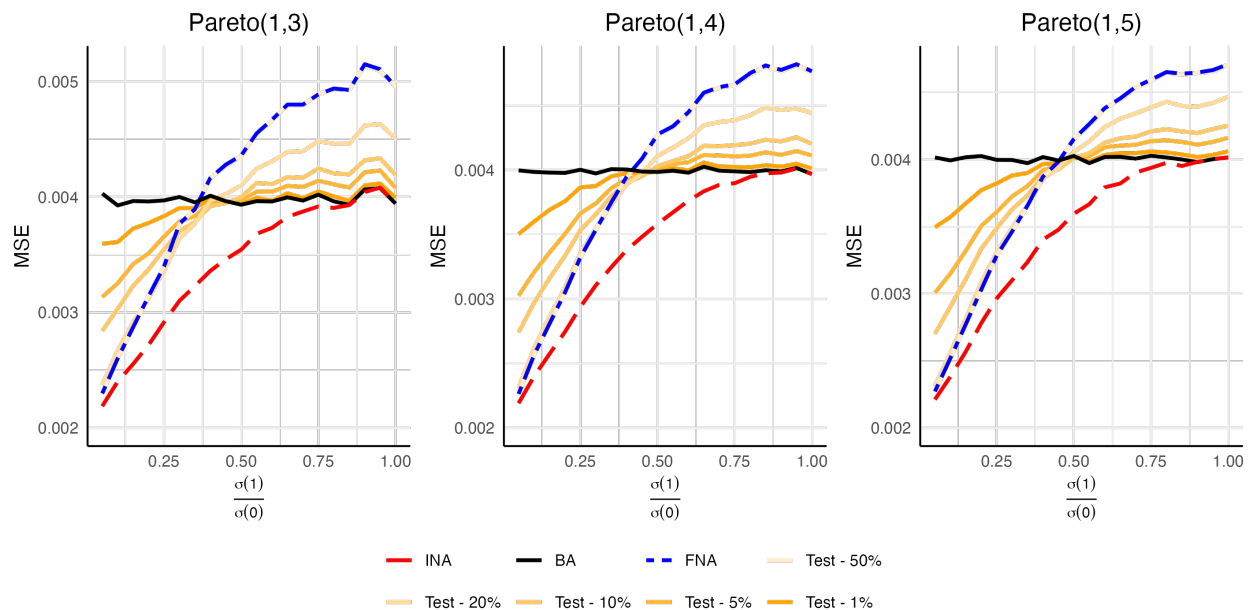


Figure 12: MSE under the Infeasible Neyman Allocation (INA), Balanced Allocation (BA) and Feasible Neyman Allocation (FNA) when $\varepsilon(1), \varepsilon(0)$ has the standardized Pareto(1, s) distribution, with $s \in \{3, 4, 5\}$.

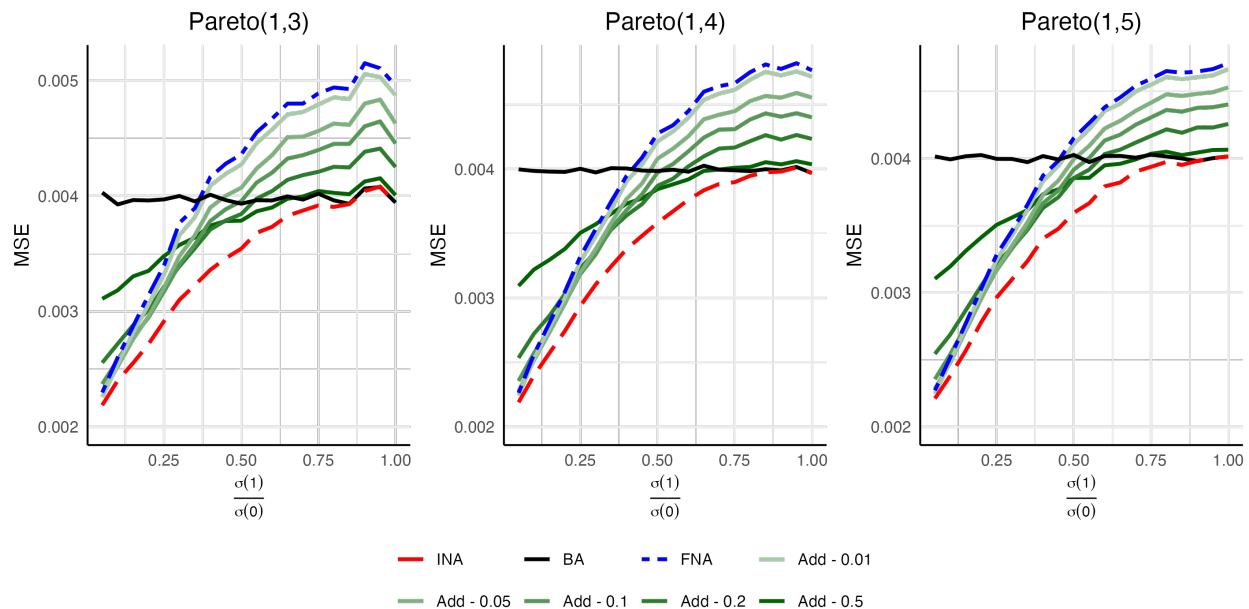


Figure 13: MSE under the Infeasible Neyman Allocation (INA), Balanced Allocation (BA) and Feasible Neyman Allocation (FNA) when $\varepsilon(1), \varepsilon(0)$ has the standardized Pareto(1, s) distribution, with $s \in \{3, 4, 5\}$.

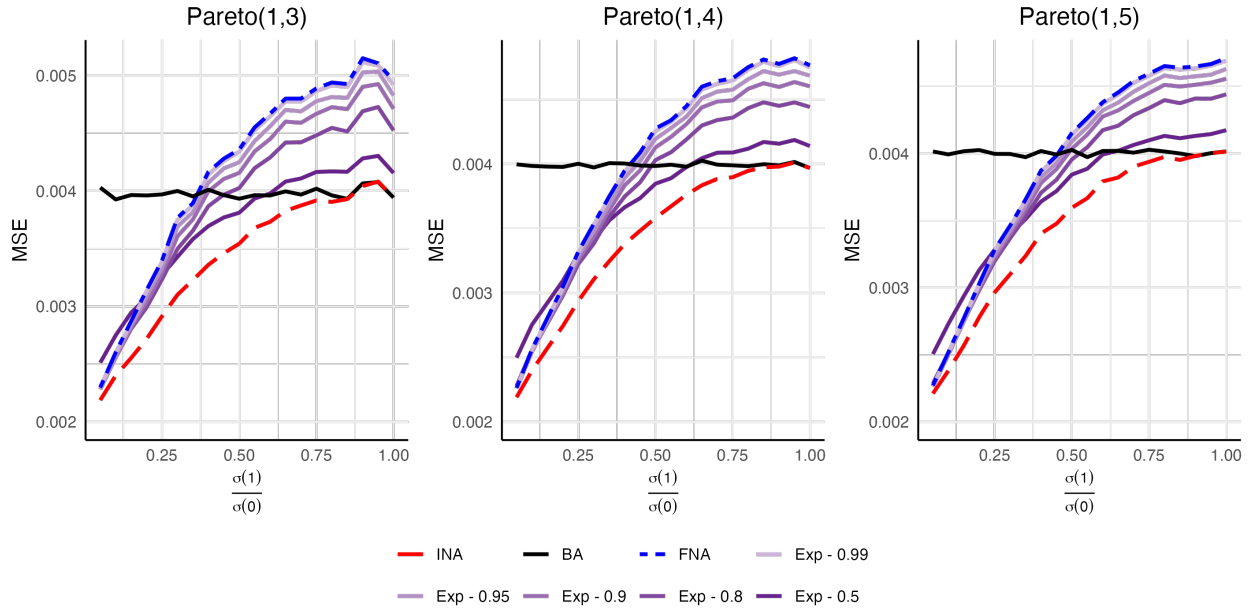


Figure 14: MSE under the Infeasible Neyman Allocation (INA), Balanced Allocation (BA) and Feasible Neyman Allocation (FNA) when $\varepsilon(1), \varepsilon(0)$ has the standardized Pareto(1, s) distribution, with $s \in \{3, 4, 5\}$.

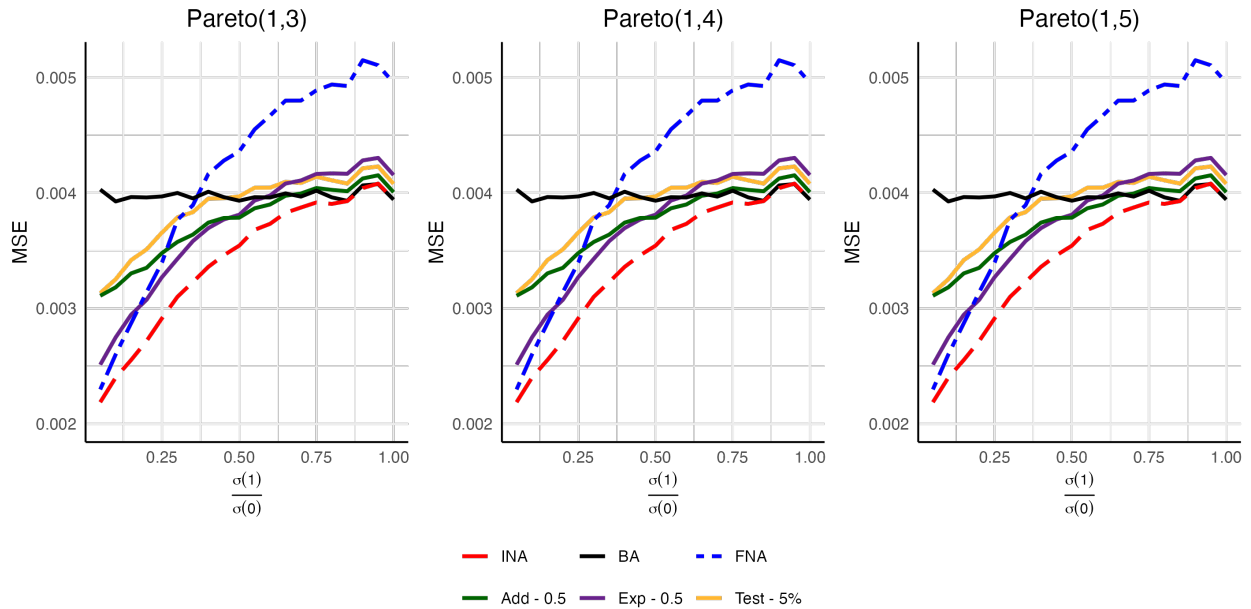


Figure 15: MSE under the Infeasible Neyman Allocation (INA), Balanced Allocation (BA) and Feasible Neyman Allocation (FNA) when $\varepsilon(1), \varepsilon(0)$ has the standardized Pareto(1, s) distribution, with $s \in \{3, 4, 5\}$.