When and How to Pilot:
Design Rules for Two-Wave Experiments

Juan C. Yamin*
Brown University

September 13, 2026

Abstract

Experimenters often run pilots, but how much a small pilot should shape the main-wave design has no settled answer. This paper shows how noisy pilot evidence should guide treatment assignment probabilities in two-wave experiments. Two canonical rules mark the extremes. Balanced assignment guards against worst cases but ignores evidence that one arm is noisier. Feasible Neyman allocation adapts, but with a finite pilot it can overreact to noise, producing arbitrarily large precision losses. I propose a Conditional Minimax Regret (CMR) rule that minimizes worst-case regret over a finite-sample confidence set for the treatment and control variances. CMR retains balance's worst-case protection with high probability, converges to the Neyman allocation as the pilot grows, and attains the minimax-regret rate up to constants. It extends to multi-arm and stratified designs, and simulations calibrated to four field experiments show it avoids feasible Neyman's severe small-pilot losses while matching its large-pilot gains.

Keywords: Experimental design; Statistical decision theory; Minimax regret.

JEL Codes: C44, C90, C93, D81.

*Brown University, Department of Economics, Robinson Hall, 64 Waterman Street, Providence, RI 02912. Email: juan_yamin_silva@brown.edu. I am grateful for the generous advice and support of Toru Kitagawa, Soonwoo Kwon, and Jonathan Roth. I thank Andrew Chesher, Yuchen Hu, Giacomo Opocher, Bobby Pakzad-Hurson, Moritz Poll, Chen Qiu, Panos Toulis, Stefan Wager, Kohei Yata, and participants at the 2025 Advances with Field Experiments (AFE) Conference and Brown's Econometric Coffee for their helpful comments. R and Python implementations of the Conditional Minimax Regret rule, covering the two-arm, multi-arm, and stratified designs, are available at https://github.com/juancyamin/cmrdesign.

1 Introduction

Nearly one in ten of the 10,905 trials registered in the AEA RCT Registry over the past decade reports running a pilot.[1] Experimenters use them to test logistics, refine survey instruments, and learn about the study population before launching the main wave. When a pilot measures outcomes under both treatment and control, its data can also inform a consequential design choice, the treatment assignment probability. By determining how many observations each arm receives, this probability directly affects the variance of the average treatment effect estimator and hence the experiment's power. The promise of pilot-based assignment is that it can use early evidence to allocate more observations where they are most valuable, and hence deliver more precise estimates from the same sample size.

If the potential-outcome variances under treatment and control were known, the variance-minimizing assignment rule would be the Neyman allocation (Neyman, 1934). It assigns a larger share of the experimental sample to the noisier arm, reflecting the principle that precision requires more observations where outcomes are more variable. In practice, these variances are unknown when the main-wave design is chosen. A natural rule estimates them from the pilot and substitutes the resulting sample variances into the Neyman formula, yielding the feasible Neyman allocation (FNA). When the pilot is large enough for the estimates to be reliable, this plug-in logic is well founded. For example, Hahn et al. (2011) show that pilot-based estimates of conditional variances can deliver asymptotically efficient assignment probabilities when both pilot and main-wave samples grow large.

When the pilot is small, however, the premise of this plug-in logic fails, and the experimenter must decide how much the design should respond to variance estimates that are not necessarily trustworthy. Keeping the main wave evenly split ensures that neither arm ends up with too few observations, which caps the possible precision loss whatever the true variances. The cost is that it forgoes the gains the pilot could deliver. Following the pilot estimates as if they were true removes this protection, and can tilt the main wave too far, or toward the wrong arm, precisely when the pilot is least informative.

These small-sample concerns are empirically relevant, since in the AEA RCT Registry the median reported pilot has about 300 observations and roughly a tenth have fewer than 30, sizes at which variance estimates can remain very noisy. In fact, Cai and Rafi (2024) show that a fixed pilot can leave the feasible Neyman allocation less precise than balance even when the main wave is large. This makes pilot-based assignment a finite-sample design

[1] Between 2016 and 2025, 991 of the 10,905 first-registered trials (9.1%) report a pilot, but only 24 (0.22% of all trials) report using it to set assignment probabilities. Because informal piloting often goes unmentioned in registrations, both counts are plausibly lower bounds on actual use. Online Appendix F provides details.

problem, not a plug-in calculation.

I formalize the experimenter's choice as a finite-sample statistical decision problem. In a two-wave experiment with binary treatment, bounded potential outcomes, and no parametric restrictions, the experimenter observes a pilot and then chooses the main-wave assignment probability. The goal is to minimize the variance of the average treatment effect estimator. I evaluate decision rules by their regret, the excess variance relative to the infeasible Neyman allocation that would be chosen if the variance pair were known. Although the setting is deliberately simple, the question it isolates is general. Design rules should be optimized for the information actually available at the moment of design, not imported from a plug-in analysis that treats pilot estimates as accurate enough to act on. Pilot-based assignment is the tractable model case in which that principle can be developed exactly.

In this setup, the two canonical assignment rules mark the endpoints of the finite-pilot design problem. Balanced assignment, the no-data rule, is optimal under minimax risk and is uniquely minimax-regret optimal without a pilot. It therefore has a formal foundation for caution, but it gives the realized pilot no authority over the assignment. Feasible Neyman gives the pilot's point estimates full authority. This plug-in logic attains the large-pilot efficiency benchmark (Armstrong, 2026), but with a finite pilot nothing limits how far noisy estimates can tilt the assignment, so the precision loss can be arbitrarily large—in the extreme, an arm whose pilot observations all coincide is estimated to have zero variance and receives no main-wave observations at all. In short, balance refuses to move, while feasible Neyman can move too far. The finite-pilot question is how far the realized evidence should move the assignment from balance toward the feasible Neyman allocation.

Exact minimax regret is the natural ex ante criterion for finite-pilot adaptation (Savage, 1951; Manski, 2021). It asks, before the pilot is drawn, which complete mapping from pilot realizations to assignments has the smallest worst-case expected regret. By this criterion, with as few as two observations per pilot arm, the minimax-regret value falls strictly below its no-pilot level, so any exact minimax-regret rule must move the assignment away from balance for some pilot realizations. The criterion does not, however, yield an operational design rule for this problem. The regret of an assignment depends on the population only through its treatment and control variances, but the expected regret of a rule does not, since populations with identical variances can induce different pilot distributions. The minimax problem therefore ranges over joint distributions of the potential outcomes, not variance pairs, and no finite-dimensional reduction is known. The search over decision rules is also not finite-dimensional, since each candidate is a complete mapping from pilot realizations to assignments. A conditional approach starts instead from the single pilot actually drawn and asks which variance pairs the evidence has ruled out and how much movement from balance the surviving configurations justify.[2]

The Conditional Minimax Regret (CMR) rule answers this question with a principle from the literature on inference for decision making, acting on what the data have not ruled out rather than on a point estimate (Manski, 2021; Chernozhukov et al., 2025). Once the pilot is observed, the only uncertainty that matters for the loss is which variance pair is true. CMR therefore builds a finite-sample confidence set for the treatment and control variances, the configurations still consistent with the realized pilot, and chooses the main-wave assignment that minimizes worst-case regret over that set. In the binary-treatment case the rule has a closed form, the Neyman formula applied to the midpoint of each arm's standard-deviation confidence interval. When the set is wide, CMR stays close to balance. As the set contracts around the true variances, CMR moves toward the Neyman allocation. Alongside the assignment, CMR reports a certificate, a bound on the regret the chosen assignment can incur (Andrews and Chen, 2025). Movement from balance is thus earned by what the pilot rules out, not imposed by a point estimate.

The CMR rule comes with finite-sample and large-pilot guarantees. In finite samples, the certificate bounds the realized regret of the chosen assignment with probability at least $1-\alpha$, turning a loss the experimenter cannot observe into a number computed from the pilot alone. This certificate is never larger than the guarantee available without a pilot, and it tightens as soon as the pilot rules out one of the adversarial configurations that make a large tilt dangerous. As the pilot grows, the caution that holds CMR near balance relaxes, and at interior variance pairs the rule converges to the infeasible Neyman allocation at the same rate as feasible Neyman. Thus asymptotic efficiency and finite-sample safety hold together rather than being traded off. Specifically, CMR's worst-case expected regret matches the minimax-regret rate up to constants, the best rate any pilot-based rule can achieve.

The CMR construction is flexible, and it remains simple to compute across the designs experimenters actually run. In multi-arm experiments with a shared control and in stratified designs, the assignment probability becomes an allocation vector, and CMR again moves from the no-information allocation toward the Neyman allocation as the pilot narrows the confidence set. The same recipe also adapts to the outcomes themselves. For binary outcomes, exact confidence sets sharpen the rule precisely at the small pilots where it matters most. For unbounded outcomes, a kurtosis condition can replace the assumption of a known outcome bound. The construction carries over to designs that target several primary out-

[2] Lim and Park (2026) characterize ex ante criteria under which constructing the rule realization by realization is optimal.

comes and to settings in which the pilot observes only a short-run proxy for the outcome that defines the loss.

I also ask when a pilot is worth running for design in the first place, and how large it should be. Observations spent on the pilot could have gone to the main wave instead, so adaptation has to repay that diversion. A worthwhile pilot must be large enough for CMR to move away from balance at all, yet small enough that even perfect adaptation recovers the observations it cost. Folding pilot observations into the final estimator softens this trade-off. No known method computes the worst-case optimal design, but I prove that a simple rule approximates it. For a total budget of n observations, a balanced pilot of total size roughly $n^{2/3}$—about 100 observations when $n = 1{,}000$—followed by CMR comes within a constant factor of the best worst-case performance available to any balanced pilot size paired with any pilot-based assignment rule. Larger experiments therefore run larger pilots but devote a shrinking share of the budget, of order $n^{-1/3}$, to piloting.

Simulations calibrated to four well-known field experiments (Miguel and Kremer, 2004; Bertrand and Mullainathan, 2004; Thornton, 2008; Abel et al., 2020) compare the rules at pilot sizes from 30 to 500 in two-arm, shared-control multi-arm, and stratified designs. I measure performance by the percentage increase in the estimator's variance relative to the infeasible Neyman allocation. A loss of one percent means the experimenter would need one percent more main-wave observations to reach the same precision. By this measure, feasible Neyman often does substantially worse than simply keeping the main wave balanced. With a pilot of 30 its mean loss is infinite in six of the nine calibrations and reaches 72 percent where finite. CMR, by contrast, is never substantially worse than either benchmark. Its mean loss exceeds balance's by at most 0.09 percentage points in any design at any pilot size, and in the multi-arm and stratified designs it outperforms feasible Neyman at every reported pilot size. With 500 pilot observations it captures 81 percent of the attainable gain in the multi-arm design and roughly half in the stratified design.

Those gains are modest in these calibrations because the standard deviations are never far apart. Balance lies less than one percent above the Neyman benchmark in the two-arm designs, 1.8 percent above it in the multi-arm design, and 5.6 percent above it in the stratified design. With one standard deviation twice the other, the two-arm gap would instead be eleven percent, and designs with more arms or strata leave adaptation even more room. The gains are also nearly free for the many experimenters who already run a pilot. The assignment probability must be chosen regardless, so CMR can realize some of these savings using data collected to test logistics and refine survey instruments.

The broader contribution is a template for building statistical uncertainty into experimental design itself. Everything in the paper runs on three features of the assignment problem. First, the unknown parameters enter the design loss only through a low-dimensional vector of variances. Second, the loss is linear in each variance, which makes regret convex, so worst-case regret over any rectangle of plausible values is attained at one of its corners. Third, a small pilot delivers a finite-sample confidence bound for each variance at any prespecified level. Any design problem with these features admits the same treatment, a decision paired with a finite-sample certificate for the loss it can incur, and every extension in this paper is an instance of that recipe.

Two literatures frame these contributions. The first is adaptive experimental design and pilot-based assignment. Hahn et al. (2011) set assignment probabilities from pilot data, Tabord-Meehan (2023) builds adaptive stratification rules, and Bai (2022) and Cytrynbaum (2026) design matched and stratified experiments from pre-experimental information, with Armstrong (2026) characterizing the efficiency frontier these designs target. Higbee (2025) studies a Bayes version of the two-wave problem, designing the main wave to maximize the prior-averaged welfare of a downstream policy choice. These guarantees are asymptotic. They identify the efficient design once the pilot estimates the variances reliably, but they are silent when the pilot is too small to be trusted. Building on the finite-pilot warning of Cai and Rafi (2024), I formulate pilot-based assignment as a finite-sample decision problem and derive a computable post-pilot rule that pairs the treatment share with a regret certificate.[3] The contribution is a characterization of how far noisy pilot evidence should move a design while the variances remain uncertain.

The second is inference for decision making. Manski (2021) examines as-if decisions with set estimates, which act on the parameter values a set estimate has not ruled out. Chernozhukov et al. (2025) select policies by balancing estimated welfare against estimation risk, and Andrews and Chen (2025) pair recommended decisions with high-probability bounds on their loss. In these settings, the data are already in hand and the decision is which policy to implement. The decision studied here comes one stage earlier. The action is the main-wave assignment probability, and the loss is the precision of an estimator whose data do not yet exist.[4] This structure yields a rule that is closed form in the binary case, immune to the failures of plug-in rules, and accompanied by a certificate for the realized pilot.

The rest of the paper is organized as follows. Sections 2 and 3 set up the decision problem

[3] A growing literature studies adaptive Neyman allocation in sequential or batched experiments, where the assignment is updated during the experiment (Dai et al., 2023; Zhao, 2024). The rule in this paper chooses the main-wave assignment once, after a single finite pilot.

[4] The regret criterion follows the statistical treatment choice tradition of Manski (2004). The closest design papers are Manski and Tetenov (2016), who choose trial size, and Hu et al. (2024), who choose whom to enroll in a Gaussian model with known variances.

and the canonical rules, Section 4 develops the CMR rule and its guarantees, Section 5 treats multi-arm and stratified designs, Section 6 reports the calibrated simulations, and Section 7 concludes. Online Appendix C extends the rule to binary, unbounded, multiple, and delayed outcomes, and Online Appendix D analyzes when a pilot is worth running and how large it should be.

2 The Two-Wave Design Problem

2.1 Experimental Setup

Population, Outcomes, and Estimand. An experimenter aims to estimate the average treatment effect (ATE) of a binary treatment in a target population. The population is characterized by the joint distribution F of the potential outcomes $(Y(1), Y(0))$, where $Y(1)$ denotes the outcome under treatment and $Y(0)$ denotes the outcome under control. Throughout, we assume that potential outcomes take values in the unit interval, $Y(d) \in [0,1]$ for each treatment status $d \in \{0,1\}$, so $F \in \mathcal{F} = \mathcal{P}([0,1]^2)$, the collection of all Borel probability measures on the compact set $[0,1]^2$. Normalizing outcomes to $[0,1]$ is without loss of generality whenever outcomes are supported on a finite interval, and the particular choice of zero and one is adopted only to simplify notation and exposition. Online Appendix C.2 relaxes boundedness, replacing the known outcome bound with a bound on kurtosis.

The estimand of interest is the ATE, $\mathrm{ATE}(F) = \mathbb{E}_F[Y(1) - Y(0)]$, where the expectation is taken with respect to $F \in \mathcal{F}$. To increase the precision of the ATE estimator, the experiment proceeds in two waves. A smaller pilot precedes a larger main wave and serves only to inform its design. The single design choice is the main-wave assignment probability. Lower estimator variance translates directly into higher power and a smaller minimum detectable effect at any target power, so variance-minimizing assignment makes the experiment as informative as possible about the ATE.

Pilot Wave. The pilot is a sample of M units whose potential outcomes $\{(Y_i(1), Y_i(0))\}_{i=1}^{M}$ are i.i.d. draws from F. For each unit i, the experimenter observes the treatment indicator $D_i \in \{0,1\}$ and the realized outcome $Y_i = D_i Y_i(1) + (1 - D_i) Y_i(0)$. Treatment in the pilot is assigned by a completely randomized design (CRD).[5] In the pilot CRD, exactly M_1 units are assigned to treatment and the remaining M_0 to control, with $M_1 + M_0 = M$. The assignment vector $(D_1, \ldots, D_M)$ is sampled uniformly at random from all binary vectors with exactly

[5] The same arguments extend to Bernoulli assignment conditional on both realized arm sizes being at least two.

M_1 treated units.

The pilot realization is denoted by $\omega = \{(Y_i, D_i)\}_{i=1}^M$, and the pilot sample space is $\Omega = ([0,1] \times \{0,1\})^M$, with the fixed-arm-size restriction imposed by the pilot design. Each pilot arm holds at least two units, $M_1, M_0 \geq 2$, so the within-arm sample variances are well defined. This minimum is maintained wherever the pilot is fixed in advance, and is relaxed only in Appendix D, where the pilot size is itself the object of choice. The extensions in Section 5 impose the same minimum on each relevant cell.

Main Wave. The main wave is a new sample of N units, independent of the pilot, whose potential outcomes $\{(Y_j(1), Y_j(0))\}_{j=1}^N$ are i.i.d. draws from F. The main wave also uses a completely randomized design. For a treatment assignment probability $\pi \in (0,1)$, the design assigns $N_1 = N\pi$ units to treatment and $N_0 = N(1-\pi)$ units to control.[6] The experimenter estimates the ATE using the difference in means, $\widehat{\text{ATE}} = \bar{Y}_1 - \bar{Y}_0$, where $\bar{Y}_d$ is the main-wave average outcome among units assigned to arm d.

2.2 Main-Wave Variance and the Neyman Allocation

The marginal potential-outcome variance in arm d is $\sigma_d^2(F) = \text{Var}_F(Y(d))$, abbreviated σ_d^2 when F is clear from context. Under the main-wave CRD, $\text{Var}_F(\widehat{\text{ATE}}) = \sigma_1^2/N_1 + \sigma_0^2/N_0$, with no cross-arm covariance term because treatment and control are measured on independent units. With $N_1 = N\pi$ and $N_0 = N(1-\pi)$, this equals $V(\pi, \sigma_1^2, \sigma_0^2)/N$, where $V(\pi, \sigma_1^2, \sigma_0^2) = \sigma_1^2/\pi + \sigma_0^2/(1-\pi)$ for $\pi \in (0,1)$. The main-wave sample size N does not depend on the assignment, so the sampling variance is proportional to V with fixed constant $1/N$, and minimizing the variance over π is the same as minimizing V, which we take as the variance criterion for the assignment problem. Boundary recommendations, $\pi \in \{0,1\}$, are assigned infinite variance by convention, since assigning all main-wave units to one arm leaves the other arm mean unobserved.

Minimizing $V(\cdot, \sigma_1^2, \sigma_0^2)$ over $\pi \in (0,1)$ yields the Neyman allocation (Neyman, 1934), which for strictly positive variances is $\pi^*(\sigma_1^2, \sigma_0^2) = \sigma_1/(\sigma_1 + \sigma_0)$. The Neyman allocation gives the larger share of the main wave to the arm with the more variable outcomes, because that arm's sample mean is noisier and benefits more from a larger allocation.

The benchmark is the smallest variance an interior assignment can achieve,

$$V^*(\sigma_1^2, \sigma_0^2) = \inf_{\pi \in (0,1)} V(\pi, \sigma_1^2, \sigma_0^2) = (\sigma_1 + \sigma_0)^2, \tag{2.1}$$

[6] We suppress integer constraints on $N\pi$ throughout, with π understood as the target treatment assignment probability implemented by the closest feasible fixed-arm design.

attained at the Neyman allocation when both variances are strictly positive. Writing V^* as an infimum keeps it well defined when one variance is zero, since the minimizer then lies on the boundary and is only approached from the interior.[7] Because V^* depends on the unknown variance pair, it is the infeasible benchmark against which feasible assignment rules are evaluated.

Remark 2.1 (How much adaptation can gain)**.** The benchmark caps what any pilot-based rule can gain over balance. For every variance pair with $\sigma_1+\sigma_0 > 0$, $V(1/2, \sigma_1^2, \sigma_0^2)/V^*(\sigma_1^2, \sigma_0^2) = 2(\sigma_1^2 + \sigma_0^2)/(\sigma_1 + \sigma_0)^2 \leq 2$, with equality exactly when one variance is zero. Even perfect adaptation can therefore at most halve the estimator's variance, equivalent to doubling the effective sample size, and realistic configurations deliver far less. Balance's excess variance relative to the benchmark is $(\sigma_1 - \sigma_0)^2/(\sigma_1 + \sigma_0)^2$, about eleven percent when one standard deviation is twice the other and four percent when it is fifty percent larger. The ceiling rises in the designs of Section 5, where the no-information allocation must divide the sample more finely. With K treatment arms sharing a control, perfect adaptation can cut variance by up to a factor of $K + \sqrt{K}$ rather than two. In a stratified design the factor is two divided by the smallest stratum share, so with S equally sized strata it is $2S$.

2.3 States, Actions, and Pilot-Based Rules

The two-wave design problem is a statistical decision problem. The unknown state of the world is the population distribution F. It both fixes the main-wave variance of any assignment and generates the pilot data through which the experimenter learns about that variance. Because F is left unrestricted beyond bounded support, the parameter is the full distribution and the parameter space is the nonparametric class $\mathcal{F}$. However, the main-wave variance depends on F only through the pair of marginal potential-outcome variances $\theta(F) = (\sigma_1^2(F), \sigma_0^2(F))$, which ranges over $\Theta = [0, 1/4]^2$. This pair is the payoff-relevant parameter.

Crucially, the variance pair governs the payoff but not the data. At a fixed assignment, the loss is the same under any two distributions sharing a variance pair, since it depends on F only through $\theta(F)$. The expected loss of a rule can still differ between them, because the rule chooses its assignment from the pilot, whose distribution depends on the arm-specific marginal distributions of $Y(1)$ and $Y(0)$ rather than on the variance pair alone. Under the fixed-arm pilot design each unit reveals only one potential outcome, so the pilot distribution depends on F only through these two marginals, and never on the within-unit dependence

[7] With $\sigma_0 = 0$, for instance, the Neyman formula returns the boundary value $\pi^* = 1$. When both variances are zero, every interior assignment attains the value zero, and we normalize $\pi^*(0, 0) = 1/2$.

between $Y(1)$ and $Y(0)$. The statistical experiment is thus indexed by the pair of marginal outcome distributions and not by the variance pair.

The experimenter's action is the main-wave assignment probability $\pi \in \mathcal{A} = [0, 1]$. The boundary actions $\pi \in \{0, 1\}$ are kept in $\mathcal{A}$ so that the framework can evaluate rules that would place the entire main wave in one arm, and the infinite-variance convention of Subsection 2.2 applies to them.

The experimenter does not know the variance pair and estimates it from the pilot. For each arm d, the natural estimate of σ_d^2 is the within-arm sample variance

$$\hat{\sigma}_d^2(\omega) = \frac{1}{M_d - 1} \sum_{i:D_i=d} (Y_i - \bar{Y}_d)^2, \qquad \bar{Y}_d = \frac{1}{M_d} \sum_{i:D_i=d} Y_i.$$

Under the fixed-arm pilot design, $\hat{\sigma}_d^2$ is unbiased for σ_d^2. The main analysis focuses on variance-based decision rules, meaning rules that use the pilot only through the two within-arm sample variances. Formally, the class of decision rules is

$$\mathcal{D} = \left\{ p : \Omega \to \mathcal{A} \;\middle|\; p(\omega) = g\big(\hat{\sigma}_1^2(\omega), \hat{\sigma}_0^2(\omega)\big) \text{ for some measurable } g : \mathbb{R}_+^2 \to \mathcal{A} \right\}.$$

The class includes balance, feasible Neyman, trimmed feasible Neyman, and the Conditional Minimax Regret rules developed below. This focus is deliberate, reflecting how pilot-based assignment is typically implemented in practice and keeping the finite-sample design problem low-dimensional enough to yield tractable rules. However, the restriction is substantive since, relative to the unrestricted class $\mathcal{D}_0 = \{p \in \mathcal{A}^\Omega \mid p \text{ is measurable}\}$, features of the full pilot beyond $(\hat{\sigma}_1^2, \hat{\sigma}_0^2)$ may contain additional information about (σ_1^2, σ_0^2) in the nonparametric model.

Once a decision rule is fixed, the realized pilot outcome ω remains random under F. For each $F \in \mathcal{F}$, let P_F denote the distribution of ω induced by the fixed-arm pilot CRD with arm sizes (M_1, M_0), suppressing this dependence in the notation. The statistical experiment is the family $\{P_F : F \in \mathcal{F}\}$ on Ω.

2.4 Loss and Regret

The variance criterion $V(\pi, \sigma_1^2, \sigma_0^2)$ is the loss from using treatment assignment probability π when the variance pair is $\theta = (\sigma_1^2, \sigma_0^2)$. Regret compares this loss with the infeasible benchmark in (2.1). For an interior assignment probability, $r(\pi, \theta) = V(\pi, \sigma_1^2, \sigma_0^2) - V^*(\sigma_1^2, \sigma_0^2)$. For boundary assignments, regret is infinite by the same convention used for V. For a pilot-based rule $p \in \mathcal{D}$, the realized regret after observing pilot sample ω is $r(p(\omega), \theta(F))$. Regret measures excess main-wave sampling variance relative to the infeasible Neyman benchmark, so it is always nonnegative. For nondegenerate variance pairs with strictly positive variances, regret is zero exactly when the chosen assignment equals the Neyman allocation. At degenerate variance pairs, regret remains well defined because the infeasible benchmark is an infimum over interior assignments.

For every interior assignment $\pi \in (0,1)$, regret has the equivalent representations

$$r(\pi,\theta) = (\sigma_1+\sigma_0)^2 \frac{\big(\pi-\pi^*(\theta)\big)^2}{\pi(1-\pi)} = \frac{\big((1-\pi)\sigma_1-\pi\sigma_0\big)^2}{\pi(1-\pi)}, \tag{2.2}$$

with the normalization $\pi^*(0,0)=1/2$ in the first. Expression (2.2) separates the size of the assignment mistake from the penalty attached to it. The squared distance $(\pi-\pi^*(\theta))^2$ to the Neyman target is weighted by $1/[\pi(1-\pi)]$, so a given mistake is most costly near the boundary, where one arm is barely sampled, and the factor $(\sigma_1+\sigma_0)^2$ scales the loss with the outcome noise. The second form, the square of an expression affine in the two standard deviations, is the version the analysis of Section 4 exploits.

3 Canonical Assignment Rules

This section studies the two assignment rules that current practice and the existing literature bring to the finite-pilot design problem. The first canonical rule is complete balance, which ignores the pilot and assigns treatment with probability one half. Although simple, balance has a decision-theoretic foundation as the solution to the minimax-risk problem. The second is the feasible Neyman allocation, which uses the pilot in the most direct way, replacing the unknown potential-outcome variances with their pilot estimates.

These two rules expose the central tension of the design problem. Balance is safe but does not adapt, while feasible Neyman adapts but is not safe. We then turn to exact minimax regret, the standard decision-theoretic criterion for resolving this tension. The criterion identifies the best worst-case performance any pilot-based rule can achieve, but it is not itself an operational assignment rule, because the game between the experimenter and Nature ranges over joint distributions of the potential outcomes rather than variance pairs.

3.1 Minimax Risk and Balanced Assignment

Intuitively, minimax risk asks for the safest assignment rule when the true variance configuration is unknown, even if the pilot is misleading. Formally, the experimenter may choose any pilot-based rule $p \in \mathcal{D}_0$, and each rule is evaluated by its frequentist risk,

$\mathcal{R}(p, F) = \mathbb{E}_{P_F}\big[V\big(p(\omega), \sigma_1^2, \sigma_0^2\big)\big]$, the expected variance of the main-wave ATE estimator when the pilot is generated under $F \in \mathcal{F}$ and assignment follows rule p. The minimax-risk criterion ranks rules by their worst-case risk over the model class $\mathcal{F}$. The following proposition shows that complete balance is a minimax-risk rule.

Proposition 3.1 (Minimax risk). *The minimax-risk problem* $\inf_{p \in \mathcal{D}_0} \sup_{F \in \mathcal{F}} \mathcal{R}(p, F)$ *has value* 1*, and the constant balanced rule* $p_{\mathrm{mm}}(\omega) = 1/2$ *for all* $\omega \in \Omega$ *attains this value.*

Proofs of all main-text results are collected in Appendix A. The intuition behind the proof is that the least favorable distribution makes both arms as noisy as possible. With outcomes bounded in $[0, 1]$, this happens when both potential outcomes are Bernoulli with success probability one half, so $\sigma_1^2 = \sigma_0^2 = 1/4$. Under this distribution, the two arms are equally variable, so the best assignment is the balanced split $\pi = 1/2$, and even it yields a main-wave variance of 1. A rule that lets the pilot tilt the main wave is reacting to noise, protecting one arm only by taking observations from an equally noisy other. This single distribution therefore prevents every rule from having worst-case risk below 1. Balance, for its part, never exceeds this value, because its risk is $2\sigma_1^2 + 2\sigma_0^2$ and each variance is capped at $1/4$. The two bounds meet at 1, so balance is minimax.

This optimality gives balance a decision-theoretic foundation.[8] It is not a default adopted by convention but the symmetric hedge against the worst-case distribution, which can make either arm maximally noisy. The same result, however, reveals the limits of the criterion. Its least favorable distribution is the same whatever the pilot shows, so minimax risk gives the pilot no value, and balance remains minimax-optimal however strongly the data suggest that one arm is noisier. The reason is that minimax risk evaluates total main-wave variance, most of which no assignment can avoid. The worst case is then driven by this unavoidable component rather than by the part the pilot can improve, a well-understood conservativeness of minimax criteria (Berger, 1985). Analogous no-data results arise in statistical treatment choice, where maximin welfare rules retain the status quo whatever the data show, in point-identified problems (Manski, 2004) and under partial identification (Montiel Olea et al., 2026).

[8]With fixed potential outcomes and no pilot, Hull (2026) shows that, within a broad class of affine estimators, balanced complete randomization is strictly suboptimal under worst-case mean squared error. The adversary exploits negative correlation between assignments, a force averaged out by the i.i.d. sampling assumed in this paper.

3.2 Asymptotic Efficiency and the Feasible Neyman Allocation

The second canonical rule takes the opposite view from minimax risk. Instead of asking for the safest rule when the pilot may mislead, it treats the pilot as a source of estimates for the variance pair that determines the Neyman allocation. If those estimates are accurate, the natural choice is to plug them into the Neyman formula. The feasible Neyman allocation does so regardless, treating the sample variances as if they were known even when the pilot is small and noisy. With $\hat{\sigma}_d(\omega) = \sqrt{\hat{\sigma}_d^2(\omega)}$ the pilot standard deviation in arm d, the rule is $\hat{p}(\omega) = \frac{\hat{\sigma}_1(\omega)}{\hat{\sigma}_1(\omega)+\hat{\sigma}_0(\omega)}$ whenever the two are not both zero, with $\hat{p}(\omega) = 1/2$ when they are. It assigns the larger share of the main wave to the arm whose pilot outcomes are more variable, the tilt the Neyman allocation prescribes when the true standard deviations are known. The appeal of the rule is asymptotic. Outcomes are bounded, so the pilot sample variances are consistent, and whenever $\sigma_1 + \sigma_0 > 0$ the assignment $\hat{p}(\omega)$ converges to the Neyman allocation $\pi^*(\theta)$. Feasible Neyman thus inherits the precision of the infeasible Neyman allocation, which no design can improve on in large samples (Hahn, 1998; Hahn et al., 2011; Armstrong, 2026).

This large-sample appeal, however, masks a finite-sample fragility that comes from the ratio form of the rule. On the event that both pilot standard deviations are positive, substituting $\hat{p}(\omega)$ into V gives

$$V\left(\hat{p}(\omega), \sigma_1^2, \sigma_0^2\right) = \sigma_1^2 \left(1 + \frac{\hat{\sigma}_0(\omega)}{\hat{\sigma}_1(\omega)}\right) + \sigma_0^2 \left(1 + \frac{\hat{\sigma}_1(\omega)}{\hat{\sigma}_0(\omega)}\right).$$

Each term pairs a true variance with a ratio of pilot standard deviations that scales it. When the pilot estimates are close to the truth and both true standard deviations are positive, the two ratios approach σ_0/σ_1 and σ_1/σ_0, and the expression collapses to $V^* = (\sigma_1 + \sigma_0)^2$. When one pilot standard deviation is far too small, the ratio that divides by it grows without bound and turns a true variance of at most $1/4$ into an arbitrarily large contribution to V. The loss comes from trusting a small estimate, not from any large variance in the population. When a pilot standard deviation is exactly zero, $\hat{p}(\omega)$ lands on the boundary, the main wave assigns no observations to one arm, and that arm's outcome mean cannot be estimated.

The finite-sample consequence, established in the next proposition, is an unbounded worst-case risk. The problem is not only that feasible Neyman sometimes makes a noisy adjustment. The problem is that, with a finite pilot, nothing limits how far a noisy estimate can move the assignment, up to and including assigning no main-wave observations to an arm whose true variance is positive.

Proposition 3.2 (Boundary vulnerability of feasible Neyman)**.** *For every finite pilot size, the feasible Neyman rule has infinite worst-case risk under the boundary convention for V,*

$\sup_{F \in \mathcal{F}} \mathcal{R}(\hat{p}, F) = \infty$.

The same inverse-probability structure that gives the Neyman allocation its efficiency makes the plug-in rule vulnerable at the boundary.

This failure is not driven by an exotic distribution or by a knife-edge pilot sample. It can arise in the simplest binary-outcome experiment. The next remark makes this point explicit.

Remark 3.1 (Discrete outcomes and zero pilot variance)**.** For a Bernoulli outcome $Y(d) \sim \text{Bernoulli}(q_d)$ with $q_d \in (0,1)$, the population variance $\sigma_d^2 = q_d(1-q_d)$ is strictly positive, while the pilot sample variance is zero whenever every observed outcome in that arm is equal, an event with probability $q_d^{M_d} + (1-q_d)^{M_d}$. Under independent sampling across pilot arms, the probability that exactly one arm shows zero pilot variance is therefore strictly positive for every finite $M_0, M_1 \geq 2$ and every $q_0, q_1 \in (0,1)$. On this event, FNA places the entire main wave in one arm despite both population variances being strictly positive.

The difficulty is not confined to exact zeros. For $M_d \geq 3$, the smallest positive sample variance for a Bernoulli arm occurs when a single observation differs from the rest, giving $\hat{\sigma}_d^2 = 1/M_d$ under the unbiased estimator, with probability $M_d q_d (1-q_d)^{M_d - 1} + M_d(1-q_d) q_d^{M_d - 1}$. A draw of this kind makes an estimated standard deviation as small as $M_d^{-1/2}$, which keeps the realized variance finite but, by the mechanism above, can still make it very large.

Remark 3.1 suggests an immediate repair. If the problem is that feasible Neyman can assign an arm zero probability, one can force the rule to remain away from zero and one. This is the trimmed feasible Neyman rule, which is discussed in the next remark.

Remark 3.2 (Trimmed feasible Neyman)**.** A natural fix trims feasible Neyman away from the boundary, $\hat{p}_\tau(\omega) = \min\{\max\{\hat{p}(\omega), \tau\}, 1-\tau\}$ with $\tau \in (0, 1/2)$, which caps the realized variance at $1/(2\tau)$ and removes the infinite-risk pathology. The protection, however, is fixed before any data arrive rather than set by the strength of the pilot evidence, and no single τ works well. A large τ guards against extreme assignments but stays far from Neyman even after a pilot that has all but resolved the variances, while a small τ tracks feasible Neyman but leaves only the weak guarantee $1/(2\tau)$.

The conflict persists even asymptotically. At $(\sigma_1^2, \sigma_0^2) = (1/4, 0)$ the Neyman target $\pi^*(\theta) = 1$ exceeds the cap $1 - \tau$, so regret converges to $\frac{\tau}{4(1-\tau)} > 0$ rather than to zero, and shrinking τ with the pilot removes this gap only by surrendering the finite-sample protection. Trimming is therefore a blunt safeguard against extreme assignments, not a criterion for how far the realized pilot justifies moving away from balance.

The feasible Neyman allocation therefore captures the large-pilot ideal but not the finite-pilot problem. It uses the pilot through point estimates and gives no way to distinguish a reliable variance imbalance from a noisy one.

3.3 Minimax Regret and Finite-Pilot Adaptation

Minimax risk is too conservative for pilot adaptation because it ranks rules by total main-wave variance, most of which no design can avoid. The regret of Subsection 2.4 instead charges a rule only for the excess variance its assignment creates relative to the infeasible Neyman benchmark, isolating the component a pilot can reduce. Ranking rules by worst-case regret rather than worst-case risk is the standard, less conservative alternative and the more appropriate criterion here (Savage, 1951).

A statistical decision rule $p \in \mathcal{D}_0$ is fixed before the pilot, prescribing the main-wave assignment probability $p(\omega)$ at every realization ω. Its expected regret at F, $R(p, F) := \mathbb{E}_{P_F}[r(p(\omega), \theta(F))]$ is the frequentist risk of p under regret loss. The exact minimax-regret criterion selects the rule with the smallest worst-case expected regret. For fixed pilot arm sizes M_1, M_0, with their dependence suppressed elsewhere in the notation, the minimax-regret value is

$$R^{\mathrm{mmr}}_{M_1,M_0} := \inf_{\delta\in\mathcal{D}_0} \sup_{F\in\mathcal{F}} \mathbb{E}_{P_{F,M_1,M_0}}[r(\delta(\omega), \theta(F))] = \inf_{\delta\in\mathcal{D}_0} \sup_{F\in\mathcal{F}} R(\delta, F).$$

The three operations read in order. The experimenter chooses a rule δ, Nature responds with the least favorable distribution F, and the expectation averages the resulting regret over the pilots that F generates. The infimum ranges over all measurable pilot-to-assignment rules in $\mathcal{D}_0$, not only those that use the pilot through its two sample variances, which makes this the most demanding ex ante minimax-regret criterion available. When the infimum is attained, an exact minimax-regret rule is any $\delta^{\mathrm{mmr}}_{M_1,M_0} \in \mathcal{D}_0$ achieving the value.

Proposition 3.3 records the basic properties of the minimax-regret value and shows when finite-pilot information must be used by an exact minimax-regret rule.

Proposition 3.3 (Exact minimax regret)**.** *The minimax-regret value satisfies the following properties.*

(i) *For every pilot size (M_1, M_0), $R^{\mathrm{mmr}}_{M_1,M_0} \le 1/4$. In the no-pilot problem, the minimax-regret value is $1/4$, and the unique minimax-regret action is balanced assignment.*

(ii) *If $M_1, M_0 \ge 2$, then $R^{\mathrm{mmr}}_{M_1,M_0} < 1/4$. Consequently, every exact minimax-regret rule $\delta^{\mathrm{mmr}}_{M_1,M_0}$ satisfies $\Pr_{P_{F,M_1,M_0}}\big(\delta^{\mathrm{mmr}}_{M_1,M_0}(\omega) \neq \frac{1}{2}\big) > 0$ for some $F \in \mathcal{F}$.*

(iii) *If* $M_d \to \infty$ *for each* $d \in \{0, 1\}$, *then* $R^{\text{mmr}}_{M_1,M_0} \to 0$.

Part (i) combines a universal upper bound with a no-pilot characterization. Balanced assignment is feasible at every pilot size and has worst-case regret exactly $1/4$, attained when one arm has variance $1/4$ and the other zero. With no pilot, the feasible rules are the constant assignments, and any $\pi \neq 1/2$ undersamples one arm, which Nature punishes by placing all variance there. Only $\pi = 1/2$ equalizes the two opposing worst cases, all variance in treatment versus all in control.

Part (ii) shows that the slightest pilot overturns the no-pilot optimality of balance. With two observations per arm, the smallest size at which within-arm variation can appear, the minimax-regret value already drops below $1/4$, so balance is no longer optimal. The reason is that Nature now faces two competing forces. Driving the arm variances far apart raises the regret balance suffers, because balance is optimal only when the two are equal. But that same asymmetry is what the pilot detects, so it also tells the experimenter which arm to favor. A slight tilt toward the arm with the larger pilot variance exploits this signal. Where the variances are far apart, the pilot usually identifies the noisier arm and the tilt reduces regret. Where they are close and balance is nearly optimal, the tilt costs almost nothing. Such a rule therefore has worst-case regret strictly below $1/4$, so every exact minimax-regret rule must leave $1/2$ with positive probability under some distribution.

Part (iii) describes the large-pilot limit. As both pilot arms grow, the variance estimates become accurate enough that a stabilized plug-in rule, constructed in the proof, approaches the Neyman allocation and its worst-case expected regret vanishes. Since the minimax-regret value is no larger than the worst-case regret of this rule, $R^{\text{mmr}}_{M_1,M_0} \to 0$.

Despite the desirable properties recorded in Proposition 3.3, Remark 3.3 shows that the exact minimax-regret problem is not directly operational, because it does not reduce to a finite-dimensional optimization over the variance pair.

Remark 3.3 (Why exact minimax regret is not directly operational)**.** Nature's choice of F enters the problem twice, through the loss via the variance pair $\theta(F)$ and through the data via P_{F,M_1,M_0}, and the two channels are not linked by the variance pair, as discussed in Subsection 2.3. The inner supremum therefore cannot be taken over Θ, and the problem does not collapse to the static problem $\inf_p \sup_{\theta \in \Theta} r(p, \theta)$ that settles the no-pilot case.

The usual shortcuts do not restore finite-dimensional structure. The guess-and-verify approach of Stoye (2009), which identifies a minimax-regret rule as Bayes against a least favorable prior, does not help here, because the least favorable prior would, for the same reason, have to range over $\mathcal{F}$ rather than over variance pairs, leaving no finite-dimensional family to guess within. The Bernoulli reduction for bounded mean-payoff problems (Schlag, 2006; Stoye, 2009), which replaces $Y \in [0,1]$ by $B \mid Y \sim \text{Bernoulli}(Y)$, preserves the mean but strictly increases the variance unless Y is already binary, so it maps the state $\theta(F)$ to a different variance pair and defines a different variance-based game.[9]

Exact minimax regret is best read as a theoretical standard of comparison rather than an operational rule. It becomes finite-dimensional on the experimenter's side only after one restricts the outcome distribution, discretizes the pilot sample space, or restricts the class of rules, and any such restriction defines a different game from the exact nonparametric problem. The simplest such restriction imposes binary outcomes. With $Y(d) \sim \text{Bernoulli}(q_d)$, the state $(q_1, q_0) \in [0,1]^2$ pins down both the loss and the pilot distribution, and sufficiency and convexity imply that a minimax rule can be chosen to use the pilot only through the arm success counts studied in Online Appendix C.1.[10] Such a rule is a complete contingency plan fixed before the pilot, a finite vector holding one assignment for every count pair the pilot could produce. The difficulty is that Nature answers the whole plan with the least favorable point in a continuum of states. Ranking even one candidate plan therefore means maximizing its expected regret over the square, a nonconcave problem whose global solution no local search can guarantee, and the search for the best plan must optimize thousands of entries jointly around that inner maximization.[11]

Even so, the value $R^{\text{mmr}}_{M_1,M_0}$ remains informative, since it is the smallest worst-case expected regret attainable by any pilot-based assignment rule. The next proposition shows that even this ideal value cannot vanish faster than the inverse-square-root rate in the pilot arm sizes.

Proposition 3.4 (Minimax regret lower bound)**.** *Fix pilot arm sizes $M_1, M_0 \geq 2$. There exists a universal constant $c > 0$ such that $R^{\text{mmr}}_{M_1,M_0} \geq c\left(M_1^{-1/2} + M_0^{-1/2}\right)$.*

The bound reflects a limit on what a finite pilot can reveal about the variance pair. Even the best possible rule faces variance configurations that generate nearly indistinguishable pilot data but call for different Neyman allocations. Because the pilot cannot reliably tell such configurations apart, any rule must make similar recommendations in states where different assignments would be optimal, and must pay regret in at least one of them. Since the minimax-regret value already optimizes over all measurable pilot-to-assignment rules,

[9] The marginal variance is $\text{Var}(B) = \mathbb{E}[Y]\,(1-\mathbb{E}[Y]) = \text{Var}(Y) + \mathbb{E}[Y(1-Y)] \geq \text{Var}(Y)$.

[10] Averaging a rule over the equally likely pilot orderings with the same counts, and over the success–failure relabeling, weakly lowers worst-case expected regret because regret is convex in the assignment (Berger, 1985).

[11] Confining Nature to a finite grid of states turns the search for a minimax-regret rule into a finite convex program, but the problem stays hard. The whole rule must be optimized jointly against every grid state, and every pilot allocation defines a separate game. Even then, a rule computed this way is a candidate rather than an exact solution, since its worst state may lie off the grid.

this cost applies to every pilot-based design. No rule can therefore improve the worst-case expected regret beyond order $M_1^{-1/2} + M_0^{-1/2}$. Any tractable rule that attains this order matches the minimax-regret rate. In sum, exact minimax regret may be out of reach computationally, but matching its rate is enough for worst-case optimality up to constants.

4 Conditional Minimax Regret Rule

Section 3 leaves unanswered how far the main-wave assignment should move from balance toward the feasible Neyman allocation given the evidence in the realized pilot. A rule that answers this question must be computable from the realized pilot and must move from balance only as far as the evidence justifies.[12]

The regret of an assignment depends on the population distribution only through the marginal potential-outcome variances, so once the pilot is observed, the uncertainty that matters for the main-wave loss is which variance pair is true. The CMR procedure summarizes the pilot's design-relevant information by a finite-sample confidence set, the variance pairs the evidence has not ruled out, and selects the assignment with the smallest worst-case regret over that realized set. The worst case thus runs over the configurations consistent with the realized pilot, not over the full model class and not averaged over pilot realizations. In this sense, CMR is the minimax-regret principle applied to the problem the experimenter faces after the pilot, with the surviving variance pairs in the role of the state space.

Alongside the assignment, CMR reports a certificate in the spirit of Andrews and Chen (2025), the largest regret the chosen assignment can incur over the configurations the set still allows. Because the set contains the true variance pair with probability at least $1 - \alpha$, the realized regret exceeds the certificate with probability at most α. The certificate is large when the pilot leaves the variances uncertain and small when the pilot has nearly pinned them down.

4.1 The Conditional Minimax Regret Procedure

The CMR rule depends on the pilot only through a confidence set for the variance pair. Because each pilot unit is observed in only one arm, treatment observations inform σ_1^2 and control observations inform σ_0^2. The set therefore combines two separate inferences, an interval for σ_1^2 and an interval for σ_0^2, pairing every treatment variance in the first with every control variance in the second. This product structure makes the confidence set a rectangle.

[12] A Bayes rule reduces to the Neyman allocation evaluated at posterior mean variances. With a small pilot these posterior means sit close to their subjective prior values, so the prior tilts the assignment exactly when the data say least.

For each arm $d \in \{0,1\}$ and one-sided error level $b \in (0,1)$, let $\underline{\sigma}_d^2(b;\omega)$ and $\overline{\sigma}_d^2(b;\omega)$ denote a lower and an upper confidence bound for $\sigma_d^2(F)$, computed from the pilot and taking values in $[0,1/4]$. They must satisfy $\Pr_{P_F}(\underline{\sigma}_d^2(b;\omega) \le \sigma_d^2(F)) \ge 1-b$ and $\Pr_{P_F}(\sigma_d^2(F) \le \overline{\sigma}_d^2(b;\omega)) \ge 1-b$ for every $F \in \mathcal{F}$. Each bound is read as one reads a standard confidence bound, with the pilot as the sample and $\sigma_d^2(F)$ the unknown parameter. Subsection 4.3 constructs such bounds from finite-sample concentration inequalities. The general CMR procedure described below needs only the coverage property.

The rectangle contains the true variance pair $\theta(F)$ exactly when all four one-sided bounds hold at once. To keep the overall miss probability at most α, the construction divides that allowance equally among the four sides and forms each at one-sided error level $\alpha/4$. The realized confidence rectangle is

$$\widehat{\Theta}_\alpha(\omega) = \left[\underline{\sigma}_1^2(\alpha/4;\omega),\, \overline{\sigma}_1^2(\alpha/4;\omega)\right] \times \left[\underline{\sigma}_0^2(\alpha/4;\omega),\, \overline{\sigma}_0^2(\alpha/4;\omega)\right]. \tag{4.1}$$

By Bonferroni's inequality (Bonferroni, 1936), the probability of a miss is at most the sum of the four failure probabilities.[13] Each is at most $\alpha/4$, so the miss probability is at most α, and $\Pr_{P_F}\left(\theta(F) \in \widehat{\Theta}_\alpha(\omega)\right) \ge 1-\alpha$ for every $F \in \mathcal{F}$. The guarantee is frequentist in the usual sense, with the random rectangle covering the fixed pair $\theta(F)$ in at least a fraction $1-\alpha$ of repeated pilot samples.

Given the realized rectangle, CMR chooses the main-wave treatment assignment probability that minimizes worst-case regret over the variance pairs the pilot has not ruled out:

$$p_{\mathrm{CMR}}(\omega) \in \arg\min_{p \in (0,1)} \sup_{\theta \in \widehat{\Theta}_\alpha(\omega)} r(p,\theta). \tag{4.2}$$

The inner supremum is the largest regret p could incur over the surviving variance pairs, and the outer minimization selects the treatment assignment probability that makes this worst case smallest. The key simplification is that, after the pilot is translated into the rectangle, the decision problem no longer ranges over joint distributions of the potential outcomes.

The certificate reported with the assignment is the worst-case regret at the chosen CMR assignment,

$$U_{\mathrm{CMR}}(\omega) = \sup_{\theta \in \widehat{\Theta}_\alpha(\omega)} r(p_{\mathrm{CMR}}(\omega),\theta)\,. \tag{4.3}$$

Whenever the rectangle contains the true variance pair, the realized regret $r(p_{\mathrm{CMR}}(\omega),\theta(F))$ is therefore at most $U_{\mathrm{CMR}}(\omega)$.

[13] Because the pilot arm samples are independent, a multiplicative, Šidák-type split of the error budget across arms is also valid (Šidák, 1967). I use the additive split for simplicity.

4.2 Geometry and Closed-Form Solution

The CMR assignment and certificate can be computed in closed form once the realized rectangle is in hand. The dependence on the pilot realization ω is left implicit in what follows. By the second form of the regret identity (2.2), regret is the square of $(1-\pi)\sigma_1 - \pi\sigma_0$, which is affine in the two standard deviations, so for fixed π the worst case over the rectangle is attained at a corner. The relevant corners are the two off-diagonal ones, the corner where treatment is as variable as the rectangle allows and control as stable as it allows, and the corner where the roles are reversed. The proposition below uses this corner structure to solve the assignment problem in closed form.

Proposition 4.1 (Closed-form CMR rule)**.** *Let $\underline{\sigma}_d = \sqrt{\underline{\sigma}_d^2}$ and $\overline{\sigma}_d = \sqrt{\overline{\sigma}_d^2}$, and suppose $\overline{\sigma}_1 > 0$ and $\overline{\sigma}_0 > 0$. The CMR assignment problem (4.2) has a unique solution, given by*

$$p_{\mathrm{CMR}} = \frac{\overline{\sigma}_1 + \underline{\sigma}_1}{\overline{\sigma}_1 + \underline{\sigma}_1 + \overline{\sigma}_0 + \underline{\sigma}_0} \in (0,1). \tag{4.4}$$

The corresponding certificate is

$$U_{\mathrm{CMR}} = \frac{\left(\overline{\sigma}_1\overline{\sigma}_0 - \underline{\sigma}_1\underline{\sigma}_0\right)^2}{\left(\overline{\sigma}_1 + \underline{\sigma}_1\right)\left(\overline{\sigma}_0 + \underline{\sigma}_0\right)}. \tag{4.5}$$

To interpret the assignment formula (4.4), multiply its numerator and denominator by one half, which shows that CMR applies the Neyman formula to the midpoint of each standard-deviation confidence interval.[14] Treatment receives more than half the main wave exactly when its confidence-interval midpoint exceeds control's. For example, if $\sigma_1 \in [0.20, 0.40]$ and $\sigma_0 \in [0.10, 0.30]$, the midpoints are 0.30 and 0.20, and the treatment assignment probability prescribed by CMR is $0.30/(0.30 + 0.20) = 0.60$.

A further feature of the closed form is that its assignment is always interior, whereas feasible Neyman can collapse to the boundary. A zero pilot sample variance is not proof that the population variance is zero, and the finite-sample constructions of the next subsection keep the upper endpoint $\overline{\sigma}_d$ strictly positive in that case. That alone keeps both arms randomized, without the ad hoc trim feasible Neyman requires.

The certificate has the same corner structure. The two off-diagonal corners pull the assignment in opposite directions, CMR equalizes their regrets, and U_{CMR} is their common value. The certificate is therefore large when the rectangle still contains variance pairs that

[14] If neither endpoint of the symmetric Maurer–Pontil interval is projected onto $[0, 1/2]$, the midpoint equals the pilot standard deviation and CMR coincides with feasible Neyman. At $\alpha = 0.05$ this requires pilot arms of at least 142 observations.

point to substantially different Neyman allocations, and it vanishes exactly when the two corners imply the same allocation, which occurs when $\overline{\sigma}_1\overline{\sigma}_0 = \underline{\sigma}_1\underline{\sigma}_0$. It measures the spread in implied Neyman allocations, not the raw width of the rectangle.

The rule and its certificate reduce to familiar values at the two extremes. When the rectangle contains no information beyond the maintained bounds and each standard-deviation interval is $[0, 1/2]$, the two midpoints coincide at $1/4$, the rule returns balance, and the certificate equals $1/4$, the worst-case regret of the no-pilot problem. When the rectangle collapses to a single variance pair with both standard deviations positive, each confidence-interval midpoint equals the corresponding true standard deviation, the rule returns the Neyman allocation $\sigma_1/(\sigma_1 + \sigma_0)$, and the certificate is zero.

4.3 Constructing the Confidence Rectangle

The geometry of the problem also implies that the CMR assignment and certificate depend on the pilot only through the four rectangle endpoints. The baseline construction of $\widehat{\Theta}_\alpha(\omega)$ is distribution-free and finite-sample, using empirical Bernstein bounds for the arm-specific standard deviations and converting them into one-sided variance bounds valid uniformly over the bounded-outcome model. When outcomes are binary, the pilot distribution is discrete and exactly tractable, so the rectangle can be tightened by inverting the folded-binomial distribution of the pilot statistic, as developed in Online Appendix C.1. Either way, the assignment and certificate are computed from the resulting rectangle exactly as in Proposition 4.1, and any endpoints satisfying the one-sided coverage requirements of Subsection 4.1 deliver the same guarantees.

The relevant concentration inequality is the empirical Bernstein bound of Maurer and Pontil (2009) for the sample standard deviation of bounded random variables.[15] For each arm d and one-sided error level $b \in (0, 1)$, write $\eta_d(b) = \sqrt{2\log(1/b)/(M_d - 1)}$ for the finite-sample concentration radius on the standard-deviation scale. Then

$$\Pr_{P_F}(\sigma_d(F) \leq \hat{\sigma}_d(\omega) + \eta_d(b)) \geq 1 - b \quad \text{and} \quad \Pr_{P_F}(\sigma_d(F) \geq \hat{\sigma}_d(\omega) - \eta_d(b)) \geq 1 - b$$

uniformly over $F \in \mathcal{F}$. The radius $\eta_d(b)$ depends only on the pilot size and the error level, shrinking at rate $1/\sqrt{M_d}$ and growing only as $\sqrt{\log(1/b)}$ as the error level falls.

The variance bounds follow by squaring these standard-deviation statements and projecting onto the maintained variance interval $[0, 1/4]$. For each arm $d \in \{0, 1\}$ and error

[15] Sharper, first-order optimal empirical Bernstein intervals for the variance are available (Martinez-Taboada and Ramdas, 2025). The theoretical guarantees extend, at the cost of more complicated formulas.

level $b \in (0,1)$, the lower and upper variance bounds are

$$\underline{\sigma}_d^2(b;\omega) = \min\Big\{\tfrac{1}{4},\ \big(\hat{\sigma}_d(\omega) - \eta_d(b)\big)_+^2\Big\}, \qquad \overline{\sigma}_d^2(b;\omega) = \min\Big\{\tfrac{1}{4},\ \big(\hat{\sigma}_d(\omega) + \eta_d(b)\big)^2\Big\}. \tag{4.6}$$

Both endpoints are projected onto $[0, 1/4]$. Because the true variance lies in that interval, the projection preserves the one-sided coverage inequalities.[16]

The four endpoints, each at error level $\alpha/4$, assemble into the rectangle of (4.1), and the union bound delivers its $1-\alpha$ coverage, as the following lemma shows.

Lemma 4.1 (Coverage of the Maurer–Pontil rectangle)**.** *Fix $\alpha \in (0,1)$. For each arm $d \in \{0,1\}$ and each one-sided error level $b \in (0,1)$, $\Pr_{P_F}(\underline{\sigma}_d^2(b;\omega) \le \sigma_d^2(F)) \ge 1-b$ and $\Pr_{P_F}(\sigma_d^2(F) \le \overline{\sigma}_d^2(b;\omega)) \ge 1-b$ for every $F \in \mathcal{F}$. Consequently, the rectangle $\widehat{\Theta}_\alpha(\omega)$ satisfies $\Pr_{P_F}\Big(\theta(F) \in \widehat{\Theta}_\alpha(\omega)\Big) \ge 1-\alpha$ for every $F \in \mathcal{F}$.*

The proof is in Appendix E.

4.4 Statistical Properties

This subsection establishes three guarantees for the CMR assignment. First, in finite samples, the reported certificate bounds realized regret with probability at least $1-\alpha$. Second, the caution built into the rule disappears as the pilot grows. At interior variance pairs, CMR converges to the Neyman allocation, with assignment error shrinking at the inverse-square-root rate and realized regret and the certificate at the inverse-pilot rate. Third, I derive an upper bound on CMR's worst-case expected regret that matches the minimax-regret lower bound up to constants.

Finite-sample validity. The link from U_{CMR} to a finite-sample guarantee is coverage, the event that the realized rectangle contains the true variance pair. On that event, the true pair is among the configurations over which the certificate takes the worst case, so the realized regret of the CMR assignment cannot exceed U_{CMR}. Lemma 4.1 shows that coverage holds with probability at least $1-\alpha$. The next theorem combines these two observations and adds that the certificate never exceeds $1/4$.

Theorem 4.1 (Finite-sample regret certificate)**.** *Fix $\alpha \in (0,1)$. Then the following statements hold.*

(i) For every $F \in \mathcal{F}$, $\Pr_{P_F}\Big(r(p_{\mathrm{CMR}}(\omega), \theta(F)) \le U_{\mathrm{CMR}}(\omega)\Big) \ge 1-\alpha$.

[16] At a degenerate pilot $\hat{\sigma}_d^2(\omega) = 0$, $\overline{\sigma}_d^2(b;\omega) = \min\{1/4, \eta_d(b)^2\} > 0$, so the rectangle never treats it as proof of zero variance.

(ii) For every pilot realization ω, $U_{\text{CMR}}(\omega) \leq \frac{1}{4}$. Moreover, equality holds if and only if $(1/4, 0) \in \widehat{\Theta}_\alpha(\omega)$ and $(0, 1/4) \in \widehat{\Theta}_\alpha(\omega)$.

Part (i) shows that, for every distribution, the certificate $U_{\text{CMR}}(\omega)$ bounds the regret of $p_{\text{CMR}}(\omega)$ with probability at least $1 - \alpha$. The guarantee is finite-sample, holding at the realized pilot sizes rather than only in a large-pilot approximation.[17] The bound is not only valid but optimal in the sense that $U_{\text{CMR}}(\omega)$ is the smallest worst-case regret bound that any assignment can attain over the variance pairs still plausible after the pilot. In particular, for every alternative assignment $\widetilde{p}(\omega) \in (0, 1)$, $U_{\text{CMR}}(\omega) \leq \sup_{\theta \in \widehat{\Theta}_\alpha(\omega)} r(\widetilde{p}(\omega), \theta)$. This makes CMR a certified decision in the framework of Andrews and Chen (2025). In their terms, $U_{\text{CMR}}(\omega)$ is a P-certificate for the loss $r(p_{\text{CMR}}(\omega), \theta(F))$.

Part (ii) bounds the certificate itself. Balance is always among the candidates and its regret is at most $1/4$ at every configuration, so the certificate never exceeds $1/4$, the guarantee available with no pilot.[18] Equality holds exactly when the rectangle still contains the two adversarial corners $(1/4, 0)$ and $(0, 1/4)$, one placing all variability in treatment and the other all in control. Because these are opposite vertices of $[0, 1/4]^2$, only the full variance space contains both, so any pilot that rules out anything at all earns a certificate strictly below $1/4$.

CMR convergence rates. Now we ask how that guarantee evolves as the pilot becomes informative. In the interior of the variance space, where both arms have positive variance, the confidence rectangle collapses around the truth. The next theorem shows that, as both pilot arms grow large, the CMR rule converges to the Neyman allocation, and both realized regret and the certificate shrink to zero.

Theorem 4.2 (Neyman convergence and interior rates)**.** *Fix $F \in \mathcal{F}$, and suppose $\sigma_1(F) > 0$ and $\sigma_0(F) > 0$. Suppose also that $\alpha \in (0, 1)$ is fixed and that $M_0, M_1 \to \infty$, with M_0/M_1 bounded away from zero and infinity. Then the CMR rule satisfies the following statements.*

(i) $p_{\text{CMR}} \xrightarrow{p} \pi^*(\theta(F))$, $|p_{\text{CMR}} - \pi^*(\theta(F))| = O_p\Big(M_1^{-1/2} + M_0^{-1/2}\Big)$.

(ii) $r(p_{\text{CMR}}, \theta(F)) = O_p\big(M_1^{-1} + M_0^{-1}\big)$.

(iii) $U_{\text{CMR}} = O_p\big(M_1^{-1} + M_0^{-1}\big)$.

[17] Part (i) also implies quantile domination: for every F and every $t \leq 1 - \alpha$, the t-quantile of realized regret is bounded by the $(t + \alpha)$-quantile of the certificate.

[18] CMR can also be certified to be no worse than balance with probability at least $1 - \alpha$. It suffices to restrict the rule to assignments whose variance is below balance's at every pair in the rectangle.

All three rates come from a single source, the contraction of the confidence rectangle around the truth. The Maurer–Pontil endpoints learn each arm's standard deviation at the inverse-square-root rate in the pilot arm sizes, shrinking the rectangle to the variance pair at rate $M_1^{-1/2}+M_0^{-1/2}$. At an interior truth the rule varies smoothly with the rectangle and inherits this rate, which is statement (i). The other two rates are faster, and the reason is that, by the regret identity (2.2), regret is quadratic in the gap between the assignment and the Neyman allocation, with a coefficient that is finite whenever both arms have positive variance. An inverse-square-root error in the assignment then enters regret squared, at rate $M_1^{-1}+M_0^{-1}$, which is statement (ii). The certificate obeys the same bound through its closed form, which squares a product gap of inverse-square-root order over a denominator bounded away from zero in the interior, giving statement (iii).

The key implication of the theorem is that, at an interior truth, the finite-sample safety of the CMR rule does not slow the rate at which it converges to the Neyman allocation. In the corresponding large-main-wave limit, this convergence implies that the resulting assignment approaches the optimized Hahn-bound allocation, which Armstrong (2026) identifies as the first-order efficiency frontier across all designs that may adapt to covariates and past outcomes. The CMR rule approaches this frontier while certifying its own regret at every pilot size with probability at least $1-\alpha$.[19] First-order asymptotic efficiency and finite-sample safety hold together, with no trade-off between them.

The interior rates assume both arms have positive variance. The following remark shows that when one arm is degenerate, realized regret and the certificate converge at the slower rate $M_d^{-1/2}$ rather than M_d^{-1}.

Remark 4.1 (Boundary rates)**.** The interior rates rely on the Neyman allocation lying strictly between zero and one. There the denominator $\pi(1-\pi)$ in the regret identity (2.2) is bounded away from zero, so regret scales with the squared assignment error, and an error of order $M_d^{-1/2}$ produces regret of order M_d^{-1}. A degenerate arm moves the Neyman allocation to a corner and breaks this scaling, slowing realized regret and the certificate to the assignment-error rate itself.

With $\sigma_1(F)>0$ and $\sigma_0(F)=0$, the Neyman allocation is $\pi^*=1$. The upper confidence endpoint for σ_0 remains strictly positive at every finite pilot size, shrinking at rate $M_0^{-1/2}$ under the Maurer–Pontil construction, so CMR keeps a small control share $1-p_{\mathrm{CMR}}$ that vanishes at the same rate. At this truth the regret identity reduces to $r(\pi,\sigma_1^2(F),0) =$

[19] Feasible Neyman approaches the frontier at the same inverse-square-root rate, because the pilot estimates the standard deviations at that rate and the Neyman formula responds smoothly to small errors at an interior truth.

$\frac{\sigma_1^2(F)\,(1-\pi)}{\pi}$, linear in the leftover control share rather than quadratic. The hedge that keeps the rule off the corner therefore costs order $M_0^{-1/2}$ in realized regret, and the certificate is of the same order. The case $\sigma_1(F) = 0$ and $\sigma_0(F) > 0$ is symmetric.

Matching the Minimax-Regret Rate. The convergence rates just established are pointwise in F, describing how the rule and its realized regret behave at a fixed population as the pilot grows. The minimax-regret criterion of Subsection 3.3 instead judges a rule by its worst-case expected regret across all populations. As the next theorem shows, the CMR rule's worst-case expected regret shrinks at the inverse-square-root rate in the pilot arm sizes, and at the inverse-pilot rate once both arms are bounded away from degeneracy.

Theorem 4.3 (Uniform expected regret)**.** *Fix $\alpha \in (0, 1)$. There exists a constant $C_\alpha < \infty$, depending only on α, such that, for all pilot arm sizes $M_1, M_0 \geq 2$,*

$$\sup_{F\in\mathcal{F}} \mathbb{E}_{P_F}[r(p_{\mathrm{CMR}}(\omega), \theta(F))] \leq C_\alpha \left(M_1^{-1/2} + M_0^{-1/2}\right).$$

Moreover, for every $\kappa > 0$, there exists a constant $C_{\alpha,\kappa} < \infty$, depending only on α and κ, such that

$$\sup_{\substack{F\in\mathcal{F}:\\ \sigma_1(F)\wedge\sigma_0(F)\geq\kappa}} \mathbb{E}_{P_F}[r(p_{\mathrm{CMR}}(\omega), \theta(F))] \leq C_{\alpha,\kappa} \left(M_1^{-1} + M_0^{-1}\right).$$

Theorem 4.3 evaluates CMR by the ex ante criterion of Section 3, worst-case expected regret. Even though the assignment adapts to the realized pilot, its worst-case expected regret vanishes at the inverse-square-root rate in the pilot arm sizes. On any class of populations with both standard deviations bounded away from zero, the rate improves to the inverse-pilot rate. The slower uniform rate has the same source as the boundary rate of Remark 4.1.

The key reason why the upper bound vanishes is that the assignment and its regret do not blow up when the rectangle $\widehat{\Theta}_\alpha(\omega)$ misses the truth $\theta(F)$. Both depend continuously on how far the pilot's variance estimates fall from the truth, not on whether the rectangle contains it, so a small estimation error keeps regret small, covered or not. The bound is therefore governed by the size of this estimation error, which shrinks at the inverse-square-root rate.

Combining Theorem 4.3 with Proposition 3.4 yields $\sup_{F\in\mathcal{F}} \mathbb{E}_{P_F}[r(p_{\mathrm{CMR}}(\omega), \theta(F))] \leq \frac{C_\alpha}{c} R^{\mathrm{mmr}}_{M_1,M_0}$. The theorem gives the CMR upper bound, of order $M_1^{-1/2} + M_0^{-1/2}$, and the proposition the matching lower bound for the exact minimax-regret value. CMR's worst-case expected regret is therefore at most a constant, depending only on α, times $R^{\mathrm{mmr}}_{M_1,M_0}$, the smallest any pilot-based rule can attain. Strikingly, CMR achieves this although it optimizes only against the realized rectangle and never solves the ex ante minimax problem. Unlike an exact minimax-regret rule, CMR is closed form and reports a finite-sample certificate for the assignment it chooses.

5 Extensions: Multiple Treatments and Stratification

5.1 Multiple Treatments with a Shared Control

Setup. Many experiments compare several treatments against a common control, testing versions of a program, prices, information treatments, or implementation modes. The design question is how to divide the main wave across all the treatments and the shared control at once. The control arm is indexed by $k = 0$ and the treatments by $k = 1, \ldots, K$. A unit's potential outcomes $(Y(0), \ldots, Y(K))$, with $Y(k) \in [0,1]$, are drawn from a population distribution F. The mean and variance of $Y(k)$ are μ_k and σ_k^2, and $\theta(F) = (\sigma_0^2, \sigma_1^2, \ldots, \sigma_K^2) \in [0, 1/4]^{K+1}$ collects the arm variances. The main wave is set by the allocation vector $\pi = (\pi_0, \pi_1, \ldots, \pi_K)$, where $\pi_k > 0$ is the fraction allocated to arm k and $\sum_{k=0}^{K} \pi_k = 1$. The estimands of interest are the K treatment-control contrasts $\mathrm{ATE}_k = \mu_k - \mu_0$.

For a main wave of size N with N_k units on arm k, the difference-in-means estimator of contrast k has variance $\sigma_k^2/N_k + \sigma_0^2/N_0$. With $N_k = N\pi_k$, the sum of the marginal variances of the K contrasts equals $V(\pi,\theta)/N$, where $V(\pi,\theta) = \frac{K\sigma_0^2}{\pi_0} + \sum_{k=1}^{K} \frac{\sigma_k^2}{\pi_k}$. Every term is an arm variance divided by the share on that arm, so the allocation problem retains the inverse-share structure of the binary case, with the control variance now counted K times. This is an A-optimality criterion for the vector of treatment-control contrasts, since it minimizes the trace of their covariance matrix (Pukelsheim, 2006). It therefore targets average marginal variance rather than the power of a particular joint test. A pilot of fixed size must now be spread over $K + 1$ arms rather than two, so each arm variance is estimated with fewer observations.

Neyman Allocation. For known positive variances, minimizing $V(\pi,\theta)$ over the simplex $\Delta_K = \{\pi \in \mathbb{R}_+^{K+1} : \sum_{k=0}^{K} \pi_k = 1\}$ yields the multi-arm Neyman allocation:

$$\pi_0^*(\theta) = \frac{\sqrt{K}\,\sigma_0}{\sqrt{K}\,\sigma_0 + \sum_{j=1}^{K} \sigma_j}, \qquad \pi_k^*(\theta) = \frac{\sigma_k}{\sqrt{K}\,\sigma_0 + \sum_{j=1}^{K} \sigma_j}, \quad k = 1, \ldots, K. \tag{5.1}$$

The benchmark value is $V^*(\theta) = \inf_\pi V(\pi,\theta) = \left(\sqrt{K}\,\sigma_0 + \sum_{k=1}^{K} \sigma_k\right)^2$, attained at the Neyman allocation (5.1) when all arm variances are strictly positive, and regret at a feasible π

is $r(\pi, \theta) = V(\pi, \theta) - V^*(\theta)$.

CMR. The CMR construction extends by replacing the scalar assignment probability with the allocation vector. Each arm's variance receives a one-sided empirical Bernstein interval as in Section 4.3, with the error budget now split equally across the $2(K+1)$ one-sided endpoints, and stacking the intervals gives the hyperrectangle $\widehat{\Theta}_\alpha(\omega) = \prod_{k=0}^{K} \left[\underline{\sigma}_k^2(\omega), \overline{\sigma}_k^2(\omega)\right] \subseteq [0, 1/4]^{K+1}$, which covers the true variance vector with probability at least $1-\alpha$ for every F. The CMR allocation $p_{\mathrm{CMR}}(\omega) \in \Delta_K$ minimizes worst-case regret over this hyperrectangle, and the reported certificate $U_{\mathrm{CMR}}(\omega)$ is again the worst-case regret at the chosen allocation. On the event that the confidence rectangle contains the true variance vector, realized regret is bounded by the supremum defining $U_{\mathrm{CMR}}(\omega)$. The rectangle's coverage guarantee therefore yields the same $1-\alpha$ finite-sample guarantee as in Theorem 4.1(i). The next proposition traces the CMR allocation from the safe default, when the hyperrectangle is the full $[0, 1/4]^{K+1}$, to the Neyman allocation, when it collapses to a point as the pilot grows.

Proposition 5.1 (Multi-arm CMR between shared-control balance and Neyman)**.** *Consider the multi-arm shared-control design.*

(i) If the pilot is uninformative, with $\widehat{\Theta}_\alpha(\omega) = [0, 1/4]^{K+1}$, the CMR allocation is

$$p_{0,\mathrm{CMR}} = \frac{1}{1+\sqrt{K}}, \qquad p_{k,\mathrm{CMR}} = \frac{1}{\sqrt{K}\,(1+\sqrt{K})}, \quad k = 1, \ldots, K.$$

(ii) Fix $\sigma_k(F) > 0$ for all k. With α fixed and $M_{\min} = \min_{0 \le k \le K} M_k \to \infty$, every measurable selection $p_{\mathrm{CMR}}(\omega)$ from the CMR solution set satisfies $p_{\mathrm{CMR}} \xrightarrow{p} \pi^*(\theta(F))$, the multi-arm Neyman allocation (5.1).

Part (i) is the safe default. Unable to treat any arm as noisier than another, CMR plays the equal-variance Neyman allocation, placing $1/(1+\sqrt{K})$ of the main wave on the control and $1/(\sqrt{K}(1+\sqrt{K}))$ on each treatment. For example, with $K = 4$, the control receives a third of the main-wave sample and each treatment a sixth. Part (ii) is the multi-arm analogue of the interior convergence in Theorem 4.2, driven by the same contraction of the confidence set around the truth. The worst case also changes shape relative to the binary case. The binding configurations are vertices of the hyperrectangle, each making some subset of arms as noisy as the pilot allows and the rest as stable as it allows, and CMR hedges against the worst of these patterns rather than against one arm's variance at a time. Away from the uninformative rectangle, neither this allocation nor its stratified counterpart has a closed form. The vertex argument used in the proofs of Propositions 5.1 and 5.2 reduces both problems to finite convex programs, which are implemented in the replication code.

5.2 Stratified Experiments

Setup. Many target populations divide into strata, such as schools, regions, or gender, known before the main wave is designed. A stratified design involves two decisions, how many observations each stratum receives and how each stratum's observations are split between treatment and control. The pilot informs both decisions.

Strata are indexed by $x = 1, \dots, S$, a unit's pre-specified stratum is $X \in \{1, \dots, S\}$, and the target-population shares $s_x = \Pr_F(X = x) > 0$ sum to one. These shares are known and fixed before the main-wave design is chosen. Within stratum x, the potential outcome under arm $d \in \{0, 1\}$ has conditional mean $\mu_{dx} = \mathbb{E}_F[Y(d) \mid X = x]$ and variance $\sigma^2_{dx} = \mathrm{Var}_F(Y(d) \mid X = x)$. As before, the estimand is the population average treatment effect, $\mathrm{ATE}(F) = \sum_{x=1}^S s_x(\mu_{1x} - \mu_{0x})$, the average of the within-stratum effects $\mu_{1x} - \mu_{0x}$ weighted by the population shares.

The two decisions combine into a single design object, the treatment-by-stratum cell shares $\pi = \{\pi_{dx} : d \in \{0, 1\},\ x = 1, \dots, S\}$, where π_{dx} is the fraction of the main wave drawn from stratum x and assigned to arm d. These shares are nonnegative fractions of one fixed main wave, so they sum to one and π ranges over the simplex $\Delta_{2S-1} = \{\pi : \pi_{dx} \geq 0,\ \sum_{x=1}^S(\pi_{1x} + \pi_{0x}) = 1\}$. Two summaries of π recover the two decisions. The total share collected from stratum x, $\pi_{\cdot x} = \pi_{1x} + \pi_{0x}$, is the sampling margin, and the within-stratum treatment probability, $\pi_{1|x} = \pi_{1x}/(\pi_{1x} + \pi_{0x})$, is the assignment margin.[20]

The stratified difference-in-means estimator $\widehat{\mathrm{ATE}} = \sum_{x=1}^S s_x(\bar{Y}_{1x} - \bar{Y}_{0x})$ estimates each within-stratum effect by the cell contrast $\bar{Y}_{1x} - \bar{Y}_{0x}$ and aggregates these contrasts using the target-population shares s_x, where $\bar{Y}_{dx}$ is the sample mean in cell (d, x) (Imbens and Rubin, 2015). For positive cell shares, independent sampling across cells gives an estimator variance of $V(\pi, \theta)/N$, where

$$V(\pi, \theta) = \sum_{x=1}^S s_x^2 \left(\frac{\sigma^2_{1x}}{\pi_{1x}} + \frac{\sigma^2_{0x}}{\pi_{0x}} \right), \qquad \theta = \{\sigma^2_{dx} : d \in \{0, 1\},\ x = 1, \dots, S\}. \tag{5.2}$$

The square s_x^2 appears because the estimator multiplies each stratum contrast by s_x.

[20] The assignment margin is defined where $\pi_{\cdot x} > 0$. As in the baseline model, boundary allocations are allowed as formal actions, but any allocation that leaves a treatment-by-stratum mean entering the estimand unobserved carries infinite loss by convention.

Neyman Allocation. If the cell variances were known, the optimal stratified design would minimize $V(\pi, \theta)$ over the cell-share simplex Δ_{2S-1}. For positive variance configurations, the solution is

$$\pi^*_{dx}(\theta) = \frac{s_x \sigma_{dx}}{\sum_{z=1}^{S} s_z(\sigma_{1z} + \sigma_{0z})}, \qquad d \in \{0, 1\}, \quad x = 1, \ldots, S. \tag{5.3}$$

The optimal design assigns observations to cells in proportion to $s_x \sigma_{dx}$, how much the cell matters for the target ATE times how noisy its mean is to estimate. A large stratum with a nearly deterministic outcome needs few observations, and a noisy cell receives little weight if it represents a small part of the population.

Summing (5.3) over arms gives the Neyman sampling margin, $\pi^*_{\cdot x}(\theta) = \frac{s_x(\sigma_{1x}+\sigma_{0x})}{\sum_{z=1}^{S} s_z(\sigma_{1z}+\sigma_{0z})}$, which samples a stratum more than proportionally to its population share exactly when its total standard deviation $\sigma_{1x} + \sigma_{0x}$ exceeds the population-weighted average of total standard deviations across strata. The Neyman assignment margin is the standard two-arm Neyman allocation, $\pi^*_{1|x} = \sigma_{1x}/(\sigma_{1x} + \sigma_{0x})$. Substituting (5.3) into (5.2) gives $V^*(\theta) = \left[\sum_{x=1}^{S} s_x(\sigma_{1x} + \sigma_{0x})\right]^2$, the smallest variance any feasible design attains. Regret takes the same form as in the baseline model, $r(\pi, \theta) = V(\pi, \theta) - V^*(\theta)$.

CMR. From the $M_{dx} \geq 2$ pilot observations in cell (d, x), a one-sided empirical Bernstein interval of Maurer and Pontil (2009) is built for each cell variance σ^2_{dx}, with the error split equally across the $4S$ one-sided endpoints at level $\alpha/(4S)$. Stacking the $2S$ cell intervals gives the hyperrectangle $\widehat{\Theta}_\alpha(\omega) = \prod_{d\in\{0,1\}} \prod_{x=1}^{S} [\underline{\sigma}^2_{dx}(\omega), \overline{\sigma}^2_{dx}(\omega)] \subseteq [0, 1/4]^{2S}$, which covers the true variance vector with probability at least $1 - \alpha$ by the union bound. The CMR rule and its certificate take the same form as before, now over the cell-share simplex Δ_{2S-1}, with $p_{\mathrm{CMR}}(\omega)$ collecting the cell shares $p_{dx,\mathrm{CMR}}(\omega)$.

Proposition 5.2 (Stratified CMR between representative balance and Neyman)**.** *Consider the stratified design.*

(i) If the pilot is uninformative, with $\widehat{\Theta}_\alpha(\omega) = [0, 1/4]^{2S}$, the CMR allocation is $p_{1x,\mathrm{CMR}}(\omega) = p_{0x,\mathrm{CMR}}(\omega) = \frac{s_x}{2}$, $x = 1, \ldots, S$. Equivalently, $p_{\cdot x,\mathrm{CMR}}(\omega) = s_x$ and $p_{1|x,\mathrm{CMR}}(\omega) = 1/2$ for every stratum x.

(ii) Suppose $\sigma_{dx}(F) > 0$ for all $d \in \{0, 1\}$ and $x = 1, \ldots, S$. With α fixed and $M_{\min} = \min_{d,x} M_{dx} \to \infty$, every measurable selection $p_{\mathrm{CMR}}(\omega)$ from the CMR solution set satisfies $p_{\mathrm{CMR}}(\omega) \xrightarrow{p} \pi^*(\theta(F))$, the stratified Neyman allocation in (5.3).

With an uninformative pilot the rule has no reason to treat any stratum or arm as noisier than another, so the allocation is representative sampling with an even split within each stratum. As the pilot shrinks the confidence set, a stratum revealed to be noisier receives more than its population share, the split within each stratum moves toward its noisier arm, and in the limit the allocation becomes the stratified Neyman allocation (5.3).

6 Calibrated Simulations

This section stress-tests the allocation rules at practice-relevant pilot sizes, between 30 and 500 total observations. I calibrate the data-generating processes to public microdata from four published field experiments: the deworming program of Miguel and Kremer (2004) (MK), the resume audit of Bertrand and Mullainathan (2004) (BM), the experiment on incentives to learn HIV results of Thornton (2008), and the reference-letter experiment of Abel et al. (2020). The two-arm designs are built from MK, BM, and Thornton, the shared-control multi-arm design from Thornton's randomized incentive amounts, and the stratified design from the gender stratification in Abel et al.

Each data-generating process fixes the outcome distributions that the original experiment induced in its randomized arms. Binary outcomes are Bernoulli draws with the arm means estimated from the microdata, and continuous and count outcomes are drawn from the empirical distribution of the corresponding arm, rescaled to the unit interval using the observed range. The multi-arm design keeps Thornton's incentive levels and their shared control as five separate arms, and the stratified design applies the same construction within gender strata at their sample shares.

Table B.1 reports the calibrated arm means, variances, and implied Neyman allocations for each design. In the two-arm designs, the infeasible Neyman allocation assigns between 46 and 55 percent of the sample to treatment, so balance is never far from optimal. Balance's excess variance over the infeasible allocation is at most 0.84 percent, which also caps the improvement any pilot-based rule can deliver, the ceiling of Remark 2.1 materializing in calibrated data.

Each replication draws a pilot of total size $M \in \{30, 100, 250, 500\}$ from the calibrated distribution, assigned evenly across arms, and evenly within each stratum at representative stratum shares in the stratified design. The performance metric is the percentage efficiency loss, $100 \times [V(p,\theta) - V^*(\theta)]/V^*(\theta)$, where p is the main-wave assignment the rule chooses from the realized pilot. An entry of 1 means the assignment raises the estimator's variance by one percent, or equivalently that the experimenter would need one percent more main-wave observations, ten additional subjects per 1,000, to reach the same precision. Each entry

Table 1: Mean efficiency loss in the calibrated two-arm simulations

Study/outcome	Rule	$M = 30$	$M = 100$	$M = 250$	$M = 500$
MK: attendance	Balance		0.02		
	CMR	0.02	0.02	0.08	0.07
	FNA	3.33	0.71	0.27	0.12
MK: test score	Balance		0.07		
	CMR	0.07	0.07	0.05	0.05
	FNA	2.19	0.58	0.25	0.12
MK: mod.-heavy infection	Balance		0.45		
	CMR	0.45	0.07	0.08	0.08
	FNA	∞	0.20	0.06	0.03
Thornton: HIV result	Balance		0.59		
	CMR	0.59	0.18	0.13	0.12
	FNA	∞	0.36	0.14	0.06
Thornton: condom purchase	Balance		0.28		
	CMR	0.28	0.18	0.10	0.07
	FNA	∞	0.49	0.18	0.08
Thornton: condom count	Balance		0.18		
	CMR	0.18	0.22	0.27	0.22
	FNA	∞	5.93	2.14	1.09
BM: callback	Balance		0.84		
	CMR	0.84	0.84	0.47	0.44
	FNA	∞	∞	1.14	0.48

Notes: Entries are percent mean efficiency losses relative to the infeasible Neyman allocation: $100\,\mathbb{E}_\omega[V(\hat{p}_a(\omega), F) - V(\pi^*(F), F)]/V(\pi^*(F), F)$. Columns $M = 30, 100, 250, 500$ index the pilot size; the Balance row reports a single value because it does not use the pilot. Each entry uses 500 pilot replications per DGP and pilot size. Because every calibrated outcome distribution is discrete, FNA discards an arm with positive probability at every pilot size, so its unconditional mean efficiency loss is infinite in every cell of its rows. ∞ marks cells in which such a boundary draw occurred among the 500 replications, and finite FNA entries average the draws in which none occurred.

averages 500 independent pilot draws.[21]

6.1 Two-Arm Results

Table 1 compares the mean efficiency losses of balance, FNA, and CMR across the seven two-arm designs. Balance has a single entry per design because it does not use the pilot.

The FNA rows quantify the downside of trusting a small pilot. At $M = 30$, the loss is

[21] Because every calibrated outcome distribution is discrete, a pilot arm can take a single value with positive probability, in which case FNA assigns the entire main wave to one arm and its loss is infinite under the boundary convention. The unconditional mean FNA loss is therefore infinite in every design and at every pilot size. In the tables, ∞ marks a cell when at least one such draw occurs among the 500 replications, while a finite FNA entry averages the draws in a cell in which none occurs.

infinite in five of the seven designs and equals 3.33 and 2.19 percent in the other two, several times the 0.84 percent ceiling on what adaptation can gain, or roughly 33 and 22 additional subjects per 1,000. The boundary failure of Proposition 3.2 persists through $M = 100$ in the callback design.[22] CMR instead returns balance in every design at $M = 30$, at a cost of at most 8.4 subjects per 1,000, because the confidence rectangles remain uninformative. Its mean loss still equals balance's at $M = 100$ in the callback design, where the rectangle is rarely informative. Once the rectangle becomes informative, the gains arrive where the variances are unequal, with the loss falling to 0.07 percent in the MK infection design by $M = 100$ and to 0.47 percent in the callback design by $M = 250$. Where the variances are nearly equal there is nothing to find, and CMR's premium over balance is at most 0.09 percentage points, in the condom-count design in which FNA loses 5.93 percent at $M = 100$. By $M = 500$, FNA is modestly ahead of CMR in two of the seven designs. The cost of overreacting to a small pilot is measured in dozens of subjects. The cost of CMR's caution is measured in a handful.

In the reported draws, the confidence rectangle always covers the true variance pair and the certificate always bounds the realized regret, with coverage well above its nominal level, reflecting the conservativeness of the distribution-free Maurer–Pontil bounds. The certificate also tightens as the pilot grows, its median falling from the uninformative one quarter at $M = 30$ to between 0.05 and 0.14 at $M = 500$, a cut in the certified worst case of roughly one half to four fifths relative to the no-pilot guarantee; the full diagnostics are in the replication files.

6.2 Multi-Arm and Stratified Results

The extension designs raise both the value and the risk of adaptation, since the allocation is now a vector rather than a single treatment share. Table 2 reports the results for the two extensions of Section 5. In the table, the balance row is each design's no-information allocation, from Propositions 5.1 and 5.2.

The extensions differ from the two-arm designs in two ways. The infeasible benchmark adapts along more margins, so balance loses 1.76 percent in Panel A and 5.59 percent in Panel B, well above the 0.84 percent maximum gain in the two-arm designs, and the same pilot must now feed one variance estimate per arm or cell. FNA's mean loss is accordingly infinite in Panel A through $M = 250$, since a single zero-variance binary arm pushes it to the

[22] Common repairs of FNA do not resolve this. Table B.2 evaluates FNA with its assignment trimmed away from the boundary, together with the pre-test and regularized rules studied by Cai and Rafi (2024). Trimming removes the infinite entries, but at $M = 30$ all four alternatives have mean losses above the corresponding balance loss in every calibration.

Table 2: Mean efficiency loss in the calibrated extension simulations

Rule	$M = 30$	$M = 100$	$M = 250$	$M = 500$
Panel A: Thornton (2008): learned HIV result, multi-arm shared-control				
Balance		1.76		
CMR	1.76	1.76	1.48	0.34
CMR (Bernoulli)	1.82	1.21	0.75	0.36
FNA	∞	∞	∞	0.41
Panel B: Abel et al. (2020): job applications, gender-stratified				
Balance		5.59		
CMR	5.59	5.59	5.38	2.94
FNA	72.13	26.01	9.27	3.86

Notes: Entries are percent mean efficiency losses relative to the infeasible Neyman allocation: $100\,\mathbb{E}_\omega[V(\hat{p}_a(\omega), F) - V(\pi^*(F), F)]/V(\pi^*(F), F)$. Columns $M = 30, 100, 250, 500$ index total pilot size; each entry uses 500 pilot replications. Balance is the appropriate non-adaptive benchmark for each design: shared-control balance-equivalent in Panel A and half treatment within each stratum in Panel B. CMR is the Maurer–Pontil bounded-outcome rule; CMR (Bernoulli) uses Bernoulli variance confidence sets for the binary multi-arm outcome. ∞ indicates infinite unconditional mean efficiency loss, which occurs when FNA assigns zero probability to a positive-variance arm or treatment-by-stratum cell in at least one pilot draw.

boundary, and severe in Panel B, 72.13 percent at $M = 30$, or roughly 721 additional subjects per 1,000 against 56 under balance. The CMR loss equals the balance loss at $M = 30$ and $M = 100$ in both panels, since the error budget now splits across more arm- or cell-level bounds, each resting on fewer observations. Once the rectangle becomes informative, the rule adapts, to 1.48 percent at $M = 250$ and 0.34 at $M = 500$ in Panel A, and to 5.38 and then 2.94 in Panel B, about 29 additional subjects per 1,000. Panel A includes a second CMR row because the outcome is binary, and it shows that the rule's caution resides in the confidence set rather than in the minimax-regret logic. With the exact-inversion Bernoulli sets of Online Appendix C.1, CMR pays a small early premium, 1.82 percent against 1.76 at $M = 30$, but adapts earlier, reaching 1.21 and 0.75 percent at $M = 100$ and $M = 250$, against 1.76 and 1.48 for the Maurer–Pontil version.

The extensions sharpen the paper's main message. Exactly where adaptation has the most to buy and plug-in rules are most dangerous, CMR captures a growing share of the attainable gain as the pilot informs, never pays more than a tenth of a percentage point for its caution, and is the only adaptive rule in these tables that arrives with a finite-sample guarantee for the assignment it selects.

7 Conclusions

The central question in pilot-based design is how far the realized evidence should move the main-wave assignment. Balanced assignment does not respond to the pilot, while feasible Neyman follows its point estimates wherever they lead. Rather than acting on point estimates or forgoing adaptation altogether, the Conditional Minimax Regret rule acts on what the pilot has ruled out, staying at balance when the evidence is weak and approaching the Neyman allocation as the evidence accumulates. Before committing the main wave, the experimenter holds a certificate that, with high probability, bounds the precision lost to the realized design and never exceeds balance's no-pilot guarantee.

The same CMR logic applies to any design choice whose loss depends on parameters a small pilot estimates imprecisely. Two directions for future research seem most valuable. The first is to develop such rules for richer designs, including cluster-randomized, sequential, and imperfect-compliance experiments. The second is to treat the confidence set itself as part of the decision. CMR splits its error budget symmetrically and takes the coverage level as given, and neither choice is necessarily optimal. Sharper sets would let the rule earn the right to adapt sooner.

Much of the theory of adaptive experimentation justifies pilot-based designs by letting the pilot grow large. Real pilots are small, and at the sizes experimenters actually run, plug-in adaptation has modest upside and unbounded downside. The practical lesson of this paper is that the choice between ignoring the pilot and trusting it is a false one. A design can move exactly as far as the pilot's evidence warrants and no further, with a guarantee in hand before the main wave begins.

References

Martin Abel, Rulof Burger, and Patrizio Piraino. The value of reference letters: Experimental evidence from South Africa. *American Economic Journal: Applied Economics*, 12(3):40–71, 2020. doi: 10.1257/app.20180666.

Charalambos D. Aliprantis and Kim C. Border. *Infinite Dimensional Analysis: A Hitchhiker's Guide*. Springer, 3rd edition, 2006.

Isaiah Andrews and Jiafeng Chen. Certified decisions, 2025. URL https://arxiv.org/abs/2502.17830.

Timothy B. Armstrong. Asymptotic efficiency bounds for a class of experimental designs, 2026. URL https://arxiv.org/abs/2205.02726.

Yuehao Bai. Optimality of matched-pair designs in randomized controlled trials. *American Economic Review*, 112(12):3911–3940, 2022. doi: 10.1257/aer.20201856.

James O. Berger. *Statistical Decision Theory and Bayesian Analysis*. Springer Series in Statistics. Springer, New York, NY, 2 edition, 1985. ISBN 978-0-387-96098-2. doi: 10.1007/978-1-4757-4286-2.

Marianne Bertrand and Sendhil Mullainathan. Are Emily and Greg more employable than Lakisha and Jamal? A field experiment on labor market discrimination. *American Economic Review*, 94(4):991–1013, 2004. doi: 10.1257/0002828042002561.

Carlo Bonferroni. Teoria statistica delle classi e calcolo delle probabilità. *Pubblicazioni del R. Istituto Superiore di Scienze Economiche e Commerciali di Firenze*, 8:3–62, 1936.

Yong Cai and Ahnaf Rafi. On the performance of the Neyman allocation with small pilots. *Journal of Econometrics*, 242(1):105793, 2024. doi: 10.1016/j.jeconom.2024.105793.

Victor Chernozhukov, Sokbae Lee, Adam M Rosen, and Liyang Sun. Policy learning with confidence, 2025. URL https://arxiv.org/abs/2502.10653.

Max Cytrynbaum. Fine stratification of survey experiments, 2026. URL https://arxiv.org/abs/2111.08157.

Jessica Dai, Paula Gradu, and Christopher Harshaw. Clip-OGD: An experimental design for adaptive Neyman allocation in sequential experiments. In *Advances in Neural Information Processing Systems*, volume 36, 2023. arXiv:2305.17187.

Jinyong Hahn. On the role of the propensity score in efficient semiparametric estimation of average treatment effects. *Econometrica*, 66(2):315–331, 1998. doi: 10.2307/2998560.

Jinyong Hahn, Keisuke Hirano, and Dean Karlan. Adaptive experimental design using the propensity score. *Journal of Business & Economic Statistics*, 29(1):96–108, 2011. doi: 10.1198/jbes.2009.08161.

Samuel D. Higbee. Experimental design for policy choice, 2025. URL https://samuelhigbee.github.io/assets/papers/higbee_samuel_jmp.pdf. Working paper.

Yuchen Hu, Henry Zhu, Emma Brunskill, and Stefan Wager. Minimax-regret sample selection in randomized experiments. In *Proceedings of the 25th ACM Conference on Economics and Computation*, pages 1209–1235, 2024.

Peter Hull. Optimal experimental design and estimation when potential outcomes are bounded, 2026. URL https://arxiv.org/abs/2608.09812.

Guido W Imbens and Donald B Rubin. *Causal inference in statistics, social, and biomedical* *sciences*. Cambridge university press, 2015.

Cheaheon Lim and Yechan Park. Dynamically consistent statistical decisions, 2026. URL https://arxiv.org/abs/2607.10519.

Charles F Manski. Statistical treatment rules for heterogeneous populations. *Econometrica*, 72(4): 1221–1246, 2004. doi: 10.1111/j.1468-0262.2004.00530.x.

Charles F Manski. Econometrics for decision making: Building foundations sketched by Haavelmo and Wald. *Econometrica*, 89(6):2827–2853, 2021. doi: 10.3982/ecta17985.

Charles F Manski and Aleksey Tetenov. Sufficient trial size to inform clinical practice. *Proceedings of the National Academy of Sciences*, 113(38):10518–10523, 2016. doi: 10.1073/pnas.1612174113.

Diego Martinez-Taboada and Aaditya Ramdas. Sharp empirical Bernstein bounds for the variance of bounded random variables, 2025. URL https://arxiv.org/abs/2505.01987.

Andreas Maurer and Massimiliano Pontil. Empirical Bernstein bounds and sample variance penalization. In *Proceedings of the 22nd Annual Conference on Learning Theory*, 2009. URL https://arxiv.org/abs/0907.3740.

Colin McDiarmid. On the method of bounded differences. In Johannes Siemons, editor, *Surveys in Combinatorics, 1989*, volume 141 of *London Mathematical Society Lecture Note Series*, pages 148–188. Cambridge University Press, Cambridge, 1989.

Edward Miguel and Michael Kremer. Worms: Identifying impacts on education and health in the presence of treatment externalities. *Econometrica*, 72(1):159–217, 2004. doi: 10.1111/j.1468-0262.2004.00481.x.

José Luis Montiel Olea, Chen Qiu, and Jörg Stoye. Decision theory for treatment choice problems with partial identification. *Review of Economic Studies*, 2026. doi: 10.1093/restud/rdag015. Advance article.

Jerzy Neyman. On the two different aspects of the representative method: The method of stratified sampling and the method of purposive selection. *Journal of the Royal Statistical Society*, 97(4): 558–625, 1934. doi: 10.2307/2342192.

Friedrich Pukelsheim. *Optimal Design of Experiments*, volume 50 of *Classics in Applied Mathematics*. Society for Industrial and Applied Mathematics, Philadelphia, 2006. doi: 10.1137/1.9780898719109.

Leonard J Savage. The theory of statistical decision. *Journal of the American Statistical Association*, 46(253):55–67, 1951. doi: 10.1080/01621459.1951.10500768.

Karl H. Schlag. ELEVEN: Tests needed for a recommendation. Economics Working Paper ECO 2006/2, European University Institute, 2006.

Zbyněk Šidák. Rectangular confidence regions for the means of multivariate normal distributions. *Journal of the American Statistical Association*, 62(318):626–633, 1967.

Jörg Stoye. Minimax regret treatment choice with finite samples. *Journal of Econometrics*, 151(1): 70–81, 2009. doi: 10.1016/j.jeconom.2009.02.013.

Max Tabord-Meehan. Stratification trees for adaptive randomisation in randomised controlled trials. *Review of Economic Studies*, 90(5):2646–2673, 2023. doi: 10.1093/restud/rdac083.

Rebecca L. Thornton. The demand for, and impact of, learning HIV status. *American Economic Review*, 98(5):1829–1863, 2008. doi: 10.1257/aer.98.5.1829.

Alexandre B. Tsybakov. *Introduction to Nonparametric Estimation*. Springer, 2009.

Jinglong Zhao. Adaptive Neyman allocation. In *Proceedings of the 25th ACM Conference on Economics and Computation*, 2024. doi: 10.1145/3670865.3673535. arXiv:2309.08808.

A Proofs of Main-Text Results

Proof of Proposition 3.1. We establish an upper bound of 1 attained by the balanced rule and a matching lower bound for every decision rule.

For the constant rule $p_{\mathrm{mm}}(\omega) \equiv 1/2$ and any $F \in \mathcal{F}$, $\mathcal{R}(p_{\mathrm{mm}}, F) = 2\sigma_1^2(F) + 2\sigma_0^2(F)$. Since $\sigma_d^2(F) \leq 1/4$ for $d \in \{0,1\}$, it follows that $\mathcal{R}(p_{\mathrm{mm}}, F) \leq 1$ for all $F \in \mathcal{F}$. Hence $\sup_{F\in\mathcal{F}} \mathcal{R}(p_{\mathrm{mm}}, F) \leq 1$.

Now fix any $p \in \mathcal{D}_0$. Choose $F^\star \in \mathcal{F}$ such that $Y(d) \sim \mathrm{Bernoulli}(1/2)$ for $d \in \{0,1\}$. Then $\sigma_1^2(F^\star) = \sigma_0^2(F^\star) = 1/4$. If $p(\omega) \in \{0,1\}$ on a set with positive probability under $P_{F^\star}$, then by the boundary-loss convention, $\mathcal{R}(p, F^\star) = +\infty \geq 1$. Otherwise, $p(\omega) \in (0,1)$ $P_{F^\star}$-almost surely, and $\mathcal{R}(p, F^\star) = \frac{1}{4}\,\mathbb{E}_{P_{F^\star}}\left[\phi\big(p(\omega)\big)\right]$, $\phi(u) := \frac{1}{u} + \frac{1}{1-u}$, $\quad u \in (0,1)$. Since $\phi(u) = 1/[u(1-u)]$ and $u(1-u) \leq 1/4$ for all $u \in (0,1)$, with equality if and only if $u = 1/2$, we have $\phi(u) \geq 4$ for all $u \in (0,1)$. Hence $\mathcal{R}(p, F^\star) \geq \frac{1}{4} \cdot 4 = 1$. Because $F^\star \in \mathcal{F}$, this implies $\sup_{F\in\mathcal{F}} \mathcal{R}(p, F) \geq 1$ for every $p \in \mathcal{D}_0$.

Combining the two bounds yields $1 \leq \inf_{p\in\mathcal{D}_0} \sup_{F\in\mathcal{F}} \mathcal{R}(p, F) \leq \sup_{F\in\mathcal{F}} \mathcal{R}(p_{\mathrm{mm}}, F) \leq 1$. Therefore $\inf_{p\in\mathcal{D}_0} \sup_{F\in\mathcal{F}} \mathcal{R}(p, F) = 1$, and the constant balanced rule p_{mm} is minimax. □

Proof of Proposition 3.2. Fix any finite pilot sizes $M_0, M_1 \geq 2$. It is enough to exhibit one distribution $F \in \mathcal{F}$ for which $\mathcal{R}(\hat{p}, F) = +\infty$.

Let F be such that $Y(1)$ and $Y(0)$ are independent Bernoulli random variables with success probability 1/2. Then both marginal variances are strictly positive, with $\sigma_1^2 = \sigma_0^2 = 1/4$. Consider the event A that all treatment-arm pilot outcomes are equal and the control-arm pilot outcomes are not all equal. On this event, the treatment-arm sample variance is zero while the control-arm sample variance is strictly positive, so $\hat{\sigma}_1(\omega) = 0$ and $\hat{\sigma}_0(\omega) > 0$. Hence the untrimmed feasible Neyman rule sets $\hat{p}(\omega) = 0$.

The event A has strictly positive probability. Conditional on the fixed pilot assignment, the treatment-arm outcomes are i.i.d. Bernoulli(1/2), so the probability that they are all equal is $2(1/2)^{M_1} = 2^{1-M_1}$. Similarly, the probability that the control-arm outcomes are not all equal is $1 - 2^{1-M_0}$. Since these events depend on disjoint sets of independent units, $\Pr_{P_F}(A) = 2^{1-M_1}\big(1 - 2^{1-M_0}\big) > 0$.

Under the boundary convention for V, assigning zero probability to treatment yields infinite main-wave variance whenever $\sigma_1^2 > 0$. Therefore, on A, $V(\hat{p}(\omega), \sigma_1^2, \sigma_0^2) = V(0, 1/4, 1/4) = +\infty$. Thus the nonnegative extended random variable $V(\hat{p}(\omega), \sigma_1^2, \sigma_0^2)$ is infinite on an event with positive probability. Its expectation is therefore infinite, so $\mathcal{R}(\hat{p}, F) = +\infty$. Since this F belongs to $\mathcal{F}$, it follows that $\sup_{F\in\mathcal{F}} \mathcal{R}(\hat{p}, F) = +\infty$. □

Proof of Proposition 3.3. For any constant action $p \in (0,1)$, pointwise regret is given by

(2.2). Boundary actions $p \in \{0,1\}$ have infinite loss by convention.

Proof of (i). Consider the constant balanced rule $p(\omega) \equiv 1/2$. By (2.2), $R(1/2, F) = (\sigma_1 - \sigma_0)^2$. Since outcomes lie in $[0,1]$, each marginal variance is at most $1/4$, so $0 \le \sigma_1, \sigma_0 \le 1/2$. Therefore $R(1/2, F) \le \frac{1}{4}$ for every $F \in \mathcal{F}$. Since the balanced rule is feasible for every pilot size, $R^{\mathrm{mmr}} \le \sup_{F \in \mathcal{F}} R(1/2, F) \le \frac{1}{4}$.

Now consider the no-pilot problem. The sample space is a singleton, so every rule is a constant action $p \in [0,1]$. Boundary actions have infinite loss, so it is enough to consider $p \in (0,1)$. For fixed p, define $r_p(\sigma_1, \sigma_0) := \frac{\sigma_1^2}{p} + \frac{\sigma_0^2}{1-p} - (\sigma_1 + \sigma_0)^2 = \frac{\big((1-p)\sigma_1 - p\sigma_0\big)^2}{p(1-p)}$. The function r_p is convex in (σ_1, σ_0), so its maximum over $[0, 1/2]^2$ is attained at a corner. These corners are attainable by distributions in $\mathcal{F}$. Evaluating,

$$r_p(0,0) = 0, \quad r_p(1/2, 0) = \frac{1-p}{4p}, \quad r_p(0, 1/2) = \frac{p}{4(1-p)}, \quad r_p(1/2, 1/2) = \frac{(1-2p)^2}{4p(1-p)}.$$

If $p \le 1/2$, then $r_p(1/2,1/2) \le r_p(1/2, 0)$. If $p \ge 1/2$, then $r_p(1/2,1/2) \le r_p(0,1/2)$. Hence $\sup_{F \in \mathcal{F}} R(p, F) = \max\left\{\frac{1-p}{4p}, \frac{p}{4(1-p)}\right\}$. The first term is strictly decreasing in p, and the second is strictly increasing in p. Their maximum is uniquely minimized at $p = 1/2$, where its value is $1/4$. Therefore, in the no-pilot problem, $R^{\mathrm{mmr}} = 1/4$, and the unique minimax-regret action is $p = 1/2$.

Proof of (ii). It is enough to construct one feasible rule whose worst-case regret is strictly below the worst-case regret of balanced assignment. Let $B = \max_{d \in \{0,1\}} \frac{M_d}{4(M_d - 1)}$. Fix $\eta \in (0, (4B)^{-1})$, to be chosen below, and define $p_\eta(\omega) = \frac{1}{2} + \eta\big(\hat{\sigma}_1^2(\omega) - \hat{\sigma}_0^2(\omega)\big)$. Since the usual unbiased within-arm sample variance satisfies $0 \le \hat{\sigma}_d^2(\omega) \le \frac{M_d}{4(M_d-1)}$, this rule takes values in $[1/2 - \eta B, 1/2 + \eta B] \subset (0,1)$. Hence $p_\eta \in \mathcal{D}$.

Fix $F \in \mathcal{F}$. For any realized pilot sample, $p_\eta(\omega) = \frac{1}{2} + \eta\big(\hat{\sigma}_1^2(\omega) - \hat{\sigma}_0^2(\omega)\big)$. Using $r(p_\eta(\omega), \theta) = \sigma_1^2 / p_\eta(\omega) + \sigma_0^2/(1 - p_\eta(\omega)) - (\sigma_1 + \sigma_0)^2$, we have

$$\begin{aligned} r(p_\eta(\omega), \theta(F)) - r(1/2, \theta(F)) &= 4\eta\big(\hat{\sigma}_1^2 - \hat{\sigma}_0^2\big)\big(\sigma_0^2 - \sigma_1^2\big) \\ &\quad + 8\eta^2\big(\hat{\sigma}_1^2 - \hat{\sigma}_0^2\big)^2 \left\{\frac{\sigma_1^2}{1 + 2\eta(\hat{\sigma}_1^2 - \hat{\sigma}_0^2)} + \frac{\sigma_0^2}{1 - 2\eta(\hat{\sigma}_1^2 - \hat{\sigma}_0^2)}\right\}. \end{aligned}$$

Because $\eta < (4B)^{-1}$ and $|\hat{\sigma}_1^2 - \hat{\sigma}_0^2| \le B$, both denominators in the second line are at least $1/2$. Since $\sigma_1^2, \sigma_0^2 \le 1/4$, the second line is bounded above by $8B^2\eta^2$. Therefore

$$R(p_\eta, F) \le (\sigma_1 - \sigma_0)^2 + 4\eta(\sigma_0^2 - \sigma_1^2)\mathbb{E}_{P_F}\big[\hat{\sigma}_1^2 - \hat{\sigma}_0^2\big] + 8B^2\eta^2. \tag{A.1}$$

We now show that η can be chosen so that the right-hand side is uniformly below $1/4$. Because $\hat{\sigma}_d^2$ is the usual unbiased within-arm sample variance, $\mathbb{E}_{P_F}[\hat{\sigma}_d^2] = \sigma_d^2$. Choose constants $c > 0$ and $c' > 0$, depending only on (M_1, M_0), such that $c' < c/2$, $1/4 - 2c' > 0$, and $\sigma_d^2 \geq \frac{1}{4} - c' \Longrightarrow \mathbb{E}_{P_F}[\hat{\sigma}_d^2] \geq c$.

Choose $\gamma > 0$ small enough that $(\sigma_1 - \sigma_0)^2 > \frac{1}{4} - \gamma$ implies either $\sigma_1^2 \geq 1/4 - c'$ and $\sigma_0^2 \leq c'$, or $\sigma_0^2 \geq 1/4 - c'$ and $\sigma_1^2 \leq c'$. This is possible because $\sigma_1, \sigma_0 \in [0, 1/2]$, and the equality $(\sigma_1 - \sigma_0)^2 = 1/4$ can occur only at the two corners $(\sigma_1, \sigma_0) = (1/2, 0)$ and $(0, 1/2)$.

First suppose $(\sigma_1 - \sigma_0)^2 \leq 1/4 - \gamma$. Since $|\sigma_0^2 - \sigma_1^2| \leq 1/4$ and $|\mathbb{E}_{P_F}[\hat{\sigma}_1^2 - \hat{\sigma}_0^2]| \leq 1/4$, equation (A.1) gives $R(p_\eta, F) \leq \frac{1}{4} - \gamma + \frac{\eta}{4} + 8B^2\eta^2$. For all sufficiently small η, this is strictly below $1/4$.

Now suppose $(\sigma_1 - \sigma_0)^2 > 1/4 - \gamma$. If $\sigma_1^2 \geq 1/4 - c'$ and $\sigma_0^2 \leq c'$, then $\mathbb{E}_{P_F}[\hat{\sigma}_1^2 - \hat{\sigma}_0^2] \geq c - c' \geq c/2$ and $\sigma_1^2 - \sigma_0^2 \geq 1/4 - 2c'$. Thus (A.1) implies $R(p_\eta, F) \leq \frac{1}{4} - 2\eta c\Big(\frac{1}{4} - 2c'\Big) + 8B^2\eta^2$. For all sufficiently small $\eta > 0$, this is strictly below $1/4$. The case $\sigma_0^2 \geq 1/4 - c'$ and $\sigma_1^2 \leq c'$ is identical, with the roles of the two arms reversed.

Combining the two cases, there exists $\eta > 0$ such that $\sup_{F \in \mathcal{F}} R(p_\eta, F) < \frac{1}{4}$. Since $p_\eta \in \mathcal{D} \subseteq \mathcal{D}_0$ and R^{mmr} is the infimum over $\mathcal{D}_0$, $R^{\mathrm{mmr}} \leq \sup_{F \in \mathcal{F}} R(p_\eta, F) < \frac{1}{4}$.

It remains to prove the final claim. Let δ be any exact minimax-regret rule. If, for every $F \in \mathcal{F}$, $\Pr_{P_{F,M_1,M_0}}(\delta(\omega) \neq 1/2) = 0$, then δ has the same expected regret as balanced assignment under every F. Hence $\sup_{F \in \mathcal{F}} R(\delta, F) = \sup_{F \in \mathcal{F}} R(1/2, F) = \frac{1}{4}$, contradicting $R^{\mathrm{mmr}} < 1/4$. Therefore every exact minimax-regret rule must satisfy $\Pr_{P_{F,M_1,M_0}}\big(\delta(\omega) \neq \frac{1}{2}\big) > 0$ for some $F \in \mathcal{F}$.

Proof of (iii). We construct a feasible rule whose worst-case regret converges to zero. For each arm $d \in \{0, 1\}$, write $\hat{v}_d = \hat{\sigma}_d^2$ and $v_d = \sigma_d^2$. Since $\hat{v}_d$ is the usual unbiased within-arm sample variance and $Y(d) \in [0, 1]$, the standard mean-square formula for the sample variance, together with bounded fourth moments, gives a universal constant $C_0 < \infty$ such that $\sup_{F \in \mathcal{F}} \mathbb{E}_{P_F}[(\hat{v}_d - v_d)^2] \leq C_0 M_d^{-1}$. Using $(\sqrt{x} - \sqrt{y})^2 \leq |x - y|$ for $x, y \geq 0$ and Cauchy–Schwarz, there is a universal constant $C_1 < \infty$ such that $\sup_{F \in \mathcal{F}} \mathbb{E}_{P_F}[(\hat{\sigma}_d - \sigma_d)^2] \leq C_1 M_d^{-1/2}$. Let $M = (M_1, M_0)$ and $\rho_M := C_1 \max\{M_0^{-1/2}, M_1^{-1/2}\}$, so $\rho_M \to 0$ whenever $M_0, M_1 \to \infty$.

Set $\kappa_M = \rho_M^{1/4}$ and $\delta_M = \rho_M^{1/4}$. Define the stabilized plug-in rule

$$p_M(\omega) := \begin{cases} 1/2, & \hat{\sigma}_0 + \hat{\sigma}_1 \leq \kappa_M, \\ \Pi_{[\delta_M,\, 1-\delta_M]}\Big(\frac{\hat{\sigma}_1}{\hat{\sigma}_0 + \hat{\sigma}_1}\Big), & \hat{\sigma}_0 + \hat{\sigma}_1 > \kappa_M. \end{cases}$$

For all sufficiently large pilot sizes, $p_M(\omega) \in [\delta_M, 1 - \delta_M]$, and hence $p_M(\omega)(1 - p_M(\omega)) \geq$

$\delta_M/2$.

Fix $F \in \mathcal{F}$, let $s = \sigma_0 + \sigma_1$, and note that if $s = 0$, regret is zero. Suppose $s > 0$, so that $\pi^*(\theta(F)) = \sigma_1/s$. Then

$$R(p_M, F) = s^2 \mathbb{E}_{P_F}\left[\frac{(p_M - \pi^*(\theta(F)))^2}{p_M(1-p_M)}\right] \le \frac{2s^2}{\delta_M}\mathbb{E}_{P_F}[(p_M - \pi^*(\theta(F)))^2].$$

If $s \le 2\kappa_M$, then $(p_M - \pi^*(\theta(F)))^2 \le 1$, so $R(p_M, F) \le \frac{8\kappa_M^2}{\delta_M}$.

Now suppose $s > 2\kappa_M$. Let $\Delta_d = \hat{\sigma}_d - \sigma_d$, $\hat{s} = \hat{\sigma}_0 + \hat{\sigma}_1$, and $E_M = \{\hat{s} \le \kappa_M\}$. On E_M, $|\Delta_0| + |\Delta_1| \ge s - \hat{s} > \kappa_M$. Hence, by Markov's inequality, $\Pr_{P_F}(E_M) \le \frac{\mathbb{E}_{P_F}[(|\Delta_0|+|\Delta_1|)^2]}{\kappa_M^2} \le \frac{4\rho_M}{\kappa_M^2}$. Since the regret integrand is bounded above by $2/\delta_M$, the contribution from E_M is at most $8\rho_M/(\delta_M\kappa_M^2)$.

On E_M^c, define $\tilde{p}_M = \hat{\sigma}_1/\hat{s}$. A direct calculation gives $|\tilde{p}_M - \pi^*(\theta(F))| = \left|\frac{\hat{\sigma}_1}{\hat{s}} - \frac{\sigma_1}{s}\right| \le \frac{|\Delta_0|+|\Delta_1|}{\kappa_M}$. Since p_M is obtained by clipping $\tilde{p}_M$ to $[\delta_M, 1-\delta_M]$, $|p_M - \pi^*(\theta(F))| \le \frac{|\Delta_0|+|\Delta_1|}{\kappa_M} + \delta_M$. Therefore, using $(x+y)^2 \le 2x^2 + 2y^2$, the contribution from E_M^c to $R(p_M, F)$ is at most $\frac{16\rho_M}{\delta_M\kappa_M^2} + 4\delta_M$. Combining the two cases, $R(p_M, F) \le \max\left\{\frac{8\kappa_M^2}{\delta_M}, \frac{24\rho_M}{\delta_M\kappa_M^2} + 4\delta_M\right\}$. With $\kappa_M = \delta_M = \rho_M^{1/4}$, this implies $\sup_{F\in\mathcal{F}} R(p_M, F) \le C\rho_M^{1/4}$ for a universal constant $C < \infty$. Since $\rho_M \to 0$, we have $\sup_{F\in\mathcal{F}} R(p_M, F) \to 0$. Finally, because $p_M \in \mathcal{D} \subseteq \mathcal{D}_0$ and R^{mmr} is the infimum over $\mathcal{D}_0$, $0 \le R^{\text{mmr}} \le \sup_{F\in\mathcal{F}} R(p_M, F) \to 0$. Therefore $R^{\text{mmr}} \to 0$. $\square$

Proof of Proposition 3.4. The proof works by embedding two simple two-state problems inside the full bounded-outcome model, following Le Cam's two-point method (Tsybakov, 2009, Section 2.2). Since the adversary in the definition of $R^{\text{mmr}}_{M_1,M_0}$ can choose any distribution in $\mathcal{F}$, the minimax value over $\mathcal{F}$ is at least as large as the minimax value over any submodel.

Fix an arbitrary pilot-to-assignment rule δ. We first construct a hard subproblem in which the control variance is difficult to distinguish from zero. In both states, let $Y(1)$ be Bernoulli with success probability 1/2. The two states differ only in the control arm. Under the first state, let $Y(0) = 0$ almost surely. Under the second state, let $Y(0)$ be Bernoulli with success probability $1/M_0$. Both states belong to $\mathcal{F}$.

In the first state, the variance pair is $(1/4, 0)$. Thus the infeasible Neyman assignment is the boundary allocation that puts all main-wave mass on treatment. If a realized rule chooses treatment assignment probability $p \in (0,1)$, its regret in this state is $r(p, (1/4, 0)) = \frac{1-p}{4p} \ge \frac{1-p}{4}$. In the second state, the treatment standard deviation is 1/2, while the control standard deviation is $\sqrt{M_0^{-1}(1 - M_0^{-1})}$. For any realized treatment assignment probability

p, the regret in this second state is

$$r\big(p, \big(1/4,\ M_0^{-1}(1-M_0^{-1})\big)\big) = \frac{1}{4p} + \frac{M_0^{-1}(1-M_0^{-1})}{1-p} - \left(\frac{1}{2} + \sqrt{M_0^{-1}(1-M_0^{-1})}\right)^2.$$

Now consider the event that all M_0 control-pilot observations are equal to zero. Under the first state this event has probability one. Under the second state it has probability $(1-M_0^{-1})^{M_0}$, which is at least $1/4$ for every $M_0 \geq 2$. Conditional on this event, the observed control pilot is the same in the two states, and the treatment pilot has the same distribution in the two states. Hence, on this event, the rule faces the same pilot evidence in both states.

We next record the deterministic trade-off faced by any assignment on this common pilot event. Let $p \in (0,1)$ be any treatment assignment probability, and write the control share as $1-p$. If $1-p \geq \sqrt{M_0^{-1}(1-M_0^{-1})}/2$, then the regret in the first state is at least $\sqrt{M_0^{-1}(1-M_0^{-1})}/8$. If instead $1-p < \sqrt{M_0^{-1}(1-M_0^{-1})}/2$, then the regret in the second state is at least

$$\frac{1}{4} + 2\sqrt{M_0^{-1}(1-M_0^{-1})} - \left(\frac{1}{2} + \sqrt{M_0^{-1}(1-M_0^{-1})}\right)^2 = \sqrt{M_0^{-1}(1-M_0^{-1})} - M_0^{-1}(1-M_0^{-1}).$$

The last expression is at least $\sqrt{M_0^{-1}(1-M_0^{-1})}/2$, because a Bernoulli standard deviation is at most $1/2$. Therefore, on the common all-zero control-pilot event, every realized assignment has regret at least $\sqrt{M_0^{-1}(1-M_0^{-1})}/8$ in one of the two states.

Averaging over the common treatment-pilot distribution gives the same trade-off in expectation on the all-zero control-pilot event. Therefore, for the rule δ, either its expected regret in the first state is at least $\sqrt{M_0^{-1}(1-M_0^{-1})}/16$, or its conditional expected regret in the second state, given the all-zero control-pilot event, is at least $\sqrt{M_0^{-1}(1-M_0^{-1})}/16$. In the latter case, since the all-zero control-pilot event has probability at least $1/4$ under the second state, the unconditional expected regret in the second state is at least $\sqrt{M_0^{-1}(1-M_0^{-1})}/64$. Thus, for every rule δ, $\sup_{F\in\mathcal{F}} \mathbb{E}_{P_F}[r(\delta(\omega), \theta(F))] \geq \frac{1}{64}\sqrt{M_0^{-1}(1-M_0^{-1})}$. Since $M_0 \geq 2$, the right-hand side is at least $(64\sqrt{2})^{-1}M_0^{-1/2}$.

The same argument with the treatment and control labels reversed gives, for every rule δ, $\sup_{F\in\mathcal{F}} \mathbb{E}_{P_F}[r(\delta(\omega), \theta(F))] \geq \frac{1}{64\sqrt{2}}M_1^{-1/2}$. Combining the two arm-specific lower bounds gives $\sup_{F\in\mathcal{F}} \mathbb{E}_{P_F}[r(\delta(\omega), \theta(F))] \geq (64\sqrt{2})^{-1}\max\{M_1^{-1/2}, M_0^{-1/2}\}$. Since $\max\{M_1^{-1/2}, M_0^{-1/2}\} \geq \frac{1}{2}\left(M_1^{-1/2} + M_0^{-1/2}\right)$, we obtain

$$\sup_{F\in\mathcal{F}} \mathbb{E}_{P_F}[r(\delta(\omega), \theta(F))] \geq \frac{1}{128\sqrt{2}}\left(M_1^{-1/2} + M_0^{-1/2}\right).$$

The rule δ was arbitrary. Taking the infimum over all measurable pilot-to-assignment rules proves the result, with $c = 1/(128\sqrt{2})$ under the $[0,1]$ normalization. $\square$

Proof of Proposition 4.1. Fix $\pi \in (0,1)$. By the square representation in (2.2), maximizing regret over the rectangle is equivalent to maximizing $\left|(1-\pi)\sigma_1 - \pi\sigma_0\right|$ over the standard-deviation rectangle $[\underline{\sigma}_1, \overline{\sigma}_1] \times [\underline{\sigma}_0, \overline{\sigma}_0]$. The affine function $(1-\pi)\sigma_1 - \pi\sigma_0$ is increasing in σ_1 and decreasing in σ_0. Its maximum over the rectangle is therefore attained at $(\overline{\sigma}_1, \underline{\sigma}_0)$, and its minimum is attained at $(\underline{\sigma}_1, \overline{\sigma}_0)$. Since the maximum of the square of a function over a set is attained either where the function is largest or where it is smallest, the supremum of regret over the rectangle is the maximum of the regrets at these two off-diagonal corners. This proves the two-corner reduction.

It remains to minimize this maximum over $\pi \in (0,1)$. Let p_{CMR} denote the value in the displayed formula of the proposition. Since $\overline{\sigma}_1^2 > 0$ and $\overline{\sigma}_0^2 > 0$, we have $p_{\mathrm{CMR}} \in (0,1)$. If the rectangle collapses to a single variance pair, then $\overline{\sigma}_d = \underline{\sigma}_d$ for $d \in \{0,1\}$. The objective is the regret from that single pair. Its derivative is $-\overline{\sigma}_1^2/\pi^2 + \overline{\sigma}_0^2/(1-\pi)^2$, which vanishes only at $\overline{\sigma}_1/(\overline{\sigma}_1 + \overline{\sigma}_0)$. This is the unique minimizer, and it equals p_{CMR} because the lower and upper endpoints coincide. The certificate is then zero, which agrees with the formula in the proposition.

Now suppose the rectangle does not collapse to a point. Let $r_+(\pi) = r(\pi, \overline{\sigma}_1^2, \underline{\sigma}_0^2)$ and $r_-(\pi) = r(\pi, \underline{\sigma}_1^2, \overline{\sigma}_0^2)$. At $\pi = p_{\mathrm{CMR}}$, the two affine terms inside the squared regret are equal in magnitude and opposite in sign, since $(1-p_{\mathrm{CMR}})\overline{\sigma}_1 - p_{\mathrm{CMR}}\underline{\sigma}_0 = -\left\{(1-p_{\mathrm{CMR}})\underline{\sigma}_1 - p_{\mathrm{CMR}}\overline{\sigma}_0\right\}$. Hence $r_+(p_{\mathrm{CMR}}) = r_-(p_{\mathrm{CMR}})$. The difference between the two corner regrets is strictly decreasing. Indeed, $\frac{d}{d\pi}\{r_+(\pi) - r_-(\pi)\} = -\frac{\overline{\sigma}_1^2 - \underline{\sigma}_1^2}{\pi^2} - \frac{\overline{\sigma}_0^2 - \underline{\sigma}_0^2}{(1-\pi)^2} < 0$, where strict negativity follows because at least one side of the rectangle has positive length. Since the two corner regrets are equal at p_{CMR}, this implies $r_+(\pi) > r_-(\pi)$ for $\pi < p_{\mathrm{CMR}}$, and $r_-(\pi) > r_+(\pi)$ for $\pi > p_{\mathrm{CMR}}$. Therefore the realized worst-case regret is $r_+(\pi)$ to the left of p_{CMR}, and $r_-(\pi)$ to the right of p_{CMR}.

Finally, p_{CMR} lies between the two corner-specific Neyman allocations. The inequalities $\frac{\underline{\sigma}_1}{\underline{\sigma}_1 + \overline{\sigma}_0} \le p_{\mathrm{CMR}} \le \frac{\overline{\sigma}_1}{\overline{\sigma}_1 + \underline{\sigma}_0}$ both reduce to $\underline{\sigma}_1\underline{\sigma}_0 \le \overline{\sigma}_1\overline{\sigma}_0$. The function $r_+(\pi)$ is strictly decreasing to the left of $\overline{\sigma}_1/(\overline{\sigma}_1 + \underline{\sigma}_0)$, and $r_-(\pi)$ is strictly increasing to the right of $\underline{\sigma}_1/(\underline{\sigma}_1 + \overline{\sigma}_0)$. Thus moving π upward toward p_{CMR} strictly lowers the realized worst-case regret whenever $\pi < p_{\mathrm{CMR}}$, while moving π downward toward p_{CMR} strictly lowers the realized worst-case regret whenever $\pi > p_{\mathrm{CMR}}$. Hence p_{CMR} is the unique minimizer of the realized worst-case regret over $\pi \in (0,1)$.

It remains only to evaluate the minimized worst-case regret. Substituting (4.4) into either binding corner regret and simplifying gives the certificate formula (4.5). This also proves that both off-diagonal corners bind at the optimum. $\square$

Proof of Theorem 4.1. We prove the two claims in turn. By Lemma 4.1, the realized rectangle has uniform coverage: $\Pr_{P_F}\left\{\theta(F) \in \widehat{\Theta}_\alpha(\omega)\right\} \geq 1-\alpha$ for every $F \in \mathcal{F}$. Now fix $F \in \mathcal{F}$. On the event $\theta(F) \in \widehat{\Theta}_\alpha(\omega)$, the realized regret of the CMR assignment is bounded by the worst-case regret over the realized rectangle: $r(p_{\text{CMR}}(\omega), \theta(F)) \leq \sup_{\theta \in \widehat{\Theta}_\alpha(\omega)} r(p_{\text{CMR}}(\omega), \theta)$. By definition of the CMR certificate, the right-hand side is $U_{\text{CMR}}(\omega)$. Therefore $\left\{\theta(F) \in \widehat{\Theta}_\alpha(\omega)\right\} \subseteq \{r(p_{\text{CMR}}(\omega), \theta(F)) \leq U_{\text{CMR}}(\omega)\}$. Taking probabilities under P_F and using the coverage lemma proves part (i).

For part (ii), fix a pilot realization ω, and write the realized standard-deviation rectangle as $[\underline{\sigma}_1, \overline{\sigma}_1] \times [\underline{\sigma}_0, \overline{\sigma}_0] \subseteq [0, 1/2]^2$. The CMR value is no larger than the worst-case regret from balanced assignment. At $p = 1/2$, the square representation of regret gives $r(1/2, \theta) = (\sigma_1 - \sigma_0)^2$. Because $\sigma_1, \sigma_0 \in [0, 1/2]$, this quantity is at most $1/4$. Therefore

$$U_{\text{CMR}}(\omega) = \inf_{p \in (0,1)} \sup_{\theta \in \widehat{\Theta}_\alpha(\omega)} r(p, \theta) \leq \sup_{\theta \in \widehat{\Theta}_\alpha(\omega)} r(1/2, \theta) \leq \frac{1}{4}.$$

We now characterize equality. By Proposition 4.1, the CMR certificate on the realized rectangle is $U_{\text{CMR}}(\omega) = \frac{(\overline{\sigma}_1\overline{\sigma}_0 - \underline{\sigma}_1\underline{\sigma}_0)^2}{(\overline{\sigma}_1 + \underline{\sigma}_1)(\overline{\sigma}_0 + \underline{\sigma}_0)}$. Since $0 \leq \underline{\sigma}_d \leq \overline{\sigma}_d \leq 1/2$ for $d \in \{0, 1\}$, we have $U_{\text{CMR}}(\omega) \leq \frac{\overline{\sigma}_1^2\overline{\sigma}_0^2}{\overline{\sigma}_1\overline{\sigma}_0} = \overline{\sigma}_1\overline{\sigma}_0 \leq \frac{1}{4}$. The first inequality uses $\overline{\sigma}_1\overline{\sigma}_0 - \underline{\sigma}_1\underline{\sigma}_0 \leq \overline{\sigma}_1\overline{\sigma}_0$ and $(\overline{\sigma}_1 + \underline{\sigma}_1)(\overline{\sigma}_0 + \underline{\sigma}_0) \geq \overline{\sigma}_1\overline{\sigma}_0$. Thus equality can hold only if both inequalities are equalities. This requires $\underline{\sigma}_1 = \underline{\sigma}_0 = 0$ and $\overline{\sigma}_1 = \overline{\sigma}_0 = 1/2$. In variance coordinates, this is exactly the condition that $(1/4, 0) \in \widehat{\Theta}_\alpha(\omega)$ and $(0, 1/4) \in \widehat{\Theta}_\alpha(\omega)$. Conversely, suppose the realized rectangle contains both adversarial corners $(1/4, 0)$ and $(0, 1/4)$. Then, for any $p \in (0, 1)$, $\sup_{\theta \in \widehat{\Theta}_\alpha(\omega)} r(p, \theta) \geq \max\{r(p, (1/4, 0)), r(p, (0, 1/4))\}$. The two terms on the right are $r(p, (1/4, 0)) = \frac{1-p}{4p}$ and $r(p, (0, 1/4)) = \frac{p}{4(1-p)}$. Their maximum is minimized when they are equal, which occurs at $p = 1/2$, and the minimized value is $1/4$. Therefore every assignment has worst-case regret at least $1/4$ on such a rectangle. Since we already proved $U_{\text{CMR}}(\omega) \leq 1/4$, it follows that $U_{\text{CMR}}(\omega) = 1/4$. This proves the equality statement and completes the proof. $\square$

Proof of Theorem 4.2. Fix $F \in \mathcal{F}$. We work under P_F, suppress dependence on ω, and write $\sigma_d = \sigma_d(F)$. By assumption, $\sigma_1, \sigma_0 > 0$.

We begin with the contraction of the rectangle endpoints. For each arm $d \in \{0, 1\}$, the raw pilot variance satisfies $\hat{\sigma}_d^2 - \sigma_d^2 = O_p(M_d^{-1/2})$. Indeed, the usual within-arm sample variance is a smooth function of the sample mean of $Y(d)$ and the sample mean of $Y(d)^2$. Since $Y(d) \in [0, 1]$, both sample means converge to their expectations at rate $M_d^{-1/2}$.

Since $\sigma_d > 0$, the square-root map is Lipschitz in a neighborhood of σ_d^2. Hence $\hat{\sigma}_d - \sigma_d = O_p(M_d^{-1/2})$. The Maurer–Pontil standard-deviation radius used at one-sided error level $\alpha/4$ is $\sqrt{\frac{2\log(4/\alpha)}{M_d-1}}$, which is $O(M_d^{-1/2})$ because α is fixed. The lower endpoint is obtained by subtracting this radius from $\hat{\sigma}_d$, taking a positive part, and capping the endpoint at $1/2$, while the upper endpoint is obtained by adding the same radius and capping the endpoint at $1/2$. These operations are nonexpansive on the standard-deviation scale. Therefore, for each $d \in \{0,1\}$, $|\underline{\sigma}_d - \sigma_d| + |\overline{\sigma}_d - \sigma_d| = O_p(M_d^{-1/2})$.

We now prove part (i). By Proposition 4.1, $p_{\mathrm{CMR}} = \frac{\overline{\sigma}_1+\underline{\sigma}_1}{\overline{\sigma}_1+\underline{\sigma}_1+\overline{\sigma}_0+\underline{\sigma}_0}$. The endpoint contraction just shown implies that $\overline{\sigma}_d + \underline{\sigma}_d = 2\sigma_d + O_p(M_d^{-1/2})$ for each arm d. Since $\sigma_1 + \sigma_0 > 0$, the denominator in the CMR formula converges in probability to $2(\sigma_1+\sigma_0)$, which is strictly positive. Hence it is bounded away from zero with probability approaching one.

The map $(x_1, x_0) \mapsto x_1/(x_1+x_0)$ is Lipschitz on any compact set where x_1+x_0 is bounded away from zero. Applying this fact to $x_1 = \frac{\overline{\sigma}_1+\underline{\sigma}_1}{2}$, $x_0 = \frac{\overline{\sigma}_0+\underline{\sigma}_0}{2}$, and comparing with (σ_1, σ_0), gives $|p_{\mathrm{CMR}} - \pi^*(\theta(F))| = O_p\left(M_1^{-1/2} + M_0^{-1/2}\right)$. In particular, $p_{\mathrm{CMR}} \xrightarrow{p} \pi^*(\theta(F))$. This proves part (i).

We next prove part (ii). The square representation of regret gives, at the true variance pair, $r(p_{\mathrm{CMR}}, \theta(F)) = \frac{(\sigma_1+\sigma_0)^2(p_{\mathrm{CMR}}-\pi^*(\theta(F)))^2}{p_{\mathrm{CMR}}(1-p_{\mathrm{CMR}})}$. Because $\sigma_1 > 0$ and $\sigma_0 > 0$, the infeasible Neyman allocation $\pi^*(\theta(F))$ lies in $(0,1)$. Since part (i) shows that $p_{\mathrm{CMR}} \xrightarrow{p} \pi^*(\theta(F))$, the denominator $p_{\mathrm{CMR}}(1-p_{\mathrm{CMR}})$ is bounded away from zero with probability approaching one. The factor $(\sigma_1+\sigma_0)^2$ is fixed. Therefore the regret is of the same stochastic order as $(p_{\mathrm{CMR}} - \pi^*(\theta(F)))^2$. Using part (i), $r(p_{\mathrm{CMR}}, \theta(F)) = O_p\left(\left(M_1^{-1/2} + M_0^{-1/2}\right)^2\right)$. Finally, $\left(M_1^{-1/2} + M_0^{-1/2}\right)^2 = M_1^{-1} + 2(M_1M_0)^{-1/2} + M_0^{-1} \le 2(M_1^{-1} + M_0^{-1})$, so $r(p_{\mathrm{CMR}}, \theta(F)) = O_p(M_1^{-1} + M_0^{-1})$. This proves part (ii).

It remains to prove part (iii). By Proposition 4.1, the CMR certificate is $U_{\mathrm{CMR}} = \frac{(\overline{\sigma}_1\overline{\sigma}_0 - \underline{\sigma}_1\underline{\sigma}_0)^2}{(\overline{\sigma}_1+\underline{\sigma}_1)(\overline{\sigma}_0+\underline{\sigma}_0)}$. The endpoint contraction implies that $\overline{\sigma}_d - \underline{\sigma}_d = O_p(M_d^{-1/2})$ for each arm. Since all standard-deviation endpoints lie in $[0, 1/2]$, $|\overline{\sigma}_1\overline{\sigma}_0 - \underline{\sigma}_1\underline{\sigma}_0| \le \frac{1}{2}|\overline{\sigma}_1 - \underline{\sigma}_1| + \frac{1}{2}|\overline{\sigma}_0 - \underline{\sigma}_0|$. Hence the numerator of the certificate is $O_p\left(\left(M_1^{-1/2} + M_0^{-1/2}\right)^2\right)$. The denominator converges in probability to $(2\sigma_1)(2\sigma_0)$, which is strictly positive because both true standard deviations are positive. Therefore the denominator is bounded away from zero with probability approaching one. Combining the numerator and denominator bounds gives $U_{\mathrm{CMR}} = O_p\left(\left(M_1^{-1/2} + M_0^{-1/2}\right)^2\right)$. As in part (ii), this is also $O_p(M_1^{-1} + M_0^{-1})$. This proves part (iii) and completes the proof. $\square$

Proof of Theorem 4.3. We prove the global square-root bound first and then the faster interior bound. Fix $F \in \mathcal{F}$. Throughout the proof, all expectations and probabilities are under P_F, and we write σ_d for $\sigma_d(F)$. Let $\eta_d = \sqrt{2\log(4/\alpha)/(M_d-1)}$ be the Maurer–Pontil radius on the standard-deviation scale. Write $\hat{\sigma}_d = \sqrt{\hat{\sigma}_d^2}$ for the raw empirical standard deviation in arm d, and write $S_d = \overline{\sigma}_d + \underline{\sigma}_d$ for the sum of the standard-deviation endpoints in arm d.

We first record the one-arm estimates used throughout the proof. Since $Y(d) \in [0,1]$, the fourth central moment of $Y(d)$ is bounded by its variance, $\mathbb{E}_F[(Y(d)-\mathbb{E}_F[Y(d)])^4] \le \sigma_d^2$. The usual formula for the variance of the sample variance therefore implies that the raw empirical variance has mean squared error bounded by $C\sigma_d^2/M_d$, uniformly over $F \in \mathcal{F}$. Since $(\sqrt{x}-\sqrt{y})^2 \le (x-y)^2/y$ for $x \ge 0$ and $y > 0$, and the case $y = 0$ is trivial because the empirical variance is then zero almost surely, it follows that $\mathbb{E}_{P_F}[(\hat{\sigma}_d - \sigma_d)^2] \le \frac{C}{M_d}$ for each $d \in \{0,1\}$. This uniform bound is the step that prevents the square-root map from creating a singularity near zero variance.

We now use the endpoint construction. On the standard-deviation scale, $\underline{\sigma}_d = \min\{\frac{1}{2}, (\hat{\sigma}_d - \eta_d)_+\}$, $\overline{\sigma}_d = \min\{\frac{1}{2}, \hat{\sigma}_d + \eta_d\}$. The empirical standard deviation is raw and may exceed $1/2$, but it is uniformly bounded. Since outcomes lie in $[0,1]$, $0 \le \hat{\sigma}_d^2 \le \frac{M_d}{4(M_d-1)}$ and $0 \le \hat{\sigma}_d \le \sqrt{\frac{M_d}{4(M_d-1)}} \le \frac{1}{\sqrt{2}}$ for $M_d \ge 2$. Since also $\eta_d \le \sqrt{2\log(4/\alpha)}$ for all $M_d \ge 2$, there exists a constant $c_\alpha > 0$, depending only on α, such that $S_d \ge c_\alpha(\hat{\sigma}_d + \eta_d)$ for all $d \in \{0,1\}$. Indeed, $S_d \ge \overline{\sigma}_d = \min\{1/2, \hat{\sigma}_d + \eta_d\}$, and $\hat{\sigma}_d + \eta_d$ is uniformly bounded above.

The endpoint sum is also close to $2\sigma_d$. The map $t \mapsto \min\{\frac{1}{2}, (t-\eta_d)_+\} + \min\{\frac{1}{2}, t+\eta_d\}$ is 2-Lipschitz. Evaluated at $t = \sigma_d \in [0,1/2]$, its value differs from $2\sigma_d$ by at most η_d. Hence $|S_d - 2\sigma_d| \le 2\,|\hat{\sigma}_d - \sigma_d| + \eta_d$.

These deterministic inequalities imply the two one-arm bounds needed below. First,

$$\mathbb{E}_{P_F}\left[\frac{(S_d - 2\sigma_d)^2}{S_d}\right] \le C_\alpha \mathbb{E}_{P_F}\left[\frac{(\hat{\sigma}_d - \sigma_d)^2 + \eta_d^2}{\hat{\sigma}_d + \eta_d}\right] \le C_\alpha \frac{\mathbb{E}_{P_F}[(\hat{\sigma}_d - \sigma_d)^2] + \eta_d^2}{\eta_d} \le C_\alpha M_d^{-1/2}.$$

Second, using $\sigma_d \le \hat{\sigma}_d + |\hat{\sigma}_d - \sigma_d|$,

$$\begin{aligned}\mathbb{E}_{P_F}\left[\frac{\sigma_d^2}{S_d}\right] \le C_\alpha \mathbb{E}_{P_F}\left[\frac{\sigma_d^2}{\hat{\sigma}_d + \eta_d}\right] &\le C_\alpha \mathbb{E}_{P_F}\left[\frac{\hat{\sigma}_d^2}{\hat{\sigma}_d + \eta_d} + \frac{(\hat{\sigma}_d - \sigma_d)^2}{\hat{\sigma}_d + \eta_d}\right] \\ &\le C_\alpha\left\{1 + \frac{\mathbb{E}_{P_F}[(\hat{\sigma}_d - \sigma_d)^2]}{\eta_d}\right\} \le C_\alpha.\end{aligned}$$

The last inequalities use the uniform bound on $\hat{\sigma}_d$, the preceding mean-square bound, and $\eta_d \asymp_\alpha M_d^{-1/2}$.

We now turn to the CMR regret. By Proposition 4.1, $p_{\mathrm{CMR}} = S_1/(S_1 + S_0)$. Substituting this expression into the square representation of regret gives $r(p_{\mathrm{CMR}}, \theta(F)) = \frac{(S_0\sigma_1 - S_1\sigma_0)^2}{S_1 S_0}$.

The numerator can be rewritten as $S_0\sigma_1 - S_1\sigma_0 = \sigma_1(S_0 - 2\sigma_0) - \sigma_0(S_1 - 2\sigma_1)$. Therefore, $r(p_{\mathrm{CMR}}, \theta(F)) \le 2\frac{\sigma_1^2}{S_1}\frac{(S_0-2\sigma_0)^2}{S_0} + 2\frac{\sigma_0^2}{S_0}\frac{(S_1-2\sigma_1)^2}{S_1}$. The two pilot arms are independent. Taking expectations and applying the two one-arm bounds gives $\mathbb{E}_{P_F}[r(p_{\mathrm{CMR}}, \theta(F))] \le C_\alpha \left(M_1^{-1/2} + M_0^{-1/2}\right)$, uniformly over $F \in \mathcal{F}$. Taking the supremum over F proves the global bound.

We now prove the interior bound. Fix $\kappa > 0$ and restrict attention to $F \in \mathcal{F}$ satisfying $\sigma_1 \wedge \sigma_0 \ge \kappa$. On the event $\{\hat\sigma_1 \ge \kappa/2,\ \hat\sigma_0 \ge \kappa/2\}$, the lower bound on S_d implies that both S_1 and S_0 are bounded below by a positive constant depending only on α and κ. Hence, on this event, $r(p_{\mathrm{CMR}}, \theta(F)) \le C_{\alpha,\kappa}\left\{(S_1 - 2\sigma_1)^2 + (S_0 - 2\sigma_0)^2\right\}$. Using $|S_d - 2\sigma_d| \le 2|\hat\sigma_d - \sigma_d| + \eta_d$, the mean-square bound above, and $\eta_d^2 \asymp_\alpha M_d^{-1}$, the expected regret on this event is bounded by $C_{\alpha,\kappa}(M_1^{-1} + M_0^{-1})$.

It remains to control the exceptional event. Let $B_d = \{\hat\sigma_d < \kappa/2\}$. Since $\sigma_d \ge \kappa$, bounded-difference concentration for the empirical variance gives constants $C_\kappa, c_\kappa > 0$, depending only on κ, such that $\Pr_{P_F}(B_d) \le C_\kappa e^{-c_\kappa M_d}$. To see this, write the unbiased sample variance as $\hat\sigma_d^2 = \frac{1}{M_d(M_d-1)}\sum_{1\le i<j\le M_d}(Y_{di} - Y_{dj})^2$. Changing one observation changes this statistic by at most $1/M_d$. Since $\mathbb{E}_{P_F}[\hat\sigma_d^2] = \sigma_d^2 \ge \kappa^2$, the event B_d implies $\hat\sigma_d^2 - \sigma_d^2 \le -3\kappa^2/4$, and McDiarmid's bounded-differences inequality (McDiarmid, 1989) gives the exponential bound. We cannot simply multiply the probability of $B_1 \cup B_0$ by the largest two-arm bound, because that would pair a two-arm blowup with a one-arm failure probability. Instead, we split the exceptional event by which arm is bad.

Consider $B_1 \cap B_0^c$. On this event, $S_0 \ge c_{\alpha,\kappa}$, so $1 - p_{\mathrm{CMR}} = S_0/(S_1 + S_0) \ge c_{\alpha,\kappa}$. Also, the lower bound on S_d gives $S_1 \ge c_\alpha\eta_1$, while $S_1 + S_0 \le 2$. Hence $p_{\mathrm{CMR}} \ge c_\alpha\eta_1$, after reducing the constant if needed. Therefore, on $B_1 \cap B_0^c$, $r(p_{\mathrm{CMR}}, \theta(F)) \le \frac{\sigma_1^2}{p_{\mathrm{CMR}}} + \frac{\sigma_0^2}{1-p_{\mathrm{CMR}}} \le C_{\alpha,\kappa}M_1^{1/2}$. Multiplying by $\Pr_{P_F}(B_1) \le C_\kappa e^{-c_\kappa M_1}$ gives a contribution bounded by $C_{\alpha,\kappa}M_1^{-1}$, after increasing the constant. The case $B_0 \cap B_1^c$ is symmetric and contributes at most $C_{\alpha,\kappa}M_0^{-1}$.

On $B_1 \cap B_0$, the lower bounds $S_d \ge c_\alpha\eta_d$ for both arms imply $p_{\mathrm{CMR}} \ge c_\alpha\eta_1$ and $1 - p_{\mathrm{CMR}} \ge c_\alpha\eta_0$. Thus $r(p_{\mathrm{CMR}}, \theta(F)) \le C_\alpha\left(M_1^{1/2} + M_0^{1/2}\right)$. By independence of the two pilot arms, the probability of $B_1 \cap B_0$ is at most $C_\kappa e^{-c_\kappa M_1}e^{-c_\kappa M_0}$. Hence the contribution of this event is bounded by $C_{\alpha,\kappa}\left(M_1^{1/2} + M_0^{1/2}\right)e^{-c_\kappa M_1}e^{-c_\kappa M_0} \le C_{\alpha,\kappa}\left(M_1^{-1} + M_0^{-1}\right)$. Combining the good-event and exceptional-event bounds gives $\mathbb{E}_{P_F}[r(p_{\mathrm{CMR}}, \theta(F))] \le C_{\alpha,\kappa}\left(M_1^{-1} + M_0^{-1}\right)$, uniformly over all $F \in \mathcal{F}$ satisfying $\sigma_1(F) \wedge \sigma_0(F) \ge \kappa$. Taking the supremum over this class proves the interior bound and completes the proof. $\square$

A preliminary observation. For the multi-arm shared-control extension, let $\mathcal{P} = \Delta_K$, the simplex over the $K+1$ arms. For the stratified extension, let $\mathcal{P} = \Delta_{2S-1}$, the simplex over the $2S$ treatment-by-stratum cells. In either design, write $\mathcal{V}_\alpha(\omega)$ for the finite set of vertices of $\widehat{\Theta}_\alpha(\omega)$, obtained by setting every variance coordinate to either endpoint of its confidence interval. For every $p \in \mathcal{P}$, $V(p,\theta)$ is affine in θ. Therefore $V^*(\theta) = \inf_{\widetilde{p}\in\mathcal{P}} V(\widetilde{p},\theta)$ is concave and $r(p,\theta) = V(p,\theta) - V^*(\theta)$ is convex in θ. It follows that worst-case regret over the confidence hyperrectangle is attained at a vertex. Moreover, at every fixed vertex the inverse-share criterion is convex in p, so the CMR allocation solves the finite convex problem $\min_{p\in\mathcal{P}} \max_{\theta\in\mathcal{V}_\alpha(\omega)} r(p,\theta)$.

Proof of Proposition 5.1. Part (i). When $\widehat{\Theta}_\alpha(\omega) = [0, 1/4]^{K+1}$, the preliminary observation reduces the inner supremum to the vertices of the box, where each σ_k^2 is 0 or $1/4$. For a nonempty $A \subseteq \{0, 1, \ldots, K\}$, let θ_A denote the vertex with $\sigma_j^2 = 1/4$ for $j \in A$ and $\sigma_j^2 = 0$ otherwise. The full-box CMR problem is to minimize $\max_{A\neq\emptyset} r(p, \theta_A)$ over the positive simplex.

At θ_A, with $\sigma_j = 1/2$ for $j \in A$,

$$V(p,\theta_A) = \frac{K\sigma_0^2}{p_0}\,\mathbf{1}\{0\in A\} + \sum_{k\in A,\,k\geq 1} \frac{\sigma_k^2}{p_k}, \qquad \sqrt{V^*(\theta_A)} = \sqrt{K}\,\sigma_0\,\mathbf{1}\{0\in A\} + \sum_{k\in A,\,k\geq 1} \sigma_k.$$

Fix a nonempty proper subset A. Under any allocation in the positive simplex, both A and A^c receive positive total allocation. Cauchy–Schwarz applied to the terms of $\sqrt{V^*(\theta_A)}$ against $(\sqrt{p_j})_{j\in A}$ gives $V^*(\theta_A) \leq V(p,\theta_A)\sum_{j\in A} p_j$. Writing $a = \sum_{j\in A} p_j \in (0,1)$, $r(p,\theta_A) = V(p,\theta_A) - V^*(\theta_A) \geq V^*(\theta_A)\,\frac{1-a}{a}$. The complement carries mass $1-a$, so likewise $r(p,\theta_{A^c}) \geq V^*(\theta_{A^c})\,\frac{a}{1-a}$. The two right-hand sides multiply to $V^*(\theta_A)V^*(\theta_{A^c})$, so their maximum is at least $\sqrt{V^*(\theta_A)V^*(\theta_{A^c})}$. Maximizing over nonempty proper subsets,

$$\max_{B\neq\emptyset} r(p,\theta_B) \geq \max_{\emptyset\neq A\subsetneq\{0,\ldots,K\}} \sqrt{V^*(\theta_A)V^*(\theta_{A^c})} \qquad \text{for every allocation } p. \tag{A.2}$$

The allocation p_{CMR} attains (A.2). Its components are $p_{0,\mathrm{CMR}} = \sqrt{K}/(\sqrt{K}+K)$ and $p_{k,\mathrm{CMR}} = 1/(\sqrt{K}+K)$, so

$$V(p_{\mathrm{CMR}},\theta_A) = \frac{\sqrt{K}+K}{4}\Big[\sqrt{K}\,\mathbf{1}\{0\in A\} + |A\cap\{1,\ldots,K\}|\Big] = \frac{\sqrt{K}+K}{2}\sqrt{V^*(\theta_A)}.$$

Since $\sqrt{V^*(\theta_A)} + \sqrt{V^*(\theta_{A^c})} = \frac{1}{2}(\sqrt{K}+K)$, counting $\frac{\sqrt{K}}{2}$ for the control and $\frac{1}{2}$ for each treatment, the last display equals $\big(\sqrt{V^*(\theta_A)} + \sqrt{V^*(\theta_{A^c})}\big)\sqrt{V^*(\theta_A)}$, hence $r(p_{\mathrm{CMR}},\theta_A) =$

$V(p_{\text{CMR}},\theta_A) - V^*(\theta_A) = \sqrt{V^*(\theta_A)V^*(\theta_{A^c})}$ for every nonempty proper A. At the full-set vertex all arm variances are equal and p_{CMR} is the corresponding Neyman allocation, so $r(p_{\text{CMR}},\theta_{\{0,\dots,K\}}) = 0$. Taking the maximum over nonempty proper subsets therefore matches (A.2), so p_{CMR} solves the full-box problem. With $\sqrt{K}/(\sqrt{K}+K) = 1/(1+\sqrt{K})$ and $1/(\sqrt{K}+K) = 1/(\sqrt{K}(1+\sqrt{K}))$, this is the allocation in part (i).

The minimizer is unique. Let A attain the maximum in (A.2) and let p be any minimizer. Equality in (A.2) forces $\max\{V^*(\theta_A)\frac{1-a}{a},\ V^*(\theta_{A^c})\frac{a}{1-a}\} = \sqrt{V^*(\theta_A)V^*(\theta_{A^c})}$, which holds only at $a = \sqrt{V^*(\theta_A)}/(\sqrt{V^*(\theta_A)} + \sqrt{V^*(\theta_{A^c})})$, and it forces equality in both Cauchy–Schwarz steps. Equality holds only when p_j is proportional to the standard-deviation score, $\sqrt{K}$ for the control and 1 for a treatment, on A and on A^c; with the value of a, the constant is $1/(\sqrt{K}+K)$ on each block, so $p = p_{\text{CMR}}$.

Part (ii). We first record that measurable selections from the CMR solution set exist. For each ω, the objective $p \mapsto \sup_{\theta\in\widehat{\Theta}_\alpha(\omega)} r(p,\theta)$ is continuous on the open simplex and diverges as any share vanishes, because every upper variance endpoint of $\widehat{\Theta}_\alpha(\omega)$ is strictly positive, so its minimization can be restricted to a compact set. The objective is measurable in ω because the rectangle enters through finitely many endpoint variables, and the measurable maximum theorem (Aliprantis and Border, 2006, Theorem 18.19) yields a measurable minimizer.

Write $\theta^0 = \theta(F)$, with all components strictly positive, and let $\pi^*(\theta^0)$ be the Neyman allocation (5.1) at θ^0. Since every $\sigma_k(F) > 0$, $\pi^*(\theta^0)$ has strictly positive components and is the unique minimizer of $V(p,\theta^0)$, equivalently of $r(p,\theta^0)$, over the positive simplex, with minimum value zero.

Since outcomes are bounded and $M_{\min} \to \infty$, the arm-level variance intervals forming $\widehat{\Theta}_\alpha(\omega)$ shrink around the true variances, so $\sup_{\theta\in\widehat{\Theta}_\alpha(\omega)} \|\theta - \theta^0\| \xrightarrow{p} 0$. We first prove a deterministic statement: along any rectangles P shrinking to θ^0, every CMR minimizer converges to $\pi^*(\theta^0)$.

Fix $\varepsilon > 0$ and let $\sigma^2_{\min} = \min_k \sigma_k^2 > 0$. Choose $\delta > 0$ so small that $\|\theta - \theta^0\| \le \delta$ implies every component of θ is at least $\sigma^2_{\min}/2$. For such θ and any allocation p, letting some $p_j \downarrow 0$ sends $K\theta_0/p_0$ or θ_j/p_j to infinity, uniformly over $\|\theta-\theta^0\| \le \delta$, since the relevant variance stays above $\sigma^2_{\min}/2$. Because $\pi^*(\theta^0)$ has strictly positive components and $\sup_{\|\theta-\theta^0\|\le\delta} r(\pi^*(\theta^0),\theta)$ is finite, there is $\eta \in (0, \min_k \pi_k^*(\theta^0)]$ such that, for every rectangle $P \subseteq \{\|\theta - \theta^0\| \le \delta\}$, every minimizer of $p \mapsto \sup_{\theta\in P} r(p,\theta)$ has all components at least η. Both these minimizers and $\pi^*(\theta^0)$ then lie in the compact set $\Delta_{K,\eta} = \{p : p_k \ge \eta,\ \sum_k p_k = 1\}$.

On $\Delta_{K,\eta} \times \{\|\theta - \theta^0\| \le \delta\}$ the map r is uniformly continuous, so as P shrinks to θ^0, $\sup_{p\in\Delta_{K,\eta}} \left| \sup_{\theta\in P} r(p,\theta) - r(p,\theta^0) \right| \longrightarrow 0$. Since $\pi^*(\theta^0)$ is the unique minimizer of $r(\cdot,\theta^0)$ on $\Delta_{K,\eta}$, with value zero, compactness gives $\gamma = \inf_{\substack{p\in\Delta_{K,\eta} \\ \|p-\pi^*(\theta^0)\|\ge\varepsilon}} r(p,\theta^0) > 0$. For all sufficiently small P, the uniform convergence yields $\sup_{\theta \in P} r(\pi^*(\theta^0), \theta) < \gamma/3$, and $\sup_{\theta \in P} r(p, \theta) > 2\gamma/3$ for every $p \in \Delta_{K,\eta}$ with $\|p - \pi^*(\theta^0)\| \geq \varepsilon$. No CMR minimizer over such a P can then lie at distance $\geq \varepsilon$ from $\pi^*(\theta^0)$.

Finally, apply this to the random rectangles. By this convergence, for every $\delta > 0$ the event $\{\widehat{\Theta}_\alpha(\omega) \subseteq \{\|\theta - \theta^0\| \leq \delta\}\}$ has probability tending to one, and on it every CMR selection lies within ε of $\pi^*(\theta^0)$ for all sufficiently large $M_{\min}$. Since $\varepsilon > 0$ was arbitrary, $p_{\mathrm{CMR}} \xrightarrow{p} \pi^*(\theta^0) = \pi^*(\theta(F))$. $\square$

Proof of Proposition 5.2. Part (i). With $\widehat{\Theta}_\alpha(\omega) = [0, 1/4]^{2S}$, the same preliminary observation places the inner supremum at a vertex of the box, where each $\sigma_{dx}^2 \in \{0, 1/4\}$. Index the cells by $\mathcal{C} = \{0,1\} \times \{1, \ldots, S\}$; for $A \subseteq \mathcal{C}$, let θ_A be the vertex with $\sigma_{dx}^2 = 1/4$ on A and 0 elsewhere, and set $w(A) = \sum_{(d,x) \in A} s_x \in [0, 2]$.

At the balanced allocation $\bar{p}_{1x} = \bar{p}_{0x} = s_x/2$, active cells have standard deviation $1/2$, so $V(\bar{p}, \theta_A) = w(A)/2$ and $\sqrt{V^*(\theta_A)} = w(A)/2$, giving $r(\bar{p}, \theta_A) = \frac{w(A)}{2} - \frac{w(A)^2}{4} = \frac{w(A)\,(2 - w(A))}{4} \leq \frac{1}{4}$, with equality at $w(A) = 1$, for instance the all-treatment vertex $A = \{(1, x) : x\}$. Hence $\sup_\theta r(\bar{p}, \theta) = 1/4$.

No allocation does better. Boundary allocations carry infinite worst-case regret under the boundary-loss convention, so take a positive allocation p and write $a = \sum_x p_{1x}$ and $1 - a = \sum_x p_{0x}$. At the all-treatment vertex $A_1 = \{(1, x) : x\}$, where $V^*(\theta_{A_1}) = 1/4$, Cauchy–Schwarz gives $1 = (\sum_x s_x)^2 \leq a \sum_x s_x^2/p_{1x}$, so $V(p, \theta_{A_1}) = \frac{1}{4} \sum_x s_x^2/p_{1x} \geq 1/(4a)$ and $r(p, \theta_{A_1}) \geq (1 - a)/(4a)$. The same step at the all-control vertex A_0 gives $r(p, \theta_{A_0}) \geq a/(4(1 - a))$, so $\sup_\theta r(p, \theta) \geq \max\left\{\frac{1-a}{4a}, \frac{a}{4(1-a)}\right\} \geq \frac{1}{4}$. Equality forces $a = 1/2$ and, by the Cauchy–Schwarz equality condition, $p_{1x} \propto s_x$, hence $p_{1x} = s_x/2$; the control vertex gives $p_{0x} = s_x/2$. The balanced allocation is therefore the unique CMR solution, with certificate $U_{\mathrm{CMR}}(\omega) = 1/4$.

Part (ii). Measurable selections from the CMR solution set exist by the same argument as in the proof of Proposition 5.1. Write $\theta^0 = \theta(F)$ and $\pi^0 = \pi^*(\theta^0)$ for the stratified Neyman allocation (5.3), with strictly positive components since $s_x, \sigma_{dx}(F) > 0$. By the same Cauchy–Schwarz argument used to derive the stratified Neyman value, π^0 is the unique minimizer of $r(\cdot, \theta^0)$ over the positive simplex, attained only at cell shares proportional to $s_x \sigma_{dx}$. Since outcomes are bounded and $M_{\min} \to \infty$, the cell intervals shrink, $\sup_{\theta \in \widehat{\Theta}_\alpha(\omega)} \|\theta - \theta^0\| \xrightarrow{p} 0$.

The remainder is the compactness and uniform-continuity argument of Proposition 5.1(ii), relabeled to cells (d, x). Boundary coercivity comes from the terms $s_x^2 \sigma_{dx}^2/p_{dx}$, which diverge uniformly near θ^0 as any $p_{dx} \downarrow 0$, so for rectangles close to θ^0 every CMR minimizer lies in a fixed compact subset of the positive simplex. There r is uniformly continuous and $r(\cdot, \theta^0)$ has unique minimizer π^0, so every CMR selection over rectangles shrinking to θ^0 converges to π^0. Applied to the random rectangles $\widehat{\Theta}_\alpha(\omega)$, $p_{\text{CMR}}(\omega) \xrightarrow{p} \pi^0 = \pi^*(\theta(F))$. $\square$

Online Appendix

B Additional Simulation Results

Table B.1: Empirical DGP calibration for the two-arm simulations

Paper	Treatment	Outcome	Type	$\mathbb{E}[Y(1)]$	$\mathbb{E}[Y(0)]$	$\text{Var}(Y(1))$	$\text{Var}(Y(0))$	π^*
Miguel & Kremer (2004)	Deworming treatment	School attendance	Continuous	0.766	0.730	0.1019	0.0965	0.507
		Academic test score	Continuous	0.542	0.552	0.0075	0.0083	0.486
	First-phase deworming	Any moderate-heavy infection	Binary	0.257	0.524	0.1909	0.2494	0.467
Thornton (2008)	Positive incentive	Learned HIV result	Binary	0.791	0.341	0.1654	0.2249	0.462
		Condom purchase	Binary	0.274	0.202	0.1990	0.1611	0.526
		Number of condoms	Count	0.055	0.043	0.0114	0.0136	0.479
Bertrand & Mullainathan (2004)	White-sounding name	Employer callback	Binary	0.097	0.064	0.0872	0.0603	0.546

Notes: The table reports the calibrated treatment-control DGPs used in the two-arm simulation exercises. $Y(1)$ denotes the treatment condition listed in the Treatment column and $Y(0)$ denotes the corresponding comparison condition. For binary outcomes, means are event probabilities and variances are Bernoulli variances implied by the calibrated arm probabilities. For non-binary outcomes, cleaned outcomes are normalized to $[0,1]$ and moments are computed from the empirical calibration distribution, using sample weights when present. $\pi^* = \sigma_1/(\sigma_1 + \sigma_0)$ is the infeasible Neyman treatment share. *Treatment definitions:* Deworming treatment compares observations in treated school-years to not-yet-treated or comparison school-years. First-phase deworming compares first-phase deworming schools to second-phase schools. Positive incentive compares any positive incentive to collect HIV test results to zero incentive. White-sounding name compares resumes assigned white-sounding names to otherwise similar resumes assigned Black-sounding names.

Outcome definitions: School attendance is school participation/attendance. Academic test score is the cleaned test-score outcome. Any moderate-heavy infection is an indicator for any moderate-heavy worm infection. Learned HIV result is an indicator for collecting or learning HIV test results. Condom purchase is an indicator for purchasing any condoms at follow-up. Number of condoms is the number of condoms purchased at follow-up. Employer callback is an indicator for receiving an employer callback.

Table B.2: Mean efficiency loss for additional two-arm allocation rules

Rule	MK attend.	MK test	MK infection	Thor. HIV	Thor. condom	Thor. cond. count	BM callback
Panel A: $M = 30$							
Balance	0.02	0.07	0.45	0.59	0.28	0.18	0.84
CMR (MP)	0.02	0.07	0.45	0.59	0.28	0.18	0.84
FNA	3.33	2.19	∞	∞	∞	∞	∞
Trimmed FNA	3.33	2.19	3.87	5.76	7.11	23.25	65.44
CMR (MTR)	0.02	0.07	0.45	0.59	0.28	0.18	0.84
CMR (Bernoulli)	–	–	0.83	1.12	1.22	–	1.52
Cai-Rafi test	2.10	1.18	0.97	1.38	1.07	5.62	0.96
Cai-Rafi add.	1.39	0.92	1.70	2.52	3.19	9.69	29.21
Cai-Rafi exp.	2.68	1.76	0.72	1.14	1.29	10.28	1.40
Panel B: $M = 100$							
Balance	0.02	0.07	0.45	0.59	0.28	0.18	0.84
CMR (MP)	0.02	0.07	0.07	0.18	0.18	0.22	0.84
FNA	0.71	0.58	0.20	0.36	0.49	5.93	∞
Trimmed FNA	0.71	0.58	0.20	0.36	0.49	5.93	7.89
CMR (MTR)	0.02	0.07	0.44	0.58	0.28	0.18	0.74
CMR (Bernoulli)	–	–	0.20	0.31	0.42	–	2.10
Cai-Rafi test	0.38	0.35	0.33	0.53	0.51	2.44	2.15
Cai-Rafi add.	0.31	0.25	0.13	0.20	0.24	2.55	3.45
Cai-Rafi exp.	0.58	0.46	0.16	0.29	0.40	4.80	2.39
Panel C: $M = 250$							
Balance	0.02	0.07	0.45	0.59	0.28	0.18	0.84
CMR (MP)	0.08	0.05	0.08	0.13	0.10	0.27	0.47
FNA	0.27	0.25	0.06	0.14	0.18	2.14	1.14
Trimmed FNA	0.27	0.25	0.06	0.14	0.18	2.14	1.14
CMR (MTR)	0.40	0.06	0.08	0.32	0.44	0.16	0.48
CMR (Bernoulli)	–	–	0.06	0.14	0.17	–	0.95
Cai-Rafi test	0.13	0.18	0.07	0.25	0.27	1.03	1.31
Cai-Rafi add.	0.12	0.11	0.08	0.13	0.11	0.97	0.57
Cai-Rafi exp.	0.22	0.20	0.05	0.12	0.14	1.74	0.92
Panel D: $M = 500$							
Balance	0.02	0.07	0.45	0.59	0.28	0.18	0.84
CMR (MP)	0.07	0.05	0.08	0.12	0.07	0.22	0.44
FNA	0.12	0.12	0.03	0.06	0.08	1.09	0.48
Trimmed FNA	0.12	0.12	0.03	0.06	0.08	1.09	0.48
CMR (MTR)	0.19	0.05	0.05	0.08	0.09	0.16	1.42
CMR (Bernoulli)	–	–	0.03	0.06	0.08	–	0.44
Cai-Rafi test	0.07	0.12	0.03	0.07	0.14	0.63	0.78
Cai-Rafi add.	0.06	0.06	0.06	0.09	0.06	0.50	0.29
Cai-Rafi exp.	0.10	0.10	0.03	0.05	0.07	0.89	0.39

Notes: Entries are percent mean efficiency losses relative to the infeasible Neyman allocation, $100\,\mathbb{E}_\omega[V(\hat{p}_a(\omega), F) - V(\pi^*(F), F)]/V(\pi^*(F), F)$. The study/outcome columns are the calibrated two-arm DGPs reported in Table B.1. Column abbreviations are MK for Miguel–Kremer, Thor. for Thornton, and BM for Bertrand–Mullainathan. Each entry uses 500 pilot replications per DGP and pilot size. CMR (MP) is the Maurer–Pontil CMR rule; CMR (MTR) is the Martinez–Taboada–Ramdas CMR rule. CMR (Bernoulli) is reported only for binary outcomes. Trimmed FNA uses a 10 percent lower and upper bound on treatment shares. The Cai–Rafi rules use one fixed tuning value from the grids studied in Cai and Rafi (2024): $\alpha = 0.10$ for the homoskedasticity test, $\nu = 0.20$ for additive regularization, and $\tau = 0.90$ for exponential regularization. The tuning values are fixed across DGPs and pilot sizes. – denotes that the rule is not applicable or was not computed for that DGP. Because every calibrated outcome distribution is discrete, any rule that can assign zero probability to an arm does so with positive probability at every pilot size, making its unconditional mean loss infinite. ∞ marks cells in which such a boundary draw occurred among the 500 replications.

C Additional Extensions

C.1 Binary Outcomes

Binary outcomes change only how the confidence rectangle is constructed, not the CMR assignment problem. In the binary-outcome model, each potential outcome satisfies $Y(d) \sim \text{Bernoulli}(q_d)$ for $d \in \{0,1\}$, with variance $\sigma_d^2 = q_d(1-q_d)$, and $\mathcal{F}^{\text{bin}} = \{F \in \mathcal{F} : F(\{0,1\}^2) = 1\}$ denotes this model. As (q_1, q_0) ranges over $[0,1]^2$, the variance pair ranges over the full parameter space $[0,1/4]^2$, so the specialization preserves the state space of the bounded-outcome problem while making the pilot distribution finite.[23] Pilot arm sizes are fixed at $M_d \geq 2$, with $M_1 = M_0 = M/2$ in the balanced design, and $X_d = \sum_{i:D_i=d} Y_i$ counts the successes in arm d, so that $X_d \sim \text{Binomial}(M_d, q_d)$, independently across arms.

The unbiased pilot sample variance in arm d is $\hat{\sigma}_d^2 = X_d(M_d - X_d)/[M_d(M_d-1)]$, which depends on the count only through its distance from an even split. This motivates the folded count $J_d = \min\{X_d,\, M_d - X_d\}$, the size of the smaller outcome category. Relabeling the two outcome categories leaves $\hat{\sigma}_d^2$, σ_d^2, and J_d unchanged, so the folded family is indexed directly by the variance $\sigma_d^2 = q_d(1-q_d)$, the object the design loss depends on.[24] The folded count takes only the $\lfloor M_d/2 \rfloor + 1$ values in $\mathcal{J}_{M_d} = \{0, 1, \ldots, \lfloor M_d/2 \rfloor\}$, and $\hat{\sigma}_d^2 = J_d(M_d - J_d)/[M_d(M_d-1)]$ is increasing in J_d on this support. Although stated for potential outcomes, the construction below uses only that arm d contributes M_d i.i.d. binary observations.

Each candidate variance $\sigma^2 \in [0,1/4]$ corresponds to the two success probabilities $\xi(\sigma^2) = \left(1 - \sqrt{1-4\sigma^2}\right)/2 \in [0,1/2]$ and $1 - \xi(\sigma^2)$. For $j < M_d/2$, the event $\{J_d = j\}$ is the union of the disjoint raw events $\{X_d = j\}$ and $\{X_d = M_d - j\}$, so the folded probability mass function under candidate variance σ^2 is

$$h_{M_d}(j;\sigma^2) = \binom{M_d}{j}\left[\xi(\sigma^2)^j\{1-\xi(\sigma^2)\}^{M_d-j} + \xi(\sigma^2)^{M_d-j}\{1-\xi(\sigma^2)\}^j\right],$$

while for even M_d and $j = M_d/2$ the two raw events coincide and $h_{M_d}(M_d/2;\sigma^2) = \binom{M_d}{M_d/2}\left[\xi(\sigma^2)\{1-\xi(\sigma^2)\}\right]^{M_d/2}$. The inversion uses the two tails of this distribution,

$$G_{\leq}(j;\sigma^2) = \Pr_{\sigma^2}(J_d \leq j) = \sum_{k=0}^{j} h_{M_d}(k;\sigma^2), \quad G_{\geq}(j;\sigma^2) = \sum_{k=j}^{\lfloor M_d/2 \rfloor} h_{M_d}(k;\sigma^2),$$

[23] "Exact" refers to finite-sample inversion of the pilot distribution, in contrast to the concentration-inequality rectangle of Subsection 4.3.

[24] If $X \sim \text{Binomial}(M, \pi)$, then $\min\{X, M-X\}$ has a folded binomial distribution (Gart, 1970; Mantel, 1970; Porzio and Ragozini, 2009).

where $\Pr_{\sigma^2}$ denotes the folded-binomial distribution indexed by σ^2, so $G_{\geq}(j;\sigma^2) = \Pr_{\sigma^2}(J_d \geq j)$. Larger variances shift the folded count upward in the sense of first-order stochastic dominance. For every $j \in \mathcal{J}_{M_d}$, the tail $G_{\geq}(j;\cdot)$ is continuous and nondecreasing on $[0,1/4]$ and $G_{\leq}(j;\cdot)$ is continuous and nonincreasing, an ordering established as Step 1 of the proof of Proposition C.1.

The two tails then separate the two ways a candidate variance can be inconsistent with the pilot. A small folded count is evidence against large variances, so the lower tail delivers the upper endpoint. A large folded count is evidence against small variances, so the upper tail delivers the lower endpoint. This is the classical test-inversion construction of Neyman (1937), with the Clopper and Pearson (1934) interval as the canonical binomial example. I invert the folded family rather than transforming an interval for q_d because J_d is the maximal invariant under relabeling and its distribution is indexed by σ_d^2 itself, so no orientation nuisance remains.

Fix a one-sided error level $b \in (0,1/2)$; the rectangle below uses $b = \alpha/4$. For a realized folded count j, the endpoints are

$$\overline{\sigma}^2_{M_d}(j;b) = \sup\{\sigma^2 \in [0,1/4] : G_{\leq}(j;\sigma^2) > b\},$$

$$\underline{\sigma}^2_{M_d}(j;b) = \begin{cases} \inf\{\sigma^2 \in [0,1/4] : G_{\geq}(j;\sigma^2) > b\}, & \text{if the set is nonempty,} \\ 1/4, & \text{otherwise.} \end{cases}$$

By the stochastic ordering, each acceptance set is an interval, the upper endpoint is strictly positive for every j, and the reported endpoints never cross. The lower acceptance set can be empty only when even the maximal variance $1/4$ assigns upper-tail probability at most b to the observed count, and reporting the boundary value $1/4$ then preserves the lower-bound coverage established below.[25] Applying the armwise bounds to both arms at level $b = \alpha/4$ gives the binary rectangle

$$\widehat{\Theta}^{\text{bin}}_{\alpha}(J_1,J_0) = \left[\underline{\sigma}^2_{M_1}(J_1;\alpha/4), \overline{\sigma}^2_{M_1}(J_1;\alpha/4)\right] \times \left[\underline{\sigma}^2_{M_0}(J_0;\alpha/4), \overline{\sigma}^2_{M_0}(J_0;\alpha/4)\right].$$

Proposition C.1 (Exact-inversion rectangle for binary outcomes). *Fix pilot arm sizes $M_d \geq 2$ for $d \in \{0,1\}$.*

[25] In computation, each endpoint is a one-dimensional root of a monotone continuous tail function. Since M_d, b, and the support are known before the pilot, the endpoint functions form a lookup table indexed by j and need not be recomputed after the pilot.

(i) *For every* $b \in (0, 1/2)$*, each arm* d*, and every* $F \in \mathcal{F}^{\mathrm{bin}}$*,*

$$\Pr_{P_F}\bigl\{\underline{\sigma}^2_{M_d}(J_d; b) \le \sigma^2_d(F)\bigr\} \ge 1 - b \qquad \textit{and} \qquad \Pr_{P_F}\bigl\{\sigma^2_d(F) \le \overline{\sigma}^2_{M_d}(J_d; b)\bigr\} \ge 1 - b,$$

so the endpoints satisfy the one-sided coverage requirement of Subsection 4.1 over $\mathcal{F}^{\mathrm{bin}}$*.*

(ii) *For every* $\alpha \in (0, 1)$ *and every* $F \in \mathcal{F}^{\mathrm{bin}}$*, the rectangle* $\widehat{\Theta}^{\mathrm{bin}}_\alpha(J_1, J_0)$ *contains* $\theta(F)$ *with probability at least* $1-\alpha$*, both upper endpoints are strictly positive, and Proposition 4.1 applied to* $\widehat{\Theta}^{\mathrm{bin}}_\alpha$ *yields a unique interior assignment* $p^{\mathrm{bin}}_{\mathrm{CMR}}$ *and a certificate* $U^{\mathrm{bin}}_{\mathrm{CMR}} \le 1/4$ *with* $\Pr_{P_F}\bigl\{r\bigl(p^{\mathrm{bin}}_{\mathrm{CMR}}, \theta(F)\bigr) \le U^{\mathrm{bin}}_{\mathrm{CMR}}\bigr\} \ge 1 - \alpha$*.*

The rectangle's coverage is at least, not exactly, $1 - \alpha$. The union bound across the four one-sided bounds and the discreteness of the folded distribution both make the construction conservative. What is exact is the inversion itself, which uses the finite-sample distribution of the pilot statistic rather than a concentration inequality or an asymptotic approximation. The gain is concrete at the smallest possible pilot.[26]

The construction supplies the CMR (Bernoulli) rules of the calibrated simulations in Section 6 and the binary activation threshold of Appendix D.

C.2 Unbounded Outcomes

Many outcomes in economics, such as earnings or expenditures, have no known bound. The bounded-outcome assumption does two jobs in the main analysis. It makes the variance space compact, so that the no-pilot minimax values of Section 3 and the universal certificate cap of Theorem 4.1(ii) are worst cases computed over the compact set $\Theta = [0, 1/4]^2$, and it delivers the finite-sample confidence rectangle of Subsection 4.3, whose coverage holds without distributional assumptions. The variance criterion, the Neyman allocation, the regret identity (2.2), and the closed form of Proposition 4.1 require only finite arm variances, so only one object must be rebuilt, the confidence rectangle $\widehat{\Theta}_\alpha(\omega)$.

Finite variances alone, however, are too weak a restriction to support any rectangle, by a Bahadur–Savage-type nonexistence argument (Bahadur and Savage, 1956; Lehmann and Romano, 2005). The obstacle is that a distribution can hide an arbitrarily large variance in

[26] With $M_d = 2$, $J_d \in \{0, 1\}$ and $\Pr_{\sigma^2}(J_d = 1) = 2\sigma^2$, so the inversion returns $[0, 1/4]$ when $J_d = 0$ and $[b/2, 1/4]$ when $J_d = 1$. Two identical observations leave the variance undetermined, treating a zero pilot variance as uncertainty rather than as proof of degeneracy, while one success and one failure already certify $\sigma^2_d \ge b/2$, so if one arm splits and the other is homogeneous the corresponding standard-deviation interval midpoints differ and CMR moves toward the split arm. Since $M_d = 2$ is the smallest arm size at which the sample variance exists, the first balanced pilot at which CMR can react to binary outcomes is $m = 4$, two observations per arm, for every $\alpha \in (0, 1)$.

a value it rarely produces. Consider a distribution that places mass $1-\varepsilon$ at a single point and mass ε at a point a units away. Its variance is $\varepsilon(1-\varepsilon)a^2$, which grows without limit in a, yet with probability $(1-\varepsilon)^{M_d}$ every pilot observation lands on the common point and the sample is indistinguishable from one drawn from a constant outcome. An upper confidence bound that is finite after such a quiet pilot therefore fails to cover, and the failure survives every pilot size, since ε can shrink as M_d grows. Certified adaptation from an unbounded outcome accordingly requires a restriction on the tails, which I impose through a known bound on the kurtosis.

Assumption C.1 (Known kurtosis bound)**.** For each arm $d \in \{0,1\}$, the marginal distribution of the potential outcome $Y(d)$ under F has mean $\mu_d(F)$, variance $\sigma_d^2(F) \in (0,\infty)$, and kurtosis $\mathbb{E}_F\left[(Y(d)-\mu_d(F))^4\right]/\sigma_d^4(F)$ at most ψ_d, for a known constant $\psi_d \in [1,\infty)$.

The lower limit $\psi_d \geq 1$ is without loss, since kurtosis is at least one for every nondegenerate distribution. The bound plays the role that bounded support plays in the main text. It is scale-free and purely moment-based, so it covers continuous and discrete outcomes alike, and it rules out degenerate arms, since kurtosis is undefined at zero variance, so the degenerate configurations that are the central adversarial states of the bounded model are absent here. The assumption converts the impossibility into an explicit pilot-size requirement.

Proposition C.2 (A necessary pilot-size condition)**.** *Fix an arm d and an error level $\alpha \in (0,1)$, and let $\overline{\sigma}_d^2$ be a measurable function of the arm-d pilot outcomes satisfying $\Pr_{P_F}(\sigma_d^2(F) \leq \overline{\sigma}_d^2) \geq 1-\alpha$ for every F obeying Assumption C.1. If $M_d < \frac{(\psi_d+3)\log(1/\alpha)}{4}$, then $\overline{\sigma}_d^2 = \infty$ at every realization in which all arm-d pilot outcomes are equal.*

The threshold quantifies how much certification a quiet pilot can support. Under the kurtosis bound the adversary can still hide the variance in a tail of probability of order $1/\psi_d$, which a pilot arm of size M_d misses entirely with probability of order e^{-M_d/ψ_d}. Until M_d reaches order $\psi_d \log(1/\alpha)$ this miss probability exceeds the allowed error, so a weaker tail restriction, meaning a larger ψ_d, has a price denominated in pilot observations.

Without a range bound the sample variance is fragile, since a single draw from the rare tail can move it arbitrarily far. I therefore estimate each arm variance with a median-of-means estimator $\widehat{v}_d(\omega)$, a classical construction whose finite-sample accuracy depends on the outcome distribution only through the kurtosis (Lugosi and Mendelson, 2019). The estimator splits the arm-d pilot outcomes into $k_d = \lceil 8\log(2/\alpha)\rceil$ blocks, computes within each block an unbiased variance estimate from b_d paired differences, and takes the median of the k_d block estimates.[27] The median across blocks moves only if half the blocks move, so no small number of extreme draws can distort the estimate. The construction requires at least one pair per block, so it needs $M_d \geq 2k_d$, a pilot arm size of order $\log(1/\alpha)$.

Lemma C.1 (Median-of-means variance estimator)**.** *Let Assumption C.1 hold, fix $\alpha \in (0,1)$, and suppose $M_d \geq 2k_d$, so that every block contains at least one pair. Then $\widehat{v}_d(\omega) \geq 0$ at every realization and $\Pr_{P_F}(|\widehat{v}_d(\omega) - \sigma_d^2(F)| \leq \rho_d\, \sigma_d^2(F)) \geq 1 - \frac{\alpha}{2}$, where $\rho_d = \sqrt{\frac{2(\psi_d+1)}{b_d}}$.*

The relative error ρ_d is known before the pilot is run, since it depends only on the pilot size and the kurtosis bound, it shrinks with the pilot at rate $M_d^{-1/2}$, and it grows with ψ_d, so weaker tail restrictions buy less precision from the same pilot. The guarantee becomes informative once $\rho_d < 1$, at pilot arm sizes of order $\psi_d \log(1/\alpha)$, matching the necessary order of Proposition C.2. Given $\rho_d < 1$, rearranging the inequality of Lemma C.1 shows that with probability at least $1 - \alpha/2$ the unknown variance lies between the endpoints

$$\underline{\sigma}_d^2(\omega) = \frac{\widehat{v}_d(\omega)}{1+\rho_d}, \qquad \overline{\sigma}_d^2(\omega) = \frac{\widehat{v}_d(\omega)}{1-\rho_d}, \tag{C.1}$$

and the confidence rectangle $\widehat{\Theta}_\alpha(\omega) = \left[\underline{\sigma}_1^2(\omega), \overline{\sigma}_1^2(\omega)\right] \times \left[\underline{\sigma}_0^2(\omega), \overline{\sigma}_0^2(\omega)\right]$ contains $\theta(F)$ with probability at least $1-\alpha$ by the union bound over the two arms.

The CMR rule for unbounded outcomes is then defined as follows. If $\rho_d \geq 1$ for some arm, the planned pilot is too small for the lemma to carry information at any realization, and the rule keeps balance and reports no finite certificate. This is a planning condition, known before any data are collected. If instead $\rho_d < 1$ for both arms but the realized pilot returns $\widehat{v}_d(\omega) = 0$ for some arm, the rule again keeps balance and reports no certificate, since a zero estimate is not evidence of zero variance, and acting on it would repeat the overreaction to quiet pilots that Remark 3.1 documents for feasible Neyman. Under Assumption C.1 a zero estimate lies inside the deviation-failure event of the lemma, so this occurs with probability at most α. In every other case the rule activates and computes the assignment and certificate of Proposition 4.1 on $\widehat{\Theta}_\alpha(\omega)$, exactly as in the main text.

Remark C.1 (Relation to feasible Neyman)**.** By (C.1), both standard-deviation endpoints for arm d are proportional to $\sqrt{\widehat{v}_d}$. When $\rho_1 = \rho_0$, as under a balanced pilot with a common kurtosis bound, the proportionality factors cancel in (4.4) and every activated assignment is exactly the feasible Neyman allocation computed from the median-of-means estimates.

[27] Concretely, form the halved squared differences $\frac{1}{2}(Y_{d,2i-1} - Y_{d,2i})^2$ of consecutive arm-d pilot outcomes, each of which has mean exactly $\sigma_d^2(F)$, group them into k_d blocks of $b_d = \lfloor \lfloor M_d/2 \rfloor / k_d \rfloor$ pairs each, discarding any remainder, and let $\widehat{v}_d(\omega)$ be the median of the k_d block averages.

Common radii therefore change only the safeguards around feasible Neyman. The variance estimator carries a finite-sample bound on its relative error, the rule declines to act when the pilot shows no dispersion, and every activated assignment is accompanied by a certificate. Unequal radii, which arise when the arms differ in pilot size or kurtosis bound, tilt the assignment toward the arm with the larger radius, relative to the plug-in feasible Neyman assignment.

The CMR guarantees of the main text carry over to this rule as follows.

Proposition C.3 (Unbounded-outcome CMR)**.** *Let Assumption C.1 hold and let* $(p_{\mathrm{CMR}}, U_{\mathrm{CMR}})$ *be the rule just defined. Parts* (i) *and* (ii) *additionally assume* $\rho_d < 1$ *for each arm* $d \in \{0,1\}$*.*

(i) *For every such* F*,*

$$\Pr_{P_F}\left(\widehat{v}_1(\omega) > 0,\ \widehat{v}_0(\omega) > 0,\ r(p_{\mathrm{CMR}}(\omega), \theta(F)) \le U_{\mathrm{CMR}}(\omega) < \infty\right) \ge 1 - \alpha.$$

(ii) *There exist constants* $c_\theta, C_\theta \in (0, \infty)$*, depending only on* $\theta(F)$*, such that whenever* $\rho_1 + \rho_0 \le c_\theta$*, with probability at least* $1 - \alpha$*,*

$$\begin{gathered} |p_{\mathrm{CMR}}(\omega) - \pi^*(\theta(F))| \le \rho_1 + \rho_0, \\ r(p_{\mathrm{CMR}}(\omega), \theta(F)) \le C_\theta(\rho_1 + \rho_0)^2, \qquad U_{\mathrm{CMR}}(\omega) \le C_\theta(\rho_1 + \rho_0)^2. \end{gathered}$$

(iii) *Fix* $\alpha \in (0,1)$ *and* F*, and let* $M_1, M_0 \to \infty$ *with* M_0/M_1 *bounded away from zero and infinity. Then the rule activates with probability tending to one, and* $|p_{\mathrm{CMR}} - \pi^*(\theta(F))| = O_p\left(M_1^{-1/2} + M_0^{-1/2}\right)$*,* $r(p_{\mathrm{CMR}}, \theta(F)) = O_p\left(M_1^{-1} + M_0^{-1}\right)$*,* $U_{\mathrm{CMR}} = O_p\left(M_1^{-1} + M_0^{-1}\right)$*.*

Part (i) extends the certificate guarantee of Theorem 4.1(i), whose argument requires only that the rectangle cover the true variance pair. With probability at least $1 - \alpha$ the rule activates and its realized regret is bounded by the finite reported certificate, so the guarantee is not met by keeping balance alone. Part (ii) states that on an event of probability at least $1 - \alpha$, uniformly in the pilot sizes, the assignment misses the Neyman allocation by at most $\rho_1 + \rho_0$, of order $\sqrt{\psi_d \log(1/\alpha)/M_d}$, while realized regret and the certificate are bounded by its square. Part (iii) turns these bounds into convergence, recovering the rates of Theorem 4.2 and letting the probability that the rule withholds activation vanish.

Beyond the tail restriction itself, what does not carry over is everything that rests on a compact variance space. Without a scale bound, every fixed assignment $p \in (0,1)$ has infinite worst-case regret over the positive finite-variance space, since a common rescaling of both standard deviations scales regret without moving the Neyman allocation, so the no-pilot minimax values of Section 3, the universal cap $U_{\mathrm{CMR}} \leq 1/4$ of Theorem 4.1(ii), and the uniform expected-regret rate of Theorem 4.3 have no analog here. In the unbounded model, absolute guarantees must be earned by the pilot. When the pilot certifies nothing, the rule holds the default of balance and reports no certificate.

C.3 Multiple Outcomes

Experimenters often care about the impact of the treatment on several outcomes. Stack the potential outcomes into the vector $Y(d) = (Y_1(d), \dots, Y_K(d))'$, where each component takes values in $[0, 1]$ as in the main text and Σ_d denotes the covariance matrix of $Y(d)$. Under a main-wave treatment share π, the vector of difference-in-means estimators has covariance matrix proportional to $\Sigma_1/\pi + \Sigma_0/(1-\pi)$, and different summaries of this matrix favor different assignments, so the CMR construction applies once the researcher commits to one summary as the design criterion. In practice the pre-analysis plan makes this commitment when it designates either a summary index or a family of co-primary outcomes.

When the plan designates a summary index, the estimand is the treatment effect on $G(d) = w'Y(d)$, where $w = (w_1, \dots, w_K)'$ collects pre-specified nonnegative weights summing to one, in the spirit of the summary indices of Kling et al. (2007). Since $G(d)$ is a convex combination of outcomes in $[0, 1]$, it is itself an outcome satisfying the assumptions of the main text, and each pilot unit contributes an observation of it, computed from its K recorded outcomes. CMR therefore applies verbatim at the variance pair $(w'\Sigma_1 w,\ w'\Sigma_0 w)$, including the closed-form rule, the certificate and its $1/4$ cap, and the convergence guarantees.

When the plan designates K co-primary outcomes whose effects are reported separately, the weights w_k, again nonnegative and summing to one, express how the researcher values precision across the K effects rather than defining an estimand, and the design criterion is the w-weighted sum of the outcome-specific variance criteria, the weighted sum of the diagonal entries of $\Sigma_1/\pi + \Sigma_0/(1-\pi)$. Criteria that cannot be written this way, such as the maximum worst-case regret across the K outcomes or the volume of a confidence ellipsoid for the K effects, either lose the closed form or depend on the cross-outcome covariances and fall outside this framework. Because V is linear in the two variances, $\sum_{k=1}^{K} w_k\, V(\pi, \sigma_{1k}^2, \sigma_{0k}^2) = V(\pi, \sigma_{1,w}^2, \sigma_{0,w}^2)$ with $\sigma_{d,w}^2 = \sum_{k=1}^{K} w_k\, \sigma_{dk}^2$ and $\sigma_{dk}^2 = \mathrm{Var}(Y_k(d))$, the weighted problem is the main-text problem at the effective variance pair $\theta_w = (\sigma_{1,w}^2, \sigma_{0,w}^2)$, which lies in $[0, 1/4]^2$ because each σ_{dk}^2 lies in $[0, 1/4]$ and the weights sum to one. Regret is measured against the best single assignment for the weighted criterion.

The confidence rectangle is built from per-outcome bounds. For each arm and outcome, the Maurer–Pontil construction of Subsection 4.3 at one-sided error level $\alpha/(4K)$ gives endpoints $\underline{\sigma}^2_{dk}$ and $\overline{\sigma}^2_{dk}$, and the weighted sums $\underline{\sigma}^2_{d,w} = \sum_k w_k \underline{\sigma}^2_{dk}$ and $\overline{\sigma}^2_{d,w} = \sum_k w_k \overline{\sigma}^2_{dk}$ bracket the effective variances whenever every per-outcome bound holds.

Proposition C.4 (Multiple-outcome CMR)**.** *For each arm $d \in \{0,1\}$ and outcome $k = 1,\dots,K$, let $\underline{\sigma}^2_{dk}(\omega)$ and $\overline{\sigma}^2_{dk}(\omega)$ be the Maurer–Pontil endpoints of Subsection 4.3 at one-sided error level $\alpha/(4K)$, and let $p_{\mathrm{CMR}}(\omega)$ and $U_{\mathrm{CMR}}(\omega)$ be computed as in Proposition 4.1 from the rectangle $\widehat{\Theta}_\alpha(\omega) = \big[\underline{\sigma}^2_{1,w}(\omega), \overline{\sigma}^2_{1,w}(\omega)\big] \times \big[\underline{\sigma}^2_{0,w}(\omega), \overline{\sigma}^2_{0,w}(\omega)\big]$. Then, for every joint distribution F of $\{(Y_k(1), Y_k(0))\}_{k=1}^K$ on $[0,1]^{2K}$, $\Pr_{P_F}(r(p_{\mathrm{CMR}}(\omega), \theta_w) \le U_{\mathrm{CMR}}(\omega)) \ge 1-\alpha$.*

The certificate never exceeds $1/4$, and the convergence arguments of Theorem 4.2 apply at the effective pair with only the constants changed.

C.4 Delayed Primary Outcomes

Many primary outcomes are observed only with long delay, as when a job training program is judged by earnings years later or an education intervention by adult outcomes (Chetty et al., 2016; Athey et al., 2026), so a pilot cannot wait for these horizons and records short-run outcomes only. Suppose the pilot observes a short-run outcome $S(d) \in [0,1]$ in each assigned arm, while the primary outcome $Y(d)$ that defines the main-wave loss is not observed at the pilot stage. The state of the world is the joint distribution F of $(S(1), S(0), Y(1), Y(0))$ on $[0,1]^4$, and the design target is unchanged, the variance pair $\theta(F) = (\sigma_1^2, \sigma_0^2)$ of the primary outcome. The pilot distribution depends on F only through the arm marginals of the short-run outcome, the loss depends on F only through the variances of the primary outcome, and the unrestricted joint distribution places no link between the two. The next proposition shows the consequence.

Proposition C.5 (Short-run pilots require a bridge)**.** *Let $\widehat{\Theta}_\alpha(\omega) = \big[\underline{\sigma}^2_1(\omega), \overline{\sigma}^2_1(\omega)\big] \times \big[\underline{\sigma}^2_0(\omega), \overline{\sigma}^2_0(\omega)\big]$ be a confidence rectangle for $\theta(F)$, computed from the pilot, whose endpoints satisfy the one-sided coverage requirements of Subsection 4.1 at level $\alpha/4$ for every distribution F of $(S(1), S(0), Y(1), Y(0))$ on $[0,1]^4$. Then, for every such F, $\Pr_{P_F}\big(\widehat{\Theta}_\alpha(\omega) = [0, 1/4]^2\big) \ge 1-\alpha$.*

However large the pilot, no movement away from balance can be certified from a short-run pilot alone. Any link between the observed and the loss-relevant outcomes is therefore a maintained assumption rather than something the pilot can verify, and I impose one as a bound on how far the two standard deviations can sit apart, stated in terms of the short-run arm variances $\sigma^2_{d,S}(F) = \mathrm{Var}_F(S(d))$.

Assumption C.2 (Standard-deviation bridge)**.** For each arm $d \in \{0, 1\}$, there is a known constant $\zeta_d \geq 0$ such that $|\sigma_d(F) - \sigma_{d,S}(F)| \leq \zeta_d$.

The bridge links the two outcomes only through their arm-specific dispersions and imposes no relationship between their means or treatment effects. In each arm the two standard deviations may differ by at most ζ_d, so a pilot that pins down the short-run standard deviation locates the primary standard deviation within a band of width $2\zeta_d$. The bound is implied by the stronger and more interpretable condition $\mathrm{Var}_F(Y(d) - S(d))^{1/2} \leq \zeta_d$, and in turn by $|Y(d) - S(d)| \leq \zeta_d$ almost surely. Under Assumption C.2, a confidence interval for the short-run standard deviation converts into one for the primary standard deviation by widening each endpoint by ζ_d. Apply the Maurer–Pontil construction of Subsection 4.3 to the pilot's short-run outcomes at one-sided level $\alpha/4$, obtaining standard-deviation endpoints $\underline{\sigma}_{d,S}(\omega)$ and $\overline{\sigma}_{d,S}(\omega)$ for each arm, and set $\underline{\sigma}_d^2(\omega) = \big(\underline{\sigma}_{d,S}(\omega) - \zeta_d\big)_+^2$, $\overline{\sigma}_d^2(\omega) = \big(\min\{\tfrac{1}{2}, \overline{\sigma}_{d,S}(\omega) + \zeta_d\}\big)^2$. The rectangle $\widehat{\Theta}_\alpha(\omega)$ assembles the two intervals, and the assignment and certificate follow from Proposition 4.1 exactly as in the main text.

Proposition C.6 (Delayed-outcome CMR)**.** *For every F satisfying Assumption C.2, the assignment and certificate satisfy* $\Pr_{P_F}(r(p_{\mathrm{CMR}}(\omega), \theta(F)) \leq U_{\mathrm{CMR}}(\omega)) \geq 1 - \alpha$*.*

The certificate is for the variance pair of the primary outcome, computed from a pilot that never observes it. Widening the short-run interval by ζ_d turns the informal practice of treating short-run dispersion as a stand-in for long-run dispersion into a design with a finite-sample guarantee.

The bridge has a price that does not vanish with the pilot. As the pilot grows, the arm-d standard-deviation interval converges to $\big[(\sigma_{d,S} - \zeta_d)_+, \min\{\tfrac{1}{2}, \sigma_{d,S} + \zeta_d\}\big]$, which retains width $2\zeta_d$ when neither truncation binds. When neither truncation binds in either arm, the interval midpoint equals $\sigma_{d,S}$, so the assignment converges to the Neyman allocation of the short-run outcome, generally different from the Neyman allocation for $\theta(F)$. The certificate converges to a positive bridge-dependent limit whenever some $\zeta_d > 0$, unless some arm has both a zero short-run standard deviation and a zero bridge constant. This is the cost of not waiting for the primary outcome, and the certificate states it as a number.

D Choosing the Pilot Size

The main text takes the pilot as given and asks how its evidence should guide the main-wave assignment. This section asks whether such a pilot is worth running in the first place, and how large it should be, a choice that must be made before any observation exists. How the pilot's cost is measured depends on whether its observations can enter the final ATE estimator, so I develop both accountings, show that no pilot plan dominates running no pilot, and derive an explicit balanced pilot of size roughly $n^{2/3}$ that, followed by CMR, is approximately optimal under worst-case evaluation.

D.1 When Can a Pilot Improve Precision?

The admissible pilot sizes form the set $\mathcal{M} = \{0\} \cup \{m \in 2\mathbb{N} : 4 \leq m \leq n-2\}$, where n is the total budget of observations. When $m > 0$ the pilot is balanced, with $m/2 \geq 2$ observations in each arm, and p_m^{CMR} denotes the CMR assignment rule of Section 4 constructed from the size-m pilot. For $m = 0$, no pilot is run and p_0^{CMR} is the constant rule equal to $1/2$. Write $R_m\big(p_m^{\mathrm{CMR}}, F\big) = \mathbb{E}_{P_{F,m}}\big[r\big(p_m^{\mathrm{CMR}}(\omega), \theta(F)\big)\big]$ for the rule's expected assignment regret after a size-m pilot. The benchmark is an infeasible experimenter who knows the variance pair $\theta(F)$, needs no pilot, spends all n observations in the main wave, and attains the smallest possible sampling variance $V^*(\theta(F))/n$, with $V^*(\theta(F)) = (\sigma_1 + \sigma_0)^2$.

Design-only accounting. Suppose first that the pilot informs the main-wave design and is set aside for estimation. The ATE estimator is then based on the $n - m$ main-wave observations alone, and averaging its conditional variance over pilot realizations, the pilot-size loss is the excess over the benchmark,

$$\mathcal{L}_m(F) = \frac{\mathbb{E}_{P_{F,m}}\big[V\big(p_m^{\mathrm{CMR}}(\omega), \theta(F)\big)\big]}{n-m} - \frac{V^*(\theta(F))}{n} = \frac{R_m\big(p_m^{\mathrm{CMR}}, F\big)}{n-m} + \frac{V^*(\theta(F))\, m}{n(n-m)}, \tag{D.1}$$

where the second equality uses the regret identity inside the expectation. The first term is the assignment error that remains after learning from the pilot. The second is the opportunity cost of diverting m observations from estimation, which remains even if the pilot leads CMR to the Neyman allocation in every realization. Without a pilot, the loss is $\mathcal{L}_0(F) = (\sigma_1 - \sigma_0)^2/n$, the bar any pilot must clear, and clearing denominators in (D.1) shows that a size-m pilot weakly improves on no pilot if and only if $(\sigma_1 - \sigma_0)^2 - R_m\big(p_m^{\mathrm{CMR}}, F\big) \geq \frac{2m}{n}\,(\sigma_1^2 + \sigma_0^2)$. Even if the pilot eliminates assignment regret completely, it can therefore pay for itself only if

$$\frac{m}{n} \leq \frac{(\sigma_1 - \sigma_0)^2}{2(\sigma_1^2 + \sigma_0^2)}. \tag{D.2}$$

This necessary cap on the pilot share is zero when the arms are equally variable, rises as their standard deviations separate, and approaches one half only as one arm's variability becomes negligible beside the other's.

Pooled accounting. When the pilot is drawn from the same target population, uses the same treatment arms, and measures the same outcome as the main wave, its observations can instead enter the final ATE estimator.[28] Combining the wave-specific estimators in this way, the excess variance over the same benchmark becomes

$$\mathcal{L}_m^{\text{pool}}(F) = \frac{m(\sigma_1 - \sigma_0)^2 + (n-m)R_m\big(p_m^{\text{CMR}}, F\big)}{n^2}, \tag{D.3}$$

which charges each pilot observation the balance regret $(\sigma_1 - \sigma_0)^2$, the price of assigning it before anything has been learned, and each main-wave observation the expected regret $R_m(p_m^{\text{CMR}}, F)$, with no sampling cost for the pilot itself. A pooled pilot of any size $m < n$ therefore weakly improves on no pilot if and only if

$$R_m\big(p_m^{\text{CMR}}, F\big) \leq (\sigma_1 - \sigma_0)^2. \tag{D.4}$$

Pooling removes the design-only share cap, although a larger pilot still leaves fewer observations to which the learned assignment can be applied. It is the accounting most favorable to piloting.

D.2 A Rate-Optimal Default

For $m \in \mathcal{M} \setminus \{0\}$, write $\mathcal{D}_0(m)$ for the rule class $\mathcal{D}_0$ of Subsection 2.3 built on the balanced size-m pilot, and let $\mathcal{D}_0(0)$ be the singleton containing the constant balance rule. Extend the pooled loss to any $\delta \in \mathcal{D}_0(m)$ as $\mathcal{L}_m^{\text{pool}}(\delta, F) = [m(\sigma_1 - \sigma_0)^2 + (n-m)R_m(\delta, F)]/n^2$, with $\mathcal{L}_0^{\text{pool}}(\delta, F) = \mathcal{L}_0(F)$ when no pilot is run. A pilot plan, a pilot size together with an assignment rule built on it, must be fixed before any outcome is observed, and no such plan dominates running no pilot.[29] I therefore rank every plan by its worst-case pooled loss, and the smallest worst-case loss any plan can secure is

$$\mathcal{L}_n^\star = \min_{m \in \mathcal{M}} \inf_{\delta \in \mathcal{D}_0(m)} \sup_{F \in \mathcal{F}} \mathcal{L}_m^{\text{pool}}(\delta, F).$$

[28] Because the main-wave assignment depends on the pilot outcomes, mechanically pooling all treated and all control observations yields a biased estimator. A wave-stratified estimator with weights fixed in advance avoids this (Hahn et al., 2011; Tabord-Meehan, 2023).

[29] Reacting to a pilot is a bet that the arms differ, and the bet is lost where they do not. Take any plan that leaves balance after some pilot realization, and a population with identical, nondegenerate arm distributions in which that realization has positive probability. Balance there is exactly optimal, so the plan is strictly worse than running no pilot. A plan that never leaves balance reproduces the no-pilot experiment and so cannot dominate it either.

This value is a target rather than a procedure, because attaining it requires optimizing the pilot size jointly with the complete mapping from pilot realizations to assignments, and no finite-dimensional reduction of the inner fixed-m problem is known, as Remark 3.3 notes. The following directly computable plan nearly attains it. Run a balanced pilot of size m_n° and assign the remaining $n - m_n^\circ$ observations by the CMR rule, where for $n \geq 8$

$$m_n^\circ = \min\{m \in \mathcal{M} \setminus \{0\} : m \geq n^{2/3}\} = 2\lceil n^{2/3}/2 \rceil, \tag{D.5}$$

the budget's two-thirds power rounded upward to the nearest admissible even integer, so that $m_n^\circ = 100$ at $n = 1{,}000$ and the pilot share falls at rate $n^{-1/3}$. The next result evaluates the plan at every budget and under both accountings of the pilot's cost, with the design-only loss extended to any rule as $\mathcal{L}_m(\delta, F) = R_m(\delta, F)/(n-m) + V^*(\theta(F))\, m/[n(n-m)]$.

Theorem D.1 (An explicit rate-optimal pilot plan)**.** *Fix $\alpha \in (0,1)$, let $c > 0$ be the universal constant of Proposition 3.4, and let $C_\alpha < \infty$ be the constant of Theorem 4.3. For every $n \geq 8$,*

$$c_1\, n^{-4/3} \;\leq\; \mathcal{L}_n^\star \;\leq\; \sup_{F \in \mathcal{F}} \mathcal{L}^{\mathrm{pool}}_{m_n^\circ}\big(p^{\mathrm{CMR}}_{m_n^\circ}, F\big) \;\leq\; C_{1,\alpha}\, n^{-4/3}, \tag{D.6}$$

with $c_1 = \min\{1/4, \sqrt{2}\, c\}$ and $C_{1,\alpha} = 1/2 + 2\sqrt{2}\, C_\alpha$. The same inequalities hold under the design-only accounting, with $\mathcal{L}^{\mathrm{pool}}_m$ replaced by $\mathcal{L}_m$ throughout, including in the definition of $\mathcal{L}_n^\star$, and $C_{1,\alpha}$ by $C_{2,\alpha} = 8 + 8\sqrt{2}\, C_\alpha$. Moreover, under either accounting, if a sequence of plans (m_n, δ_n) has worst-case loss at most a fixed multiple of $n^{-4/3}$ for all large n, then $m_n \asymp n^{2/3}$, with implied constants depending only on that multiple.

The rate follows from balancing two unavoidable costs, the pilot's exploration cost of order m/n^2 in (D.3) against the worst-case assignment regret it leaves behind, of order $m^{-1/2}$ and hence $1/(n\sqrt{m})$ in the loss. Equating the two gives $m \asymp n^{2/3}$ and a loss of order $n^{-4/3}$. This rules out the common defaults, since a pilot of fixed size eventually stops learning and a pilot of fixed budget share overpays for exploration, so both remain at the no-pilot order n^{-1}. The constants are conservative byproducts of the concentration bounds behind Theorem 4.3, so the theorem pins down rates rather than loss levels, and when credible planning information is available the operational checks below can be more informative than the worst-case rule.[30]

[30] Kawato and Sakaguchi (2026) study the analogous sample-size question for test-and-roll experiments, in which the experiment selects the treatment deployed to the remaining population and the objective is total welfare rather than estimator precision. Their worst-case marginal-benefit criterion sizes the experiment at roughly one third of the population, a constant share rather than the vanishing share implied by Theorem D.1.

D.3 Operational Planning Checks

Activation. A pilot can improve the assignment only if CMR can respond to it. With $m/2$ observations per arm and outcomes in $[0,1]$, the interval half-width and the largest reportable standard deviation are

$$\eta_m(\alpha) = \left\{ \frac{2\log(4/\alpha)}{m/2-1} \right\}^{1/2}, \qquad \hat{\sigma}_{\max}(M) = \left\{ \frac{\lfloor M^2/4 \rfloor}{M(M-1)} \right\}^{1/2},$$

and the first pilot size at which CMR can possibly leave balance is

$$m^{\text{act}}(\alpha) = \min\left\{ m \in \mathcal{M} : \eta_m(\alpha) < \hat{\sigma}_{\max}(m/2) \right\}, \tag{D.7}$$

with $m^{\text{act}}(\alpha) = \infty$ if the set is empty. Below the threshold, every realized confidence interval covers all standard deviations the model allows, the rectangle equals the full variance space $[0, 1/4]^2$, and CMR returns balance with certificate $1/4$, exactly as with no pilot. With all even pilot sizes admissible, $m^{\text{act}}(0.05) = 72$. With Bernoulli outcomes, Appendix C.1 shows that the exact-inversion folded-binomial rectangle lowers the threshold to $m = 4$, two observations per arm, for every $\alpha \in (0, 1)$. Below the threshold a pilot is harmless under the pooled design, where it reproduces the no-pilot experiment, and strictly wasteful under design-only accounting, so at small budgets with continuous outcomes balance remains the default and the binary construction is how small pilots earn their keep.

Viability. Under design-only accounting, combining the activation threshold with the share cap (D.2) confines defensible pilots to the band

$$m^{\text{act}}(\alpha) \;\le\; m \;\le\; n \cdot \frac{(\sigma_1 - \sigma_0)^2}{2(\sigma_1^2 + \sigma_0^2)},$$

and even a nonempty band is necessary rather than sufficient, since the cap credits the pilot with reaching the Neyman allocation exactly. The band is often narrow, and can be empty. With $n = 1{,}000$, $\alpha = 0.05$, and a two-to-one ratio of planning standard deviations, the band is $72 \le m \le 100$ and the default $m_n^\circ = 100$ sits at its upper endpoint, while at $n = 600$ the same band is empty and $m_n^\circ = 72$ exceeds the design-only cap. Under the pooled design the upper limit disappears. A pilot must still clear the activation threshold, and it improves precision only when its expected assignment regret falls below the balance regret, as in (D.4), a comparison that depends on the entire distribution of pilot realizations and cannot be guaranteed population by population.

Three questions. Three questions organize the planning decision. First, are the pilot and main wave comparable enough, drawn from the same population with the same arms and outcome, for a wave-stratified estimator to use both? A yes selects the pooled accounting, a no the conservative one. Second, does the planned pilot clear the relevant activation threshold? Third, given plausible planning values for the outcome variances, can the learned assignment reasonably be expected to improve on balance by enough to repay the pilot's cost? When these checks support piloting but no sharper loss calculation is credible, $2\lceil n^{2/3}/2 \rceil$ is the default, checked against the band above. Otherwise balance remains the safer assignment choice, even when a pilot is valuable for testing instruments, logistics, and procedures.

E Additional Proofs

E.1 Proof of the Coverage Lemma

Proof of Lemma 4.1. Fix $F \in \mathcal{F}$ and an arm $d \in \{0, 1\}$. We first prove the two one-sided statements for $b \in (0, 1)$. Conditional on the pilot assignment vector, the M_d outcomes observed in arm d are i.i.d. draws from the marginal distribution of $Y(d)$ under F. Denote these observations by $Y_{d1}, \ldots, Y_{dM_d}$.

Maurer and Pontil's empirical Bernstein inequality is stated in terms of the pairwise variance statistic $V_{M_d} = \frac{1}{M_d(M_d-1)} \sum_{1 \le i < j \le M_d} (Y_{di} - Y_{dj})^2$. This statistic is exactly the usual unbiased sample variance. Indeed, for any real numbers $x_1, \ldots, x_M$ with $\overline{x} = M^{-1} \sum_{i=1}^{M} x_i$, $\sum_{1 \le i < j \le M} (x_i - x_j)^2 = M \sum_{i=1}^{M} (x_i - \overline{x})^2$. Applying this identity with $M = M_d$ gives $V_{M_d} = \frac{1}{M_d - 1} \sum_{i=1}^{M_d} (Y_{di} - \overline{Y}_d)^2 = \hat{\sigma}_d^2(\omega)$. Moreover, $\mathbb{E}_{P_F}[V_{M_d}] = \sigma_d^2(F)$, because the unbiased sample variance has expectation equal to the population variance.

By Theorem 10 of Maurer and Pontil (2009), for every $b \in (0, 1)$,

$$\Pr_{P_F}\left(\sqrt{\mathbb{E}_{P_F}[V_{M_d}]} > \sqrt{V_{M_d}} + \sqrt{\frac{2\log(1/b)}{M_d - 1}} \right) \le b,$$

and

$$\Pr_{P_F}\left(\sqrt{V_{M_d}} > \sqrt{\mathbb{E}_{P_F}[V_{M_d}]} + \sqrt{\frac{2\log(1/b)}{M_d - 1}} \right) \le b.$$

Since $\sqrt{\mathbb{E}_{P_F}[V_{M_d}]} = \sigma_d(F)$ and $\sqrt{V_{M_d}} = \sqrt{\hat{\sigma}_d^2(\omega)}$, the first inequality is

$$\Pr_{P_F}\left(\sigma_d(F) > \sqrt{\hat{\sigma}_d^2(\omega)} + \sqrt{\frac{2\log(1/b)}{M_d - 1}}\right) \le b.$$

The only additional step is the projection of the reported variance endpoints onto the maintained variance space $[0, 1/4]$. Because $\sigma_d^2(F) \le 1/4$, the event $\overline{\sigma}_d^2(b;\omega) < \sigma_d^2(F)$ can occur only if the unprojected upper endpoint also falls below $\sigma_d^2(F)$, that is, only if $\left(\sqrt{\hat{\sigma}_d^2(\omega)} + \sqrt{\frac{2\log(1/b)}{M_d-1}}\right)^2 < \sigma_d^2(F)$. All quantities are nonnegative, so this last event is equivalent to $\sqrt{\hat{\sigma}_d^2(\omega)} + \sqrt{\frac{2\log(1/b)}{M_d-1}} < \sigma_d(F)$. Therefore, $\Pr_{P_F}\left(\overline{\sigma}_d^2(b;\omega) < \sigma_d^2(F)\right) \le b$. Next, the second Maurer–Pontil inequality gives $\Pr_{P_F}\left(\sqrt{\hat{\sigma}_d^2(\omega)} > \sigma_d(F) + \sqrt{\frac{2\log(1/b)}{M_d-1}}\right) \le b$. If $\underline{\sigma}_d^2(b;\omega) > \sigma_d^2(F)$, then projection of the reported lower endpoint cannot be responsible for the failure. Since $\underline{\sigma}_d^2(b;\omega) \le \left(\sqrt{\hat{\sigma}_d^2(\omega)} - \sqrt{2\log(1/b)/(M_d-1)}\right)_+^2$, we must have $\left(\sqrt{\hat{\sigma}_d^2(\omega)} - \sqrt{\frac{2\log(1/b)}{M_d-1}}\right)_+^2 > \sigma_d^2(F)$. Again all quantities are nonnegative, so this implies $\left(\sqrt{\hat{\sigma}_d^2(\omega)} - \sqrt{\frac{2\log(1/b)}{M_d-1}}\right)_+ > \sigma_d(F)$. Since $\sigma_d(F) \ge 0$, it follows that $\sqrt{\hat{\sigma}_d^2(\omega)} > \sigma_d(F) + \sqrt{\frac{2\log(1/b)}{M_d-1}}$. Thus, $\Pr_{P_F}\left(\underline{\sigma}_d^2(b;\omega) > \sigma_d^2(F)\right) \le b$.

For $b = 0$, the conventions $\underline{\sigma}_d^2(0;\omega) = 0$ and $\overline{\sigma}_d^2(0;\omega) = \frac{1}{4}$ make both one-sided failures impossible, because $\sigma_d^2(F) \in [0, 1/4]$. Hence the two inequalities also hold for $b = 0$.

It remains to prove the joint coverage statement for $\widehat{\Theta}_\alpha(\omega)$. Since $\alpha \in (0,1)$, the one-sided inequalities above apply with $b = \alpha/4$. The event that the true variance pair is not contained in $\widehat{\Theta}_\alpha(\omega)$ is contained in the union of the four one-sided failure events:

$$\begin{aligned}\left\{\theta(F) \notin \widehat{\Theta}_\alpha(\omega)\right\} \subseteq\ & \left\{\underline{\sigma}_1^2(\alpha/4;\omega) > \sigma_1^2(F)\right\} \cup \left\{\overline{\sigma}_1^2(\alpha/4;\omega) < \sigma_1^2(F)\right\} \\ & \cup \left\{\underline{\sigma}_0^2(\alpha/4;\omega) > \sigma_0^2(F)\right\} \cup \left\{\overline{\sigma}_0^2(\alpha/4;\omega) < \sigma_0^2(F)\right\}.\end{aligned}$$

By the union bound, $\Pr_{P_F}\left(\theta(F) \notin \widehat{\Theta}_\alpha(\omega)\right) \le \frac{\alpha}{4}+\frac{\alpha}{4}+\frac{\alpha}{4}+\frac{\alpha}{4} = \alpha$. Therefore, $\Pr_{P_F}\left(\theta(F) \in \widehat{\Theta}_\alpha(\omega)\right) \ge 1-\alpha$. $\square$

E.2 Proofs for Additional Extensions

Proof of Proposition C.1. Step 1: stochastic ordering of the folded count. The map $\xi : [0, 1/4] \to [0, 1/2]$, $\xi(\sigma^2) = \left(1 - \sqrt{1-4\sigma^2}\right)/2$, is continuous and strictly increasing, and under success probability $q = \xi(\sigma^2)$ the folded count satisfies $\{J_d \ge j\} = \{j \le X_d \le M_d - j\}$ with $X_d \sim \text{Binomial}(M_d, q)$. It therefore suffices to show that, for each fixed $j \in \mathcal{J}_{M_d}$, the function $\varphi_j(q) = \Pr_q(j \le X_d \le M_d - j)$ is nondecreasing on $[0, 1/2]$; continuity is immediate because φ_j is a polynomial in q. For $j = 0$, $\varphi_0 \equiv 1$. Fix $1 \le j \le \lfloor M_d/2 \rfloor$ and write $M = M_d$. Differentiating the binomial survival function term by term and telescoping gives, for $k \ge 1$, $\frac{d}{dq} \Pr_q(X_d \ge k) = M\binom{M-1}{k-1} q^{k-1}(1-q)^{M-k}$. Since $\varphi_j(q) = \Pr_q(X_d \ge j) - \Pr_q(X_d \ge M-j+1)$ and $\binom{M-1}{M-j} = \binom{M-1}{j-1}$, $\varphi_j'(q) = M\binom{M-1}{j-1}[q(1-q)]^{j-1}\left[(1-q)^{M-2j+1} - q^{M-2j+1}\right]$. For $q \in [0, 1/2]$ and $j \le M/2$, the exponent satisfies $M-2j+1 \ge 1$ and $1-q \ge q$, so the bracket is nonnegative and φ_j is nondecreasing on $[0, 1/2]$. Composing with ξ shows that $G_{\ge}(j; \sigma^2) = \varphi_j(\xi(\sigma^2))$ is continuous and nondecreasing in σ^2, and $G_{\le}(j; \sigma^2) = 1 - G_{\ge}(j+1; \sigma^2)$, with the convention $G_{\ge}(\lfloor M_d/2 \rfloor + 1; \cdot) \equiv 0$, is continuous and nonincreasing. Finally, since $G_{\le}(j; \sigma^2) + G_{\ge}(j; \sigma^2) = 1 + \Pr_{\sigma^2}(J_d = j) \ge 1 > 2b$, no candidate variance is rejected by both tails, and the reported endpoints satisfy $\underline{\sigma}^2_{M_d}(j; b) \le \overline{\sigma}^2_{M_d}(j; b)$ in every case.

Step 2: one-sided coverage. Fix $F \in \mathcal{F}^{\text{bin}}$ with arm-d success probability q_d and variance $v_d = \sigma_d^2(F) = q_d(1-q_d) \in [0, 1/4]$. Under P_F, the folded count J_d has the folded-binomial distribution indexed by v_d, because q_d and $1 - q_d$ induce the same distribution for J_d.

Upper endpoint. By Step 1, $G_{\le}(j; \cdot)$ is nonincreasing, so the acceptance set $A_j = \{\sigma^2 \in [0, 1/4] : G_{\le}(j; \sigma^2) > b\}$ is an interval containing 0, with supremum $\overline{\sigma}^2_{M_d}(j; b)$. If $\overline{\sigma}^2_{M_d}(J_d; b) < v_d$, then $v_d \notin A_{J_d}$, so $G_{\le}(J_d; v_d) \le b$. It remains to bound $\Pr_{P_F}\{G_{\le}(J_d; v_d) \le b\}$. If no $j \in \mathcal{J}_{M_d}$ satisfies $G_{\le}(j; v_d) \le b$, this probability is zero. Otherwise, since $j \mapsto G_{\le}(j; v_d)$ is nondecreasing, the set of such j is $\{0, \ldots, j^*\}$ with $j^* = \max\{j : G_{\le}(j; v_d) \le b\}$, and $\Pr_{P_F}\{G_{\le}(J_d; v_d) \le b\} = \Pr_{P_F}\{J_d \le j^*\} = G_{\le}(j^*; v_d) \le b$. Hence $\Pr_{P_F}\{\sigma_d^2(F) \le \overline{\sigma}^2_{M_d}(J_d; b)\} \ge 1 - b$.

Lower endpoint. Let $B_j = \{\sigma^2 \in [0, 1/4] : G_{\ge}(j; \sigma^2) > b\}$. By Step 1, $G_{\ge}(j; \cdot)$ is nondecreasing, so B_j is an interval ending at $1/4$ when nonempty, with infimum $\underline{\sigma}^2_{M_d}(j; b)$. Suppose $\underline{\sigma}^2_{M_d}(J_d; b) > v_d$. If B_{J_d} is nonempty, then $v_d < \inf B_{J_d}$, so $v_d \notin B_{J_d}$ and $G_{\ge}(J_d; v_d) \le b$. If B_{J_d} is empty, then $G_{\ge}(J_d; \sigma^2) \le b$ for every σ^2, in particular at v_d. Either way the failure event is contained in $\{G_{\ge}(J_d; v_d) \le b\}$. Since $j \mapsto G_{\ge}(j; v_d)$ is nonincreasing, the set $\{j : G_{\ge}(j; v_d) \le b\}$, when nonempty, is $\{j^*, \ldots, \lfloor M_d/2 \rfloor\}$ with $j^* = \min\{j : G_{\ge}(j; v_d) \le b\}$, and $\Pr_{P_F}\{J_d \ge j^*\} = G_{\ge}(j^*; v_d) \le b$. Hence $\Pr_{P_F}\{\underline{\sigma}^2_{M_d}(J_d; b) \le \sigma_d^2(F)\} \ge 1 - b$. Both inequalities hold for every $q_d \in [0, 1]$ and depend on F only through the arm-d marginal, so the endpoints satisfy the one-sided coverage requirement of Subsection 4.1 restricted to $\mathcal{F}^{\text{bin}}$.

Step 3: the rectangle and the certificate. The rectangle misses $\theta(F)$ only if at least one of the four one-sided bounds at level $\alpha/4$ fails, so by Step 2 and the union bound, which does not require independence across arms, the miss probability is at most α. By Step 1, $G_{\le}(j; \cdot)$ is continuous with $G_{\le}(j; 0) = 1 > b$, so A_j contains an interval $[0, \varepsilon)$ with $\varepsilon > 0$

and $\overline{\sigma}^2_{M_d}(j;b) > 0$ for every $j \in \mathcal{J}_{M_d}$. Both upper endpoints of $\widehat{\Theta}^{\mathrm{bin}}_{\alpha}$ are therefore strictly positive, and Proposition 4.1 yields the unique interior assignment $p^{\mathrm{bin}}_{\mathrm{CMR}}$ and the certificate $U^{\mathrm{bin}}_{\mathrm{CMR}}$. The certificate guarantee follows from the coverage just established by the same coverage-event argument as in the proof of Theorem 4.1(i). Finally, $U^{\mathrm{bin}}_{\mathrm{CMR}} \le 1/4$ follows as in the proof of Theorem 4.1(ii), since $\widehat{\Theta}^{\mathrm{bin}}_{\alpha} \subseteq [0,1/4]^2$. □

Proof of Proposition C.2. Under the fixed-arm pilot design, the arm-d outcomes are M_d i.i.d. draws from the arm-d marginal of F (Subsection 2.3). Let $\varepsilon_d \in (0,1/2]$ solve $\varepsilon(1-\varepsilon) = 1/(\psi_d+3)$, which exists since $\psi_d \ge 1$ implies $1/(\psi_d+3) \le 1/4$. We first show that $M_d < (\psi_d+3)\log(1/\alpha)/4$ implies $(1-\varepsilon_d)^{M_d} > \alpha$. Since $\varepsilon_d \le 1/2$, we have $1-\varepsilon_d \ge 1/2$ and hence $\varepsilon_d = 1/[(\psi_d+3)(1-\varepsilon_d)] \le 2/(\psi_d+3)$. Combining with $\log(1/(1-\varepsilon)) \le 2\varepsilon$ for $\varepsilon \in (0,1/2]$ gives $M_d \log\left(\frac{1}{1-\varepsilon_d}\right) < \frac{(\psi_d+3)\log(1/\alpha)}{4} \cdot \frac{4}{\psi_d+3} = \log(1/\alpha)$, which is the claim.

Fix a constant $c \in \mathbb{R}$ and suppose instead that $\overline{\sigma}^2_d = B < \infty$ at the realization in which all arm-d pilot outcomes equal c. For $a > 0$, let G^c_a place mass $1-\varepsilon_d$ at c and mass ε_d at $c+a$, so that its variance is $\varepsilon_d(1-\varepsilon_d)a^2 > 0$. Direct computation gives its kurtosis as $\frac{\varepsilon_d^3+(1-\varepsilon_d)^3}{\varepsilon_d(1-\varepsilon_d)} = \frac{1}{\varepsilon_d(1-\varepsilon_d)} - 3 = \psi_d$, free of a and of c, since kurtosis is invariant to location and scale. Let F have arm-d marginal G^c_a with a large enough that $\varepsilon_d(1-\varepsilon_d)a^2 > B$, and any arm-$(1-d)$ marginal satisfying Assumption C.1, so that F obeys the assumption. Under P_F, all arm-d pilot outcomes equal c with probability $(1-\varepsilon_d)^{M_d} > \alpha$, and on that event $\overline{\sigma}^2_d = B < \sigma^2_d(F)$. The miss probability therefore exceeds α, contradicting the coverage requirement. Since c was arbitrary, the bound is infinite at every constant realization. □

Proof of Lemma C.1. Write $W_{di} = \frac{1}{2}\left(Y_{d,2i-1} - Y_{d,2i}\right)^2$, $i = 1,\ldots,\lfloor M_d/2 \rfloor$, for the halved squared differences that form the estimator, discarding at most one observation. Nonnegativity is immediate, since each W_{di} is nonnegative and the median of nonnegative block means is nonnegative. Fix F satisfying Assumption C.1 and write $\sigma^2 = \sigma^2_d(F)$ and $\mu_4 = \mathbb{E}_F[(Y(d)-\mu_d(F))^4] \le \psi_d \sigma^4$. Under the fixed-arm pilot design, the paired differences are built from disjoint pairs of i.i.d. draws, so the W_{di} are i.i.d. across i.

First, the moments of W_{di}. Writing A and B for two independent copies of $Y(d) - \mu_d(F)$, the difference $Y_{d,2i-1} - Y_{d,2i}$ is distributed as $A - B$, so $\mathbb{E}[W_{di}] = \mathbb{E}[(A-B)^2]/2 = \sigma^2$, and, expanding and using independence with $\mathbb{E}[A] = \mathbb{E}[B] = 0$, $\mathbb{E}[(A-B)^4] = 2\mu_4 + 6\sigma^4$. Hence $\mathbb{E}[W^2_{di}] = \mu_4/2 + 3\sigma^4/2$ and $\mathrm{Var}(W_{di}) = \frac{\mu_4+\sigma^4}{2} \le \frac{(\psi_d+1)\,\sigma^4}{2}$.

Second, the block means. Each block mean $\bar{W}_j$ averages b_d i.i.d. pairs, so $\mathrm{Var}(\bar{W}_j) \le (\psi_d+1)\sigma^4/(2b_d)$. With $\rho^2_d = 2(\psi_d+1)/b_d$, Chebyshev's inequality gives $p_0 = \Pr\left(|\bar{W}_j - \sigma^2| > \rho_d\,\sigma^2\right) \le \frac{(\psi_d+1)\sigma^4/(2b_d)}{\rho^2_d\,\sigma^4} = \frac{(\psi_d+1)}{2b_d\rho^2_d} = \frac{1}{4}$.

Third, the median. At least half of the block means lie weakly above the median and at least half lie weakly below it, so if $|\widehat{v}_d - \sigma^2| > \rho_d\sigma^2$, at least $\lceil k_d/2\rceil$ block means must deviate from σ^2 by more than $\rho_d\sigma^2$. The indicators of these deviations are i.i.d. Bernoulli with success probability $p_0 \le 1/4$, so by Hoeffding's inequality (Hoeffding, 1963),

$$\Pr\left(|\widehat{v}_d - \sigma^2| > \rho_d\,\sigma^2\right) \le \Pr\left(\sum_{j=1}^{k_d} \mathbf{1}\{|\bar{W}_j - \sigma^2| > \rho_d\sigma^2\} \ge \frac{k_d}{2}\right)$$
$$\le \exp\left(-2k_d\left(\tfrac{1}{2} - \tfrac{1}{4}\right)^2\right) = e^{-k_d/8} \le \frac{\alpha}{2},$$

using $k_d = \lceil 8\log(2/\alpha)\rceil$. This is the claimed deviation bound. □

Proof of Proposition C.3. For each arm d, let $E_d = \{\omega : |\widehat{v}_d(\omega) - \sigma_d^2(F)| \le \rho_d\,\sigma_d^2(F)\}$ and $E = E_1 \cap E_0$. By Lemma C.1 and the union bound, $\Pr_{P_F}(E) \ge 1-\alpha$.

(i) Fix $\omega \in E$. Since $\sigma_d^2(F) > 0$ by Assumption C.1 and $\rho_d < 1$, the deviation inequality gives $\widehat{v}_d(\omega) \ge (1-\rho_d)\sigma_d^2(F) > 0$, so the rule activates. The same inequality inverts to $\underline{\sigma}_d^2(\omega) \le \sigma_d^2(F) \le \overline{\sigma}_d^2(\omega)$ for each arm, so $\theta(F) \in \widehat{\Theta}_\alpha(\omega)$. The endpoints in (C.1) are positive and finite, so Proposition 4.1 delivers an interior assignment and a finite certificate, and since the certificate is the supremum of regret over a set containing $\theta(F)$, the event-wise argument from the proof of Theorem 4.1(i) gives $r(p_{\mathrm{CMR}}(\omega), \theta(F)) \le U_{\mathrm{CMR}}(\omega)$. Hence E is contained in the event of the display.

(ii) Write $\pi^* = \pi^*(\theta(F))$ and set $c_\theta = \min\{1/4, \min\{\pi^*, 1-\pi^*\}/2\}$. Fix $\omega \in E$ and suppose $\rho_1 + \rho_0 \le c_\theta$, so $\rho_d \le 1/4$ for each arm. On E, $\sqrt{\widehat{v}_d(\omega)} \in \sigma_d(F)\,[\sqrt{1-\rho_d}, \sqrt{1+\rho_d}]$, so the standard-deviation endpoints satisfy

$$\underline{\sigma}_d, \overline{\sigma}_d \in \sigma_d(F)\left[\sqrt{\tfrac{1-\rho_d}{1+\rho_d}}, \sqrt{\tfrac{1+\rho_d}{1-\rho_d}}\right] \subseteq \sigma_d(F)\,[1-2\rho_d,\, 1+2\rho_d],$$

where the inclusion holds for $\rho_d \le 1/4$ by squaring. The midpoint $s_d = (\underline{\sigma}_d + \overline{\sigma}_d)/2$ therefore satisfies $s_d = \sigma_d(F)(1+e_d)$ with $|e_d| \le 2\rho_d$. Since $p_{\mathrm{CMR}} = s_1/(s_1+s_0)$, $p_{\mathrm{CMR}} - \pi^* = \frac{\sigma_1\sigma_0\,(e_1-e_0)}{(s_1+s_0)(\sigma_1+\sigma_0)}$, and combining $\sigma_1\sigma_0 \le (\sigma_1+\sigma_0)^2/4$, $|e_1 - e_0| \le 2(\rho_1+\rho_0)$, and $s_1 + s_0 \ge (\sigma_1+\sigma_0)/2$ yields $|p_{\mathrm{CMR}} - \pi^*| \le \rho_1 + \rho_0$.

For the regret bound, $\rho_1 + \rho_0 \le \min\{\pi^*, 1-\pi^*\}/2$ and the assignment bound give $p_{\mathrm{CMR}} \ge \pi^*/2$ and $1 - p_{\mathrm{CMR}} \ge (1-\pi^*)/2$, so $p_{\mathrm{CMR}}(1-p_{\mathrm{CMR}}) \ge \pi^*(1-\pi^*)/4$, and the regret identity (2.2) gives $r(p_{\mathrm{CMR}}, \theta(F)) \le 4(\sigma_1+\sigma_0)^2(\rho_1+\rho_0)^2/[\pi^*(1-\pi^*)]$. For the certificate, the endpoint bounds give $\overline{\sigma}_1\overline{\sigma}_0 - \underline{\sigma}_1\underline{\sigma}_0 \le \sigma_1\sigma_0\,[(1+2\rho_1)(1+2\rho_0) - (1-2\rho_1)(1-2\rho_0)] = 4\sigma_1\sigma_0(\rho_1+\rho_0)$, and $(\underline{\sigma}_1 + \overline{\sigma}_1)(\underline{\sigma}_0 + \overline{\sigma}_0) \ge 4\sigma_1\sigma_0(1-2\rho_1)(1-2\rho_0) \ge \sigma_1\sigma_0$, so the closed form (4.5) gives $U_{\mathrm{CMR}} \le 16\sigma_1\sigma_0(\rho_1+\rho_0)^2$. Taking $C_\theta = \max\{4(\sigma_1+\sigma_0)^2/[\pi^*(1-\pi^*)],\, 16\sigma_1\sigma_0\}$ proves part (ii).

(iii) With α fixed, the block count k_d is fixed while $b_d \to \infty$, so $\rho_d = \sqrt{2(\psi_d + 1)/b_d} = O(M_d^{-1/2})$ deterministically, and $\rho_d < 1$ eventually, so the planning condition holds for all large pilots. Each block mean satisfies $\bar{W}_j - \sigma_d^2 = O_p(b_d^{-1/2}) = O_p(M_d^{-1/2})$ by Chebyshev's inequality, and since the median lies weakly between the smallest and largest of the k_d block means and k_d is fixed, $|\widehat{v}_d - \sigma_d^2| \leq \max_{1 \leq j \leq k_d} |\bar{W}_j - \sigma_d^2| = O_p\Big(M_d^{-1/2}\Big)$. In particular $\widehat{v}_d \xrightarrow{p} \sigma_d^2 > 0$, so $\Pr_{P_F}(\widehat{v}_d = 0 \text{ for some } d) \to 0$ and the rule activates with probability tending to one. Because the rule keeps balance and reports no certificate only on this vanishing-probability event, it suffices to establish the three O_p statements when the rule activates. There, $|\sqrt{\widehat{v}_d} - \sigma_d| = |\widehat{v}_d - \sigma_d^2|/(\sqrt{\widehat{v}_d} + \sigma_d) = O_p(M_d^{-1/2})$, and since $(1 \pm \rho_d)^{-1/2} = 1 \mp \rho_d/2 + O(\rho_d^2)$ with $\rho_d = O(M_d^{-1/2})$, the endpoints (C.1) satisfy $|\underline{\sigma}_d - \sigma_d| + |\overline{\sigma}_d - \sigma_d| = O_p(M_d^{-1/2})$ for each arm. This is the same endpoint contraction as in the proof of Theorem 4.2, and the arguments given there for the assignment and regret rates use only this contraction, the closed form (4.4), the positivity of σ_1 and σ_0, and the regret identity (2.2), so they apply verbatim and give the first two displayed rates. The certificate numerator bound used there relies on endpoints lying in $[0, 1/2]$ and must be replaced. The gap in the closed form (4.5) satisfies

$$\overline{\sigma}_1 \overline{\sigma}_0 - \underline{\sigma}_1 \underline{\sigma}_0 = \overline{\sigma}_0 \left(\overline{\sigma}_1 - \underline{\sigma}_1\right) + \underline{\sigma}_1 \left(\overline{\sigma}_0 - \underline{\sigma}_0\right) = O_p\Big(M_1^{-1/2} + M_0^{-1/2}\Big),$$

since $\overline{\sigma}_0$ and $\underline{\sigma}_1$ converge in probability to σ_0 and σ_1 by the contraction, and the certificate numerator is the square of this gap. The denominator converges in probability to $4\sigma_1\sigma_0 > 0$, again by the contraction. Hence $U_{\mathrm{CMR}} = O_p(M_1^{-1} + M_0^{-1})$, completing part (iii). □

Proof of Proposition C.4. The construction involves $4K$ one-sided bounds, two per arm–outcome pair, each with failure probability at most $\alpha/(4K)$, so by the union bound they hold simultaneously with probability at least $1 - \alpha$. On that event, weighted sums with nonnegative weights preserve the per-outcome inequalities, so $\underline{\sigma}_{d,w}^2(\omega) \leq \sigma_{d,w}^2 \leq \overline{\sigma}_{d,w}^2(\omega)$ for each arm and $\theta_w \in \widehat{\Theta}_\alpha(\omega)$. The Maurer–Pontil upper endpoints are strictly positive at every realization and the weights sum to one, so $\overline{\sigma}_{d,w}^2(\omega) > 0$ for each arm and Proposition 4.1 applies. The certificate guarantee follows from the same coverage-event argument as in the proof of Theorem 4.1(i). □

Proof of Proposition C.5. Fix F and an arm d. For the upper endpoint, let F' have the same joint distribution of $(S(1), S(0))$ as F and, independently of $(S(1), S(0))$, let $Y(d)$ equal 0 or 1 with probability one half each, so $\sigma_d^2(F') = 1/4$. The pilot records only short-run outcomes, so its distribution depends on the state only through the joint distribution of $(S(1), S(0))$, and $P_{F'} = P_F$. Since the endpoints take values in $[0, 1/4]$ by the coverage requirements, one-sided coverage applied at F' gives $\Pr_{P_F}\big(\overline{\sigma}_d^2(\omega) = 1/4\big) = \Pr_{P_{F'}}\big(\sigma_d^2(F') \leq \overline{\sigma}_d^2(\omega)\big) \geq 1 - \alpha/4$. For the lower endpoint, let F'' agree with F on $(S(1), S(0))$ and set $Y(d) \equiv 1/2$, so $\sigma_d^2(F'') = 0$, and the same argument gives $\Pr_{P_F}\big(\underline{\sigma}_d^2(\omega) = 0\big) \geq 1 - \alpha/4$. The union bound over the four endpoints gives $\Pr_{P_F}\big(\widehat{\Theta}_\alpha(\omega) = [0, 1/4]^2\big) \geq 1 - \alpha$. □

Proof of Proposition C.6. On the event that the four one-sided Maurer–Pontil bounds for the short-run standard deviations hold, which has probability at least $1-\alpha$ by the union bound, $\underline{\sigma}_{d,S}(\omega) \leq \sigma_{d,S}(F) \leq \overline{\sigma}_{d,S}(\omega)$ for each arm. Assumption C.2 gives $\sigma_{d,S}(F) - \zeta_d \leq \sigma_d(F) \leq \sigma_{d,S}(F) + \zeta_d$, hence $\underline{\sigma}_{d,S}(\omega) - \zeta_d \leq \sigma_d(F) \leq \overline{\sigma}_{d,S}(\omega) + \zeta_d$. Since $\sigma_d(F) \in [0, 1/2]$, applying the positive part to the lower endpoint and the cap to the upper preserves the bracketing, and squaring gives $\theta(F) \in \widehat{\Theta}_\alpha(\omega)$. The upper endpoints are at least the Maurer–Pontil upper endpoints, hence strictly positive, so Proposition 4.1 applies, and the certificate guarantee follows from the same coverage-event argument as in the proof of Theorem 4.1(i). □

E.3 Proofs for Pilot-Size Results

Proof of Theorem D.1. Throughout, write $R_m^{\mathrm{mmr}} := R_{m/2,m/2}^{\mathrm{mmr}}$ for the minimax-regret value of the balanced size-m pilot. Proposition 3.4 with $M_1 = M_0 = m/2$ gives

$$R_m^{\mathrm{mmr}} \;\geq\; 2\sqrt{2}\, c\, m^{-1/2}, \tag{E.1}$$

and Theorem 4.3 with the same arm sizes gives

$$\sup_{F \in \mathcal{F}} R_m\big(p_m^{\mathrm{CMR}}, F\big) \;\leq\; 2\sqrt{2}\, C_\alpha\, m^{-1/2}. \tag{E.2}$$

Because outcomes lie in $[0, 1]$, $\sigma_d(F) \in [0, 1/2]$ and $(\sigma_1 - \sigma_0)^2 \leq 1/4$, with equality at any $F^A \in \mathcal{F}$ whose treated marginal is Bernoulli(1/2) and whose control marginal is degenerate, so that $\theta(F^A) = (1/4, 0)$.

Step 1: floors at a fixed m. Fix $m \in \mathcal{M} \setminus \{0\}$ and $\delta \in \mathcal{D}_0(m)$. Evaluating the pooled loss at F^A, where the balance regret equals $1/4$ and $R_m(\delta, F^A) \geq 0$, and separately dropping the nonnegative pilot term and using $\sup_{F \in \mathcal{F}} R_m(\delta, F) \geq R_m^{\mathrm{mmr}}$ together with (E.1), gives

$$\sup_{F \in \mathcal{F}} \mathcal{L}_m^{\mathrm{pool}}(\delta, F) \;\geq\; \max\left\{ \frac{m}{4n^2},\; \frac{2\sqrt{2}\, c\,(n - m)}{n^2\sqrt{m}} \right\}. \tag{E.3}$$

Step 2: CMR at a fixed m. Bounding the balance regret by $1/4$ and applying (E.2),

$$\sup_{F \in \mathcal{F}} \mathcal{L}_m^{\mathrm{pool}}\big(p_m^{\mathrm{CMR}}, F\big) \;\leq\; \frac{m}{4n^2} + \frac{2\sqrt{2}\, C_\alpha\,(n - m)}{n^2\sqrt{m}}. \tag{E.4}$$

Step 3: feasibility and size of m_n°. Fix $n \geq 8$ and set $x = n^{2/3}$, so that $x \geq 4$. The smallest even integer no smaller than x is $2\lceil x/2 \rceil$, and $x \leq m_n^\circ < x + 2 \leq 2x$. Since $n - n^{2/3} \geq 4$ for $n \geq 8$, also $m_n^\circ < x + 2 \leq n - 2$. Hence $4 \leq m_n^\circ \leq n - 2$, so $m_n^\circ \in \mathcal{M}$, which justifies the identity in (D.5). Also $n - m_n^\circ > n - x - 2 \geq n/4$, since $(3/4)n - n^{2/3} \geq 2$ for $n \geq 8$.

Step 4: upper bound for the explicit plan. Apply (E.4) at $m = m_n^\circ$. Using $m_n^\circ \leq 2n^{2/3}$, $\sqrt{m_n^\circ} \geq n^{1/3}$, and $n - m_n^\circ \leq n$,

$$\sup_{F \in \mathcal{F}} \mathcal{L}^{\text{pool}}_{m_n^\circ}\big(p^{\text{CMR}}_{m_n^\circ}, F\big) \;\leq\; \frac{2n^{2/3}}{4n^2} + \frac{2\sqrt{2}\, C_\alpha\, n}{n^2\, n^{1/3}} \;=\; \Big(\tfrac{1}{2} + 2\sqrt{2}\, C_\alpha\Big) n^{-4/3} \;=\; C_{1,\alpha}\, n^{-4/3}.$$

The plan $(m_n^\circ, p^{\text{CMR}}_{m_n^\circ})$ is feasible, so $\mathcal{L}_n^\star$ is bounded by its worst-case loss, which proves the two right-hand inequalities in (D.6).

Step 5: floor for every plan. If $m = 0$, then $\sup_{F \in \mathcal{F}} \mathcal{L}_0(F) = 1/(4n) \geq \frac{1}{4} n^{-4/3}$. If $m \geq n^{2/3}$, the first term in (E.3) gives at least $n^{2/3}/(4n^2) = \frac{1}{4} n^{-4/3}$. If $0 < m < n^{2/3}$, then $n - m \geq n - n^{2/3} \geq n/2$ for $n \geq 8$, and the second term in (E.3) gives at least $(n/2) \cdot 2\sqrt{2}\, c\, m^{-1/2}/n^2 = \sqrt{2}\, c/(n\sqrt{m}) \geq \sqrt{2}\, c\, n^{-4/3}$, where the last inequality uses $m < n^{2/3}$. Hence every plan has worst-case loss at least $c_1 n^{-4/3}$ with $c_1 = \min\{1/4, \sqrt{2}\, c\}$, and minimizing over plans proves the left-hand inequality in (D.6).

Step 6: necessity of the $n^{2/3}$ scaling. Fix $Q < \infty$ and let the plans (m_n, δ_n) satisfy $\sup_{F \in \mathcal{F}} \mathcal{L}^{\text{pool}}_{m_n}(\delta_n, F) \leq Q\, n^{-4/3}$ for all large n, and consider any $n \geq \max\{8, 512Q^3\}$ at which the bound holds, so that $n^{1/3} \geq 8Q$. Since $1/(4n) > Q\, n^{-4/3}$ whenever $n^{1/3} > 4Q$, the no-pilot plan violates the bound, so $m_n > 0$ and (E.3) applies. Its first term gives $m_n/(4n^2) \leq Q\, n^{-4/3}$, hence $m_n \leq 4Q\, n^{2/3}$, and $4Q\, n^{2/3} \leq n/2$ because $n^{1/3} \geq 8Q$, so $n - m_n \geq n/2$. The second term then gives $\sqrt{2}\, c/(n\sqrt{m_n}) \leq Q\, n^{-4/3}$, hence $m_n \geq 2c^2 Q^{-2}\, n^{2/3}$. Therefore $2c^2 Q^{-2}\, n^{2/3} \leq m_n \leq 4Q\, n^{2/3}$ for all large n, which is the claim $m_n \asymp n^{2/3}$ with implied constants depending only on Q.

Step 7: design-only accounting. Fix $m \in \mathcal{M} \setminus \{0\}$ and $\delta \in \mathcal{D}_0(m)$, and recall the design-only loss $\mathcal{L}_m(\delta, F) = R_m(\delta, F)/(n - m) + V^*(\theta(F))\, m/[n(n - m)]$, where $V^*(\theta(F)) = (\sigma_1 + \sigma_0)^2 \in [0, 1]$ and $V^*(\theta(F^A)) = 1/4$. Using $1/(n - m) \geq 1/n$, evaluating at F^A and separately using $\sup_{F \in \mathcal{F}} R_m(\delta, F) \geq R_m^{\text{mmr}}$ with (E.1) gives

$$\sup_{F \in \mathcal{F}} \mathcal{L}_m(\delta, F) \;\geq\; \max\left\{ \frac{m}{4n^2},\, \frac{2\sqrt{2}\, c}{n\sqrt{m}} \right\},$$

which is at least the right-hand side of (E.3). Repeating Step 5 with this floor, and Step 6 with the second branch $2\sqrt{2}\, c/(n\sqrt{m})$ in place of $\sqrt{2}\, c/(n\sqrt{m})$, which no longer requires $n - m \geq n/2$, yields the same constant c_1 and the same necessity bounds under the design-only accounting. For the upper bound, Step 3 gives $n - m_n^\circ > n/4$, so with $V^* \le 1$, (E.2), $m_n^\circ \le 2n^{2/3}$, and $\sqrt{m_n^\circ} \ge n^{1/3}$,

$$\sup_{F \in \mathcal{F}} \mathcal{L}_{m_n^\circ}\big(p_{m_n^\circ}^{\mathrm{CMR}}, F\big) \;\le\; \frac{2\sqrt{2}\,C_\alpha}{\sqrt{m_n^\circ}\,\big(n - m_n^\circ\big)} + \frac{2n^{2/3}}{n\big(n - m_n^\circ\big)} \;\le\; \frac{8\sqrt{2}\,C_\alpha}{n^{4/3}} + \frac{8}{n^{4/3}} \;=\; C_{2,\alpha}\, n^{-4/3},$$

which proves the design-only chain with $C_{2,\alpha} = 8 + 8\sqrt{2}\,C_\alpha$. □

F Pilot Use in Economics RCTs

This appendix documents how the pilot-use statistics in Section 1 were constructed. The frame is the April 2026 monthly snapshot of the AEA RCT Registry (American Economic Association, 2026), restricted to the 10,905 trials first registered between 2016 and 2025. The registry has no structured field for pilot use, so I screen nine free-text fields for terms indicating a pilot, feasibility study, or multi-wave design, classify the high- and medium-confidence registrations with a large language model under a fixed rubric with structured output, and inspect by hand all 24 registrations classified as using pilot evidence to set the treatment-assignment probability. None of the 24 reports a feasible-Neyman or variance-proportional calculation, and pilot size is reported for 399 of the 991 pilots. Among these 399 pilots, the median is 300 observations, the interquartile range is 100 to 1,102, and approximately one tenth report fewer than 30 observations. Screening patterns, prompts, rubric, model identifier, and trial-level classifications are in the replication files.

Two checks support the classification. In 30 high-confidence registrations held out from those used to build the rubric, I agree with the model in 29 cases, a 95% Wilson interval of 83–99%. In 45 randomly drawn low- and no-confidence registrations I find no missed pilots, bounding the false-negative rate in those groups by about 8% at 95% confidence. Because informal piloting may go unmentioned in registry text, the reported counts are lower bounds.

References for the Online Appendix

American Economic Association. Registrations in the AEA RCT Registry. Harvard Dataverse, 2026. URL `https://doi.org/10.7910/DVN/PWNOZZ`. April 2026 monthly snapshot.

Susan Athey, Raj Chetty, Guido W. Imbens, and Hyunseung Kang. The surrogate index: Combining short-term proxies to estimate long-term treatment effects more rapidly and precisely. *The Review of Economic Studies*, 93(4):2284–2312, 2026. doi: 10.1093/restud/rdaf087.

Raghu R. Bahadur and Leonard J. Savage. The nonexistence of certain statistical procedures in nonparametric problems. *The Annals of Mathematical Statistics*, 27(4):1115–1122, 1956. doi: 10.1214/aoms/1177728077.

Raj Chetty, Nathaniel Hendren, and Lawrence F. Katz. The effects of exposure to better neighborhoods on children: New evidence from the Moving to Opportunity experiment. *American Economic Review*, 106(4):855–902, 2016.

C. J. Clopper and E. S. Pearson. The use of confidence or fiducial limits illustrated in the case of the binomial. *Biometrika*, 26(4):404–413, 1934. doi: 10.1093/biomet/26.4.404.

John J. Gart. A locally most powerful test for the symmetric folded binomial distribution. *Biometrics*, 26(1):129–138, 1970. doi: 10.2307/2529049.

Jinyong Hahn, Keisuke Hirano, and Dean Karlan. Adaptive experimental design using the propensity score. *Journal of Business & Economic Statistics*, 29(1):96–108, 2011. doi: 10.1198/jbes.2009.08161.

Wassily Hoeffding. Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association*, 58(301):13–30, 1963. doi: 10.1080/01621459.1963.10500830.

Kentaro Kawato and Shosei Sakaguchi. Prior-free sample size design for test-and-roll experiments, 2026. URL https://arxiv.org/abs/2605.02414.

Jeffrey R. Kling, Jeffrey B. Liebman, and Lawrence F. Katz. Experimental analysis of neighborhood effects. *Econometrica*, 75(1):83–119, 2007.

Erich Leo Lehmann and Joseph P Romano. *Testing statistical hypotheses*. Springer, third edition, 2005.

Gábor Lugosi and Shahar Mendelson. Mean estimation and regression under heavy-tailed distributions: A survey. *Foundations of Computational Mathematics*, 19(5):1145–1190, 2019. doi: 10.1007/s10208-019-09427-x.

Nathan Mantel. An alternative test for the symmetric folded binomial distribution. *Biometrics*, 26 (4):848–851, 1970. doi: 10.2307/2528731.

Andreas Maurer and Massimiliano Pontil. Empirical Bernstein bounds and sample variance penalization. In *Proceedings of the 22nd Annual Conference on Learning Theory*, 2009. URL https://arxiv.org/abs/0907.3740.

Jerzy Neyman. Outline of a theory of statistical estimation based on the classical theory of probability. *Philosophical Transactions of the Royal Society of London. Series A, Mathematical and Physical Sciences*, 236(767):333–380, 1937. doi: 10.1098/rsta.1937.0005.

Giovanni C. Porzio and Giancarlo Ragozini. On the stochastic ordering of folded binomials. *Statistics & Probability Letters*, 79(9):1299–1304, 2009. doi: 10.1016/j.spl.2009.01.021.

Max Tabord-Meehan. Stratification trees for adaptive randomisation in randomised controlled trials. *Review of Economic Studies*, 90(5):2646–2673, 2023. doi: 10.1093/restud/rdac083.