Waiting for Balance: Covariate-Adaptive Randomization in Sequential Experiments

EEA Congress 2025

Pedro Vergara Merino

CREST - ENSAE Paris - Institut Polytechnique de Paris

August 27th, 2025

Allocation Design in Randomized Experiments

Coin-toss assignment:

- Ex-ante: treatment groups are "equal" in expectation
- Ex-post: by chance, treatment groups may differ substantially

Solutions: When pre-treatment information is available, using more sophisticated allocation mechanisms (e.g., complete randomization, stratification, matched pairs) may lead to:

- More precise ATE estimates (Bai et al., 2024)
- Less exposure to publication bias (Snyder and Zhuo, 2024):

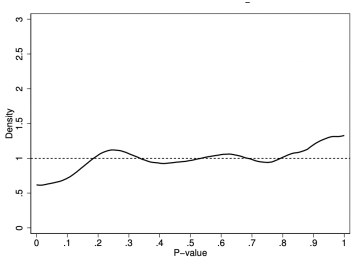

Balance in Sequential Experiments

Many experiments involve units arriving one by one:

- Field experiments with enrollment-based designs, such as unemployment benefits
- Clinical trials, such as cancer research
- Lab experiments in online platforms (Prolific or MTurk)

This poses several challenges for achieving balance:

- We do not observe all units beforehand, making classical randomization techniques infeasible
- Stratified permuted block design (SPBD): only possible on discrete variables. If one discretizes continuous covariates, one needs to choose a threshold
- Adaptive randomization and minimization methods (mostly used in clinical trials) $\rightarrow$ increased predictability of treatment assignment and/or complications in inference

This Paper

Can we improve balance in sequential experiments?

- ↪ Discuss the trade-off between balance and immediacy from the researcher's perspective
- ↪ Introduce a new randomization method achieving exact balance at the cost of short delays in assignment
- ↪ Derive statistical properties of the algorithm and the HT estimator of ATE: imbalances, waiting periods, unbiasedness, asymptotic normality

What are the practical benefits of this randomization method?

- ↪ Suitability for multi-armed bandits and/or response-adaptive methods: non-deterministic, heterogeneous probabilities

Literature Review

Data-Generating Process

I work on the superpopulation Neyman-Rubin causal framework with a binary treatment and units arriving one by one:

t	Experimental unit index by time of arrival
X_t	Vector of pre-treatment covariates of dimension p
$Y_t(d)$	Potential outcome when treated or untreated ($d = 1$ and $d = 0$), respectively

Assumption 1 (i.i.d. + moments).

$$(Y_t(0), Y_t(1), X_t) \overset{i.i.d.}{\sim} (Y(0), Y(1), X) \quad \text{with} \quad \mathbb{E}\left[Y(0)^2 + Y(1)^2 + ||X||^2\right] < \infty.$$

Assignment Mechanism

The researcher assigns individuals to treatment or control ($D_t = \mathbb{1}\{\text{unit } t \text{ is treated}\}$), with a prespecified probability π_t (propensity score).

For a given sample size of n units, an assignment mechanism defines

$$\mathbb{P}((D_1, \dots, D_n) = (d_1, \dots, d_n) | X_1, \dots, X_n) \text{ for all } (d_1, \dots, d_n) \in \{0, 1\}^n.$$

Assumption 2 (Conditional Unconfoundedness and Strict Common Support).

The assignment mechanism is such that

$$(D_1, \dots, D_n) \perp\!\!\!\perp (Y_1(0), Y_1(1), \dots Y_n(0), Y_n(1)) | (X_1, \dots, X_n)$$

and, for any $t = 1, \dots, n$,

$$\mathbb{P}(D_t | X_1, \dots, X_n) = p(X_t) =: \pi_t$$

with $c < \pi_t < 1 - c$ for some positive constant c.

After the experiment, the researcher observes $Y_t := Y_t(D_t)$, for $t = 1, \dots, n$.

ATE Estimation

I focus on the estimation of the population average treatment effect

$$\theta_0 = \mathbb{E}[Y(1) - Y(0)]$$

through the Horvitz-Thompson estimator

$$\widehat{\theta} = \frac{1}{n}\sum_{t=1}^{n} \frac{Y_t D_t}{\pi_t} - \frac{Y_t(1-D_t)}{1-\pi_t}.$$

Remark (Difference-in-Means).

If treatment probabilities are homogeneous ($\pi_t = \pi$) and the group sizes are fixed ($\sum_t^n D_t = n\pi$), one has

$$\widehat{\theta} = \frac{1}{n_1}\sum_{t=1}^{n} Y_t D_t - \frac{1}{n_0}\sum_{t=1}^{n} Y_t(1-D_t).$$

Exact Balance

Given X, the researcher specifies some vector $Z_t = g(X_t)$ of dimension q. I assume here $\mathbb{E}[||Z||^2] < \infty$.

Definition 1 (Exactly Balanced Allocation).

A treatment allocation $(d_1, \ldots, d_n) \in \{0,1\}^n$ is *exactly balanced* with respect to $(Z_1, \ldots, Z_n)$ if

$$\frac{1}{n}\sum_{t=1}^{n}\frac{Z_t d_t}{\pi_t} = \frac{1}{n}\sum_{t=1}^{n}\frac{Z_t(1-d_t)}{1-\pi_t}. \tag{1}$$

Let $\widetilde{Z}_t = \frac{Z_t}{\pi_t(1-\pi_t)}$, $\widetilde{\boldsymbol{Z}} = \left(\widetilde{Z}_1, \ldots, \widetilde{Z}_n\right)^T \in \mathbb{R}^{n\times q}$, $\boldsymbol{\pi} = (\pi_1, \ldots, \pi_n)^T \in \mathbb{R}^n$ and $\boldsymbol{d} = (d_1, \ldots, d_n)^T \in \mathbb{R}^n$, then Equation (1) becomes

$$(\boldsymbol{d} - \boldsymbol{\pi}) \in \ker\left(\widetilde{\boldsymbol{Z}}^T\right).$$

The Precision Gains from Exact Balance

Proposition 1.

Let Assumptions 1 and 2 hold. Let $f(X_t) = \mathbb{E}[Y_t(1)|X_t](1-\pi_t) + \mathbb{E}[Y_t(0)|X_t]\pi_t$. Given the vector of covariates X, and the propensity scores $\pi = p(X)$ an allocation design guaranteeing

$$\frac{1}{n}\sum_{t=1}^{n}\frac{f(X_t)D_t}{\pi_t} = \frac{1}{n}\sum_{t=1}^{n}\frac{f(X_t)(1-D_t)}{1-\pi_t}$$

minimizes $\text{MSE}\left(\widehat{\theta}\right)$.

If $f(X_t)$ is linear in Z_t, then the design is efficient. In theory, increasing q improves the approximation of $f(X_t)$ and may lead to more precise estimators.

Algorithm with One Constraint

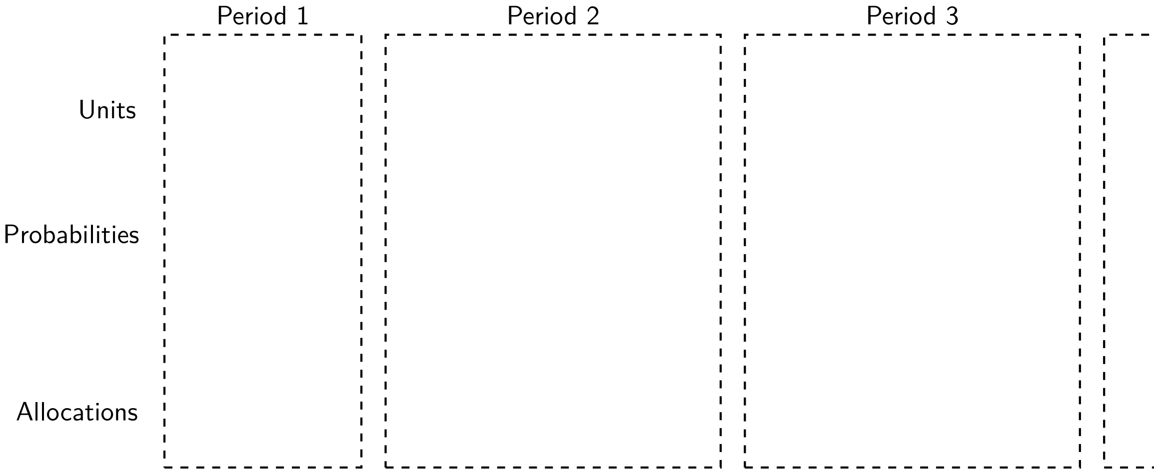

Algorithm with One Constraint

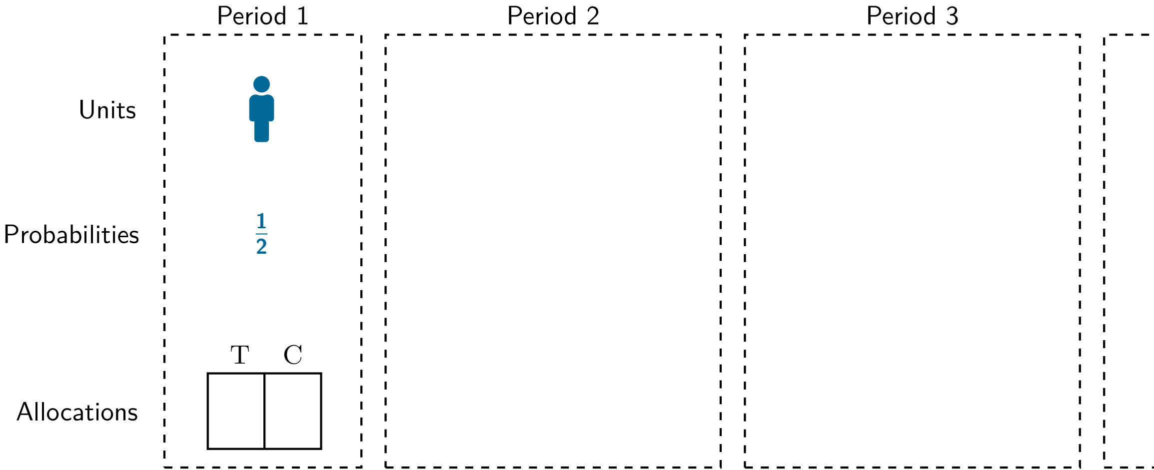

Algorithm with One Constraint

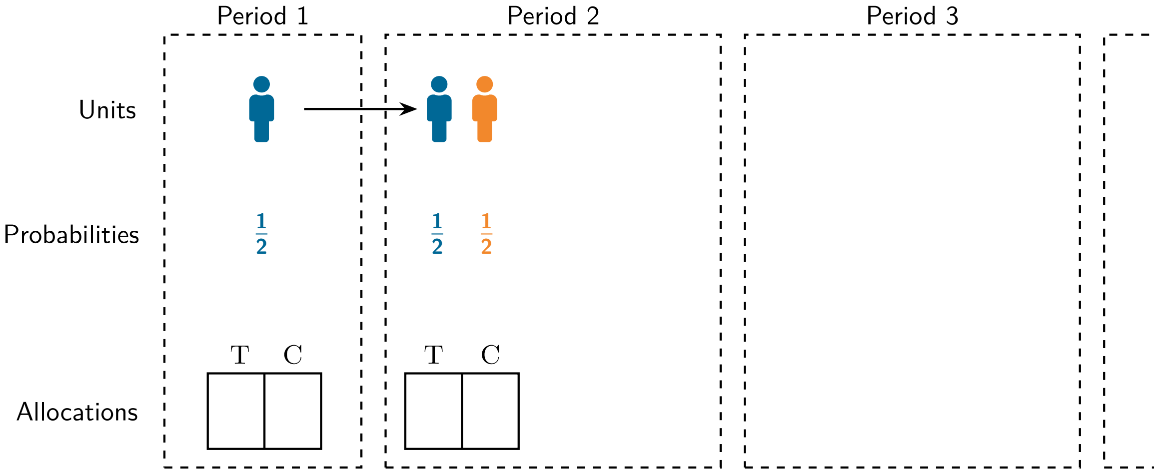

Algorithm with One Constraint

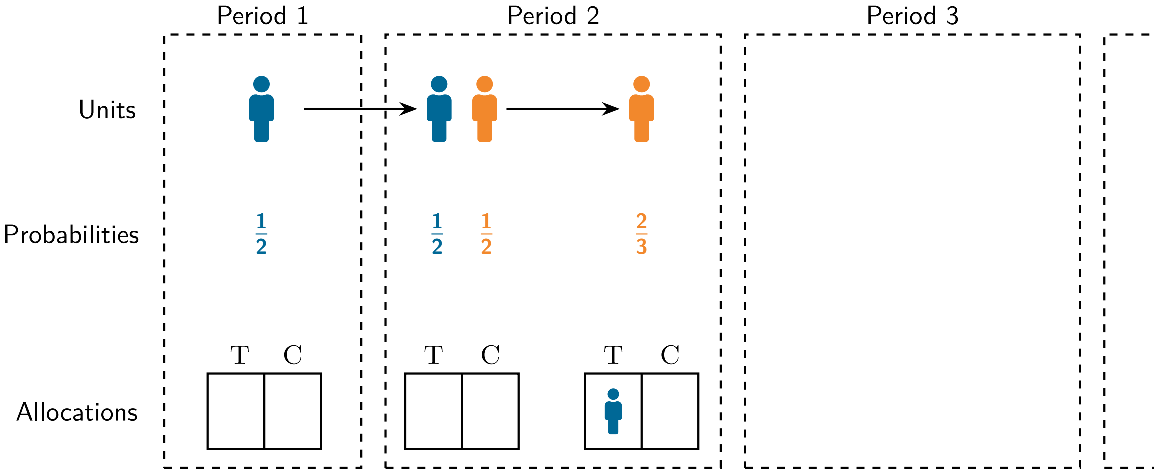

Algorithm with One Constraint

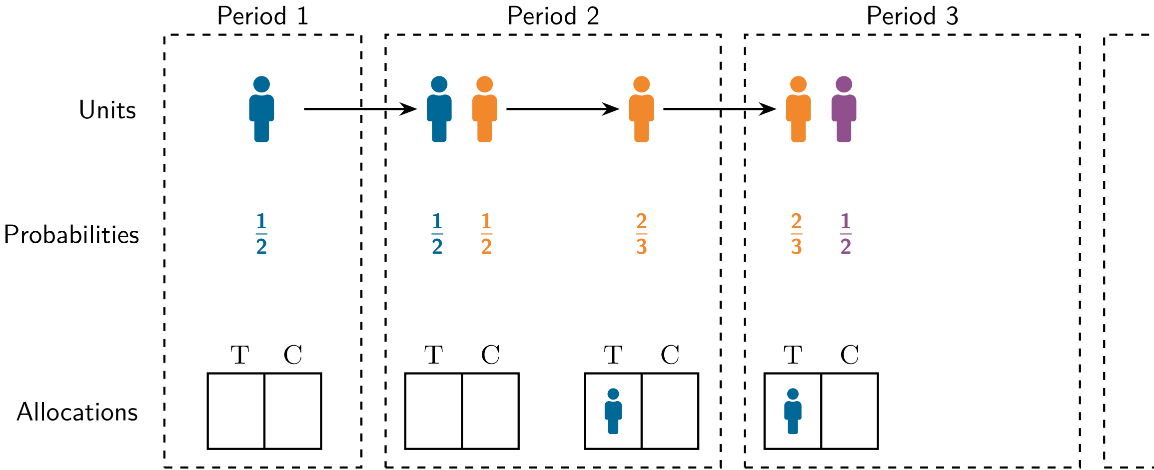

Algorithm with One Constraint

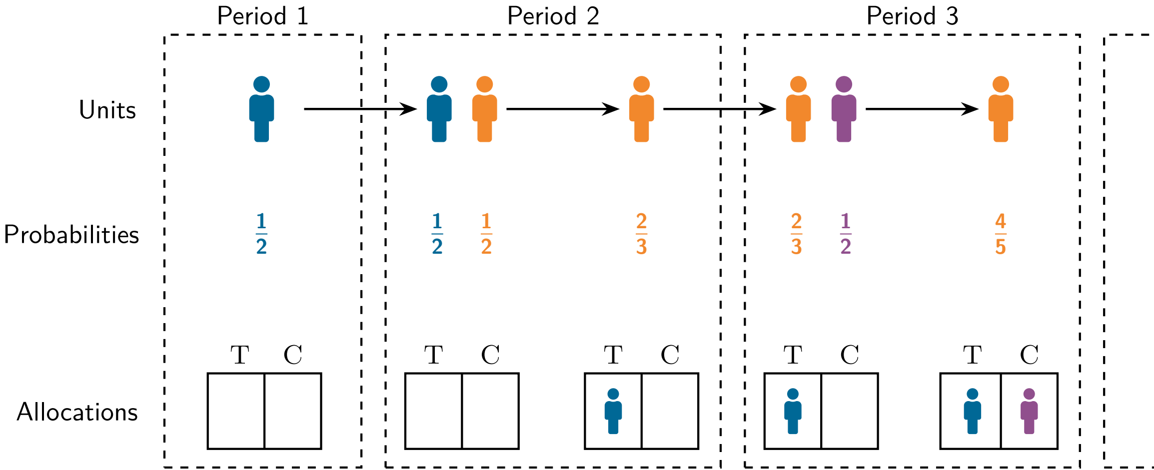

Algorithm with One Constraint

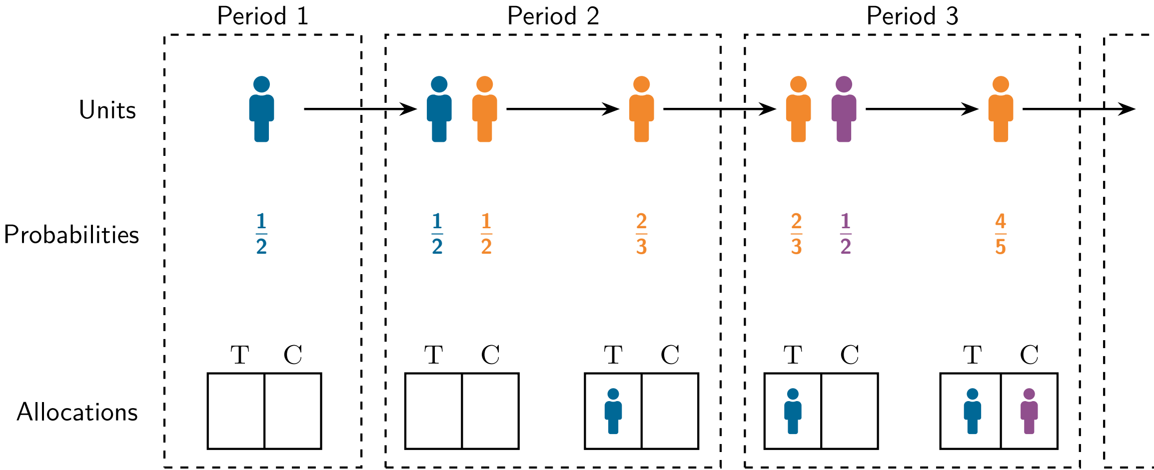

Zoom on Period 2: Balancing Constraint

$\pi_1 = \pi_2 = \frac{1}{2}$, $X_1 = -2.5$ and $X_2 = 7.5$

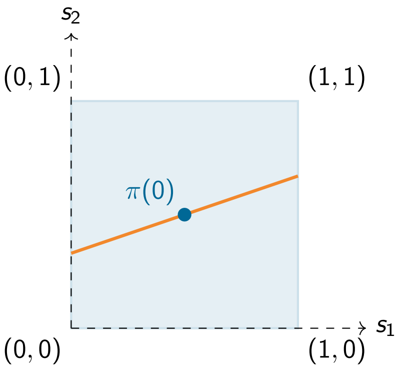

$$\left\{ (s_1, s_2) \in [0,1]^2 \middle| \frac{X_1 s_1}{\pi_1} + \frac{X_2 s_2}{\pi_2} = \frac{X_1(1-s_1)}{1-\pi_1} + \frac{X_2(1-s_2)}{1-\pi_2} \right\}$$

Zoom on Period 2: Treatment Allocation (General Algorithm)

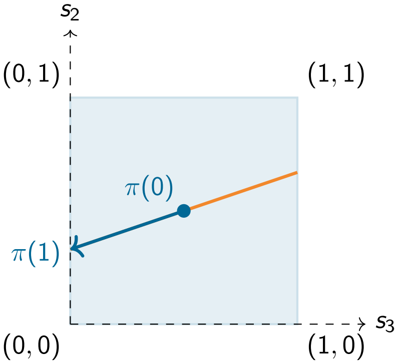

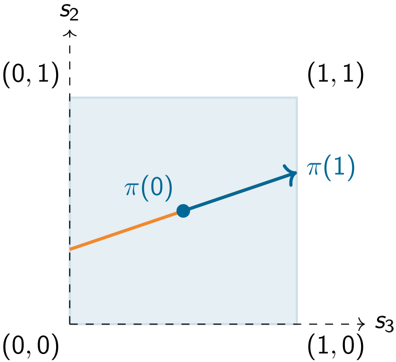

$$(\pi_1(1), \pi_2(1)) = \left(0, \frac{1}{3}\right) \quad \text{w.p.} \quad \frac{1}{2}$$

$$(\pi_1(1), \pi_2(1)) = \left(1, \frac{2}{3}\right) \quad \text{w.p.} \quad \frac{1}{2}$$

Zoom on Period 3: Balancing Constraints (General Algorithm)

$\pi_1 = \pi_2 = \pi_3 = \frac{1}{2}$, $X_1 = -2.5$, $X_2 = 7.5$ and $X_3 = 2$

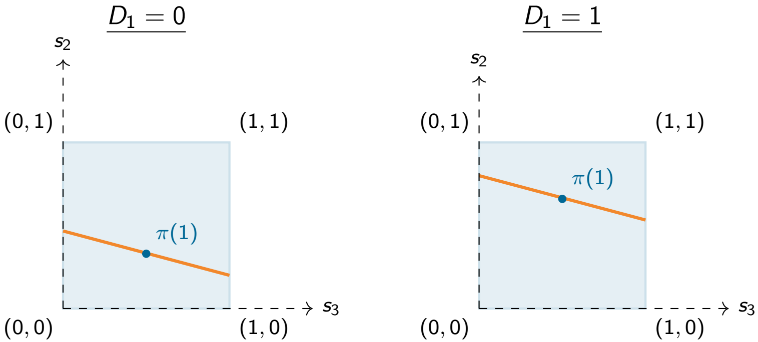

$$\left\{ (s_2, s_3) \in [0,1]^2 \middle| \frac{X_1 D_1}{\pi_1} + \frac{X_2 s_2}{\pi_2} + \frac{X_3 s_3}{\pi_3} = \frac{X_1(1-D_1)}{1-\pi_1} + \frac{X_2(1-s_2)}{1-\pi_2} + \frac{X_3(1-s_3)}{1-\pi_3} \right\}$$

Zoom on Period 3: Treatment Allocations (General Algorithm)

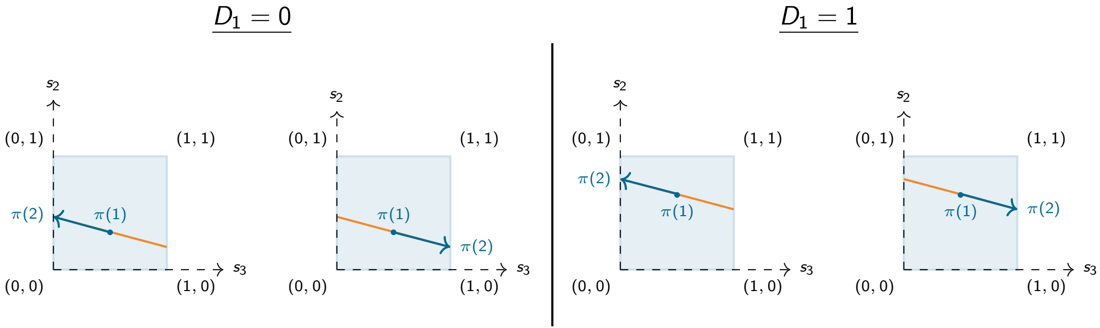

$$(\pi_1(2), \pi_2(2), \pi_3(2)) =$$

$\left(0, \frac{7}{15}, 0\right)$ w.p. $\frac{1}{4}$ $\quad$ $\left(0, \frac{1}{5}, 1\right)$ w.p. $\frac{1}{4}$ $\quad$ $\left(1, \frac{4}{5}, 0\right)$ w.p. $\frac{1}{4}$ $\quad$ $\left(1, \frac{8}{15}, 1\right)$ w.p. $\frac{1}{4}$

Unbiasedness of the HT Estimator

Proposition 2 (Unbiasedness).

Let Assumptions 1-2 hold. Under the general algorithm,

$$\mathbb{E}[D_t|X_1, \ldots, X_n] = \pi_t$$

and, therefore, the HT estimator is unbiased, i.e.,

$$\mathbb{E}\left[\widehat{\theta} - \theta_0\right] = 0.$$

Proposition 3 (Selection Bias).

Let $\pi_t = 1/2$ for $t = 1, \ldots, n$. Then, the ability to predict an assignment under the general algorithm is identical to coin toss randomization.

Imbalances

Imbalances are measured according to:

$$B_{n,q} = \left\| \frac{1}{n} \sum_{i=1}^{n} \frac{Z_t D_t}{\pi_t} - \frac{Z_t(1 - D_t)}{1 - \pi_t} \right\|^2 .$$

Proposition 4 (Imbalances).

Let Assumptions 1-2 hold. Under the coin toss assignment,

$$\mathbb{E}[B_{n,q}] = \frac{1}{n} \sum_{k=1}^{q} \mathbb{E}\left[\frac{Z_k^2}{\pi(1 - \pi)}\right] .$$

Under the general algorithm,

$$\mathbb{E}[B_{n,q}] = \frac{q}{n^2} \sum_{k=1}^{q} \mathbb{E}\left[\frac{Z_k^2}{\pi(1 - \pi)}\right] .$$

Imbalances: Simulations

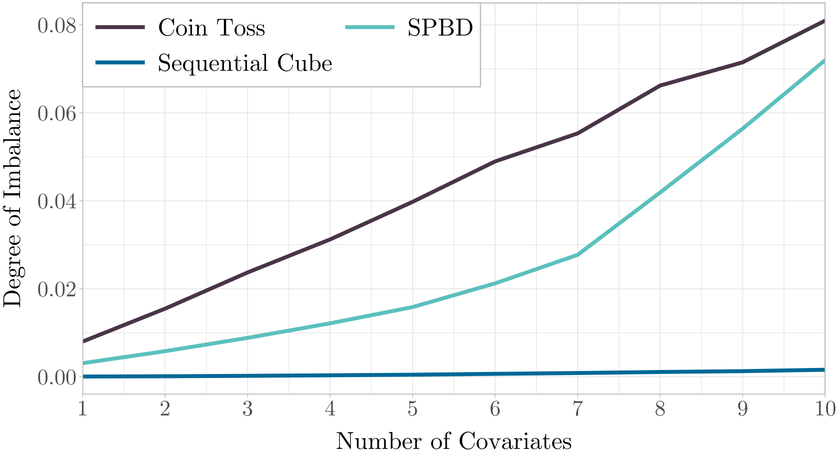

Waiting Times

Under the general algorithm units can wait for more than one period. I denote by T_t the number of periods that unit t waits. If unit t is immediately assigned, $T_t = 0$. At most, $T_t = n - t + 1$.

Proposition 5 (Sample Distribution of Waiting Times).

Under the algorithm,

$$\overline{T}_{n,q} = \frac{q(2n - q + 1)}{2n},$$

and

$$\frac{(q^2 + q)(2q + 1)}{6(n-1)} - \frac{q^2(q+1)^2}{4n(n-1)} \leq \sigma^2_{n,q} \leq \frac{(q+1)(2q^2 - 6qn + q + 6n^2)}{6(n-1)} - \frac{(q+1)^2(2n-q)^2}{4n(n-1)}.$$

Waiting Times: Simulations

# Covariates	1	2	3	4	5	6	7	8	9	10
Mean	1.00	2.00	2.99	3.99	4.98	5.97	6.96	7.94	8.93	9.91
SD	2.22	2.88	4.12	5.19	6.22	7.16	8.08	9.00	9.91	10.77
Min	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
1%	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
5%	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.01	0.06
10%	0.00	0.00	0.00	0.00	0.00	0.01	0.18	0.64	0.92	0.99
25%	0.00	0.00	0.30	0.98	1.00	1.20	1.89	2.02	2.40	2.91
50%	0.00	1.00	1.91	2.12	3.00	3.82	4.40	5.09	5.89	6.60
75%	1.00	2.61	3.99	5.20	6.67	7.97	9.25	10.58	11.91	13.21
90%	3.00	4.93	7.44	9.79	12.08	14.29	16.41	18.58	20.77	22.85
95%	4.77	7.03	10.54	13.70	16.67	19.67	22.39	25.12	28.01	30.68
99%	10.14	13.45	18.94	23.98	28.75	33.04	37.50	41.87	46.04	49.98
Max	22.11	27.03	36.08	43.76	50.99	57.45	64.15	70.84	77.06	83.64

Waiting Times: Simulations

# Covariates	1	2	3	4	5	6	7	8	9	10
Mean	1.00	2.00	2.99	3.99	4.98	5.97	6.96	7.94	8.93	9.91
SD	2.22	2.88	4.12	5.19	6.22	7.16	8.08	9.00	9.91	10.77
Min	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
1%	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
5%	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.01	0.06
10%	0.00	0.00	0.00	0.00	0.00	0.01	0.18	0.64	0.92	0.99
25%	0.00	0.00	0.30	0.98	1.00	1.20	1.89	2.02	2.40	2.91
50%	0.00	1.00	1.91	2.12	3.00	3.82	4.40	5.09	5.89	6.60
75%	1.00	2.61	3.99	5.20	6.67	7.97	9.25	10.58	11.91	13.21
90%	3.00	4.93	7.44	9.79	12.08	14.29	16.41	18.58	20.77	22.85
95%	4.77	7.03	10.54	13.70	16.67	19.67	22.39	25.12	28.01	30.68
99%	10.14	13.45	18.94	23.98	28.75	33.04	37.50	41.87	46.04	49.98
Max	22.11	27.03	36.08	43.76	50.99	57.45	64.15	70.84	77.06	83.64

Waiting Times: Experimental Data

Experimental data from Caria et al. (2024), which include exact arrival times.

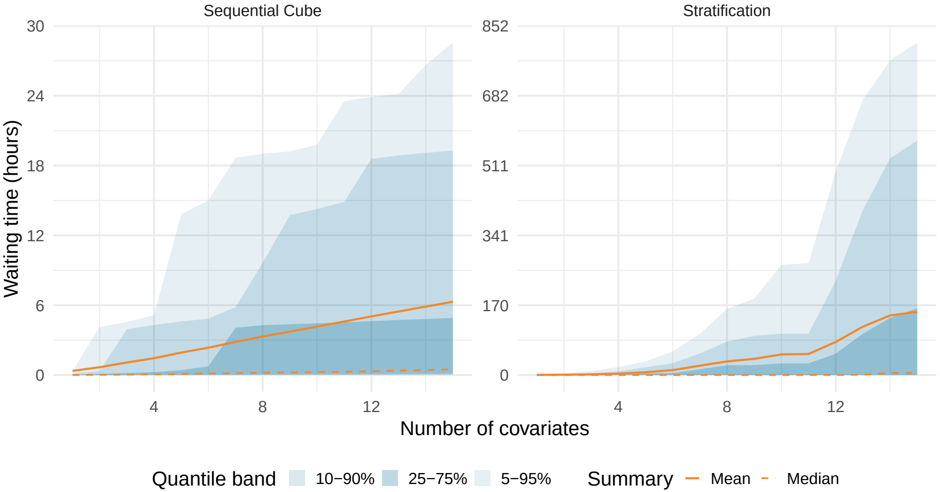

Asymptotic Normality

Definition 2 (Orthogonal Projections).

We consider the following orthogonal projections:

$$Y_t(0) = \alpha_0 + \frac{Z_t'}{\pi_t}\beta_0 + \varepsilon_t(0) \tag{2}$$

$$Y_t(1) = \alpha_1 + \frac{Z_t'}{1-\pi_t}\beta_1 + \varepsilon_t(1) \tag{3}$$

with $\mathbb{E}[\varepsilon_i(0)] = \mathbb{E}[\varepsilon_t(1)] = \mathbb{E}\left[\varepsilon_t(0)\frac{Z_t}{\pi_t}\right] = \mathbb{E}\left[\varepsilon_t(1)\frac{Z_t}{1-\pi_t}\right] = 0$.

Proposition 6 (Asymptotic Normality).

Let Assumptions 1-2 hold. Then,

$$\sqrt{n}\left(\widehat{\theta}_{HT} - \theta\right) \xrightarrow{d} \mathcal{N}\left(0, V\right),$$

with $V = \mathbb{E}[\pi_1]\mathbb{E}\left[\frac{\sigma_1^2(X_1)}{\pi_1^2}\right] + \mathbb{E}[1-\pi_1]\mathbb{E}\left[\frac{\sigma_0^2(X_1)}{(1-\pi_1)^2}\right] + V\left(\frac{Z_1'\beta_1}{1-\pi_1} - \frac{Z_1'\beta_0}{\pi_1}\right)$.

Simulations – DGPs

I try two different specifications. In all of them, $n = 500$, $p = 10$, $\pi_t = 1/2$, $X_t \sim \mathcal{N}(0,1)^p$, $\varepsilon_t(d) \sim \mathcal{N}(0,1)$ for $d = 0, 1$ and $\theta_0 = 0$.

1. Linear Model: $Y_t(0) = 1 + X_{1,t} + \ldots + X_{10,t} + \varepsilon_t(0)$ and $Y_t(1) = Y_t(0) - \varepsilon_t(0) + \varepsilon_t(1)$
2. Misspecified Model: $Y_t(0) = \frac{1}{\sqrt{10}}(X_{1,t}^2 + \ldots + X_{10,t}^2) + \varepsilon_t(0)$ and $Y_t(1) = Y_t(0) - \varepsilon_t(0) + \varepsilon_t(1)$
3. Noise Model: $Y_t(0) = \varepsilon_t(0)$ and $Y_t(1) = \varepsilon_t(1)$.

I first compare the precision of three allocation mechanisms: the sequential cube method, coin toss procedure, and SPBD. I look at the ratio

$$\frac{\mathbb{V}(\widehat{\theta}^{\Pi})}{\mathbb{V}(\widehat{\theta}^{CT})}$$

Precision Gains: Linear Model Misspecified Noise

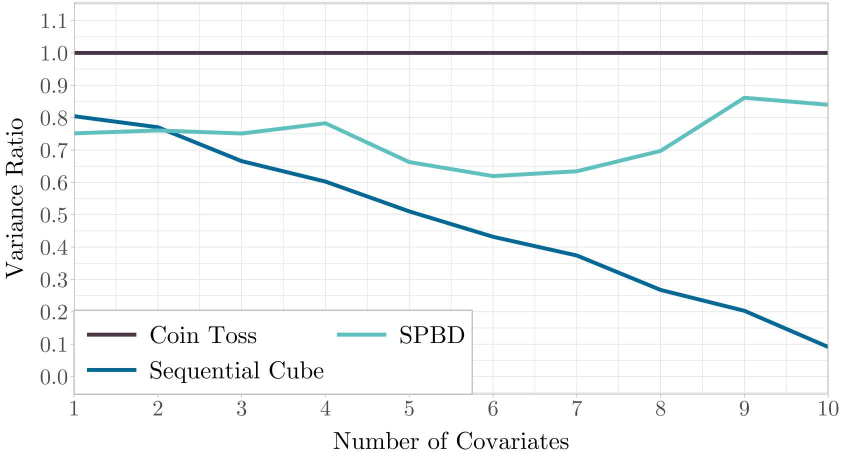

Inference: Coverage Rates

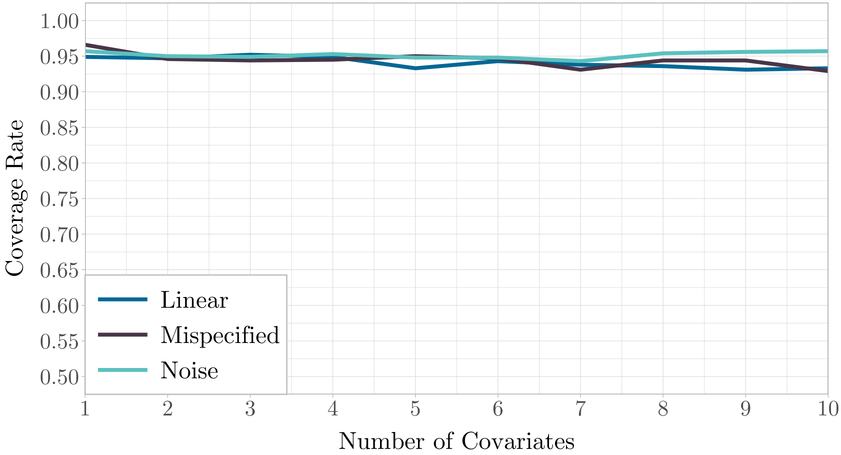

Conclusion

- Motivation: Sequential experiments often struggle to achieve covariate balance.
- Contribution: Introduced a randomization algorithm that:
  - Achieves *almost exact covariate balance*
  - Maintains *non-deterministic* assignment
  - Requires limited *waiting periods*
- Statistical Results:
  - HT estimator is *unbiased* and *asymptotically normal*
  - Substantially reduces imbalances and variance
- Practical Implications:
  - Applicable in *field*, *lab*, and *online* experiments
  - Adaptable to *multi-armed bandits* and *response-adaptive designs*
- Future Work: Response-adaptiveness and anytime-valid inference

References I

Athey, Susan et al. (Nov. 2022). *Contextual Bandits in a Survey Experiment on Charitable Giving: Within-Experiment Outcomes versus Policy Learning*. arXiv:2211.12004 [cs, econ, stat]. DOI: 10.48550/arXiv.2211.12004.

Atkinson, A. C. (1982). "Optimum Biased Coin Designs for Sequential Clinical Trials with Prognostic factors". *Biometrika* 69.1. Publisher: [Oxford University Press, Biometrika Trust], pp. 61–67. ISSN: 0006-3444. DOI: 10.2307/2335853.

Bai, Yuehao, Azeem M. Shaikh, and Max Tabord-Meehan (May 2024). *A Primer on the Analysis of Randomized Experiments and a Survey of some Recent Advances*. arXiv:2405.03910 [econ, stat]. DOI: 10.48550/arXiv.2405.03910.

Caria, A Stefano et al. (Apr. 2024). "An Adaptive Targeted Field Experiment: Job Search Assistance for Refugees in Jordan". *Journal of the European Economic Association* 22.2, pp. 781–836. ISSN: 1542-4766. DOI: 10.1093/jeea/jvad067.

References II

Chipman, Jonathan J., Lindsay Mayberry, and Robert A. Greevy Jr. (2023). "Rematching on-the-fly: Sequential matched randomization and a case for covariate-adjusted randomization". en. *Statistics in Medicine* 42.22. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/sim.9843, pp. 3981–3995. ISSN: 1097-0258. DOI: 10.1002/sim.9843.

Davezies, Laurent, Guillaume Hollard, and Pedro Vergara Merino (July 2024). *Revisiting Randomization with the Cube Method*. arXiv:2407.13613. DOI: 10.48550/arXiv.2407.13613.

Deville, Jean-Claude and Yves Tillé (Dec. 2004). "Efficient balanced sampling: The cube method". *Biometrika* 91.4, pp. 893–912. ISSN: 0006-3444. DOI: 10.1093/biomet/91.4.893.

Fisher, Sir Ronald Aylmer (1935). *The Design of Experiments*. en. Oliver and Boyd.

Hu, Yanqing and Feifang Hu (2012). "Asymptotic Properties of Covariate-Adaptive Randomization". *The Annals of Statistics* 40.3. Publisher: Institute of Mathematical Statistics, pp. 1794–1815. ISSN: 0090-5364.

References III

Jauslin, Raphaël, Bardia Panahbehagh, and Yves Tillé (2022). "Sequential spatially balanced sampling". en. *Environmetrics* 33.8. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/env.2776, e2776. ISSN: 1099-095X. DOI: 10.1002/env.2776.

Kapelner, Adam and Abba Krieger (2014). "Matching On-the-Fly: Sequential Allocation with Higher Power and Efficiency". *Biometrics* 70.2. Publisher: [Wiley, International Biometric Society], pp. 378–388. ISSN: 0006-341X.

Ma, Wei et al. (Jan. 2024). "A New and Unified Family of Covariate Adaptive Randomization Procedures and Their Properties". *Journal of the American Statistical Association* 119.545. Publisher: ASA Website _eprint: https://doi.org/10.1080/01621459.2022.2102986, pp. 151–162. ISSN: 0162-1459. DOI: 10.1080/01621459.2022.2102986.

Pocock, Stuart J. and Richard Simon (1975). "Sequential Treatment Assignment with Balancing for Prognostic Factors in the Controlled Clinical Trial". *Biometrics* 31.1. Publisher: International Biometric Society, pp. 103–115. ISSN: 0006-341X. DOI: 10.2307/2529712.

Qin, Yichen et al. (July 2018). *Pairwise Sequential Randomization and Its Properties*. arXiv:1611.02802 [stat]. DOI: 10.48550/arXiv.1611.02802.

References IV

Snyder, Christopher M. and Ran Zhuo (Feb. 2024). "Examining Selection Pressures in the Publication Process through the Lens of Sniff Tests". *The Review of Economics and Statistics*, pp. 1–45. ISSN: 0034-6535. DOI: 10.1162/rest_a_01410.

Zhou, Quan et al. (Sept. 2018). "Sequential rerandomization". *Biometrika* 105.3, pp. 745–752. ISSN: 0006-3444. DOI: 10.1093/biomet/asy031.

Related Literature

Balancing in Sequential Trials:

- Stratified-permuted-block design (Fisher, 1935)
- Matching-on-the-flight (Chipman et al., 2023; Kapelner and Krieger, 2014)
- Biased coin-designs (Atkinson, 1982; Y. Hu and F. Hu, 2012; Ma et al., 2024; Pocock and Simon, 1975)

↪ Proposed method: non-deterministic, no covariate assumptions, and does not alter initial propensity scores. Cost: waiting periods

Sequential Methods with Waiting Periods:

- Stream sampling (Jauslin et al., 2022)
- Use of batches (Qin et al., 2018; Zhou et al., 2018)

Cube method (Davezies et al., 2024; Deville and Tillé, 2004)

◂ Back

General Algorithm

1. Before the experiment: the empiricist commits on a sample size n, a function $g\colon \mathbb{R}^p \to \mathbb{R}^q$ selecting the moments of X to balance, and a function $p(.)$ defining the propensity scores.

2. Initial waiting periods: Units $t = 1, \dots, q$ arrive and wait. We denote $\boldsymbol{\pi}(\boldsymbol{t}) = (\pi_1, \dots, \pi_t)'$.

3. Sequential assignment: At the arrival of units $t = q+1, \dots, n$

   3.1 Redefine $\boldsymbol{\pi}(\boldsymbol{t}-\mathbf{1}) = (\boldsymbol{\pi}(\boldsymbol{t}-\mathbf{1})', \pi_t)'$.

   3.2 Let $\widetilde{\boldsymbol{Z}}(\boldsymbol{t}) = \left(\frac{g(X_1)}{\pi_1(1-\pi_1)}, \dots, \frac{g(X_t)}{\pi_t(1-\pi_t)}\right)'$, $U(t) = \{s \in \{1, \dots, t\} | \pi_s(t-1) \in (0,1)\}$ and $W(t) = \mathrm{diag}(a_1, \dots, a_t)$, with $a_s = 1$ if $s \in U(t)$, 0 otherwise.

   3.3 Draw $v(t) \in \mathbb{R}^t$, compute $u(t) = W(t)v(t) - W(t)\widetilde{\boldsymbol{Z}}(\boldsymbol{t})\left(\widetilde{\boldsymbol{Z}}(\boldsymbol{t})^T W(t)\widetilde{\boldsymbol{Z}}(\boldsymbol{t})\right)^{-1}\widetilde{\boldsymbol{Z}}(\boldsymbol{t})W(t)v(t)$, $\lambda_1(t) = \min_{s\in U(t)} \frac{\mathbb{1}\{u_s(t)>0\} - \pi_s(t-1)}{u_s(t)}$, $\lambda_2(t) = \min_{s\in U(t)} \frac{\pi_s(t-1) - \mathbb{1}\{u_s(t)<0\}}{u_s(t)}$, $r(t) = \frac{\lambda_2(t)}{\lambda_1(t)+\lambda_2(t)}$.

   3.4 Update probabilities $\pi(t) = \begin{cases} \pi(t-1) + \lambda_1 u(t), & \text{w.p. } r(t) \\ \pi(t-1) - \lambda_2 u(t), & \text{w.p. } 1 - r(t) \end{cases}$

4. After period n: Allocate unassigned units (at most q) using a coin toss with probabilities $\pi_s(n)$.

Back

General Algorithm

1. Before the experiment: the empiricist commits on a sample size n, a function $g: \mathbb{R}^p \to \mathbb{R}^q$ selecting the moments of X to balance, and a function $p(.)$ defining the propensity scores.

2. Initial waiting periods: Units $t = 1, \ldots, q$ arrive and wait. We denote $\boldsymbol{\pi(t)} = (\pi_1, \ldots, \pi_t)'$.

3. Sequential assignment: At the arrival of units $t = q+1, \ldots, n$

   3.1 Redefine $\boldsymbol{\pi(t-1)} = (\boldsymbol{\pi(t-1)}', \pi_t)'$.

   3.2 Let $\widetilde{\boldsymbol{Z}}\boldsymbol{(t)} = \left(\frac{g(X_1)}{\pi_1(1-\pi_1)}, \ldots, \frac{g(X_t)}{\pi_t(1-\pi_t)}\right)'$, $U(t) = \{s \in \{1, \ldots, t\} | \pi_s(t-1) \in (0,1)\}$ and $W(t) = \text{diag}(a_1, \ldots, a_t)$, with $a_s = 1$ if $s \in U(t)$, 0 otherwise.

   3.3 Draw $v(t) \in \mathbb{R}^t$, compute $u(t) = W(t)v(t) - W(t)\widetilde{\boldsymbol{Z}}\boldsymbol{(t)}\left(\widetilde{\boldsymbol{Z}}\boldsymbol{(t)}^T W(t)\widetilde{\boldsymbol{Z}}\boldsymbol{(t)}\right)^{-1}\widetilde{\boldsymbol{Z}}\boldsymbol{(t)}W(t)v(t)$, $\lambda_1(t) = \min_{s \in U(t)} \frac{\mathbb{1}\{u_s(t)>0\} - \pi_s(t-1)}{u_s(t)}$, $\lambda_2(t) = \min_{s \in U(t)} \frac{\pi_s(t-1) - \mathbb{1}\{u_s(t)<0\}}{u_s(t)}$, $r(t) = \frac{\lambda_2(t)}{\lambda_1(t)+\lambda_2(t)}$.

   3.4 Update probabilities $\pi(t) = \begin{cases} \pi(t-1) + \lambda_1 u(t), & \text{w.p. } r(t) \\ \pi(t-1) - \lambda_2 u(t), & \text{w.p. } 1 - r(t) \end{cases}$

4. After period n: Allocate unassigned units (at most q) using a coin toss with probabilities $\pi_s(n)$.

Back

General Algorithm

1. Before the experiment: the empiricist commits on a sample size n, a function $g: \mathbb{R}^p \to \mathbb{R}^q$ selecting the moments of X to balance, and a function $p(.)$ defining the propensity scores.

2. Initial waiting periods: Units $t = 1, \dots, q$ arrive and wait. We denote $\boldsymbol{\pi}(\boldsymbol{t}) = (\pi_1, \dots, \pi_t)'$.

3. Sequential assignment: At the arrival of units $t = q + 1, \dots, n$

   3.1 Redefine $\boldsymbol{\pi}(\boldsymbol{t-1}) = (\boldsymbol{\pi}(\boldsymbol{t-1})', \pi_t)'$.

   3.2 Let $\widetilde{\boldsymbol{Z}}(\boldsymbol{t}) = \left(\frac{g(X_1)}{\pi_1(1-\pi_1)}, \dots, \frac{g(X_t)}{\pi_t(1-\pi_t)}\right)'$, $U(t) = \{s \in \{1, \dots, t\} | \pi_s(t-1) \in (0,1)\}$ and $W(t) = \text{diag}(a_1, \dots, a_t)$, with $a_s = 1$ if $s \in U(t)$, 0 otherwise.

   3.3 Draw $v(t) \in \mathbb{R}^t$, compute $u(t) = W(t)v(t) - W(t)\widetilde{\boldsymbol{Z}}(\boldsymbol{t})\left(\widetilde{\boldsymbol{Z}}(\boldsymbol{t})^T W(t)\widetilde{\boldsymbol{Z}}(\boldsymbol{t})\right)^{-1}\widetilde{\boldsymbol{Z}}(\boldsymbol{t})W(t)v(t)$,
   $\lambda_1(t) = \min_{s\in U(t)} \frac{\mathbb{1}\{u_s(t)>0\}-\pi_s(t-1)}{u_s(t)}$, $\lambda_2(t) = \min_{s\in U(t)} \frac{\pi_s(t-1)-\mathbb{1}\{u_s(t)<0\}}{u_s(t)}$, $r(t) = \frac{\lambda_2(t)}{\lambda_1(t)+\lambda_2(t)}$.

   3.4 Update probabilities $\boldsymbol{\pi}(t) = \begin{cases} \boldsymbol{\pi}(t-1) + \lambda_1 u(t), & \text{w.p. } r(t) \\ \boldsymbol{\pi}(t-1) - \lambda_2 u(t), & \text{w.p. } 1 - r(t) \end{cases}$

4. After period n: Allocate unassigned units (at most q) using a coin toss with probabilities $\pi_s(n)$.

Back

General Algorithm

1. Before the experiment: the empiricist commits on a sample size n, a function $g : \mathbb{R}^p \to \mathbb{R}^q$ selecting the moments of X to balance, and a function $p(.)$ defining the propensity scores.

2. Initial waiting periods: Units $t = 1, \ldots, q$ arrive and wait. We denote $\boldsymbol{\pi(t)} = (\pi_1, \ldots, \pi_t)'$.

3. Sequential assignment: At the arrival of units $t = q+1, \ldots, n$

   3.1 Redefine $\boldsymbol{\pi(t-1)} = (\boldsymbol{\pi(t-1)}', \pi_t)'$.

   3.2 Let $\widetilde{\boldsymbol{Z}}\boldsymbol{(t)} = \left(\frac{g(X_1)}{\pi_1(1-\pi_1)}, \ldots, \frac{g(X_t)}{\pi_t(1-\pi_t)}\right)'$, $U(t) = \{s \in \{1, \ldots, t\} | \pi_s(t-1) \in (0,1)\}$ and $W(t) = \text{diag}(a_1, \ldots, a_t)$, with $a_s = 1$ if $s \in U(t)$, 0 otherwise.

   3.3 Draw $v(t) \in \mathbb{R}^t$, compute $u(t) = W(t)v(t) - W(t)\widetilde{\boldsymbol{Z}}\boldsymbol{(t)} \left(\widetilde{\boldsymbol{Z}}\boldsymbol{(t)}^T W(t) \widetilde{\boldsymbol{Z}}\boldsymbol{(t)}\right)^{-1} \widetilde{\boldsymbol{Z}}\boldsymbol{(t)} W(t)v(t)$, $\lambda_1(t) = \min_{s \in U(t)} \frac{\mathbb{1}\{u_s(t)>0\} - \pi_s(t-1)}{u_s(t)}$, $\lambda_2(t) = \min_{s \in U(t)} \frac{\pi_s(t-1) - \mathbb{1}\{u_s(t)<0\}}{u_s(t)}$, $r(t) = \frac{\lambda_2(t)}{\lambda_1(t)+\lambda_2(t)}$.

   3.4 Update probabilities $\boldsymbol{\pi}(t) = \begin{cases} \boldsymbol{\pi}(t-1) + \lambda_1 u(t), & \text{w.p. } r(t) \\ \boldsymbol{\pi}(t-1) - \lambda_2 u(t), & \text{w.p. } 1 - r(t) \end{cases}$

4. After period n: Allocate unassigned units (at most q) using a coin toss with probabilities $\pi_s(n)$.

Back

Link to the "Cube Method"

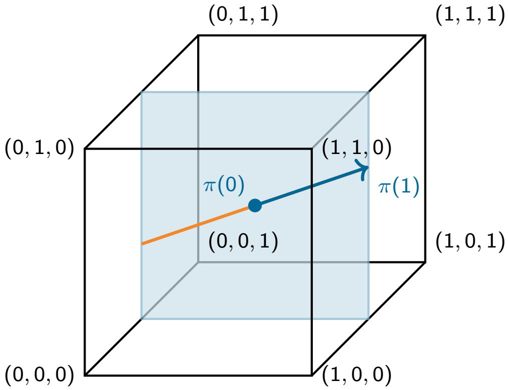

Back

Link to the "Cube Method"

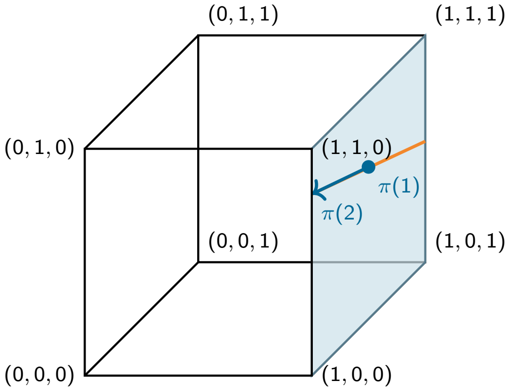

Back

Link to the "Cube Method"

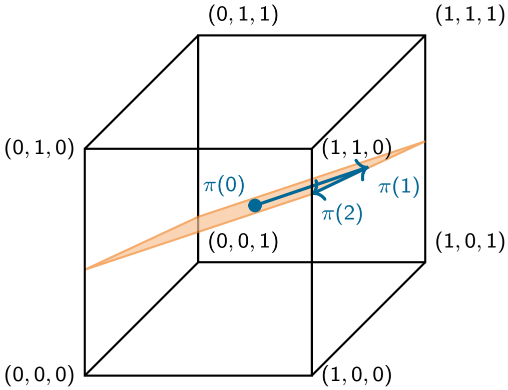

$$\left\{ (s_1, s_2, s_3) \in [0,1]^3 \middle| \frac{X_1 s_1}{\pi_1} + \frac{X_2 s_2}{\pi_2} + \frac{X_3 s_3}{\pi_3} = \frac{X_1(1-s_1)}{1-\pi_1} + \frac{X_2(1-s_2)}{1-\pi_2} + \frac{X_3(1-s_3)}{1-\pi_3} \right\}$$

Back

Possible Modifications of the Algorithm

In practice, the empiricist might want to slightly modify the algorithm

1. X with discrete support: It becomes necessary to reintroduce waiting periods if more than one unit are allocated in the same period.

2. Maximal waiting rule: If at some period t, a unit has waited more than T periods, allocate the unassigned units $s \in U(t)$ using independent coin tosses with probabilities $\pi_s(t)$.

3. Batches: Units are allowed to wait more before being assigned, and we assign multiple units at the same time $\rightarrow$ useful for batched multi-arm bandits.

$\hookrightarrow$ This paper focuses on the general algorithm for statistical properties.

Back

Precision Gains: Misspecified Model

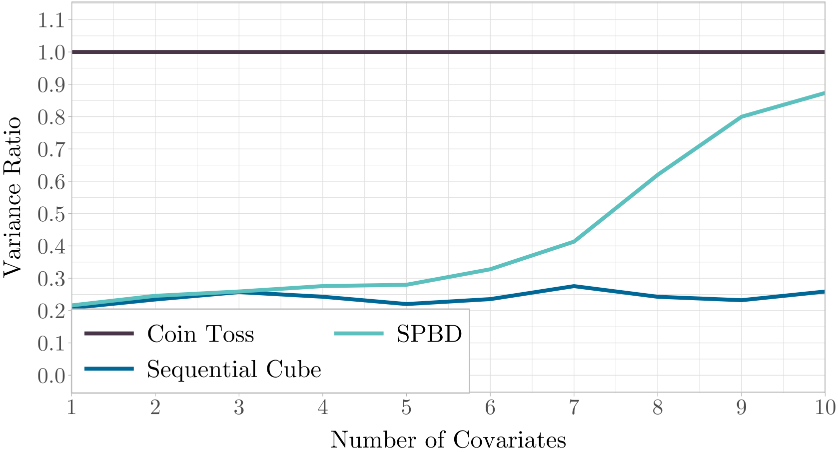

Figure 1: Variance Ratio in Misspecified Model

◂ Back

Precision Gains: Noise Model

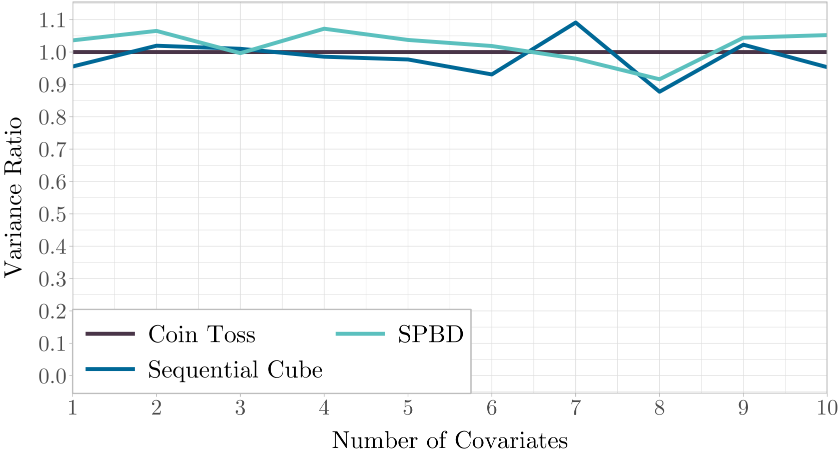

Figure 2: Variance Ratio in Noise Model

◂ Back

Multiple Treatments

I first extend the method to the case of multiple treatments. I do so by nesting the algorithm. This introduces extra constraints and thus further waiting periods.

For simplicity consider, $D_t \in \{0, 1, 2\}$ and propensity scores π_{t0}, π_{t1} and π_{t2}, with $\pi_{t0} + \pi_{t1} + \pi_{t2} = 1$. We start by allocating units to $D = 2$, so we update the propensity scores π_{t2} and $1 - \pi_{t2} = \pi_{t1} + \pi_{t0}$. At a period s, the update vector $u(s)$, must satisfy

$$\boldsymbol{u(s)} \in \ker\left(\frac{\boldsymbol{X^T}}{\boldsymbol{\pi_2(\pi_1 + \pi_0)}}\right).$$

Some units will be excluded from treatment 2 at a period s ($\pi_{2t}(s) = 0$). These units must be allocated between treatments 0 and 1. This assignment requires two constraints:

$$\boldsymbol{u(s)} \in \ker\left(\frac{\boldsymbol{X^T}}{\boldsymbol{\pi_0}}, \frac{\boldsymbol{X^T}}{\boldsymbol{\pi_1}}\right).$$

Graphical Intuition

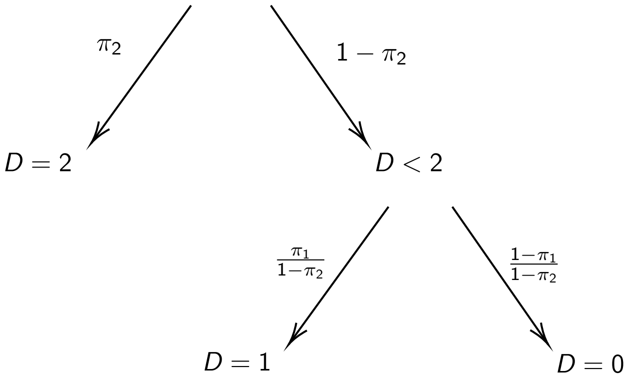

Remark (Homogeneous Probabilities).

For homogeneous π_i, one constraint at each step is sufficient.

Multi-Armed Bandits

Several aspects of the algorithm could be interesting for bandit algorithms:

1. ETC: Use a balancing algorithm in the exploratory phase.

→ More precise estimates should lead to less regret.

2. Bootstrap Linear Thompson Sampling (Athey et al., 2022; Caria et al., 2024): The treatment probabilities are a convex combination between the estimated probability of success of a treatment and a positive constant.

→ I can exploit the suitability for heterogeneous probabilities in my algorithm.

Targeting Precision

Ongoing research on other adaptive features of the method that could increase precision:

1. Adapt probabilities with respect to the feasible Neyman allocation: $\pi_t = \frac{\widehat{\sigma_1}}{\widehat{\sigma_1}+\widehat{\sigma_0}}$
2. Contextual Neyman Allocation: $\pi_t = \frac{\hat{\sigma}_1(X_t)}{\hat{\sigma}_1(X_t)+\hat{\sigma}_0(X_t)}$
3. Use of responses to adapt the covariates used for balancing. E.g., estimate the function $f(.)$ non-parametrically and balance on $\widehat{f}(X_t)$ for future units.