

Supplementary Materials of “Fair Adaptive Experiments”

Waverly Wei

Division of Biostatistics

University of California, Berkeley

Xinwei Ma

Department of Economics

University of California, San Diego

Jingshen Wang *

Division of Biostatistics

University of California, Berkeley

Contents

A Proof of Results	2
A.1 Proof of Theorem 1	2
A.2 Proof of Theorem 2	3
A.3 Proof of Theorem 3	4
B Additional Details on Implementation	5
B.1 Stage 1	5
B.2 The Welfare Constraint and $\delta(\cdot)$	6
B.3 The Envy-freeness Constraint	6
C Additional Simulation Results	7
D Synthetic Data Analyses	8
D.1 Effect of Bundling Health Insurance with Microfinancing on Loan Renewal	8
D.2 Effect of Genetically-guided Therapy on Treating Major Depressive Disorder Patients	10

*Corresponding author: jingshenwang@berkeley.edu

A Proof of Results

A.1 Proof of Theorem 1

Proof. To start, we note that the convergence of \widehat{p}_{tj} follows from the standard law of large numbers, as the covariates X_{it} are iid. To be precise, we have that $\widehat{p}_{tj} - \mathbb{P}[X_{it} \in \mathcal{S}_j] = O_p(1/\sqrt{N_t})$.

Next, we show that moments of the potential outcomes are consistently estimated. For example, consider

$$\frac{\sum_{s=1}^t \sum_{i=1}^{n_s} \mathbb{1}_{(X_{is} \in \mathcal{S}_j)} D_{is} Y_{is}}{\sum_{s=1}^t \sum_{i=1}^{n_s} \mathbb{1}_{(X_{is} \in \mathcal{S}_j)} D_{is}} = \frac{\sum_{s=1}^t \sum_{i=1}^{n_s} \mathbb{1}_{(X_{is} \in \mathcal{S}_j)} D_{is} (Y_{is}(1) - \mu_j(1))}{\sum_{s=1}^t \sum_{i=1}^{n_s} \mathbb{1}_{(X_{is} \in \mathcal{S}_j)} D_{is}} + \mu_j(1),$$

where $\mu_j(1) = \mathbb{E}[Y_{is}(1)|X_{is} \in \mathcal{S}_j]$. We further notice that by the law of iterated expectation,

$$\begin{aligned} \mathbb{E} \left[\mathbb{1}_{(X_{is} \in \mathcal{S}_j)} D_{is} (Y_{is}(1) - \mu_j(1)) \middle| \mathcal{F}_{s-1} \right] &= \mathbb{E} \left[\mathbb{1}_{(X_{is} \in \mathcal{S}_j)} \widehat{e}_{sj}^* \left(\mu(X_{is}, 1) - \mu_j(1) \right) \middle| \mathcal{F}_{s-1} \right] \\ &= \widehat{e}_{sj}^* \mathbb{E} \left[\mathbb{1}_{(X_{is} \in \mathcal{S}_j)} \left(\mu(X_{is}, 1) - \mu_j(1) \right) \middle| \mathcal{F}_{s-1} \right] = 0, \end{aligned}$$

where we use $\mu(x, 1) = \mathbb{E}[Y_{is}(1)|X_{is} = x]$. Here and throughout this Supplementary Materials, we use \mathcal{F}_{s-1} to denote the sigma-algebra formed by all information up to time $s - 1$. As a result, straightforward variance computation suggests that

$$\sum_{s=1}^t \sum_{i=1}^{n_s} \mathbb{1}_{(X_{is} \in \mathcal{S}_j)} D_{is} (Y_{is}(1) - \mu_j(1)) = O_p(\sqrt{N_t}).$$

Finally, we notice that the denominator satisfies

$$\sum_{s=1}^t \sum_{i=1}^{n_s} \mathbb{1}_{(X_{is} \in \mathcal{S}_j)} D_{is} = \underbrace{\sum_{s=1}^t \sum_{i=1}^{n_s} p_j \widehat{e}_{sj}^*}_{\gtrsim N_t} + \underbrace{\sum_{s=1}^t \sum_{i=1}^{n_s} \left(\mathbb{1}_{(X_{is} \in \mathcal{S}_j)} D_{is} - p_j \widehat{e}_{sj}^* \right)}_{= O_p(\sqrt{N_t})}.$$

The first term on the rhs is bounded below by the order of N_t due to our assumption on p_j and the feasibility constraint, while the second term on the rhs is a sum of martingale difference sequence, and its probability order follows from standard variance calculations. As a result, we have that

$$\sum_{s=1}^t \sum_{i=1}^{n_s} \mathbb{1}_{(X_{is} \in \mathcal{S}_j)} D_{is} \gtrsim_p N_t.$$

The same strategy can be used to prove the consistency of other estimated moments of the potential outcomes. \square

A.2 Proof of Theorem 2

Proof. The proof has two parts. **Part I: consistency of \widehat{e}_{tj}^* .** From Theorem 1, we have shown that moments of the potential outcomes are consistently estimated, which implies that the objective function in Problem B is uniformly consistent:

$$\begin{aligned} & \sup_{c_2 \leq e_j \leq 1 - c_2; j=1,2,\dots,m} \left| \sum_{j=1}^m \widehat{p}_{t-1,j} \left(\frac{\widehat{\sigma}_{t-1,j}^2(1)}{e_j} + \frac{\widehat{\sigma}_{t-1,j}^2(0)}{1 - e_j} \right) - \sum_{j=1}^m p_j \left(\frac{\sigma_j^2(1)}{e_j} + \frac{\sigma_j^2(0)}{1 - e_j} \right) \right| \\ & \leq \max_{j=1,2,\dots,m} |\widehat{p}_{t-1,j} - p_j| O_p(1) + \frac{1}{c_2} \max_{j=1,2,\dots,m} |\widehat{\sigma}_{t-1,j}^2(1) - \sigma_j^2(1)| + \frac{1}{c_2} \max_{j=1,2,\dots,m} |\widehat{\sigma}_{t-1,j}^2(0) - \sigma_j^2(0)| \\ & = O_p(1/\sqrt{N_{t-1}}). \end{aligned}$$

In addition, it is routine to check that the solution to Problem A is unique.

We now discuss the welfare constraint, since it involves the estimated subgroup treatment effect. First assume $\tau_j > 0$, then the welfare constraint in Problem A for subgroup j reduces to $e_j \geq 0.5$. In this case, $\mathbb{P}[\widehat{\tau}_{t-1,j} > 0] \rightarrow 1$, which means with probability approaching 1, the welfare constraint in Problem B corresponds to

$$\log \left(\frac{e_j}{1 - e_j} \right) \geq -\frac{\sqrt{N_{t-1}}\delta(N_{t-1})}{\sqrt{N_{t-1}}\tau_j + O_p(1)} \rightarrow_p 0,$$

where \rightarrow_p denotes convergence in probability. Similarly, if $\tau_j < 0$, then with probability approaching 1, the welfare constraint takes the form

$$\log \left(\frac{e_j}{1 - e_j} \right) \leq \frac{\sqrt{N_{t-1}}\delta(N_{t-1})}{\sqrt{N_{t-1}}(-\tau_j) + O_p(1)} \rightarrow_p 0.$$

When $\tau_j = 0$, the constraint becomes

$$\begin{aligned} \widehat{\tau}_{t-1,j} > 0 : \quad & \log \left(\frac{e_j}{1 - e_j} \right) \geq -\frac{\sqrt{N_{t-1}}\delta(N_{t-1})}{\sqrt{N_{t-1}}\widehat{\tau}_{t-1,j}} \\ \widehat{\tau}_{t-1,j} < 0 : \quad & \log \left(\frac{e_j}{1 - e_j} \right) \leq \frac{\sqrt{N_{t-1}}\delta(N_{t-1})}{\sqrt{N_{t-1}}(-\widehat{\tau}_{t-1,j})}. \end{aligned}$$

In either case, the constraint is nonbinding, and hence reduces to the welfare constraint in Problem A. To summarize, we showed that both the objective function and the feasible set of treatment allocation rules converge to their oracle counterparts. The consistency of \widehat{e}_{tj}^* then follows from standard M-estimation arguments (such as Theorem 2.1 in [Newey and McFadden 1994](#)).

Part II: convergence of the actual treatment allocation. Now consider the actual

treatment allocation after stage T :

$$\frac{\sum_{s=1}^T \sum_{i=1}^{n_s} \mathbb{1}_{(X_{is} \in \mathcal{S}_j)} D_{is}}{\sum_{s=1}^T \sum_{i=1}^{n_s} \mathbb{1}_{(X_{is} \in \mathcal{S}_j)}} = \underbrace{\frac{\sum_{s=1}^T \sum_{i=1}^{n_s} \mathbb{1}_{(X_{is} \in \mathcal{S}_j)} (D_{is} - \widehat{e}_{sj}^*)}{\sum_{s=1}^T \sum_{i=1}^{n_s} \mathbb{1}_{(X_{is} \in \mathcal{S}_j)}}}_{= O_p(1/\sqrt{N})} + \frac{\sum_{s=1}^T \sum_{i=1}^{n_s} \mathbb{1}_{(X_{is} \in \mathcal{S}_j)} \widehat{e}_{sj}^*}{\sum_{s=1}^T \sum_{i=1}^{n_s} \mathbb{1}_{(X_{is} \in \mathcal{S}_j)}},$$

where the probabilistic order of the first term on the rhs follows from mean and variance calculations.

To analyze the second term, we notice that it can be further written as

$$\underbrace{\frac{\sum_{s=1}^T \sum_{i=1}^{n_s} (\mathbb{1}_{(X_{is} \in \mathcal{S}_j)} - p_j) \widehat{e}_{sj}^*}{\sum_{s=1}^T \sum_{i=1}^{n_s} \mathbb{1}_{(X_{is} \in \mathcal{S}_j)}}}_{= O_p(1/\sqrt{N})} + \frac{\sum_{s=1}^T n_s \widehat{e}_{sj}^*}{N} \underbrace{\frac{p_j}{\frac{1}{N} \sum_{s=1}^T \sum_{i=1}^{n_s} \mathbb{1}_{(X_{is} \in \mathcal{S}_j)}}}_{\rightarrow_p 1}.$$

Again the probabilistic order of the first term follows from mean and variance calculations.

Finally, given the consistency of \widehat{e}_{sj}^* as we discussed in Part I and our Assumption 3 on n_t ,

$$\frac{\sum_{s=1}^T n_s \widehat{e}_{sj}^*}{N} - e_j^* = o_p(1),$$

which closes the proof. \square

A.3 Proof of Theorem 3

Proof. Part I: asymptotic normality of $\widehat{\tau}_j$. We write

$$\sqrt{N}(\widehat{\tau}_j - \tau_j) = \frac{\sqrt{N} \sum_{t=1}^T \sum_{i=1}^{n_s} \mathbb{1}_{(X_{it} \in \mathcal{S}_j)} D_{it} (Y_{it} - \mu_j(1))}{\sum_{t=1}^T \sum_{i=1}^{n_s} \mathbb{1}_{(X_{it} \in \mathcal{S}_j)} D_{it}} - \frac{\sqrt{N} \sum_{t=1}^T \sum_{i=1}^{n_s} \mathbb{1}_{(X_{it} \in \mathcal{S}_j)} (1 - D_{it}) (Y_{it} - \mu_j(0))}{\sum_{t=1}^T \sum_{i=1}^{n_s} \mathbb{1}_{(X_{it} \in \mathcal{S}_j)} (1 - D_{it})}.$$

By Theorem 2,

$$\frac{1}{N} \sum_{t=1}^T \sum_{i=1}^{n_s} \mathbb{1}_{(X_{is} \in \mathcal{S}_j)} D_{is} - p_j e_j^* = o_p(1), \quad \frac{1}{N} \sum_{t=1}^T \sum_{i=1}^{n_s} (1 - D_{is}) \mathbb{1}_{(X_{is} \in \mathcal{S}_j)} - p_j (1 - e_j^*) = o_p(1).$$

Therefore, to prove the asymptotic normality of $\widehat{\tau}_j$, we can study

$$\frac{1}{\sqrt{N}} \frac{\sum_{t=1}^T \sum_{i=1}^{n_s} \mathbb{1}_{(X_{it} \in \mathcal{S}_j)} D_{it} (Y_{it} - \mu_j(1))}{p_j e_j^*} - \frac{1}{\sqrt{N}} \frac{\sum_{t=1}^T \sum_{i=1}^{n_s} \mathbb{1}_{(X_{it} \in \mathcal{S}_j)} (1 - D_{it}) (Y_{it} - \mu_j(0))}{p_j (1 - e_j^*)},$$

which is a martingale difference sequence with respect to the natural filtration.

We can derive the conditional variance as

$$\frac{1}{N} \sum_{t=1}^T \sum_{i=1}^{n_t} \mathbb{V} \left[\frac{\mathbb{1}_{(X_{it} \in \mathcal{S}_j)} D_{it} (Y_{it} - \mu_j(1))}{p_j e_j^*} - \frac{\mathbb{1}_{(X_{it} \in \mathcal{S}_j)} (1 - D_{it}) (Y_{it} - \mu_j(0))}{p_j (1 - e_j^*)} \middle| \mathcal{F}_{t-1} \right]$$

$$= \frac{1}{N} \sum_{t=1}^T \sum_{i=1}^{n_t} \left\{ \frac{\widehat{e}_{tj}^* \sigma_j^2(1)}{p_j e_j^{*2}} + \frac{(1 - \widehat{e}_{tj}^*) \sigma_j^2(0)}{p_j (1 - e_j^*)^2} \right\}.$$

By Theorem 2, the above converges to $v_j^2(e_j^*)$. Given our assumptions on the moments of the potential outcomes, it is straightforward to verify Lindeberg's condition. The asymptotic normality then follows from the martingale central limit theorem (see, for example, [Hall and Heyde 2014](#)).

Part II: asymptotic normality of $\widehat{\tau}$. We write

$$\sqrt{N}(\widehat{\tau} - \tau) = \sqrt{N} \left(\sum_{j=1}^m \widehat{p}_j (\widehat{\tau}_j - \tau_j) \right) + \sqrt{N} \left(\sum_{j=1}^m \widehat{p}_j \tau_j - \tau \right).$$

From the discussion in Part I, we have that

$$\sqrt{N} \left(\sum_{j=1}^m \widehat{p}_j (\widehat{\tau}_j - \tau_j) \right) = \frac{1}{\sqrt{N}} \sum_{t=1}^T \sum_{i=1}^{n_t} \sum_{j=1}^m \mathbf{1}_{(X_{it} \in \mathcal{S}_j)} \left(\frac{D_{it}(Y_{it} - \mu_j(1))}{e_j^*} - \frac{(1 - D_{it})(Y_{it} - \mu_j(0))}{1 - e_j^*} \right) + o_p(1).$$

In addition,

$$\sqrt{N} \left(\sum_{j=1}^m \widehat{p}_j \tau_j - \tau \right) = \frac{1}{\sqrt{N}} \sum_{t=1}^T \sum_{i=1}^{n_t} \sum_{j=1}^m \mathbf{1}_{(X_{it} \in \mathcal{S}_j)} (\tau_j - \tau).$$

The final asymptotic normality result then follows from the martingale central limit theorem and (conditional) variance calculation. \square

B Additional Details on Implementation

B.1 Stage 1

In this section, we first provide implementation details of the Stage 1 experiment as mentioned in Algorithm 1 (lines 1-2) in the main manuscript. In Stage 1, we assign treatment with probability $e_{1j} = 1/2$, for $j = 1, \dots, m$. After collecting the outcome information, we update the subgroup treatment effect estimates and the associated variances as

$$\begin{aligned} \widehat{p}_{1j} &= \frac{\sum_{i=1}^{n_1} \mathbf{1}_{(X_{i1} \in \mathcal{S}_j)}}{n_1}, \\ \bar{Y}_{1j}(1) &= \frac{\sum_{i=1}^{n_1} \mathbf{1}_{(X_{i1} \in \mathcal{S}_j)} D_{i1} Y_{i1}}{\sum_{i=1}^{n_1} \mathbf{1}_{(X_{i1} \in \mathcal{S}_j)} D_{i1}}, \quad \bar{Y}_{1j}(0) = \frac{\sum_{i=1}^{n_1} \mathbf{1}_{(X_{i1} \in \mathcal{S}_j)} (1 - D_{i1}) Y_{i1}}{\sum_{i=1}^{n_1} \mathbf{1}_{(X_{i1} \in \mathcal{S}_j)} (1 - D_{i1})}, \\ \widehat{\tau}_{1j} &= \bar{Y}_{1j}(1) - \bar{Y}_{1j}(0), \\ \widehat{\sigma}_{1j}^2(1) &= \frac{\sum_{i=1}^{n_1} \mathbf{1}_{(X_{i1} \in \mathcal{S}_j)} D_{i1} (Y_{i1} - \bar{Y}_{1j}(1))^2}{\sum_{i=1}^{n_1} \mathbf{1}_{(X_{i1} \in \mathcal{S}_j)} D_{i1}}, \end{aligned}$$

$$\widehat{\sigma}_{1j}^2(0) = \frac{\sum_{i=1}^{n_1} \mathbb{1}_{(X_{i1} \in \mathcal{S}_j)} (1 - D_{i1}) (Y_{i1} - \bar{Y}_{1j}(0))^2}{\sum_{i=1}^{n_1} \mathbb{1}_{(X_{i1} \in \mathcal{S}_j)} (1 - D_{i1})}.$$

B.2 The Welfare Constraint and $\delta(\cdot)$

We first discuss why the naive approach (i.e., setting $\delta(\cdot) = 0$ in the welfare constraint of Problem B) may lead to inconsistent treatment allocation. Consider the scenario where $\tau_j = 0$. The welfare constraint in Problem A becomes non-binding (i.e., it does not restrict e_j). However, with $\delta(\cdot) = 0$, the welfare constraint in Problem B becomes

$$\log\left(\frac{e_j}{1 - e_j}\right) \widehat{\tau}_{t-1,j} \geq 0 \quad \Leftrightarrow \quad \log\left(\frac{e_j}{1 - e_j}\right) \sqrt{N_{t-1}} \widehat{\tau}_{t-1,j} \geq 0.$$

As we showed in Theorem 3, $\sqrt{N_{t-1}} \widehat{\tau}_{t-1,j}$ is asymptotically $\mathcal{N}(0, v_j^2(e_j^*))$ if $\tau_j = 0$. This implies that the above constraint becomes

$$\begin{aligned} \log\left(\frac{e_j}{1 - e_j}\right) \geq 0 &\Leftrightarrow e_j \geq 0.5 && \text{with probability approximately } \frac{1}{2} \\ \log\left(\frac{e_j}{1 - e_j}\right) \leq 0 &\Leftrightarrow e_j \leq 0.5 && \text{with probability approximately } \frac{1}{2}. \end{aligned}$$

In other words, the naive approach leads to a feasible set that alternates between $[0.5, 1]$ and $[0, 0.5]$ (subject to the additional envy-free and feasibility constraints). In this case, the solution to Problem B will not converge to the oracle solution (from Problem A) even asymptotically.

The relaxation we introduced can be intuitively understood as a t-test of the hypothesis that $\tau_j = 0$ with a diverging threshold/critical value. The particular choice of $\delta(N_{t-1}) = \sqrt{\log N_{t-1}/N_{t-1}}$ is related to Schwarz's minimum BIC rule. As we discussed in the main manuscript, it is possible to further scale the welfare constraint by a measure of the randomness in $\widehat{\tau}_{t-1,j}$, which takes the form

$$\widehat{v}_{t-1,j}^2 = \frac{1}{\widehat{p}_{t-1,j}} \left(\frac{\widehat{\sigma}_{t-1,j}^2(1)}{\widehat{e}_{t-1,j}} + \frac{\widehat{\sigma}_{t-1,j}^2(0)}{1 - \widehat{e}_{t-1,j}} \right), \quad \text{where } \widehat{e}_{t-1,j} = \frac{\sum_{s=1}^{t-1} \sum_{i=1}^{n_s} \mathbb{1}_{(X_{is} \in \mathcal{S}_j)} D_{is}}{\sum_{s=1}^{t-1} \sum_{i=1}^{n_s} \mathbb{1}_{(X_{is} \in \mathcal{S}_j)}}.$$

As a final remark, the welfare constraint in Problem B is not uniformly consistent for that in Problem A. In particular, if the true treatment effect is small in magnitude (say, if we model $\tau_j = a_j/\sqrt{N_{t-1}}$ for some $a_j > 0$), then the welfare constraint in Problem A and B may not coincide even asymptotically. Loosely speaking, the relaxation we introduced leads to a down weight of the welfare constraint when the treatment effect is small in magnitude or when it is imprecisely estimated.

B.3 The Envy-freeness Constraint

Lastly, we provide some guidance on selecting the fairness constraints. Figure 1 below compares estimation efficiency under different values of c_1 , demonstrating that the lowest estimation efficiency corresponds to the smallest c_1 , suggesting that the fairest treatment allocation might lead to a

compromise in estimation efficiency. We recognize that the "optimal" balance between estimation efficiency and fairness can vary based on the specific context of application. For general guidance, we recommend setting $0.1 \leq c_1 \leq 0.3$ when fairness concern is slightly prioritized over estimation efficiency. Conversely, when estimation efficiency is slightly prioritized over fairness, we recommend setting $0.3 \leq c_1 \leq 0.5$ to achieve a more balanced trade-off between estimation efficiency and fairness.

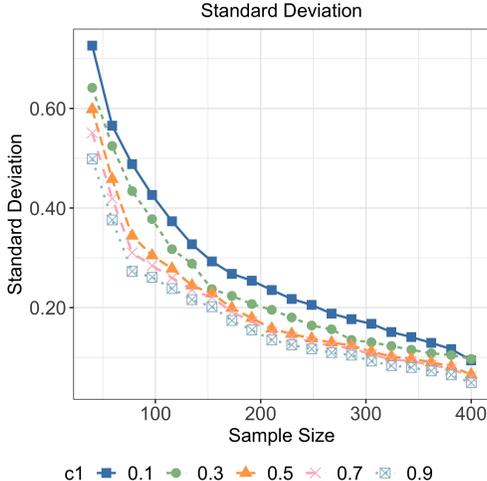


Figure 1: Standard deviation comparison with respect to various sample sizes under different values of c_1 , where c_1 is the user-specified value in the envy-freeness constraint.

C Additional Simulation Results

We provide additional simulation results in this section to further illustrate our proposed design. Figure 2 shows the empirical treatment allocation for Group 1 following DGP-1 in our manuscript. The empirical treatment allocation refers to the treatment assignment probability derived from the optimization Problem B. The dashed line indicates the oracle treatment allocation solved from Problem A. Figure 2 shows that the empirical subgroup treatment allocation converges to the oracle treatment allocation fairly quickly.

In Figure 3, we provide a comparison with the rerandomization design. The simulation setup follows DGP-1 in the main manuscript, and we further generate an additional covariate $W_{it} \sim \text{Bernoulli}(0.4)$. Given that rerandomization requires a pre-specified treatment assignment probability, we use the oracle subgroup treatment assignment probabilities from Problem A. Then, within each subgroup, we perform rerandomization to balance the covariate W . Figure 3 shows that given known oracle treatment assignment probabilities, the empirical treatment allocation using rerandomization can approximate the oracle allocation quite closely. However, when the oracle treatment assignment probabilities are unknown, our design can sequentially approach the oracle treatment allocation by learning from the collected data.

In Table 1, we provide the coverage probabilities of our proposed design, complete randomization (CR), and the doubly adaptive debiased coin design (DBCD) under the two data-generating processes (DGPs) as described in the main manuscript in the simulation studies. We observe that our proposed design reaches the nominal level coverage, which verifies the validity of our constructed confidence interval in Section 2 of the main manuscript.

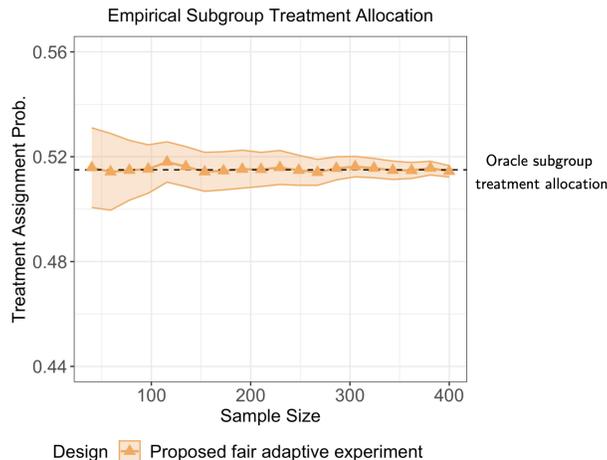


Figure 2: Empirical subgroup treatment allocation with respect to various sample sizes. The blue dashed line indicates the oracle treatment allocation.

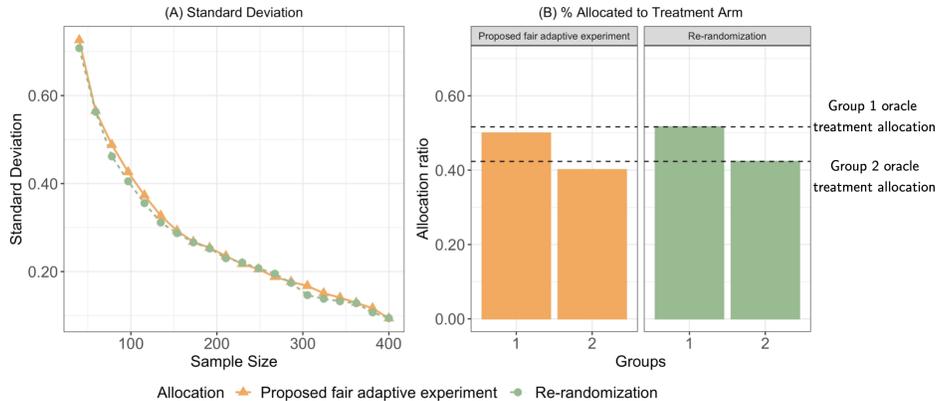


Figure 3: Comparison of our proposed design with the rerandomization design. The dashed lines indicate oracle treatment allocations derived from solving Problem A in our manuscript.

D Synthetic Data Analyses

D.1 Effect of Bundling Health Insurance with Microfinancing on Loan Renewal

In a field experiment, Banerjee et al. [Banerjee et al. \(2014\)](#) studied the effect of bundling health insurance with microfinancing on loan renewal. Their original data consists of household clients of SKS Microfinance, the largest microfinance institution in India in 2006. The study randomly

Table 1: Coverage probabilities and the associated standard errors of three designs under two data-generating processes (DGPs) described in the main manuscript. “CR” refers to complete randomization. “DBCD” refers to doubly adaptive biased coin design.

		Proposed design	CR	DBCD
DGP 1	$n = 100$	0.93 (0.02)	0.97 (0.02)	0.92 (0.03)
	$n = 400$	0.95 (0.01)	0.96 (0.01)	0.93 (0.02)
DGP 2	$n = 100$	0.96 (0.02)	0.97 (0.02)	0.94 (0.03)
	$n = 400$	0.94 (0.01)	0.95 (0.01)	0.92 (0.03)

assigned clients to either the treatment arm, where loan renewal was bundled with health insurance, or the control arm, where clients could renew their loans without purchasing health insurance. The binary outcome variable represents the client’s decision to renew the loan. We refer interested readers to the original paper for additional details on the institutional background and the study.

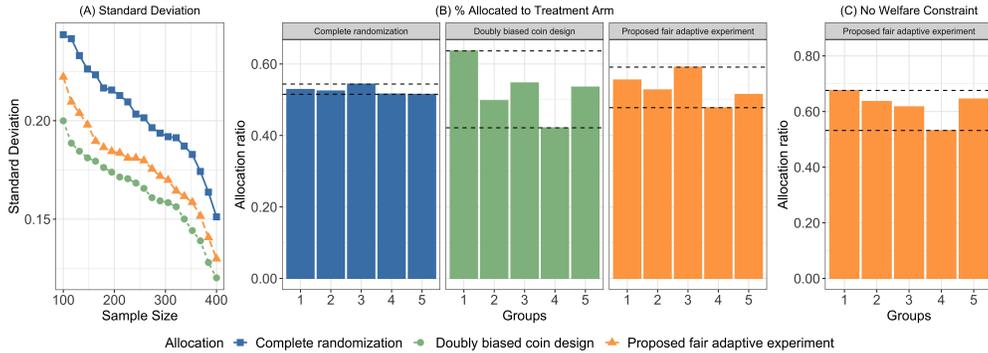


Figure 4: Comparison of the proposed adaptive experiment, the complete randomization design, and the doubly adaptive biased coin design using synthetic data. (A) shows the standard deviations. (B) shows the percentage of participants allocated to the treatment arm in each group. (C) shows the treatment allocation under our proposed experiment strategy without the welfare constraint.

In this section, we conduct a synthetic data analysis to illustrate the effectiveness of our proposed design. To be specific, we first employ the dataset in [Banerjee *et al.* \(2014\)](#) to estimate group-level mean potential outcomes and their variances, and then use these estimates to generate synthetic data. In our synthetic data analysis, we focus on estimating the log relative risk parameter and examining the estimation efficiency and fairness of treatment allocations across different assignment regimes. Fairness is crucial in this context, as lower-income households were more likely to forgo the loan due to the additional financial burden of purchasing health insurance, potentially leading to disproportionate effects. The results show that our approach significantly improves upon complete randomization regarding statistical estimation efficiency. Compared with the doubly adaptive biased coin design, treatment assignment probabilities produced by our procedure are less extreme, and they also exhibit much less between-group variation.

To generate the synthetic data, we first subset the original data in [Banerjee *et al.* \(2014\)](#) and select participants with a chronic disease. Then, we partition the data into five groups based on

annual expenditures. We then estimate the mean potential outcomes and variances thereof at the group level, which are used to generate synthetic potential outcome variables.

After generating the synthetic data, we consider a fully adaptive experiment setting where the treatment is assigned sequentially. We first sample $n_1 = 40$ participants from the synthetic data to be included in the first stage. In each subsequent stage, $t = 2, \dots, m$, we randomly sample $n_t = 1$ participants from the synthetic data. The time horizon ranges from $m \in \{40, 80, \dots, 480\}$ and compare three adaptive experiment design strategies in terms of their estimation efficiency and allocation fairness: (1) the complete randomization (CR) design, (2) the doubly adaptive biased coin design (DBCD), and (3) our proposed adaptive experiment strategy.

We summarize the synthetic data analysis results in Figure 4. Panel (A) demonstrates that the two adaptive design methods (our proposal and the DBCD) achieve higher estimation efficiency than the non-adaptive procedure (complete randomization). Compared with the DBCD design, our strategy has slightly lower efficiency. We conjecture that the slightly compromised efficiency is due to our additional fairness constraints. Figure 4 (B) shows the fractions of participants allocated to the treatment arm in the five groups. As expected, the actual treatment assignment probability is very close to 0.5 with complete randomization. Under the DBCD design, the largest treatment assignment probability is about 70% (i.e., around 70% of participants in the first group are expected to be assigned to receive treatment). In addition, there is also a large discrepancy in treatment assignment probabilities across groups: the difference can be as large as 20%, which clearly raises fairness concerns. To compare, treatment probabilities constructed by our fair adaptive design strategy are less extreme (i.e., they are closer to 50%), and they exhibit much fewer variations across the groups. Figure 4 also suggests a natural tension between fairness and welfare constraints. Panel (C) shows the assignment probabilities using our proposed method without the welfare constraint. In this case, all groups receive higher treatment assignment probabilities (relative to panel B).

In sum, our proposed adaptive experimental design strategy achieves fair treatment allocation and accounts for participants’ welfare, with limited sacrifice in estimation efficiency.

D.2 Effect of Genetically-guided Therapy on Treating Major Depressive Disorder Patients

In this section, we provide an additional synthetic case study to showcase the performance of our proposed design. Here, we consider a clinical dataset investigating the treatment effect of genetically-guided therapy on treating major depressive disorder patients (Ruaño *et al.*, 2021). As genetically-guided therapy potentially benefits patients, we hope to have fair treatment allocations in different subgroups. In addition, as patient subgroups may respond differently to the therapy, it is natural to incorporate the welfare constraint to maximize patients’ welfare. We aim to design adaptive experiments to efficiently estimate the treatment effect of genetically guided therapy on major depressive disorder (MDD) patients under fairness constraints. The original trial was conducted at the Institute of Living at Hartford Hospital, consisting of 1459 patients (Tortora *et al.*, 2020). There are two considered therapies: (1) the standard therapy (control arm), and (2) the

genetically-guided therapy (treatment arm). The outcome is the length of stay in the hospital, measured by hours. The shorter the length of stay, the more beneficial the therapy is. The patient subgroups are defined by age: (1) 18-20; (2) 21-30; (3) 31-40; (4) 41-50; (5) 51-60; (6) > 60. We summarize the results under our proposed design, the doubly adaptive debiased coin (DBCD) design, and the complete randomization (CR) design in Figure 5.

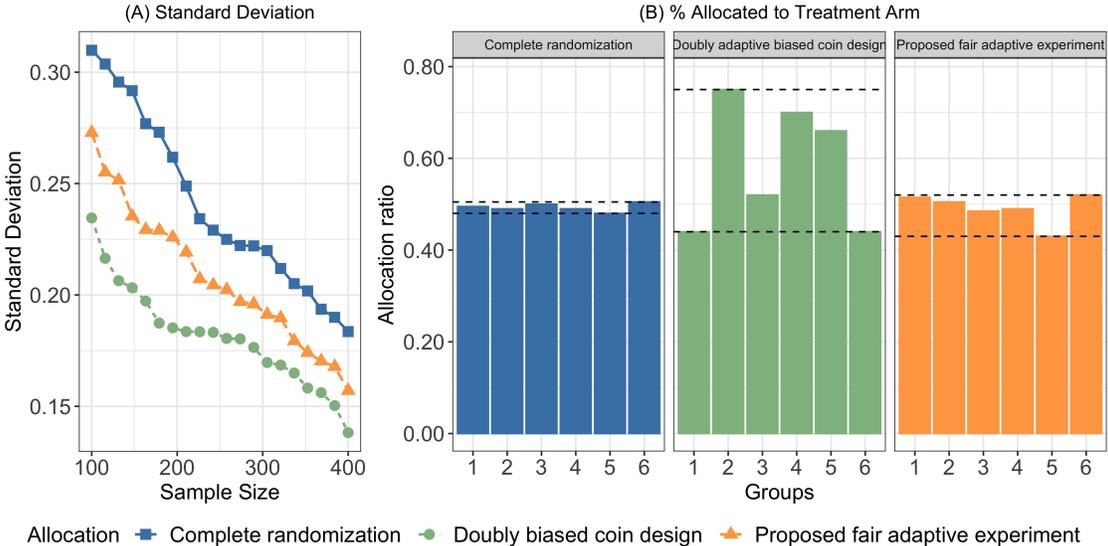


Figure 5: Comparison of the proposed adaptive experiment, the complete randomization design, and the doubly adaptive biased coin design in our case study. (A) shows the standard deviation comparison. (B) shows the percentage of participants allocated to the treatment arm in each group

Figure 5 (A) shows that compared with the DBCD design, our strategy is slightly comprised in estimation efficiency. Figure 5 (B) shows that the treatment assignment probabilities are roughly equal under the CR design. Under the DBCD design, there is a rather large discrepancy in treatment assignment probabilities across the six groups, while the treatment probabilities under our fair adaptive design strategy are less extreme.

References

- [1] Banerjee, A., Duflo, E., and Hornbeck, R. (2014). “Bundling health insurance and microfinance in India: There cannot be adverse selection if there is no demand,” *American Economic Review: Papers & Proceedings*, 104(5), 291–297.
- [2] Hall, P. and Heyde, C. C. (2014). *Martingale Limit Theory and Its Application*: Academic Press.
- [3] Newey, W. K. and McFadden, D. L. (1994). “Large sample estimation and hypothesis testing,” In R. F. Engle and D. L. McFadden (eds.) *Handbook of Econometrics, Volume IV*: Elsevier Science B.V. 2111–2245.
- [4] Ruaño, G., Tortora, J., Robinson, S., Baker, S., Holford, T., Winokur, A., and Goethe, J. W. (2021). “Subanalysis of the CYP-GUIDES Trial: CYP2D6 functional stratification and operational timeline selection,” *Psychiatry Research*, 297, 113571.
- [5] Tortora, J., Robinson, S., Baker, S., and Ruaño, G. (2020). “Clinical database of the CYP-guides trial: An open data resource on psychiatric hospitalization for severe depression,” *Data in Brief*, 30, 105457.