When and How to Pilot:
Statistical Decision Theory for Two-Wave Experiments

Juan C. Yamin
Brown University

2025 Advances with Field Experiments (AFE) Conference

October 8, 2025

Treatment Assignment Probability as a Design Problem

- Experimental design requires making decisions under uncertainty
- One such decision: the treatment assignment probability $\pi = \Pr(D = 1)$
- The goal is to choose π to minimize the variance of the ATE estimator
- If outcome variances were known, the Neyman allocation (Neyman, 1934) is optimal:

$$\pi^* = \frac{\sigma_1}{\sigma_1 + \sigma_0}$$

$\Rightarrow$ Allocate more units to the group with higher outcome variance

Motivation: Using Pilot Data in Two-Wave Experiments

- Estimate variances in the pilot and then plug into Neyman's rule

$$\hat{\pi}^* = \frac{\hat{\sigma}_1}{\hat{\sigma}_1 + \hat{\sigma}_0}$$

- This Feasible Neyman Allocation (FNA) sets the main-wave assignment probability

- But pilots are small and noisy
  - Cai & Rafi (2024): FNA can perform worse than a 50/50 split

- **This paper:** How should experimenters use pilot data to choose the main-wave treatment assignment probability?
  - Model the choice of treatment assignment probability as a statistical decision problem
  - Exact finite-sample analysis (no parametric assumptions or asymptotics)

Roadmap of Talk

Setup as a Decision Problem

- Two-wave experiment (both waves under CRD):
  - Pilot of size m: balanced treatment and control
  - Main wave of size n: assign treatment with probability π

- Goal: estimate the ATE $= \mathbb{E}[Y(1) - Y(0)]$ with minimum variance

- Variance of diff-in-means (IPW) estimator:

$$V(\pi, \sigma_1^2, \sigma_0^2) = \frac{\sigma_1^2}{\pi} + \frac{\sigma_0^2}{1 - \pi}$$

- Assumption: bounded outcomes $Y(d) \in [0, 1] \Rightarrow (\sigma_1^2, \sigma_0^2) \in \Theta = [0, 1/4]^2$

- Decision problem: use pilot data to choose π to minimize $V(\pi, \sigma_1^2, \sigma_0^2)$

Roadmap of Talk

Regret and Minimax-regret

- Expected regret of a rule $p(\cdot)$. Regret is the expected excess variance from using $p(\cdot)$ instead of π^*:

$$R(p, \sigma_1^2, \sigma_0^2) \;=\; \mathbb{E}\big[\, V(p, \sigma_1^2, \sigma_0^2) - V(\pi^*, \sigma_1^2, \sigma_0^2) \,\big].$$

- Following Manski (2004) and others, find rule that minimizes worst-case regret:

$$p_{\text{MR}} \;\in\; \arg\min_{p \in \mathcal{D}} \; \max_{(\sigma_1^2, \sigma_0^2) \in \Theta} R(p, \sigma_1^2, \sigma_0^2).$$

  $\Rightarrow$ Rule that keeps worst-case excess variance as small as possible.

- Game-theoretic view:
  - **Experimenter:** chooses p to minimize regret.
  - **Nature:** chooses $(\sigma_1^2, \sigma_0^2) \in \Theta$ to maximize regret.

Nash Equilibrium

Proposition (Minimax Regret Rule)

In the simultaneous zero-sum game between the Experimenter and Nature:

1. *The Experimenter's unique best reply is*

$$p_{MR} \equiv \frac{1}{2}.$$

2. *Nature's best replies are mixtures of the two extreme points* $(\frac{1}{4}, 0)$ *and* $(0, \frac{1}{4})$.
3. *The value of the game is* $1/4$.

- Any imbalance can be exploited by Nature $\Rightarrow$ the safe design is always 50/50
- Regret is uniformly bounded under p_{MR}: the maximum regret is at most $1/4$
- The pilot data is uninformative!

Why the Feasible Neyman Allocation (FNA) is not *Optimal*

- FNA plugs pilot estimated variances into Neyman's allocation

$$\hat{p}^* = \frac{\hat{\sigma}_1}{\hat{\sigma}_1 + \hat{\sigma}_0}.$$

- Regret of FNA at true variances (σ_1^2, σ_0^2):

$$R(\hat{p}^*, \sigma_1^2, \sigma_0^2) = \sigma_1^2\, \mathbb{E}\left[\frac{\hat{\sigma}_0}{\hat{\sigma}_1}\right] + \sigma_0^2\, \mathbb{E}\left[\frac{\hat{\sigma}_1}{\hat{\sigma}_0}\right] - 2\sigma_1\sigma_0.$$

- **Problem:** ratios of sample standard deviations can explode
  $\Rightarrow$ Worst-case regret of FNA is unbounded.

- **Why people use it?** If $\hat{\sigma}_d \xrightarrow{p} \sigma_d$ as $m \to \infty$, then $\hat{p}^* \xrightarrow{p} \pi^*$ and $R(\hat{p}^*, \theta) \xrightarrow{p} 0$.

- What to do?
  - FNA is *optimal* when pilot is large, but can perform very poorly with small or noisy pilots
  - Minimax regret is always safe but ignores pilot information entirely

Roadmap of Talk

Conditional Minimax Regret (CMR): Definition

Minimize worst-case regret, but only over *plausible* variances supported by the pilot.

1. Estimate treatment and control variances from the pilot: $(\hat{\sigma}_1^2, \hat{\sigma}_0^2)$
2. Construct confidence intervals $I_d = [\sigma^2_{\ell,d}, \sigma^2_{u,d}]$ around each $\hat{\sigma}^2_d$ ▸ Details
3. Form the confidence rectangle such that $\Pr\big((\sigma_1^2, \sigma_0^2) \in \hat{\Theta}_\alpha\big) \geq 1 - \alpha$:

$$\hat{\Theta}_\alpha = I_1 \times I_0 = [\sigma^2_{\ell,1}, \sigma^2_{u,1}] \times [\sigma^2_{\ell,0}, \sigma^2_{u,0}] \subseteq \ [0, \tfrac{1}{4}]^2$$

4. Restrict Nature to $\hat{\Theta}_\alpha$ and solve minimax regret problem: ▸ Closed-form

$$p_{\text{CMR}} \in \arg\min_{p \in \mathcal{A}} \max_{(\sigma_1^2, \sigma_0^2) \in \hat{\Theta}_\alpha} R(p, \theta)$$

CMR: Statistical Guarantees

Proposition (CMR regret on the good event)

Fix $\alpha \in (0, 1)$. On the event that the pilot confidence set $\hat{\Theta}_\alpha$ contains the truth (which occurs with probability at least $1 - \alpha$):

1. *Uniform safety.*
$$\max_{\theta \in \hat{\Theta}_\alpha} R(p_{\mathrm{CMR}}, \theta) \;\le\; \frac{1}{4}.$$
2. *Vanishing regret. If (σ_1^2, σ_0^2) lies in the interior, $R(p_{\mathrm{CMR}}, \sigma_1^2, \sigma_0^2) \;=\; O_p(m^{-1})$. If (σ_1^2, σ_0^2) lies on the boundary of $[0, \frac{1}{4}]^2$, then $R(p_{\mathrm{CMR}}, \sigma_1^2, \sigma_0^2) \;=\; O_p(m^{-1/2})$.*

High-probability guarantees:

- Regret is at most 1/4
- Regret goes to zero as the pilot grows

▸ Certified Decision

Roadmap of Talk

Model

Minimax Regret Rule

Conditional Minimax Regret Rule

Calibrated Simulations

Conclusions

Bibliography

Setup

- Follow Bai (2022); apply to multiple empirical papers (e.g., Abel et al., 2019; Herskowitz, 2021; Abel et al., 2020,...). Details

- For each paper, impute counterfactual potential outcomes using three models:
  - **Model 1:** No treatment effect $Y_i(1) = Y_i(0)$
  - **Model 2:** Nearest–neighbor (Mahalanobis distance) match
  - **Model 3:** Prediction using local linear regression

- Draw (with replacement) pilot samples from *population* and randomize treatment

- Compute treatment assignment probability of main wave with:
  1. Feasible Neyman Allocation (Plug-in)
  2. Minimax Regret
  3. Conditional Minimax Regret
  4. Neyman allocation using population variances (to compute regret)

Models

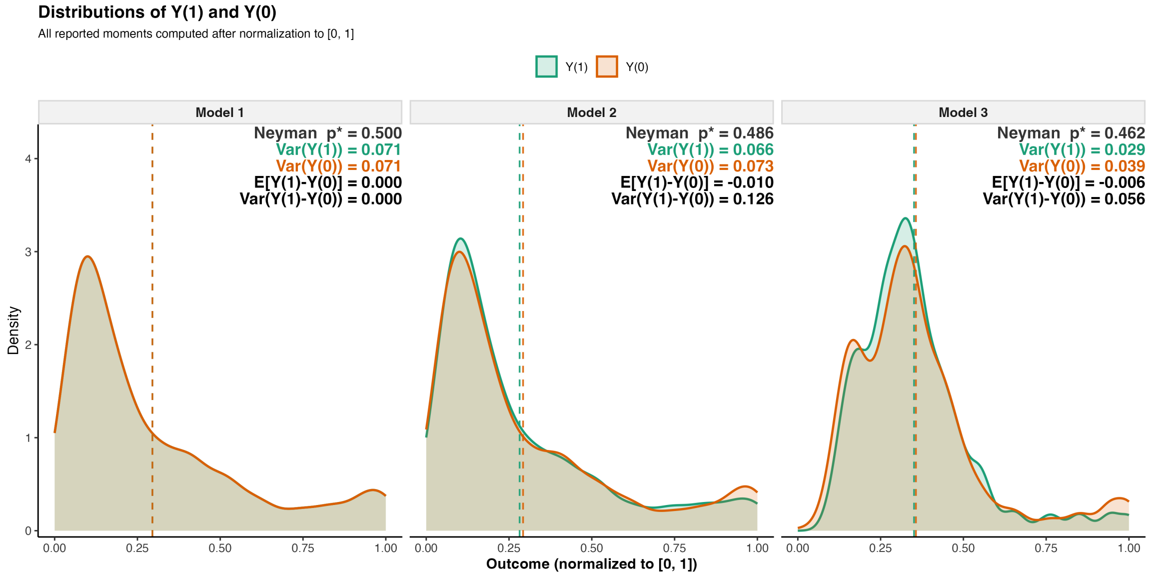

Distributions of Y(1) and Y(0)
All reported moments computed after normalization to [0, 1]
Y(1)
Y(0)
Model 1
Neyman p* = 0.500
Var(Y(1)) = 0.071
Var(Y(0)) = 0.071
E[Y(1)-Y(0)] = 0.000
Var(Y(1)-Y(0)) = 0.000
Model 2
Neyman p* = 0.486
Var(Y(1)) = 0.066
Var(Y(0)) = 0.073
E[Y(1)-Y(0)] = -0.010
Var(Y(1)-Y(0)) = 0.126
Model 3
Neyman p* = 0.462
Var(Y(1)) = 0.029
Var(Y(0)) = 0.039
E[Y(1)-Y(0)] = -0.006
Var(Y(1)-Y(0)) = 0.056
Density
0
1
2
3
4
0.00
0.25
0.50
0.75
1.00
Outcome (normalized to [0, 1])

Results

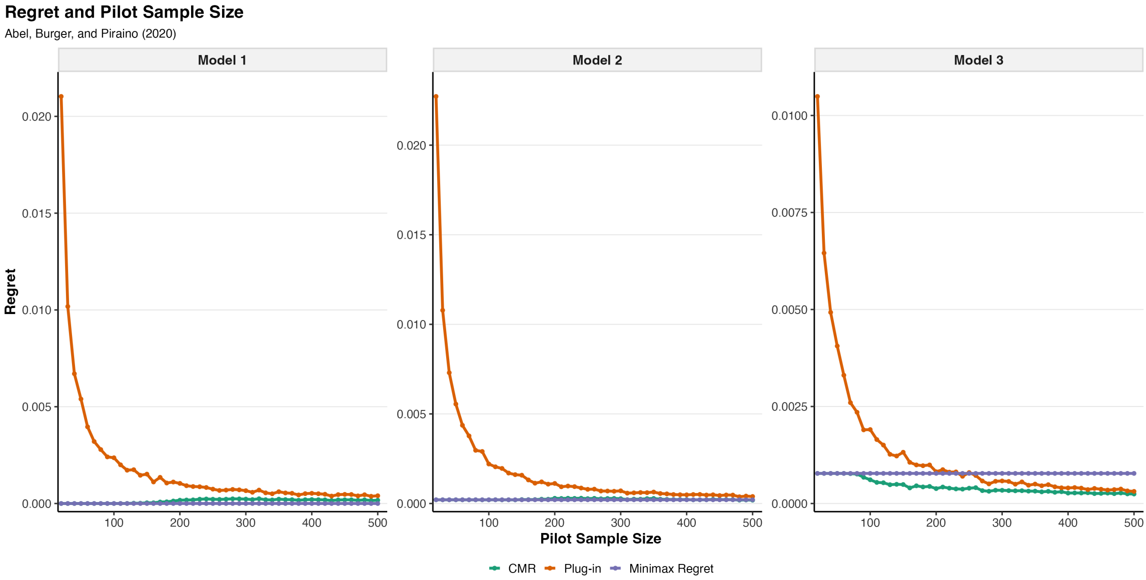
Regret and Pilot Sample Size
Abel, Burger, and Piraino (2020)
Model 1
Model 2
Model 3
Regret
0.020
0.015
0.010
0.005
0.000
0.0100
0.0075
0.0050
0.0025
0.0000
100
200
300
400
500
Pilot Sample Size
CMR
Plug-in
Minimax Regret

Roadmap of Talk

Conclusions

- Focus only on pilots used to choose the main-wave treatment assignment probability
  - Pilots have many other advantages (testing logistics, feasibility, implementation, etc.)
- When to pilot?
  - If the goal is worst-case protection (minimax regret), pilots do not help
  - With CMR or FNA, larger pilots give more information (better alignment with Neyman)
- How to pilot?
  - Be cautious with small/noisy pilots: asymptotic optimality $\neq$ finite-sample optimality
  - Conditional Minimax Regret (CMR) provides a principled middle ground:
    - Defaults to balance with small pilots
    - Converges (w.h.p.) to Neyman allocation with informative pilots

Roadmap of Talk

Bibliography I

Abel, M., Burger, R., Carranza, E., & Piraino, P. (2019). Bridging the intention-behavior gap? the effect of plan-making prompts on job search and employment. *American Economic Journal: Applied Economics*, *11*(2), 284–301.

Abel, M., Burger, R., & Piraino, P. (2020). The value of reference letters: Experimental evidence from south africa. *American Economic Journal: Applied Economics*, *12*(3), 40–71.

Bai, Y. (2022). Optimality of matched-pair designs in randomized controlled trials. *American Economic Review*, *112*(12), 3911–3940.

Cai, Y., & Rafi, A. (2024). On the performance of the neyman allocation with small pilots. *Journal of Econometrics*, *242*(1), 105793.

Herskowitz, S. (2021). Gambling, saving, and lumpy liquidity needs. *American Economic Journal: Applied Economics*, *13*(1), 72–104.

Manski, C. F. (2004). Statistical treatment rules for heterogeneous populations. *Econometrica*, *72*(4), 1221–1246.

Bibliography II

Neyman, J. (1934). On the two different aspects of the representative method: the method of stratified sampling and the method of purposive selection. In *Breakthroughs in statistics: Methodology and distribution* (pp. 123–150). Springer.

Confidence Intervals

- Apply Empirical Bernstein inequality (Maurer–Pontil, 2009) for bounded outcomes $Y(d) \in [0, 1]$: with prob. $\geq 1 - \alpha/2$,

$$\left|\hat{\sigma}_d^2 - \sigma_d^2\right| \;\leq\; 2\sqrt{\frac{2\hat{\sigma}_d^2 \ln(4/\alpha)}{m_d - 1}} + \frac{2\ln(4/\alpha)}{m_d - 1}.$$

- Define the random radius:

$$t_{d,\alpha}(\omega) = 2\sqrt{\frac{2\hat{\sigma}_d^2(\omega)\ln(4/\alpha)}{m_d - 1}} + \frac{2\ln(4/\alpha)}{m_d - 1}.$$

- Construct per-arm confidence interval (clipped to $[0, 1/4]$):

$$I_d(\omega) = \left[\max\{0, \hat{\sigma}_d^2 - t_{d,\alpha}\},\ \min\{\tfrac{1}{4}, \hat{\sigma}_d^2 + t_{d,\alpha}\}\right].$$

- Joint confidence set is the rectangle:

$$\hat{\Theta}_\alpha = I_1(\omega) \times I_0(\omega),$$

which covers the true (σ_1^2, σ_0^2) with probability at least $1 - \alpha$.

Back

What p_{CMR} does, its formula, and a simple example

- The worst case occurs at the two off-diagonal corners $A = (\sigma_{u,1}, \sigma_{\ell,0})$ and $B = (\sigma_{\ell,1}, \sigma_{u,0}) \Rightarrow$ p_{CMR} equalizes their regrets

- Closed-form solution:

$$p_{\mathrm{CMR}}(\omega) = \begin{cases} \dfrac{\Delta_1 + \Delta_0 + C - \mathrm{sign}(C)\sqrt{(\Delta_1 + \Delta_0 + C)^2 - 4C\,\Delta_1}}{2C}, & C \neq 0, \\ \dfrac{\Delta_1}{\Delta_1 + \Delta_0}, & C = 0. \end{cases}$$

 where $\Delta_d = \sigma_{u,d}^2 - \sigma_{\ell,d}^2$ (arm d width) and $C = (\sigma_{u,1}+\sigma_{\ell,0})^2 - (\sigma_{\ell,1}+\sigma_{u,0})^2$ compares the overall scale of the two corners (i.e., $C>0 \Rightarrow B$ is scarier).

- Examples:
  1. **Disjoint or shifted bands:** if $\sigma_{\ell,1}^2 > \sigma_{u,0}^2$ then treatment is known noisier $\Rightarrow$ $p_{\mathrm{CMR}} > 1/2$
  2. **High precision:** $\Delta_1, \Delta_0 \to 0 \Rightarrow$ $p_{\mathrm{CMR}} \to p^*_{\mathrm{Neyman}}$.
  3. **Maximal ambiguity:** wide and symmetric bands ($\Delta_1 \approx \Delta_0$, $C \approx 0$) $\Rightarrow$ $p_{\mathrm{CMR}} = 1/2$.

Back

Why the m^{-1} Rate in the Interior & CMR vs FNA

- The optimal allocation $p^* = \sigma_1/(\sigma_1 + \sigma_0)$ minimizes $V(\cdot)$, so $V'(p^*) = 0$.
- Variance estimates have order $m^{-1/2}$, hence $\hat{p} - p^* = O_p(m^{-1/2})$.
- Around p^*, regret is second order:

$$R(\hat{p}, \theta) \approx \frac{1}{2} V''(p^*)(\hat{p} - p^*)^2,$$

  so with $\hat{p} - p^* = O_p(m^{-1/2})$ we get $R = O_p(m^{-1})$.
- The confidence rectangle for (σ_1^2, σ_0^2) shrinks at width $O_p(m^{-1/2})$. CMR inherits the square of the width $\Rightarrow O_p(m^{-1})$.
- **FNA in the interior.** Same local logic: plug in $\hat{\sigma}_d$ gives $\hat{p}$ that is root-m consistent, hence $R_{\text{FNA}} = O_p(m^{-1})$, provided variances are bounded away from 0 and $1/4$.

Summary of interior rates:

- CMR: $O_p(m^{-1})$ and high-probability safety bound.
- FNA: $O_p(m^{-1})$ but no safety guarantee.

Why Rates Are Slower Near the Boundary

- **Boundary geometry.** When a true variance is near 0 or $1/4$, the confidence interval becomes one-sided and clipped. The symmetric "two-sided" shrinkage that powers the m^{-1} cancellation is lost.
- **First-order term no longer cancels.** In the interior, $V'(p^*) = 0$ and the leading error is quadratic in $\hat{p} - p^*$. Near the boundary, moving within the confidence set shifts the worst case mostly in one direction, so the contribution becomes **first order** in the set's width.
- **Width of the set.** Pilot uncertainty still shrinks like $m^{-1/2}$. Without the interior cancellation, worst-case regret scales like the **width itself** $\Rightarrow O_p(m^{-1/2})$.
- **Ratios become unstable.** Near $\sigma_d^2 \approx 0$, ratios such as $\hat{\sigma}_1/\hat{\sigma}_0$ can be heavy tailed. CMR guards by optimizing over a bounded set; FNA can suffer large allocations and large regret.
- **Takeaway.** Interior: errors in opposite directions offset, leaving a quadratic remainder (m^{-1}). Boundary: offset breaks, leaving a linear remainder ($m^{-1/2}$).

CMR and Certified Decisions (Andrews & Chen, 2025)

- **CMR is a $(1-\alpha)$ P-certified decision**
  - Define the certificate
  $$R_{\mathrm{CMR}}(\omega) = \sup_{\theta \in \widehat{\Theta}_\alpha(\omega)} r(p_{\mathrm{CMR}}(\omega), \theta).$$
  - By coverage, $P\{\, L(p_{\mathrm{CMR}}(\omega), \theta) \le R_{\mathrm{CMR}}(\omega) \,\} \ge 1-\alpha$. Hence $R_{\mathrm{CMR}}(\omega)$ is a valid P-certificate for regret.

- **CMR belongs to an essentially complete class of certified decisions**
  - *Section 3 result.* P-certified decisions obtained by as-if optimization over any $1-\alpha$ set form an essentially complete class. For any certified $(\tilde{\delta}, \tilde{R})$ there exists an as-if (δ, R) that weakly dominates it.
  - *Implication for CMR.* CMR is an as-if rule using $\widehat{\Theta}_\alpha(\omega)$, so it lies in this class. Focusing on CMR-type designs is without loss under this order.

▸ Back

Empirical Application: Abel, Burger & Piraino (2020)

- **Context:** High youth unemployment in South Africa; employers face information frictions about applicants' skills
- **Intervention:** Researchers designed a standardized *reference letter template* for former employers to complete
- **Treatment:** Job seekers were randomly encouraged to obtain and use a standardized reference letter from a former employer
- **Outcome:** total hours worked in the last week measured at 3 months

▸ Back to Setup