THE TRADE-OFF BETWEEN COORDINATION AND STABILITY IN RANDOMIZED DECISIONS

Anonymous authors
Paper under double-blind review

ABSTRACT

When a decision is implemented by a randomized rule, the rule usually commits to each individual's probability of an action but does not constrain how the actions are drawn together. For example, two members of one household can be served by independent coins, which often split them, or by one shared coin, which never does. Yet one coin shared across a batch also makes the number of actions taken, or a group's action rate, swing widely from one draw to the next. We ask for the best attainable balance between these two properties of the realized decisions over all joint distributions with the prescribed marginals. Stability is measured by the worst normalized variance of a bounded linear aggregate of the decisions. We show that this coordination–stability frontier is not determined by local summaries. Arrangement matters: two placements on a cycle have the same counts of each probability value and edge type, yet their optimal stability costs can differ by a factor that grows with batch size. On arbitrary graphs, equal-probability components provide a one-sided test for whether zero excess can meet a stability cap. Block and Gaussian implementations provide explicit caps and certified tails. On equalized-odds policies, dependence placed along graph edges yields lower disagreement than random placement at matched stability, and in village lotteries that select exactly K people with balanced gender counts it keeps more families together than graph-blind sampling under the same certified cap.

1 INTRODUCTION

Randomized decision rules are a standard instrument of algorithmic and institutional policy. Randomized post-processing achieves equalized odds by acting on some individuals with a probability strictly between zero and one (Hardt et al., 2016), weighted lotteries randomize prediction-informed allocation of scarce resources (Jain et al., 2024), and experimental designs assign treatment with prescribed probabilities. Operational resource-allocation policies have also been designed and evaluated using data collected in deployment, including policies for allocating scarce housing resources subject to budget and fairness constraints (Tang et al., 2023). Each such policy commits to a vector of *marginal action probabilities* $p_i = \mathbb{P}(Z_i = 1)$, where $Z_i \in \{0, 1\}$ is the realized action for individual i. These marginals are commitments of the policy rather than predictive estimates. Compared with the randomized rule, thresholding them would produce a deterministic rule with different group error rates, that is, a different decision rule. The marginals determine every expected quantity, such as the expected number of actions, a group's expected action rate or the expected accuracy. However, they leave the joint distribution of $(Z_1, \ldots, Z_n)$ unspecified, and this joint distribution is a design choice that shapes the realized actions rather than their expectations (Budish et al., 2013; Aziz et al., 2024), in two ways.

The joint distribution governs two properties of the realized decisions that are in tension. The first is local coordination. For example, suppose two individuals whose decisions should agree each have $p_i = 1/2$. Two independent coin flips disagree half of the time, whereas a single shared coin never produces disagreement, and each individual still has $p_i = 1/2$. Such pairs form a *coordination graph*, whose edges mark where joint outcomes carry value, such as shared households, clusters or complementary treatments. The second property is aggregate stability. A shared draw that coordinates many decisions also moves them together, so statistics of the realized decisions, such as the number of actions, a group action rate, a resource load or a batch loss, keep their expectations but fluctuate

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

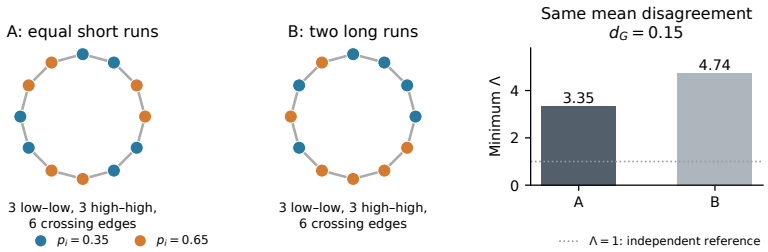


Figure 1: Left, middle: placements A and B on a 12-vertex cycle, $p_i = 0.35$ (blue) or 0.65 (orange), each with three low–low, three high–high and six crossing edges and floor $b_G = 0.15$. Right: exact minimum of Λ at zero excess over all joint distributions with these marginals; dotted: $\Lambda = 1$, independent sampling.

more. The same marginals can therefore yield stable or volatile aggregates, depending only on the joint implementation (Gur-Arieh & Lee, 2025).

Existing work studies the two properties separately or trades them off only within one family of constructions. On the coordination side, correlated school-choice lotteries preserve every assignment probability while keeping communities together (Ashlagi & Shi, 2014). On the stability side, dependent rounding correlates assignments at fixed marginals to reduce variance (Yamin et al., 2025), and the decompositions of one random allocation carry different guarantees for the realized allocation (Budish et al., 2013; Aziz et al., 2024). Where both properties appear, the trade-off is examined inside a chosen family: a shared random threshold for a stochastic classifier preserves metric fairness best but approximates its aggregate output worst, with hashing interpolating between the extremes (Wu et al., 2022), and covariance design with Gaussian rounding optimizes within Gaussian couplings (Thiyageswaran et al., 2026). Three questions therefore remain open: what the best attainable trade-off over all joint distributions is, whether the placement of prescribed probabilities changes it, and how to measure stability without a ground-truth action.

The goal of this paper is to characterize the best attainable trade-off between local coordination and aggregate stability over every joint distribution that preserves the prescribed marginals. The challenge is that the feasible set is a polytope in $\mathbb{R}^{2^n}$ and that any uniform stability criterion is a maximum over exponentially many aggregates. Because the frontier ranges over this whole set, its lower bounds must hold without assuming independent blocks, finite dependence or Gaussian generation. We therefore study the *consistency–stability frontier*, the least stability cost attainable at each level of coordination, which traces the Pareto frontier of all implementations. Figure 1 previews the central phenomenon: two placements of the same probabilities on the same cycle have identical local summaries and minimum disagreement, yet Theorem 2 shows that their optimal stability costs differ by a factor that grows with the batch size.

Our contributions are threefold. First, we formalize the frontier over all distributions with the prescribed marginals, with a stability criterion that needs no ground-truth actions and so applies to allocation as well as prediction (Section 2). Second, on cycles a computable score based on log-odds path lengths characterizes the zero-excess cost for arbitrary probabilities up to absolute constants. Two placements with the same counts of each probability and edge type stay separated at excess $n^{-1/2}$ (Section 3). Third, on arbitrary graphs a zero-excess diagnostic can rule out a selected stability cap (Section 3.3). For practitioners, Section 4 gives implementations whose stability cap is fixed before any draw: bounded blocks whose partition is chosen by its exact excess and Gaussian rounding of an optimized covariance, both with Bernstein tails for every fixed linear aggregate, and lotteries with a fixed number of places. Section 5 verifies the theorems numerically and studies an equalized-odds prediction policy and the allocation of exactly K places in villages.

2 Randomized policies and realized aggregates

A randomized policy fixes individual action probabilities but leaves their joint sampling open. Let $V = \{1, \dots, n\}$ index a batch and let $\mathbf{p} = (p_1, \dots, p_n) \in [0, 1]^n$ be the prescribed probabilities. An

implementation P is a joint distribution of decisions $Z = (Z_1, \dots, Z_n) \in \{0,1\}^n$ with $\mathbb{P}(Z_i = 1) = p_i$; $\mathcal{P}(\mathbf{p})$ is the set of all such distributions. We keep $\mathbf{p}$ fixed and choose only P.

A coordination graph marks pairs whose realized decisions should agree. Write $G = (V, E, w)$, where E contains unordered pairs $ij = \{i, j\}$ of distinct vertices and $w_{ij} > 0$ is the weight of edge ij. The graph is specified before the draw, for example from household membership or an observed network.

Definition 1 (Disagreement and excess)**.** Let $W = \sum_{ij \in E} w_{ij} > 0$ denote the total edge weight. The weighted average disagreement of P and its floor are defined as

$$d_G(P) = \frac{1}{W} \sum_{ij \in E} w_{ij} \mathbb{P}(Z_i \neq Z_j), \qquad b_G(\mathbf{p}) = \frac{1}{W} \sum_{ij \in E} w_{ij} |p_i - p_j|, \tag{1}$$

and the *excess* is $e_G(P) = d_G(P) - b_G(\mathbf{p})$.

Each edge has its own disagreement floor, determined by its two marginal probabilities. For an edge ij, relabel its endpoints so that $p_i \leq p_j$. Separating the two ways its decisions can differ gives

$$\mathbb{P}(Z_i \neq Z_j) - |p_i - p_j| = 2\mathbb{P}(Z_i = 1, Z_j = 0), \tag{2}$$

Thus the disagreement probability on ij equals $|p_i - p_j|$ exactly when $\mathbb{P}(Z_i = 1, Z_j = 0) = 0$. With the marginals fixed, this condition uniquely determines the joint distribution of (Z_i, Z_j): the event $\{Z_i = 1\}$ is contained in $\{Z_j = 1\}$ almost surely. We call these action events *nested*. Because all weights are positive, $e_G(P) = 0$ exactly when every edge attains its disagreement floor. These requirements fix the joint distribution of each edge pair, but generally leave that of the full batch open.

Zero excess is always feasible. Draw one $U \sim \mathrm{Uniform}[0,1]$ and set $Z_i = \mathbb{I}\{U \leq p_i\}$ for all i, where $\mathbb{I}\{\cdot\}$ is the indicator function. This *common-uniform* implementation nests every pair, including nonedges. If all $p_i = 1/2$, for example, its total action count is either 0 or n. On some graphs, the same edge requirements permit less aggregate fluctuation.

Next, we measure aggregate stability by how much totals formed from realized decisions fluctuate across repeated draws. For coefficients $a = (a_1, \dots, a_n) \in [-1,1]^n$ fixed before the draw, $\sum_i a_i Z_i$ has the same expectation $\sum_i a_i p_i$ under every implementation. Its centered value is $T_a = \sum_i a_i (Z_i - p_i)$. For a baseline, independent sampling draws $U_1, \dots, U_n$ independently from $\mathrm{Uniform}[0,1]$ and sets $Z_i = \mathbb{I}\{U_i \leq p_i\}$. Then $v_i = p_i(1 - p_i)$ is the variance of Z_i, and $s_{\mathbf{p}}^2 = \sum_i v_i$ is the variance of the total action count. We assume $s_{\mathbf{p}}^2 > 0$. Otherwise all decisions are fixed.

Definition 2 (Aggregate stability)**.** The *aggregate stability* of P is the worst normalized variance over bounded linear aggregates,

$$\Lambda(P) = \frac{1}{s_{\mathbf{p}}^2} \sup_{\|a\|_\infty \leq 1} \mathrm{Var}(T_a) = \frac{1}{s_{\mathbf{p}}^2} \max_{a \in \{-1,1\}^n} a^\top \mathrm{Cov}(Z)\, a, \tag{3}$$

where $\|a\|_\infty = \max_i |a_i|$, $\mathrm{Cov}(Z)$ is the covariance matrix of Z, and $^\top$ denotes transposition.

The sign-vector form holds because $a^\top \mathrm{Cov}(Z) a$ is convex in a, so its maximum over $[-1,1]^n$ occurs at a vertex. Under independent draws, $\Lambda(P) = 1$. Under any implementation, averaging the variance over all sign choices cancels cross-covariances and leaves $s_{\mathbf{p}}^2$. The maximum is at least this average, so $\Lambda(P) \geq 1$. Smaller values mean more stable aggregates.

The same criterion covers common totals and prediction loss. For $S \subseteq V$, let $\mathbf{1}_S$ be its indicator vector. The total count uses $a = \mathbf{1}_V$; a group count uses $a = \mathbf{1}_S$; a contrast between disjoint groups $S_1, S_2 \subseteq V$ uses $a = \mathbf{1}_{S_1} - \mathbf{1}_{S_2}$; and a cost total uses $a_i = c_i$ for costs $c_i \in [-1,1]$. Equalized-odds postprocessing can randomize binary predictions to meet fairness constraints (Hardt et al., 2016). For fixed labels $y = (y_1, \dots, y_n) \in \{0,1\}^n$, unknown when P is chosen, the batch error rate $L_y = n^{-1} \sum_i \mathbb{I}\{Z_i \neq y_i\}$ obeys

$$L_y = \frac{1}{n} \sum_i \big[y_i + (1 - 2y_i) Z_i\big], \qquad \mathbb{E}L_y = \frac{1}{n} \sum_i \big[y_i + (1 - 2y_i) p_i\big], \tag{4}$$

since an error is Z_i if $y_i = 0$ and $1 - Z_i$ if $y_i = 1$. Thus expected loss depends only on $\mathbf{p}$, while $n(L_y - \mathbb{E}L_y) = T_{\mathbf{1}_V - 2y}$. As y ranges over binary labels, $\mathbf{1}_V - 2y$ ranges over all sign vectors: $\Lambda(P)$ is the largest variance of the batch error count, divided by its variance $s_{\mathbf{p}}^2$ under independent draws. The theory does not require labels.

Definition 3 (Stability frontier)**.** For fixed G, $\mathbf{p}$, and excess budget $e \geq 0$, define

$$K_{G,\mathbf{p}}(e) = \min\{\Lambda(P) : P \in \mathcal{P}(\mathbf{p}),\ e_G(P) \leq e\}. \tag{5}$$

The frontier asks how small the worst aggregate variance can be while keeping excess at most e. At $e = 0$ every edge attains its own floor. A larger e cannot increase $K_{G,\mathbf{p}}(e)$. The minimum ranges over all implementations with marginals $\mathbf{p}$ that meet the budget, and is attained. Appendix A gives an equivalent linear program. The floor $b_G(\mathbf{p})$ averages the edgewise disagreement floors, but does not record how those edges connect into paths. We next ask whether placing the probabilities differently on the graph changes the best aggregate stability.

3 What determines the frontier

The frontier is defined for any coordination graph. We first show how probability placement shapes it on a cycle, then use the same path mechanism to check zero-excess stability on arbitrary graphs. Section 4 constructs implementations under a selected stability cap.

3.1 A coordination-cost score for zero excess on a cycle

Let C_n be the unit-weight cycle on vertices $1, \ldots, n$, where $n \geq 3$, with edges $(i, i+1)$ for $1 \leq i < n$ and $(n, 1)$. We study $K_{C_n,\mathbf{p}}(0)$, the best stability attainable at zero excess on this cycle. By (5), this quantity minimizes the worst normalized aggregate variance while every edge attains its disagreement floor. Here $p_i \in (0, 1)$ so the log odds below are finite. On the cycle, $p_{n+1} = p_1$ and $Z_{n+1} = Z_1$. Under a zero-excess implementation, the joint distribution of each adjacent pair is fixed, although the full decision vector can still be coupled in different ways.

The endpoint probabilities determine the correlation under a zero-excess implementation. Define the log odds $\eta_i = \log\{p_i/(1-p_i)\}$. If $p_i \leq p_{i+1}$, edgewise nesting gives $\mathrm{Cov}(Z_i, Z_{i+1}) = p_i(1-p_{i+1})$. Normalizing by $\sqrt{v_i v_{i+1}}$ gives the correlation

$$\mathrm{Corr}(Z_i, Z_{i+1}) = e^{-\ell_i}, \qquad \ell_i = \tfrac{1}{2}|\eta_{i+1} - \eta_i|, \tag{6}$$

where ℓ_i is the *log-odds length* of edge $(i, i+1)$. Thus equal endpoint probabilities give length zero and identical decisions, while a larger difference in log odds gives a smaller edge correlation. This formula fixes the correlation on an edge, not generally between distant vertices. Appendix B.2 gives the algebra.

The log-odds edge lengths, the marginal variances v_i and their sum $s^2_{\mathbf{p}}$ determine a score for the cycle's coordination cost.

Definition 4 (Coordination-cost score)**.** For vertices i and j, let $d_{\mathbf{p}}(i, j)$ be the smaller of the sums of ℓ_k along the two cycle paths between them, with $d_{\mathbf{p}}(i, i) = 0$. The coordination-cost score is

$$\Psi(\mathbf{p}) = \frac{1}{s^2_{\mathbf{p}}} \sum_{i,j \in V} \sqrt{v_i v_j}\, e^{-d_{\mathbf{p}}(i,j)}. \tag{7}$$

The score combines three components: marginal variability, path length, and normalization. The factor $\sqrt{v_i v_j}$ gives more weight to pairs whose individual decisions are more variable. The Bernoulli variance $v_i = p_i(1 - p_i)$ is largest at $p_i = 1/2$ and decreases toward zero as p_i approaches 0 or 1. The factor $e^{-d_{\mathbf{p}}(i,j)}$ is smaller when the shorter route has larger accumulated log-odds changes. Because $d_{\mathbf{p}}$ sums edge lengths along paths, the score records how those edges connect around the cycle. Dividing by $s^2_{\mathbf{p}}$ compares the score with the independent-sampling baseline. Placing the same probability values, with each value occurring the same number of times, at different vertices can change the score. The next theorem relates it to the best stability attainable at zero excess.

Theorem 1 (Coordination-cost score bounds the zero-excess cost)**.** *For every $\mathbf{p} \in (0, 1)^n$ on the cycle C_n,*

$$\max\{1, \Psi(\mathbf{p})/40\} \leq K_{C_n,\mathbf{p}}(0) \leq \min\{n, 2\Psi(\mathbf{p})\}. \tag{8}$$

The theorem turns the score $\Psi(\mathbf{p})$ into a way to assess zero-excess stability before choosing a joint decision rule: it bounds the lowest achievable worst normalized aggregate variance up to universal

constants. A high score signals unavoidable aggregate fluctuation, while a low score guarantees that an implementation with controlled fluctuation exists. The bounds hold for each finite n, with constants independent of n and the prescribed probabilities. Appendix B.2 gives the proof, also for $p_i \in \{0, 1\}$, and a copy–reset implementation.

The score sums over vertex pairs along the shortest cycle paths, measured in the log-odds lengths ℓ_i of (6). Runs of equal probabilities make this path dependence visible. A *run* is a maximal sequence of consecutive cycle vertices with the same prescribed probability. At zero excess, all decisions within a run are identical, and all vertex pairs in it have distance zero in the score. A longer run can therefore increase the lowest achievable worst normalized aggregate variance. The next subsection compares placements with the same counts of each prescribed probability and the same disagreement floor but different run lengths. Their optimal costs can differ by a factor that grows with n.

3.2 Placement matters beyond zero excess

The score $\Psi(\mathbf{p})$ characterizes the zero-excess endpoint up to constants. We now ask whether placement still affects the frontier when additional disagreement is allowed. We compare two assignments of the same two probabilities to the same cycle, keeping the counts of both probability values and of each kind of neighboring pair fixed.

For simplicity, take $n \geq 8$ divisible by four and $g \in (0, 1/2]$. Let $p = (1-g)/2$ be the low probability and $1 - p = (1 + g)/2$ the high probability, so their difference is g. Placement A repeats the pattern $p, p, 1-p, 1-p$ around the cycle. Placement B has a low-probability run of length $n/4+1$ followed by a high-probability run of the same length; the remaining $n/2 - 2$ vertices alternate between low and high probabilities, starting with low. Each placement has $n/2$ vertices of each type, $n/4$ low–low edges, $n/4$ high–high edges, and $n/2$ mixed edges. Thus both have the same disagreement floor $b_{C_n}(\mathbf{p}) = g/2$: each mixed edge contributes g and each same-probability edge contributes zero. Figure 1 shows the case $n = 12$ and $g = 0.30$. Let K_A and K_B denote the frontiers of these two probability placements. In the following rates, we interpret $1/e = \infty$ at $e = 0$.

Theorem 2 (Placement matters at a common excess budget)**.** *For the two placements above and* $0 \leq e \leq 1/8$*,*

$$K_A(e) = \Theta\left(\min\left\{n, \frac{1}{g+e}\right\}\right), \qquad K_B(e) = \Theta\left(\min\left\{n, \frac{1}{e}\right\}\right). \tag{9}$$

The rates compare optimal stability at the same allowed excess disagreement. Here $\Theta(f)$ means that each optimum is within constant factors of its displayed rate f, uniformly in n, g, and e. Both $K_A(e)$ and $K_B(e)$ minimize $\Lambda(P)$ over all implementations with their prescribed marginals and $d_{C_n}(P) \leq g/2 + e$. The lower bounds therefore apply beyond any particular implementation. Appendix B.3.2 gives matching explicit implementations.

The ratio $K_B(e)/K_A(e)$ depends on how the excess budget scales with n. Fix $g > 0$. At $e = 0$, A has $K_A(0) = \Theta(1)$ and B has $K_B(0) = \Theta(n)$, so the ratio is $\Theta(n)$: under B the worst aggregate variance is of the same order as if the whole batch shared one coin, whereas under A it stays bounded. At the positive budget $e_n = n^{-1/2}$, the ratio is $\Theta(\sqrt{n})$ as $n \to \infty$. The theorem applies to this budget when $n \geq 64$. This budget allows at most $ne_n = \sqrt{n}$ additional disagreements in expectation above the floor on the n-edge cycle. For any fixed $e \in (0, 1/8]$, both costs and their ratio are $\Theta(1)$ as n grows.

Long equal-probability runs explain why B costs more. Every edge within B's long equal-probability runs has disagreement floor zero, so a small excess budget bounds their total expected disagreement by ne. On run segments of order $\min\{n, 1/e\}$, enough agreement remains to force aggregate variance on that scale. Placement A instead disperses the same kinds of edges among runs of length two. The shared block-sign step is in Appendix B.1, and Appendix B.3 gives this theorem's lower-bound details. The same path-to-variance mechanism also supports Theorem 1 and the following general-graph diagnostic.

3.3 Checking zero excess on an arbitrary graph

On the cycle, a long run of equal probabilities moves as one unit at zero excess, which is what makes placement B costly. On a general graph the same holds for any set of vertices joined by edges with

equal endpoint probabilities, since at zero excess each such edge has identical decisions at its two ends. Let $\mathcal{C}$ be the connected components, including singletons, of the subgraph containing edges $ij \in E$ with $p_i = p_j$; all vertices in $C \in \mathcal{C}$ share a probability p_C. At zero excess, decisions within each component are therefore identical almost surely, so each component acts as one block drawn with a single coin. The block-sign bound (Lemma 3, Appendix B.1), which bounds Λ from below by the variance summed within any partition, then gives

$$K_{G,\mathbf{p}}(0) \geq \frac{\sum_{C\in\mathcal{C}} |C|^2 p_C(1-p_C)}{s_{\mathbf{p}}^2}. \tag{10}$$

The long runs in placement B make this bound of order n at zero excess. For a chosen stability cap $\kappa \geq 1$, a bound greater than κ rules out every zero-excess implementation with $\Lambda(P) \leq \kappa$, while a bound at or below κ does not establish feasibility. Since Λ compares aggregate variance with independent draws, a cap $\kappa = 4$, for instance, allows every bounded aggregate at most twice its standard deviation under independent draws. This test requires exactly equal probabilities along component edges, and Section 5.2 evaluates it on the evaluation batches of a prediction policy.

4 Implementations with explicit stability guarantees

This section gives practitioners implementations for an arbitrary supplied graph whose stability is certified before any draw. Section 5 evaluates each of them. We adapt bounded-block coupling and Gaussian rounding of an optimized covariance, give a tail bound for every fixed linear aggregate to select the cap, and extend the blocks to lotteries with a fixed number of places.

Bounded blocks. Fix a partition $\mathcal{B}$ of V into blocks $B \in \mathcal{B}$ with $|B| \leq \kappa$, draw one independent $U_B \sim \mathrm{Unif}[0,1]$ per block, and set $U_i = U_B$, so that $Z_i = \mathbb{I}\{U_B \leq p_i\}$, for $i \in B$. Every within-block edge is then nested and attains its marginal floor. Because blocks are independent, the variance of a linear form splits over blocks. Within a block of size at most κ, the Cauchy–Schwarz inequality bounds the resulting quadratic form by κ times its diagonal. Therefore, for every $a \in \mathbb{R}^n$,

$$\mathrm{Var}\left(\sum_i a_i(Z_i - p_i)\right) \leq \kappa \sum_i a_i^2 p_i(1-p_i). \tag{11}$$

The upper bound controls every linear aggregate, and $a \in \{-1,1\}^n$ gives $\Lambda(P) \leq \kappa$. A graph-based partition can keep selected edges within blocks, and its maximum block size is a stability certificate set without outcomes.

Both sides of the trade-off can be read from the partition before any draw. Write $P_{\mathcal{B}}$ for this joint distribution and $\gamma_{ij} = \min\{p_i, p_j\}(1 - \max\{p_i, p_j\})$, the covariance of a nested pair. Independence across blocks and nesting within blocks give

$$e_G(P_{\mathcal{B}}) = \frac{2}{W} \sum_{\substack{ij\in E:\\ i,j \text{ in different blocks}}} w_{ij}\gamma_{ij}, \qquad \Lambda(P_{\mathcal{B}}) = 1 + \frac{2}{s_{\mathbf{p}}^2} \sum_{B\in\mathcal{B}} \sum_{\substack{i<j\\ i,j\in B}} \gamma_{ij}, \tag{12}$$

the second because all covariances are nonnegative, so the all-ones sign attains (3). Merging two blocks A, B therefore saves the excess $2H_{AB}/W$, with H_{AB} the sum of $w_{ij}\gamma_{ij}$ over edges between them, and adds $2C_{AB}/s_{\mathbf{p}}^2$ to Λ, with C_{AB} the sum of γ_{ij} over all pairs between them. Merging greedily along edges under the size cap can rank candidate merges by edge weight, by H_{AB}, or by H_{AB}/C_{AB}. Because (12) is exact and needs only $\mathbf{p}$ and the graph, one can build all three partitions and keep the one with the smallest excess. All share the certificate κ. In Section 5.2 this choice lowers excess at the same cap in all 80 batch–cap pairs (Appendix E.8). Whether any of them reaches the frontier on a general graph remains open.

Gaussian covariance design. Let $Q \in \mathbb{R}^{n\times n}$ be symmetric with $Q \succeq 0$, unit diagonal $\mathrm{diag}\, Q = \mathbf{1}_n$, and $Q \preceq \kappa I$, that is, $\kappa I - Q$ is positive semidefinite for the identity I. This spectral cap on the latent correlation matrix Q is what bounds Λ. With Φ the standard normal distribution function, draw $H \sim N(0, Q)$ and set $Z_i = \mathbb{I}\{H_i \leq \Phi^{-1}(p_i)\}$, that is, $U_i = \Phi(H_i)$, using deterministic conventions at $p_i \in \{0,1\}$. The distribution P_Q of Z lies in $\mathcal{P}(\mathbf{p})$ because $\mathbb{P}(H_i \leq \Phi^{-1}(p_i)) = p_i$. Covariance

optimization followed by Gaussian rounding is an existing design approach (Thiyageswaran et al., 2026). We use a capped graph-Laplacian form with $\mathcal{L} \in \mathbb{R}^{n\times n}$, where $\mathcal{L}_{ij} = -w_{ij}$ for $ij \in E$, off-diagonal entries are zero otherwise, and $\mathcal{L}_{ii} = \sum_j w_{ij}$, so that $\operatorname{tr}(\mathcal{L}Q) = \sum_{ij\in E} w_{ij}(2 - 2Q_{ij})$:

$$\min_{Q\succeq 0,\, \operatorname{diag} Q=\mathbf{1}_n,\, Q\preceq \kappa I} \frac{\operatorname{tr}(\mathcal{L}Q)}{4W}. \tag{13}$$

Because $\operatorname{tr}(\mathcal{L}Q)/(4W) = W^{-1}\sum_{ij\in E} w_{ij}(1 - Q_{ij})/2$, the program seeks latent correlations near one on heavily weighted edges, while $Q \preceq \kappa I$ keeps the whole batch from moving together.

The latent objective is not the excess itself, which is hard to optimize directly, but it bounds the excess of the thresholded decisions:

$$e_G(P_Q) \le \frac{1}{W}\sum_{ij\in E} w_{ij}\frac{\arccos Q_{ij}}{\pi} \le \sqrt{\frac{\operatorname{tr}(\mathcal{L}Q)}{4W}}. \tag{14}$$

The first term is the Gaussian orthant probability $\arccos(Q_{ij})/\pi$ of opposite signs, with equality at $p_i = p_j = 1/2$. Appendix C.2 proves the bound for heterogeneous thresholds. Section 5 evaluates the thresholded decisions, not the latent objective.

Selecting a stability cap. Both constructions satisfy $\Lambda(P) \le \kappa$, and both provide exponential tails for every fixed linear aggregate. With $S_a = \sum_i a_i^2 p_i(1 - p_i)$, the variance of T_a under independent draws, and $M_a = \|a\|_\infty$, block independence and Gaussian product inequalities (Chen et al., 2015) give, for every fixed $a \in \mathbb{R}^n$ and every $u > 0$,

$$\mathbb{P}\{|T_a| > u\} \le 2\exp\left\{-\frac{u^2}{2\kappa(S_a + M_a u/3)}\right\}. \tag{15}$$

The bound is proved for these two constructions (Appendices C.1 and C.2) and does not follow from $\Lambda(P) \le \kappa$ for an arbitrary distribution. Since $S_a < s_{\mathbf{p}}^2 \|a\|_\infty^2$ for a sparse a such as a group indicator, it selects the cap from the monitored statistic itself, before any draw, more tightly than the full-batch variance bound: for a deviation u of T_a at failure probability δ, any feasible $\kappa \le u^2/\{2(S_a + M_a u/3)\log(2/\delta)\}$ suffices, provided the right side is at least one. For instance, in binary prediction, $a = \mathbf{1}_n - 2y$ and $u = nt$ give $S_a = s_{\mathbf{p}}^2$ and $M_a = 1$, so that for every fixed y

$$\mathbb{P}\{|L_y - \mathbb{E}L_y| > t\} \le 2\exp\left\{-\frac{n^2t^2}{2\kappa(s_{\mathbf{p}}^2 + nt/3)}\right\}. \tag{16}$$

Agreement can then be compared among the feasible designs.

A fixed number of places. Programs that must select exactly K people use $\mathcal{P}_K(\mathbf{p}) = \{P \in \mathcal{P}(\mathbf{p}) : \sum_i Z_i = K \text{ a.s.}\}$, which requires $\sum_i p_i = K$. On size-K selections we take the distribution proportional to $\exp\{\theta^\top z + \beta\sum_{B\in\mathcal{B}}\sum_{ij\in E:\, i,j\in B} w_{ij}\mathbb{I}\{z_i = z_j\}\}$, a second-order conditional Poisson design (Bondesson, 2012) with interactions of strength $\beta \ge 0$ only inside blocks, and θ calibrated so that $\mathbb{E}Z_i = p_i$. A recursion over blocks and the number selected, optionally also over a group's count, gives its marginals, covariance and samples exactly. Its cap is certified from that covariance (Appendix C.3).

5 Experiments

We check the cycle theorems (Section 5.1), test where to place dependence in a randomized fairness policy (Section 5.2), and allocate a fixed number of program places in villages (Section 5.3).

5.1 Checking the cycle theorems

To verify Theorems 1 and 2, we solve the linear program of Appendix A exactly, over all joint distributions with the prescribed marginals, and check the two predictions: the ratio K_B/K_A grows with n (Theorem 2), and $K_{C_n,\mathbf{p}}(0)/\Psi(\mathbf{p})$ lies in $[1/40, 2]$ (Theorem 1, (8)). For placements A and B on cycles of 8 and 12 vertices, the zero-excess ratio of B's optimum to A's rises from $n = 8$ to $n = 12$ at every gap, from 1.07 to 1.41 at $g = 0.30$. Beyond $n = 12$, Figure 3 shows B's lower bounds growing with n while A's implementation stays bounded. On cycles of 3 to 68 vertices, with probability vectors from nearly constant to highly heterogeneous, some near zero, the zero-excess optimum lies between 0.489 and 1.042 times $\Psi(\mathbf{p})$ (Appendix D lists every case).

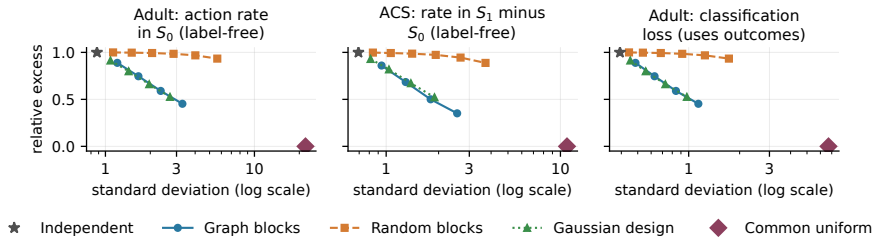


Figure 2: Horizontal: standard deviation of the statistic (log scale; rates in percentage points). Vertical: e_G relative to independent draws (1: independence; 0: floor). Smaller is better on both axes. Curves run through caps 2–16 (random blocks also 32, 64); averages over 4 Adult and 16 ACS batches.

5.2 Coordinating a randomized fairness policy

Setup. We fix the equalized-odds postprocessor of Hardt et al. (2016, Section 3.1) and change only how its decisions are drawn together. For each group and hard prediction of a trained classifier, it acts with a probability chosen by a linear program to equalize true- and false-positive rates across the two groups. We train classifiers on UCI Adult and on the ACS Income and Public Coverage tasks (Ding et al., 2021) and fit the postprocessor on a separate calibration split, so each person's p_i depends only on their group and hard prediction. People used in neither fit form the evaluation batches: four Adult batches of 2,048 and sixteen ACS batches of 1,024 across eight predictors. The postprocessor makes some cells deterministic, so 42–58% of each batch have $p_i \in \{0, 1\}$, and the designs differ only on the rest (splits and models: Appendix E.1). The graph joins ten nearest neighbours in the fitted features, weighted by $1/(1 + \text{distance})$. We compare independent draws, the common uniform, graph blocks (partition chosen by (12)), random blocks and the Gaussian design at caps $\kappa \in \{2, 4, 8, 16\}$ through the exact excess e_G and the exact standard deviation of each group's action rate and of their difference; labels enter only in the batch loss.

Results. Figure 2 shows the trade-off. At the same standard deviation of a group rate or of the difference of the two rates, graph blocks have smaller excess than random blocks in every comparison, 16 per statistic on Adult and 182 on ACS, by 0.10 to 0.85 of the independent excess. The Gaussian design is marginally better than graph blocks on Adult (57 of 60 comparisons, by at most 0.024) but not on ACS (55 of 240). The common uniform (diamond in Figure 2) removes all excess, but the standard deviation of the Adult S_0 action rate is 22.0 percentage points under it, against 0.88 under independent draws. Zero excess is out of reach under every tested cap, because the component bound (10) is at least 46 in every batch. Appendices E.3–E.5 report per-predictor results and the comparison at equal Λ.

Choosing a design before any draw. Before any draw and without outcomes, one fixes a tolerated deviation t of the monitored statistic and a confidence level; (15) then gives the largest admissible cap, and one uses graph blocks at that cap. For a batch-loss deviation of three points at 95% confidence, (16) selects $\kappa = 4$ in every batch and lowers disagreement below independent draws by 1.0 point on Adult and 0.3–0.4 on ACS (Appendix E.6). For a group's action rate, the coefficient-specific scale S_a admits up to twice the cap that the full-batch scale allows, and every selected design exceeds its target far less often than the allowed 5% (Appendix E.7).

5.3 Allocating a fixed number of program places in villages

Setup. We consider a program with K places allocated within each village. We use the village surveys and networks of Banerjee et al. (2013); Jackson et al. (2012). A logistic score for current non-employment, fitted on other villages only, sets $p_i \propto$ score with $\sum_i p_i = K = \lfloor Bn + 1/2 \rfloor$, so that a fraction $B \in \{0.1, 0.3\}$ of the village is selected. The graph links household members and named close relatives, a posited preference against splitting families whose value the data do not establish. Every design keeps each p_i, selects exactly K people and keeps each gender's selected count within one of its expectation. We compare graph-blind quota conditional Poisson sampling

(CP) with the graph-aware quota design of Section 4 on six villages of 124 to 304 people chosen by size rank before any run (implementation details: Appendix F).

Table 1: B: selected fraction; κ: certified cap; feasible: villages with a graph-aware candidate with $\bar{\Lambda} \leq \kappa$; d_G in pp, median over villages; gain: median (range) in pp and median share of CP's excess over b_G; –: no candidate in the β grid is certified.

		graph	d_G (pp)		gain over quota CP	
B	κ	feasible	quota CP	graph-aware	pp	% of excess
0.1	2	3/6	17.1	15.8	1.3 (1.1–1.4)	11%
	4	6/6	17.1	13.5	3.7 (2.3–7.3)	31%
	8	6/6	17.1	7.6	9.6 (7.4–10.8)	80%
0.3	2	0/6	41.3	–	–	–
	4	6/6	41.3	30.8	10.5 (6.9–13.8)	42%
	8	6/6	41.3	21.7	19.7 (15.3–23.3)	77%

Results. Table 1 reports, for each fraction B and certified cap κ, the villages admitting a graph-aware candidate and the disagreement of quota CP and of the best such candidate. Under $\kappa = 4$ and 8 the graph-aware lottery lowers disagreement in every village, by a median of 3.7 and 9.6 percentage points at $B = 0.1$ and 10.5 and 19.7 at $B = 0.3$, a median 31 to 80% of the CP excess over the unrestricted floor. The same interaction on random blocks gains at most 1.7 points, so the gain comes from the family edges. The method loses where the cap is tight: under $\kappa = 2$ only three of six villages admit a candidate at $B = 0.1$, gaining about one point, and none at $B = 0.3$. Appendix F adds the pivotal method (Deville & Tillé, 1998) and help-seeking edges.

6 Related work

Randomized allocation mechanisms and correlated lotteries already implement prescribed probabilities in practice (Budish et al., 2013; Aziz et al., 2024; Ashlagi & Shi, 2014), which motivates choosing their joint distribution deliberately. Simultaneous pairwise optimal couplings are characterized by Angel & Spinka (2019). Our upper bounds adapt constructions from derandomization and design, while our lower bounds assume none: hashing-based derandomization (Cotter et al., 2019), shared thresholds (Wu et al., 2022), dependent rounding (Yamin et al., 2025), cluster randomization (Ugander et al., 2013), Gaussian rounding (Thiyageswaran et al., 2026), robust joint design (Borusyak et al., 2026), spatial rates (Leung, 2022), and fixed-size sampling (Deville & Tillé, 1998; Bondesson, 2012), whose conditional Poisson family our fixed-capacity design confines to blocks. Post-processing (Hardt et al., 2016), prediction-informed lotteries (Jain et al., 2024), deployed allocation policies (Tang et al., 2023) and institutional consistency (Gur-Arieh & Lee, 2025) supply the marginals we take as given.

7 Conclusion

This paper treats the dependence structure of a randomized policy as a design choice. Arrangement along paths, which local counts do not reveal, sets the price of coordination in aggregate variability. A component bound can rule out zero excess on any graph, exact partition costs select blocks before any draw, blocks extend to exactly K places, and the cap follows from a tolerated deviation without labels. The cycle theorems do not cover general graphs or fixed capacities, the tails hold only for the stated constructions, and the exact-capacity study covers six villages.

AI USE STATEMENT

In this work, we used generative AI tools to help formulate the research problem and mathematical claims, develop proofs, design experiments, implement methods, and process data. AI-generated synthetic datasets and qualitative or thematic data analysis were not applicable to this work. We reviewed all AI-assisted work. We take responsibility for the final content of this work, including text, claims, and artifacts produced with the aid of generative AI.

REFERENCES

Omer Angel and Yinon Spinka. Pairwise optimal coupling of multiple random variables. *arXiv preprint arXiv:1903.00632*, 2019.

Itai Ashlagi and Peng Shi. Improving community cohesion in school choice via correlated-lottery implementation. *Operations Research*, 62(6):1247–1264, 2014.

Haris Aziz, Rupert Freeman, Nisarg Shah, and Rohit Vaish. Best of both worlds: Ex ante and ex post fairness in resource allocation. *Operations Research*, 72(4):1674–1688, 2024.

Abhijit Banerjee, Arun G. Chandrasekhar, Esther Duflo, and Matthew O. Jackson. The diffusion of microfinance. *Science*, 341(6144):1236498, 2013.

Lennart Bondesson. On sampling with prescribed second-order inclusion probabilities. *Scandinavian Journal of Statistics*, 39:813–829, 2012.

Kirill Borusyak, Peter Hull, and Evan Munro. Robust signal maximization in spillover experiments. *arXiv preprint arXiv:2607.18601*, 2026.

Eric Budish, Yeon-Koo Che, Fuhito Kojima, and Paul Milgrom. Designing random allocation mechanisms: Theory and applications. *American Economic Review*, 103(2):585–623, 2013.

Wei-Kuo Chen, Nikos Dafnis, and Grigoris Paouris. Improved Hölder and reverse Hölder inequalities for Gaussian random vectors. *Advances in Mathematics*, 280:643–689, 2015.

Andrew Cotter, Maya Gupta, and Harikrishna Narasimhan. On making stochastic classifiers deterministic. In *Advances in Neural Information Processing Systems*, volume 32, 2019.

Jean-Claude Deville and Yves Tillé. Unequal probability sampling without replacement through a splitting method. *Biometrika*, 85(1):89–101, 1998.

Frances Ding, Moritz Hardt, John Miller, and Ludwig Schmidt. Retiring adult: New datasets for fair machine learning. In *Advances in Neural Information Processing Systems*, volume 34, 2021.

Shira Gur-Arieh and Christina Lee. Consistently arbitrary or arbitrarily consistent: Navigating the tensions between homogenization and multiplicity in algorithmic decision-making. In *Proceedings of the 2025 ACM Conference on Fairness, Accountability, and Transparency*, pp. 3336–3349, 2025.

Moritz Hardt, Eric Price, and Nati Srebro. Equality of opportunity in supervised learning. *Advances in Neural Information Processing Systems*, 29, 2016.

Matthew O. Jackson, Tomas Rodriguez-Barraquer, and Xu Tan. Social capital and social quilts: Network patterns of favor exchange. *American Economic Review*, 102(5):1857–1897, 2012.

Shomik Jain, Kathleen Creel, and Ashia Wilson. Position: Scarce resource allocations that rely on machine learning should be randomized. In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pp. 21148–21169, 2024.

Michael P. Leung. Rate-optimal cluster-randomized designs for spatial interference. *The Annals of Statistics*, 50(5):3064–3087, 2022.

Bill Tang, Çağıl Koçyiğit, Eric Rice, and Phebe Vayanos. Learning optimal and fair policies for online allocation of scarce societal resources from data collected in deployment. *arXiv preprint arXiv:2311.13765*, 2023.

Vydhourie Thiyageswaran, Alex Kokot, Jennifer Brennan, Marina Meila, Christina Lee Yu, and Maryam Fazel. Optimal design under interference, homophily, and robustness trade-offs. *arXiv preprint arXiv:2601.17145*, 2026.

Johan Ugander, Brian Karrer, Lars Backstrom, and Jon Kleinberg. Graph cluster randomization: Network exposure to multiple universes. In *Proceedings of the 19th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 329–337, 2013.

Jimmy Wu, Yatong Chen, and Yang Liu. Metric-fair classifier derandomization. In *International Conference on Machine Learning*, pp. 23999–24016. PMLR, 2022.

Khurram Yamin, Edward Kennedy, and Bryan Wilder. Dependent randomized rounding for budget constrained experimental design. In Silvia Chiappa and Sara Magliacane (eds.), *Proceedings of the Forty-first Conference on Uncertainty in Artificial Intelligence*, volume 286 of *Proceedings of Machine Learning Research*, pp. 4681–4700. PMLR, 2025.

GUIDE TO THE APPENDIX

Each appendix section starts by naming the part of the main text it supports. The table lists where each result of the main text is proved or supported.

Main text	Statement	Appendix
Section 2	edge identity (2); the frontier (5) as a linear program	A.1
Section 3.3	Lemma 3 and the component bound (10)	B.1
Section 3	Theorem 1 (coordination-cost score)	B.2
Section 3	Theorem 2 (placement), with matching implementations	B.3
Section 4	bounded blocks: (12), tails (15), (16)	C.1
Section 4	Gaussian design: (14), tails (15), (16)	C.2
Section 4	fixed number of places: calibration, exact computation, certified cap	C.3
Section 4	choice among the three greedy partitions	E.8
Section 5.1	exact linear programs for both theorems	D
Section 5.2	setup	E.1–E.2
Section 5.2	results	E.3–E.5
Section 5.2	choosing a design before any draw	E.6–E.7
–	sensitivity to the graph (not used in the main text)	E.9
Section 5.3	data, designs and checks of the lottery	F

A NOTATION AND FACTS FROM SECTION 2

Table 2 collects the symbols of the main text. Appendix A.1 proves the two facts stated without proof in Section 2. Throughout, the graph and the action probabilities are fixed, as are outcomes where they exist, and only the implementation P varies.

A.1 FACTS STATED IN SECTION 2

The edge identity (2)**.** For binary Z_i, Z_j, $\mathbb{P}(Z_i \neq Z_j) = \mathbb{P}(Z_i = 1, Z_j = 0) + \mathbb{P}(Z_i = 0, Z_j = 1)$, and the marginals give $\mathbb{P}(Z_i = 0, Z_j = 1) = \mathbb{P}(Z_i = 1, Z_j = 0) + p_j - p_i$. For $p_i \leq p_j$ this is (2).

The frontier is a linear program and is attained. Because the means are fixed, $a^\top \operatorname{Cov}(Z)a = \mathbb{E}[(a^\top(Z - \mathbf{p}))^2]$ is linear in P, so Λ is a maximum of 2^n linear functions of P. Minimizing an auxiliary variable λ subject to $a^\top \operatorname{Cov}(Z)a \leq \lambda s_{\mathbf{p}}^2$ for every $a \in \{-1, 1\}^n$, $e_G(P) \leq e$ and $P \in \mathcal{P}(\mathbf{p})$ is therefore a linear program over a compact polytope. It is feasible for every $e \geq 0$ because the common-uniform implementation has zero excess, so its minimum is attained. Appendix D.1 gives the finite form solved in Section 5.1.

B PROOFS FOR SECTION 3

Appendix B.1 proves the block-sign bound (Lemma 3), which gives the component bound (10) of Section 3.3 and is the lower-bound tool in both theorems. Appendix B.2 proves Theorem 1 and Appendix B.3 proves Theorem 2.

B.1 THE BLOCK-SIGN BOUND AND THE COMPONENT TEST

Lemma 3 (Independent block signs)**.** *For any partition $\mathcal{B}$ of the vertices and any implementation,*

$$s_{\mathbf{p}}^2 \Lambda(P) \;\geq\; \sum_{B \in \mathcal{B}} \operatorname{Var}\left(\sum_{i \in B} Z_i \right). \tag{17}$$

Proof. Give each block $B \in \mathcal{B}$ an independent fair sign σ_B and assign it to all vertices of B, so that $a_i = \sigma_B$ for $i \in B$. Averaging the quadratic form $a^\top \operatorname{Cov}(Z)a$ over the signs cancels every covariance term between distinct blocks and leaves $\sum_B \operatorname{Var}(\sum_{i \in B} Z_i)$. The maximum over fixed signs in (3) is at least this average, which proves (17). □

Symbol	Meaning
$V = \{1, \dots, n\}$	batch of decision subjects, one per vertex
$G = (V, E, w)$	coordination graph; $E \subseteq \{\{i, j\} : i \neq j\}$, edge written ij
$w : E \to (0, \infty)$	edge weights, written w_{ij}; $W = \sum_{ij \in E} w_{ij}$
$\mathbf{p} = (p_1, \dots, p_n)$	prescribed action probabilities, $p_i \in [0, 1]$
$Z = (Z_1, \dots, Z_n)$	random decision vector in $\{0, 1\}^n$
P, $\mathcal{P}(\mathbf{p})$	an implementation, i.e. a distribution of Z with $\mathbb{E}Z_i = p_i$; the set of all
$\mathbb{E}, \mathbb{P}, \mathrm{Var}, \mathrm{Cov}$	expectation, probability, variance, covariance under P
$\mathrm{Cov}(Z)$	$n \times n$ covariance matrix of Z
$d_G(P)$, $b_G(\mathbf{p})$	weighted average disagreement; its marginal floor, (1)
$e_G(P)$	excess $d_G(P) - b_G(\mathbf{p}) \geq 0$
U	a single $\mathrm{Unif}[0, 1]$ variable; $Z_i = \mathbb{I}\{U \leq p_i\}$ gives $e_G = 0$
a, T_a	fixed coefficients in $[-1, 1]^n$; centred linear aggregate $\sum_i a_i(Z_i - p_i)$
S_a, M_a	$\sum_i a_i^2 p_i(1 - p_i)$ and $\|a\|_\infty$, (15)
$s^2_{\mathbf{p}}$	$\sum_i p_i(1 - p_i)$, assumed positive
$\Lambda(P)$	worst normalized linear-aggregate variance, (3); $\Lambda \geq 1$, $= 1$ if independent
$y \in \{0, 1\}^n$	outcome vector in binary prediction, unknown when P is chosen
L_y	batch classification loss $n^{-1} \sum_i \mathbb{I}\{Z_i \neq y_i\}$, (4)
$K_{G,\mathbf{p}}(e)$	frontier $\min\{\Lambda(P) : e_G(P) \leq e\}$ over all of $\mathcal{P}(\mathbf{p})$, (5)
κ	cap imposed on Λ, i.e. the constraint $\Lambda(P) \leq \kappa$
$\mathbf{1}$, $\mathbb{I}\{\cdot\}$, I	all-ones vector, indicator, identity matrix
$a^\top$	transpose of a
C_n, g, p, v	unit-weight cycle; probability gap; the value $(1 - g)/2$; $v = p(1 - p)$
$\mathbf{r} = (r_1, \dots, r_R)$	run lengths in cyclic order; R even, $\sum_j r_j = n$
ℓ_i, $d_{\mathbf{p}}(i, j)$	log-odds edge length, (6); shorter-arc length between i and j
$\Psi(\mathbf{p})$	coordination-cost score, (7)
K_A, K_B	frontiers under placements A and B
$\mathcal{B}$, B	a partition of V and one of its blocks
γ_{ij}	$\min\{p_i, p_j\}(1 - \max\{p_i, p_j\})$, covariance of a nested pair
H_{AB}, C_{AB}	$\sum w_{ij}\gamma_{ij}$ over edges between blocks A, B; $\sum \gamma_{ij}$ over all pairs between them
D	expected number of disagreeing edges, $\sum_{ij \in E} \mathbb{P}(Z_i \neq Z_j)$
Q, H, Φ	latent correlation matrix; $H \sim N(0, Q)$; standard normal cdf
P_Q	distribution of $Z_i = \mathbb{I}\{H_i \leq \Phi^{-1}(p_i)\}$, an element of $\mathcal{P}(\mathbf{p})$
$\mathcal{L}$	graph Laplacian, $\mathcal{L}_{ij} = -w_{ij}$ on edges, $\mathcal{L}_{ii} = \sum_j w_{ij}$
u, t, δ	target deviation of T_a; of L_y; failure probability
$\mathcal{P}_K(\mathbf{p})$, $\bar{\Lambda}$	distributions selecting exactly K; certified upper bound on Λ (App. C.3)
θ, β	calibration parameters; within-block interaction strength of that design

Table 2: Symbols of the main text, grouped by the section that introduces them.

Proof of (10)**.** At zero excess every edge attains its floor. For an edge with $p_i = p_j$ the floor is zero, so $Z_i = Z_j$ almost surely, and every component $C \in \mathcal{C}$ is constant, with $\mathrm{Var}(\sum_{i\in C} Z_i) = |C|^2 p_C(1-p_C)$. Lemma 3 with the partition $\mathcal{C}$ gives (10).

B.2 Proof of Theorem 1

Throughout, $\eta_i = \mathrm{logit}(p_i)$, $\ell_i = \frac{1}{2}|\eta_{i+1} - \eta_i|$ is the length of edge $\{i, i+1\}$ from (6), $L = \sum_i \ell_i$ is the total length of the cycle, $w_i = \sqrt{v_i}$ and $s^2 = s^2_{\mathbf{p}}$. Zero excess means $\mathbb{P}(Z_i \neq Z_{i+1}) = |p_i - p_{i+1}|$ on every edge, equivalently $Z_i \leq Z_{i+1}$ almost surely when $p_i \leq p_{i+1}$. For $p_i \leq p_{i+1}$, the identity (6) follows from $\mathrm{Cov}(Z_i, Z_{i+1}) = p_i(1-p_{i+1})$ for a nested pair and $p_i(1-p_{i+1})/\sqrt{v_i v_{i+1}} = \{p_i(1-p_{i+1})/(p_{i+1}(1-p_i))\}^{1/2}$. The bounds $1 \leq \Psi(\mathbf{p}) \leq (\sum_i w_i)^2/s^2 \leq n$ follow from the diagonal terms, from $e^{-d} \leq 1$ and from the Cauchy–Schwarz inequality.

B.2.1 Lower bound

Lemma 4 (Forced covariance at short log-odds distance)**.** *Let P have zero excess and let a path of C_n join i and j with total length $D = \sum \ell_k$ along it. Then $\mathrm{Cov}_P(Z_i, Z_j) \geq \sqrt{v_i v_j}\,(1 - De^D)$. In particular, $\mathrm{Cov}_P(Z_i, Z_j) \geq \frac{1}{2}\sqrt{v_i v_j}$ when $D \leq 1/4$.*

Proof. For any two binary variables, $\mathrm{Cov}(Z_i, Z_j) = \frac{1}{2}\{v_i + v_j + (p_i - p_j)^2 - \mathbb{P}(Z_i \neq Z_j)\}$, and by the union bound along the path and zero excess on its edges, $\mathbb{P}(Z_i \neq Z_j) \leq \sum_{uv} |p_u - p_v|$. Write $p(\eta) = (1 + e^{-\eta})^{-1}$ and $v(\eta) = p(\eta)\{1 - p(\eta)\}$, so that $dp/d\eta = v(\eta)$ and $|d \log v/d\eta| = |1 - 2p(\eta)| \leq 1$. Interpolating each edge monotonically in η gives $\sum_{uv} |p_u - p_v| = \int v(\eta)\,|d\eta|$ over a curve of total variation $2D$. At a point at variation A from i and B from j, where $A + B = 2D$, the derivative bound gives $\log v \leq \log v_i + A$ and $\log v \leq \log v_j + B$; averaging, $v \leq e^D \sqrt{v_i v_j}$. Hence $\mathbb{P}(Z_i \neq Z_j) \leq 2De^D\sqrt{v_i v_j}$, and $v_i + v_j \geq 2\sqrt{v_i v_j}$ yields the first claim. The second follows from $\frac{1}{4}e^{1/4} < \frac{1}{2}$. □

For the lower bound in (8), let P have zero excess. If $L \leq 1/4$, every pair is joined by a path of length at most $1/4$, and Lemma 3 with a single block gives $s^2\Lambda(P) \geq \mathrm{Var}(\sum_i Z_i) \geq \frac{1}{2}(\sum_i w_i)^2 \geq \frac{1}{2}s^2\Psi(\mathbf{p})$. If $L > 1/4$, represent the cycle as a circle of circumference L, on which vertex i sits at coordinate $\sum_{k<i} \ell_k$, and divide it into $m = \lceil 4L \rceil$ half-open cells of length $b = L/m$, so that $1/8 < b \leq 1/4$. Vertices at a common coordinate share a cell, and cells may be empty. With $A_k = \sum_{i \in \text{cell } k} w_i$, the lemma applies inside each cell, and Lemma 3 gives $s^2\Lambda(P) \geq \frac{1}{2}\sum_k A_k^2$. Two cells at cyclic index distance r are separated along both arcs by at least $(r-1)_+$ complete cells, so $d_{\mathbf{p}}(i,j) \geq (r-1)_+ b$ for their vertices. By $2A_kA_l \leq A_k^2 + A_l^2$ and a geometric series,

$$s^2\Psi(\mathbf{p}) \leq \sum_{k,l} A_k A_l e^{-(r(k,l)-1)_+ b} \leq \Big(1 + \frac{2}{1 - e^{-b}}\Big) \sum_k A_k^2 \leq 19 \sum_k A_k^2,$$

using $e^{-1/8} \leq 8/9$. Hence $\Lambda(P) \geq \Psi(\mathbf{p})/38 \geq \Psi(\mathbf{p})/40$ for every zero-excess P, and $\Lambda \geq 1$ holds always.

B.2.2 Upper bound: the copy–reset implementation

If $L = 0$, every probability is equal and a common uniform attains $\Lambda = n = \Psi(\mathbf{p})$. Otherwise, for each edge $i \to i+1$ set

$$\tau_i = \frac{1 - p_{i+1}}{1 - p_i} \;\; (p_{i+1} \geq p_i), \qquad \tau_i = \frac{p_{i+1}}{p_i} \;\; (p_{i+1} < p_i),$$

and draw independent random maps $F_i : \{0,1\} \to \{0,1\}$ equal to the identity with probability τ_i and otherwise to the constant 1 on an increasing edge or 0 on a decreasing edge. Then $\mathbb{E}F_i(x) = p_{i+1} + \tau_i(x - p_i)$, and a direct computation from (6) gives $\tau_i\sqrt{v_i/v_{i+1}} = e^{-\ell_i}$, so that $\alpha = \prod_i \tau_i = e^{-L} < 1$. Draw the maps conditionally on the event $\mathcal{E}$ that at least one map is constant, which has probability $1 - \alpha$, and let Z solve $Z_{i+1} = F_i(Z_i)$ for every i modulo n. The solution exists and is unique, because a constant map fixes the next state and the remaining states follow by propagation.

Every edge has zero excess: copying or resetting to 1 gives $Z_i \le Z_{i+1}$ on an increasing edge, and copying or resetting to 0 gives $Z_i \ge Z_{i+1}$ on a decreasing edge.

Let G_i be the composition of all n maps, starting and ending at vertex i. Independence of the maps gives $\mathbb{E}G_i(x) = p_i + \alpha(x - p_i)$ before conditioning. On $\mathcal{E}^c$, G_i is the identity; on $\mathcal{E}$, it is the constant map with value Z_i, since $Z_i = G_i(Z_i)$. Setting $x = 0$ gives $(1-\alpha)\mathbb{E}[Z_i \mid \mathcal{E}] = \mathbb{E}G_i(0) = (1-\alpha)p_i$, so $\mathbb{E}Z_i = p_i$ exactly. A draw requires $O(n)$ operations without rejection: in a fixed linear order, the first constant map is at edge j with probability $(\prod_{k<j} \tau_k)(1 - \tau_j)/(1 - \alpha)$, the maps after it are drawn from their unconditional distributions, and states propagate from edge j.

Fix $i \ne j$, let D_+ and $D_- = L - D_+$ be the lengths of the two arcs from i to j, and set $A = e^{-D_+}$ and $B = e^{-D_-}$, so that $AB = \alpha$. Before conditioning, the compositions along the two arcs use disjoint maps and are independent. The clockwise composition is the identity with probability $t = A\sqrt{v_j/v_i}$ and the constant 1 with probability $p_j - tp_i$; the other is the identity with probability $u = B\sqrt{v_i/v_j}$ and the constant 1 with probability $p_i - up_j$. Excluding the event that both are identities, which has probability $tu = \alpha$, and enumerating the remaining cases gives

$$(1 - \alpha)\mathbb{P}(Z_i = Z_j = 1) = (p_j - tp_i)(p_i - up_j) + t(p_i - up_j) + u(p_j - tp_i),$$

and therefore, with $r = e^{-|\eta_i - \eta_j|/2}$ and $(p_i + p_j - 2p_ip_j)/\sqrt{v_iv_j} = r + r^{-1}$,

$$\frac{\mathrm{Cov}(Z_i, Z_j)}{\sqrt{v_iv_j}} = \frac{A + B - AB(r + r^{-1})}{1 - AB} = A + \frac{B(1 - A/r)(1 - Ar)}{1 - AB}.$$

Both arcs are at least as long as $|\eta_i - \eta_j|/2$, so $A, B \le r \le 1$ and the last term is nonnegative; by symmetry the correlation is at least $\max\{A, B\} = e^{-d_{\mathbf{p}}(i,j)}$. Since $r + r^{-1} \ge 2 \ge A + B$, it is at most $(A+B-2AB)/(1-AB) \le A+B \le 2e^{-d_{\mathbf{p}}(i,j)}$. The diagonal satisfies the same comparison. Every covariance is nonnegative, so the all-ones sign vector attains the maximum in (3), and summing gives $\Psi(\mathbf{p}) \le \Lambda \le 2\Psi(\mathbf{p})$. The universal bound $\Lambda \le (\sum_i w_i)^2/s^2 \le n$ completes (8).

B.2.3 Special cases used in Section 5.1

Monotone marginals. If the vertices can be indexed cyclically so that $p_1 \le \cdots \le p_n$, zero excess forces $Z_1 \le \cdots \le Z_n$, which determines the distribution. Then $\mathrm{Cov}(Z_i, Z_j) = p_i(1 - p_j) = \sqrt{v_iv_j}e^{-|\eta_i - \eta_j|/2}$ for $i < j$, the increasing arc is a shortest arc, and $K_{C_n,\mathbf{p}}(0) = \Psi(\mathbf{p})$.

Deterministic coordinates. Let $\mathbf{p} \in [0,1]^n$ with $s^2 > 0$. A vertex with $p_i \in \{0, 1\}$ attains the floor on its incident edges under every implementation, so deleting these vertices leaves active components C that are paths, or the whole cycle. Keep the lengths ℓ_i inside components, set $d_{\mathbf{p}} = \infty$ between components and sum (7) over active vertices only. With $s_C^2 = \sum_{i \in C} v_i$, $K_{C_n,\mathbf{p}}(0) = \sum_C (s_C^2/s^2) K_{C,\mathbf{p}_C}(0)$: for $\ge$, take maximizing signs within each component and average independent signs across components, which removes cross-component covariance; for $\le$, use independent optimal implementations on the components. On a path, propagating copy–reset maps from a first Bernoulli(p) decision preserves every marginal and gives $\mathrm{Cov}(Z_i, Z_j) = \sqrt{v_iv_j}e^{-d_{\mathbf{p}}(i,j)}$ exactly, and the cell argument above with cells of length $1/4$ on a line gives the lower constant $1/21$. Theorem 1 therefore holds for all such $\mathbf{p}$.

Two complementary values. For $p_i \in \{(1 - g)/2, (1 + g)/2\}$, every v_i equals $v = (1 - g^2)/4$, edges inside a run have length zero and each run boundary has length $\log\{(1 + g)/(1 - g)\}$, which gives the run form of $\Psi(\mathbf{p})$ used to analyze placements in Section 3. With $R \ge 2$ runs, an even number, of equal length a and $\rho = (1 - g)/(1 + g)$, the sum over run distances is $a\{1 + 2\sum_{k=1}^{R/2-1} \rho^k + \rho^{R/2}\} = a(1+\rho)(1-\rho^{R/2})/(1-\rho)$, and $(1+\rho)/(1-\rho) = 1/g$. Since $1-\rho \ge g$ and $1 - x^q \ge (1 - e^{-1}) \min\{1, q(1-x)\}$ for $x \in [0, 1]$, the score lies between $(1-e^{-1}) \min\{a/g, n/2\}$ and $\min\{a/g, n\}$, which gives $K_{C_n,\mathbf{p}}(0) = \Theta(\min\{n, a/g\})$.

B.3 Proof of Theorem 2

Let $n = 4m$, $m \ge 2$, and write the marginal-type words as $A = (0011)^m$ and $B = 0^{m+1}1^{m+1}(01)^{m-1}$, where 0 denotes probability $p = (1 - g)/2$ and 1 denotes probability $1 - p$.

Both directed edge histograms contain m occurrences of each of $00, 01, 10, 11$. In particular both attainable total disagreement floors are $ng/2$. Throughout the proof $v = (1 - g^2)/4 \geq 3/16$. We prove the explicit bounds

$$\frac{9}{64} \min\{n, (g/2 + e)^{-1}\} \leq K_A(e) \leq \min\{n, 8/(g/2 + e)\}, \tag{18}$$

$$\frac{3}{1024} \min\{n, e^{-1}\} \leq K_B(e) \leq \min\{n, 2/e\}. \tag{19}$$

We interpret $1/e = \infty$ at $e = 0$, as specified before Theorem 2. Replacing $g/2 + e$ by $g + e$ changes only a universal factor. Appendix B.3.1 gives the variance lower bound used for A, Appendix B.3.2 the implementations behind both upper bounds, Appendices B.3.3 and B.3.4 the two placements, and Appendix B.3.5 the finite bounds plotted in Figure 3.

B.3.1 A VARIANCE LOWER BOUND FROM TOTAL DISAGREEMENT

For binary decisions with common marginal variance $v > 0$, the identity $\mathbb{P}(Z_i \neq Z_j) = p_i + p_j - 2\mathbb{E}Z_iZ_j$ gives

$$\mathrm{Cov}(Z_i, Z_j) = v - \tfrac{1}{2}\{\mathbb{P}(Z_i \neq Z_j) - (p_i - p_j)^2\} \geq v - \tfrac{1}{2}\mathbb{P}(Z_i \neq Z_j). \tag{20}$$

For a binary vector on a cycle with this common variance, write $d_i = \mathbb{P}(Z_i \neq Z_{i+1})$ and $D = \sum_i d_i$. The resulting lower bound is

$$\Lambda(P) \geq \frac{3}{8} \min\left\{n, \frac{2vn}{D}\right\}. \tag{21}$$

At $D = 0$, the minimum is interpreted as n. To prove it, fix $L \in \{1, \ldots, n\}$ and write $n = qL + r$ with $0 \leq r < L$. Partition the cyclic linear order into q intervals of size L and, if $r > 0$, one interval of size r. If b_B is the size of an interval, then

$$S_2 := \sum_B b_B^2 = qL^2 + r^2 = nL - r(L - r) \geq \tfrac{3}{4}nL. \tag{22}$$

For each ordered pair of vertices in an interval, use the unique path inside that interval. The edge after position t in an interval of size b is used by $2t(b - t)$ ordered paths, and the sum of these incidences is

$$A_b = \sum_{t=1}^{b-1} 2t(b - t) = \frac{b(b^2 - 1)}{3} \leq \frac{b^2(b - 1)}{2}. \tag{23}$$

The final inequality is valid also for $b = 1, 2$.

Apply Lemma 3 to the partition and use (20). Average the partition, including its paths, over all n cyclic translations. Each edge then has average incidence $n^{-1} \sum_B A_{b_B}$, while the distribution and the marginal placement stay fixed. Consequently

$$\begin{aligned} nv\Lambda(P) &\geq vS_2 - \frac{D}{2n} \sum_B A_{b_B} \\ &\geq vS_2 \left(1 - \frac{D(L - 1)}{4vn}\right). \end{aligned} \tag{24}$$

If $D > 0$, take $L = \min\{n, \lfloor 2vn/D \rfloor + 1\}$. The parenthesis is at least $1/2$, and (22) gives $\Lambda(P) \geq 3L/8 \geq (3/8) \min\{n, 2vn/D\}$. If $D = 0$, take $L = n$ and obtain $\Lambda(P) \geq n$. This proves (21).

For $0 < g \leq 1/2$ we have $v \geq 3/16$, hence

$$\Lambda(P) \geq \frac{9}{64} \min\{n, n/D\}. \tag{25}$$

B.3.2 EXPLICIT IMPLEMENTATIONS

Independent blocks at positive excess. For $e > 0$, set $s = \lceil 4v/e \rceil$. If $s \geq n$, use one common uniform across the cycle; this has zero excess and $\Lambda \leq n$. Otherwise partition into $q = \lceil n/s \rceil \geq 2$

consecutive blocks of size at most s and use independent common uniforms within them. Across a boundary, independent Bernoulli variables of means $r, t \in \{p, 1-p\}$ have excess

$$r + t - 2rt - |r - t| = 2\min(r, t)(1 - \max(r, t)) \le 2v. \tag{26}$$

There are exactly q boundary edges, including the wrap edge. Thus $D_{\mathrm{ex}} \le 2vq \le 4vn/s \le ne$. The covariance parameter is at most s, because the blocks are independent and each block has at most s variables of variance v. Since $4v \le 1$ and $e \le 1/4$, $s \le 1/e + 1 \le 2/e$. For every placement of the two values this proves $K(e) \le \min(n, 2/e)$.

Zero excess on an alternating cycle. Let n be even, let the values p and $1-p$ alternate around C_n, and write K_{alt} for the frontier. Complement the high-probability decisions,

$$X_i = Z_i \text{ if } p_i = p, \qquad X_i = 1 - Z_i \text{ if } p_i = 1 - p, \tag{27}$$

so that every X_i has mean p. By (2) the excess of an edge is $2\mathbb{P}(X_i = X_j = 1)$, so zero excess means that X is supported on independent sets of C_n, and any distribution on independent sets with means p is a zero-excess implementation. With S diagonal with entries ± 1, $\mathrm{Cov}(Z) = S\,\mathrm{Cov}(X)S$, so Λ is unchanged. The construction applies to every $0 < g < 1$.

Set $h = \lceil 1/g \rceil \ge 2$. If $n < 2h$, choose a fair global checkerboard and independently retain each selected vertex with probability $1 - g$. Every marginal is p and no adjacent ones occur. For every real a, the universal covariance Cauchy–Schwarz bound gives

$$\mathrm{Var}\Big(\sum_i a_i X_i\Big) \le v\Big(\sum_i |a_i|\Big)^2 \le nv\,\|a\|_2^2. \tag{28}$$

Evenness implies $n \le 2h - 2 < 2/g$ in this branch.

If $n \ge 2h$, let $m = \lfloor n/h \rfloor \ge 2$. Form m consecutive cells with sizes $b_k \in \{\lfloor n/m \rfloor, \lceil n/m \rceil\}$. Each size is at least h. Let $B_{\max} = \max_k b_k$. When $h = 2$, evenness makes all cells have size two. When $h \ge 3$, write $n = mh + r$ with $0 \le r < h$; then

$$B_{\max} = h + \lceil r/m \rceil \le h + \lceil (h-1)/2 \rceil \le 2(h-1) < 2/g. \tag{29}$$

Independently select a uniform deleted vertex in each cell. Also select an independent fair phase for each open clockwise arc between successive deleted vertices. On that arc select precisely the vertices whose global parity equals its phase. Deleted vertices are zero. Every edge either touches a deleted vertex or lies in one arc with a common phase; it cannot have two selected endpoints, including at the cycle's wrap edge. The preliminary selection probability of any vertex in cell k is $(1 - 1/b_k)/2$. Thin it independently with probability

$$\theta_k = \frac{1-g}{1 - 1/b_k} \le 1. \tag{30}$$

This produces marginal mean p exactly. Thinning preserves the independent-set property.

Let $Y_k = \sum_{i \in I_k} a_i (X_i - p)$ denote a centered cell statistic. It depends only on that cell's deletion, its private thinning coins, and the two adjacent arc phases. Thus nonadjacent cells have disjoint primitive randomness. The cell dependency graph has a proper coloring with $\chi \le 3$ colors (one edge when $m = 2$). Within each color class the Y_k are jointly independent. Moreover,

$$\mathrm{Var}(Y_k) \le v\Big(\sum_{i \in I_k} |a_i|\Big)^2 \le v b_k \sum_{i \in I_k} a_i^2. \tag{31}$$

Writing T_c for the centered sum of cells of color c, pointwise Cauchy–Schwarz and then color-class independence give

$$\begin{aligned} \mathrm{Var}\Big(\sum_k Y_k\Big) &\le \chi \sum_c \mathrm{Var}(T_c) = \chi \sum_k \mathrm{Var}(Y_k) \\ &\le 3B_{\max} v\,\|a\|_2^2 \le (6v/g)\,\|a\|_2^2. \end{aligned} \tag{32}$$

Together with (28), this proves a covariance parameter at most $\min(n, 6/g)$, and therefore the same bound on Λ. Combined with the block implementation, for $0 < g \le 1/2$ and $0 \le e \le 1/4$,

$$K_{\mathrm{alt}}(e) \le \min(n, 6/g, 2/e) \le \min\{n, 8/(g+e)\}. \tag{33}$$

B.3.3 PLACEMENT A

The lower bound follows from (25) with $D \le n(g/2+e)$. For the upper bound, put $N = n/2$ and take an alternating-marginal implementation W on C_N. Duplicate each W_j at the two corresponding consecutive vertices of A. Internal pair edges contribute zero disagreement; all other edges correspond one-to-one to the smaller cycle. Consequently total excess is preserved and $e_{\mathrm{lift}} = e_{\mathrm{super}}/2$.

For any original sign vector a, put $c_j = (a_{2j} + a_{2j+1})/2 \in \{-1, 0, 1\}$. Then $\sum_i a_i Z_i = 2\sum_j c_j W_j$. The positive-semidefinite quadratic form $c^\top \mathrm{Cov}(W)c$ attains its maximum over $[-1,1]^N$ at a sign vertex: maximize the convex function successively in each coordinate. Every sign vertex is attainable by choosing identical signs within a pair. It follows that $\Lambda(Z) = 2\Lambda(W)$ exactly for this lifted implementation. Using the alternating upper bound (33) at $e_{\mathrm{super}} = 2e \le 1/4$ gives

$$K_A(e) \le 2K_{\mathrm{alt}}(2e) \le \min\{n, 16/(g+2e)\}.$$

B.3.4 PLACEMENT B

Let I be either long run and $L = m + 1 \ge n/4$. Its internal disagreement D_I is at most ne, because its edges have zero marginal floor and every edgewise excess elsewhere is nonnegative. Partition I into intervals of length H and one possible remainder, where $1 \le H \le L$, and make every outside vertex a singleton. Write S_2 for the sum of squared interval lengths within I. Then $S_2 \ge 3LH/4$. Each internal edge of a block of length $b \le H$ belongs to $2t(b-t) \le H(H-1)$ ordered paths. Within I the marginals are equal, so (20) holds with equality in its first lower bound. Lemma 3 gives

$$nv\Lambda(P) \ge vS_2 - \tfrac{1}{2}H(H-1)D_I + (n-L)v.$$

For $D_I > 0$, take $H = \min\{L, \lfloor vL/(2D_I)\rfloor + 1\}$. The penalty is at most $vLH/4$, so discarding the nonnegative singleton term yields

$$\Lambda(P) \ge \frac{LH}{2n} \ge \frac{L^2}{2n}\min\{1, v/(2D_I)\} \ge \frac{3}{1024}\min\{n, 1/e\}.$$

Here we used $L \ge n/4$, $D_I \le ne$, and $v \ge 3/16$. If $D_I = 0$, all decisions in the run are equal almost surely and $\Lambda(P) \ge L^2/n + (n-L)/n$, which implies the same stated lower bound. The block implementation of Appendix B.3.2 works for any ordering of the two marginal values and gives $K_B(e) \le \min\{n, 2/e\}$. This completes (19).

B.3.5 FINITE LOWER BOUNDS USED IN FIGURE 3

Let $S_2(H) = \lfloor L/H\rfloor H^2 + (L \bmod H)^2$ and $\Gamma_H = \lfloor H^2/2\rfloor$. Partition each long run separately into intervals of length at most H, with all other vertices as singletons. The internal edge sets of the two runs are disjoint, so their combined disagreement is at most ne. The exact maximum ordered-path incidence within an interval is Γ_H. Thus

$$K_B(e) \ge \max_{1\le H\le L} \frac{2vS_2(H) + (n-2L)v - \Gamma_H ne/2}{nv}. \tag{34}$$

Taking $H = 1$ recovers the universal bound one. At $e = 0$ the two long-run decisions A_0, B_0 also satisfy $A_0 \le B_0$ almost surely because their runs are adjacent. The nested distribution has $\mathrm{Cov}(A_0, B_0) = p^2$. Treating their union as one auxiliary-sign block gives the stronger endpoint bound

$$K_B(0) \ge \frac{4L^2}{n(1+g)} + \frac{n-2L}{n}. \tag{35}$$

The remaining singleton terms ensure that these are bounds on Λ for the complete batch, including vertices outside the two long runs.

C PROOFS FOR SECTION 4

Appendix C.1 proves the excess and stability identities (12) and the tail bounds (15) and (16) for bounded blocks, Appendix C.2 proves the agreement bound (14) and the same tails for the Gaussian design, and Appendix C.3 describes how the design with a fixed number of places is calibrated, computed and certified.

C.1 BOUNDED BLOCKS

For an edge ij crossing two blocks, independence gives disagreement $p_i + p_j - 2p_ip_j$. Subtracting its floor $|p_i - p_j|$ gives $2\gamma_{ij}$; within-block edges have zero excess. Summing with the edge weights proves the first identity in (12). Within a block the common uniform gives $\mathrm{Cov}(Z_i, Z_j) = \gamma_{ij} \geq 0$, and across blocks the covariance is zero, so the all-ones sign vector attains the maximum in (3); its value $s_{\mathbf{p}}^2 + 2\sum_B \sum_{i<j\in B} \gamma_{ij}$ gives the second.

Let blocks B of size at most κ carry independent common uniforms, and write $T_a = \sum_B Y_B$ with $Y_B = \sum_{i\in B} a_i(Z_i - p_i)$. The Y_B are independent and centred. Cauchy–Schwarz on the within-block covariances gives $\mathrm{Var}\, Y_B \leq (\sum_{i\in B} |a_i|\sqrt{p_i(1-p_i)})^2 \leq \kappa \sum_{i\in B} a_i^2 p_i(1-p_i)$, and $|Y_B| \leq \kappa M_a$. The scalar Bernstein moment bound for each Y_B, multiplied over blocks, yields

$$\log \mathbb{E} e^{tT_a} \leq \frac{\kappa t^2 S_a}{2(1 - \kappa|t|M_a/3)}, \qquad |t| < \frac{3}{\kappa M_a},$$

the same form as the Gaussian bound (39) below, hence (15) and (16) with the maximum block size as κ.

C.2 GAUSSIAN DESIGN

Let $H \sim N(0, Q)$ with $\mathrm{diag}\, Q = \mathbf{1}$ and $0 \preceq Q \preceq \kappa I$. The Gaussian product inequality of Chen et al. (2015, Theorem 1(i)) states

$$\mathbb{E} \prod_i f_i(H_i) \leq \prod_i \left(\mathbb{E} f_i(H_i)^\kappa\right)^{1/\kappa} \tag{36}$$

for the nonnegative functions used below, including singular Q. Substitute $f_i(h) = \exp(ta_i(\mathbb{I}\{h \leq \Phi^{-1}(p_i)\} - p_i))$, so that $\mathbb{E} f_i(H_i)^\kappa = \mathbb{E} e^{\kappa t a_i (Z_i - p_i)}$. Scalar Hoeffding at $\kappa t a_i$ gives

$$\log \mathbb{E} e^{t\sum_i a_i(Z_i - p_i)} \leq \kappa t^2 \|a\|_2^2 / 8. \tag{37}$$

The scalar variance-sensitive bound

$$\log \mathbb{E} e^{ta_i(Z_i - p_i)} \leq \frac{t^2 a_i^2 p_i(1-p_i)}{2(1 - |ta_i|/3)} \tag{38}$$

at $\kappa t a_i$ in place of ta_i, inserted in (36), gives

$$\log \mathbb{E} e^{t\sum_i a_i(Z_i - p_i)} \leq \frac{\kappa t^2 \sum_i a_i^2 p_i(1-p_i)}{2(1 - \kappa|t|\,\|a\|_\infty/3)}, \qquad |t| < \frac{3}{\kappa \|a\|_\infty}. \tag{39}$$

Deterministic coordinates have zero centered contribution. The second-order term proves $\mathrm{Cov}(Z) \preceq \kappa\, \mathrm{diag}(p_i(1-p_i))$, which implies $\Lambda \leq \kappa$. The right side is the Bernstein moment bound with variance proxy κS_a and scale $\kappa M_a/3$, so the usual Chernoff optimization gives (15); $a = \mathbf{1} - 2y$, $S_a = s_{\mathbf{p}}^2$, $M_a = 1$ and $u = nt$ give (16).

For the Gaussian agreement claim, suppose $p_i \leq p_j$ and write $a = \Phi^{-1}(p_i) \leq b = \Phi^{-1}(p_j)$, $\rho = Q_{ij}$. By (2), edge excess equals $2\mathbb{P}(H_i \leq a, H_j > b)$, bounded by $2F(a)$ where $F(a) = \Phi(a) - \Phi_2(a, a; \rho)$. For $-1 < \rho < 1$,

$$F'(a) = \phi(a)\left[1 - 2\Phi\left(a\sqrt{\frac{1-\rho}{1+\rho}}\right)\right]. \tag{40}$$

It is nonnegative for $a \leq 0$ and nonpositive for $a \geq 0$. Thus F is maximized at zero, with $2F(0) = \arccos(\rho)/\pi$. The endpoints $\rho = \pm 1$ and $p_i, p_j \in \{0, 1\}$ follow by continuity or directly from their deterministic distributions. Finally, $\arccos(1 - 2x) \leq \pi\sqrt{x}$ for $x \in [0, 1]$, applied with $x = (1 - Q_{ij})/2$, and Cauchy–Schwarz over the edges with weights w_{ij}/W give $W^{-1}\sum_{ij\in E} w_{ij} \arccos(Q_{ij})/\pi \leq \sqrt{\mathrm{tr}(\mathcal{L}Q)/(4W)}$, which is the second inequality in (14).

C.3 THE DESIGN WITH A FIXED NUMBER OF PLACES

In the design of Section 4, coordinates with $p_i \in \{0, 1\}$ are fixed, and the support is every $z \in \{0, 1\}^n$ with $\sum_i z_i = K$ that agrees with them; the quota version also requires the selected count N_1 of

one group to lie in $\{\lfloor\mu_1\rfloor, \lceil\mu_1\rceil\}$, $\mu_1 = \sum_{i\in S_1} p_i$, which fixes the other group's count within one of its expectation as well. On this support $P_\theta(z) \propto \exp\{\theta^\top z + \beta E_{\mathcal{B}}(z)\}$ with $E_{\mathcal{B}}(z) = \sum_B \sum_{ij\in E:\, i,j\in B} w_{ij}\mathbb{I}\{z_i = z_j\}$. At $\beta = 0$ it is conditional Poisson sampling, with or without the quota. The function $\theta \mapsto \log Z(\theta) - \theta^\top \mathbf{p}$ is convex with gradient $\mathbb{E}_\theta Z - \mathbf{p}$. The constraints $0 \le x_i \le 1$, on the total and on N_1 form a laminar family, hence a totally unimodular system, so the polytope they define is the convex hull of the support. Whenever $0 < p_i < 1$ for the randomized coordinates, $\mathbf{p}$ satisfies every inequality of this system strictly except those that hold with equality on the whole support, so it lies in the relative interior and a calibrating θ exists. We minimize this function by L-BFGS with the exact gradient, followed by Newton steps that use the exact covariance as Jacobian when needed; every reported design matches each p_i to within 10^{-8}, below the tolerance 10^{-7} fixed in advance.

Each block's $2^{|B|}$ states are enumerated with their counts in total and in group 1; a recursion over blocks gives the normalizing constant, and prefix and suffix products give every marginal. Conditioning on $Z_i = 1$ and repeating gives $\mathbb{P}(Z_i = 1, Z_j = 1)$, hence the covariance C, and sampling backward through the blocks is exact. On 72 random instances with at most 12 coordinates these quantities agree with full enumeration to within 5×10^{-15}. Since $\bar\Lambda \ge \Lambda$ needs only a bound on $\max_{a\in\{-1,1\}^n} a^\top C a$, we report $\bar\Lambda = \min\{n\lambda_{\max}(C),\ \sum_{ij}|C_{ij}|,\ \sum_i u_i\}/s_{\mathbf{p}}^2$, where $\mathrm{Diag}(u) \succeq C$ is checked by its smallest eigenvalue; the last term is a feasible point of the dual of the max-cut semidefinite relaxation. On the small instances $\bar\Lambda$ is 1.00 to 1.11 times the exact Λ.

D Numerical checks for Section 5.1

This section gives the linear programs behind Section 5.1.

D.1 The exact linear program

For each binary state $z \in \{0,1\}^n$, assign a probability w_z. The exact finite program minimizes K subject to

$$\begin{aligned} &w_z \ge 0, \quad \sum_z w_z = 1, \quad \sum_z w_z z_i = p_i \quad (i \in V), \\ &\sum_z w_z \sum_{ij\in E} 2\mathbb{I}\{z_{\text{low}} = 1, z_{\text{high}} = 0\} \le ne, \\ &\sum_z w_z [a^\top (z - \mathbf{p})]^2 \le s_{\mathbf{p}}^2 K \quad \text{for every } a \in \{-1,1\}^n. \end{aligned} \tag{41}$$

For each edge the low endpoint has no larger marginal; ties use a fixed orientation, and (2) makes the middle constraint the expected excess. At $e = 0$, states violating any oriented constraint are discarded in advance. Sign constraints are added until none is violated, and the final maximum over all sign vectors and the dual objective certify each optimum.

D.2 Theorem 2

For placements A and B the grid is $n \in \{8, 12\}$, $g \in \{0.15, 0.30, 0.50\}$ and $e \in \{0, 0.02, 0.05, 0.10, 0.125\}$: 60 programs, with largest marginal residual 2.1×10^{-10} and largest primal–dual difference 2.5×10^{-9}. Figure 3 (left) shows the $n = 12$, $g = 0.30$ slice: the optima are 3.349 for A and 4.735 for B at zero excess, and 2.456 and 2.518 at $d_G = 0.20$, that is $e = 0.05$. The right panel shows how the gap grows with n, beyond the reach of the linear program.

Exact moments of the construction behind A. Consider the alternating cycle C_n of Appendix B.3.2 with $n = mb$, $m \ge 2$, even b, and $1/b \le g < 1$. Use cells of size b in the zero-excess construction and let $q = (1-g)/(1-1/b)$ be the private thinning probability. Every off-diagonal covariance of the transformed Z is nonnegative. In the same cell, vertices at distance $h \ge 1$ have covariance

$$\frac{q^2}{4}\left(\frac{b-1-h}{b} - \frac{(-1)^h}{b^2}\right).$$

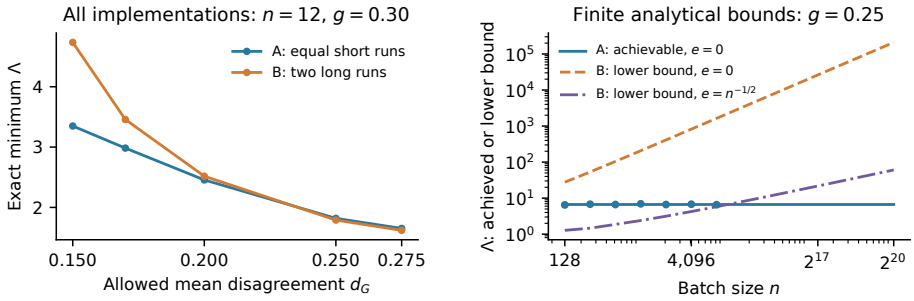


Figure 3: Left: exact minimum of Λ over all joint distributions for placements A and B ($n = 12$, $g = 0.30$) against the allowed mean disagreement d_G. Right, for $g = 0.25$ and log scales: Λ of the explicit zero-excess implementation on A (markers: variance of 5,000 draws) and lower bounds for every implementation on B at $e = 0$ and $e = n^{-1/2}$, against the batch size n (Appendices B.3.5 and D.2).

Between adjacent cells, their shared arc phase contributes $q^2 i(b-1-j)/(4b^2)$ at local positions i, j. When $m = 2$, the second shared phase contributes the same expression with i, j reversed relative to their cell ends. Other cells are independent. Summing both ordered directions gives the exact value

$$\Lambda = 1 + \frac{q^2}{bv}\left\{\frac{(b-1)(b-2)}{6} + \frac{1}{4b} + \frac{(b-1)^2}{8}\right\}. \tag{42}$$

Nonnegativity of the covariances implies that constant signs attain the worst-sign maximum. At $b = 4, g = 1/4$, this value is $10/3$. By Appendix B.3.3, the lifted implementation on A has $2n$ vertices and Λ equal to twice (42); Figure 3 (right) plots these lifted values. For $b \in \{4, 8, 16\}$, $g = 1/b$ and n from 64 to 4096, the variances of 5,000 draws agree with (42) to within 2.2 Monte Carlo standard errors.

D.3 Theorem 1

The heterogeneous programs compute the exact zero-excess optimum $K_{C_n,\mathbf{p}}(0)$ by a linear program over every binary state that satisfies the zero-excess support constraints and every sign vector modulo global reversal. For each $n \in \{3, \ldots, 10\}$, a fixed seed generates five marginal vectors, from highly heterogeneous to nearly constant: independent uniform values on $[0.02, 0.98]$; probabilities near zero, the logistic transform of normal log odds with mean -4 and standard deviation 2; values drawn from $\{0.2, 0.35, 0.7\}$; sorted uniform values on $[0.05, 0.95]$, which increase monotonically around the cycle; and nearly constant values in 0.5 ± 0.02. Two further vectors contain probabilities equal to zero and one, and the placements A and B of Figure 1 complete 44 programs with $3 \le n \le 12$. Every program satisfies (8), and $K(0)/\Psi(\mathbf{p})$ ranges over $[0.595, 1.015]$. The two placements reproduce the optima 3.349 and 4.735 of Appendix D.2, with scores $\Psi(\mathbf{p}) = 5.63$ and 6.53. In the eight monotone programs, $K(0)$ equals $\Psi(\mathbf{p})$ to within 2×10^{-15}.

The copy–reset implementation is computed exactly, by enumerating every copy and reset pattern with at least one reset, for the 41 programs whose probabilities lie in $(0, 1)$ and are not all equal. Its largest marginal error is 7×10^{-16}, its largest deviation from an edge floor is 4×10^{-16}, every covariance is nonnegative, and $\Lambda/\Psi(\mathbf{p})$ ranges over $[1.000, 1.068]$.

Larger cycles use two probability values $(1 \pm g)/2$, arranged in $R \in \{2, 4, \ldots, 12\}$ runs of equal or fixed unequal lengths, with $g \in \{1/2, 2/5, 3/10, 1/4, 1/10\}$. At zero excess each run moves as one unit, so these 60 programs are solved on the run-level quotient, which reaches 68 vertices; all residuals are below 10^{-8}. With the run form of $\Psi(\mathbf{p})$ (Appendix B.2), their optima satisfy (8), and $K(0)/\Psi(\mathbf{p})$ ranges over $[0.489, 1.042]$. Over all 104 programs the ratio therefore lies in $[0.489, 1.042]$, the interval reported in Section 5.1.

E THE FAIRNESS-POLICY EXPERIMENT OF SECTION 5.2

Appendices E.1 and E.2 give the setup. Appendices E.3–E.5 support the Results paragraph, and Appendices E.6–E.7 the paragraph on choosing a design before any draw. Appendix E.8 supports the partition choice of Section 4, and Appendix E.9 checks sensitivity to the graph.

E.1 DATA, PREDICTORS AND POSTPROCESSOR

Adult uses the official UCI training and test split; training rows are split 80/20 into fitting and calibration, and 8,192 test rows form four batches of 2,048. ACS uses the 2018 one-year person data for California and Texas through Folktables 0.0.12 (Ding et al., 2021), with the ACSIncome and ACSPublicCoverage tasks and their published filters, labels and covariates; each state and task is split 60/20/20 into fitting, calibration and test, and the first 2,048 test rows form two batches. Sex is excluded from the predictors and the graph features and serves only as the postprocessor's group. The predictors are a regularized logistic regression (Adult and ACS) and histogram gradient boosting (ACS), which gives eight ACS predictors. The equalized-odds postprocessor (Hardt et al., 2016) has one decision probability for each group and hard prediction; a linear program on the calibration split maximizes expected accuracy subject to equal true- and false-positive rates across groups, and the fitted probabilities are then frozen. All splits and seeds are fixed in advance, and nothing is chosen on the evaluation batches. Adult is distributed under CC BY 4.0, and ACS PUMS data are supplied by the U.S. Census Bureau through the Folktables loader (MIT license).

E.2 GRAPH, DESIGNS AND METRICS

The graph is the union of ten-nearest-neighbour pairs under Euclidean distance in the fitted representation, each weighted by $1/(1 + \text{distance})$. Every fitted postprocessor assigns probability zero or one to at least two of its four cells, so 51.6–53.0% of each Adult batch and 41.6–58.2% of each ACS batch are deterministic. Their edges count in d_G, and the designs differ only on the remaining subjects.

Graph blocks use the partition of Appendix E.8, computed from $\mathbf{p}$ and the graph only. Random blocks cut a fixed label-independent permutation into groups of at most the cap. The Gaussian design solves (13) on the randomized vertices by ADMM, and each solution is repaired so that it meets the cap and the unit diagonal exactly.

For blocks, d_G, Λ and the variance of every linear statistic are exact, from the within-block covariance γ_{ij} (Appendix C.1); for Gaussian designs, bivariate normal probabilities give exact edge disagreements and covariances. Exceedance rates use 2,000 draws of each design. The certified radius of (16) at confidence $1 - \delta$ is

$$r_\delta = \min\left\{1,\ \frac{\kappa x}{3n} + \sqrt{\frac{2\kappa s_{\mathbf{p}}^2 x}{n^2} + \left(\frac{\kappa x}{3n}\right)^2}\right\}, \qquad x = \log(2/\delta);$$

in every design setting at most 0.65% of draws exceed the 95% radius. Expected accuracy is computed from $\mathbf{p}$ and is the same for every design (4).

E.3 ZERO EXCESS AND RESULTS BY PREDICTOR

This subsection supports two statements of the Results paragraph: zero excess is out of reach, and under the same cap graph blocks keep less excess than the other designs. Table 3 reports, for each predictor, the zero-excess bound (10), the scale of disagreement, and the excess kept at cap 16. The bound is at least 46 in every batch, above every tested cap. At cap 16 graph blocks keep 16.2–45.3% of the independent excess, the Gaussian design 43.4–56.0% and random blocks 96.8–98.5%.

E.4 COMPARISONS AT THE SAME STANDARD DEVIATION OF A STATISTIC

This subsection supports the matched comparison of Figure 2. Table 4 uses statistics that need no outcome: the overall action rate ($a = \mathbf{1}$), the action rate of each postprocessor group (a the indicator of S_0 or S_1; S_0 is the female group) and the difference of the two rates, with classification loss

Table 3: Means over the batches of each predictor. Accuracy: expected accuracy (percent), the same for every design. Zero-excess Λ bound: range over batches of (10). d_G: disagreement (percent) under independent draws and its floor b_G, which the common uniform attains. Excess: e_G at cap 16 as a percentage of its value under independent draws. G, N, R: graph blocks, Gaussian design, random blocks.

Predictor	Accuracy	Zero-excess Λ bound	d_G (%) Indep.	Floor	Excess at cap 16 (%) G	N	R
Adult	82.5	451–482	13.96	10.01	45.3	52.7	98.5
CA Income HGB	82.0	273–288	24.76	24.11	36.7	55.1	97.8
CA Income Logistic	81.7	247–266	25.21	24.42	38.2	56.0	96.8
CA Coverage HGB	72.2	420–429	20.02	18.08	37.6	50.3	97.6
CA Coverage Logistic	68.7	312–313	16.05	14.55	36.6	51.3	97.1
TX Income HGB	79.6	291–305	25.37	24.04	39.3	54.5	97.2
TX Income Logistic	80.1	262–264	25.57	24.18	40.2	55.5	96.9
TX Coverage HGB	85.4	46–56	8.40	7.84	16.2	43.4	97.9
TX Coverage Logistic	83.1	256–267	6.74	5.33	36.4	51.5	97.3

Table 4: Standard deviation: exact standard deviation of the statistic (percentage points) under independent draws and the common uniform, mean over 4 Adult or 16 ACS batches. Matched wins: comparisons at the same standard deviation in which graph blocks have smaller excess than random blocks, one per batch and graph-block cap; points outside the random-block range are not compared. Excess gap: random minus graph-block excess, relative to the independent excess, over these comparisons.

Statistic	Standard deviation (pp) Indep.	Common	Matched wins	Excess gap random − graph
Adult, overall action rate	0.39	9.14	16/16	0.11–0.54
Adult, action rate in S_0	0.88	22.01	16/16	0.11–0.54
Adult, action rate in S_1	0.39	7.28	16/16	0.11–0.52
Adult, rate in S_1 minus S_0	0.96	22.62	16/16	0.11–0.54
Adult, classification loss	0.39	6.70	16/16	0.11–0.54
ACS, overall action rate	0.34	5.64	62/62	0.11–0.61
ACS, action rate in S_0	0.48	9.22	62/62	0.11–0.85
ACS, action rate in S_1	0.43	4.91	58/58	0.10–0.61
ACS, rate in S_1 minus S_0	0.69	10.91	62/62	0.11–0.61
ACS, classification loss	0.34	3.42	61/61	0.11–0.80

($a = \mathbf{1} - 2y$) as the labelled case. Their standard deviations are exact: within a block $\mathrm{Cov}(Z_i, Z_j) = \gamma_{ij}$, and for a Gaussian design $\mathrm{Cov}(Z_i, Z_j) = \Phi_2(\Phi^{-1}(p_i), \Phi^{-1}(p_j); Q_{ij}) - p_i p_j$. For the Adult S_0 rate the standard deviation is 0.88 percentage points under independent draws and 22.0 under the common uniform. Random blocks are also run at caps 32 and 64 so that their curve spans the standard deviations that graph blocks reach. Each graph-block point, one per batch and cap (4×4 on Adult), is compared with the random-block excess interpolated at the same standard deviation; points outside the random-block range are skipped, which leaves 58 to 62 of the 64 ACS points. Graph blocks win every comparison. Since block covariances are nonnegative, the overall action count has variance ratio exactly Λ, so for that statistic this comparison repeats the one in Appendix E.5. Against graph blocks, Gaussian designs are ahead in 57 of 60 Adult comparisons, by at most 0.024 of the independent excess, and in 55 of 240 ACS comparisons.

E.5 Comparisons at the same worst-case stability

This subsection compares the two block families at the same Λ rather than at the same cap. We mix each family's cap-$2, 4, 8, 16$ designs with independent draws and the common uniform: one design is chosen with probabilities λ_j and the decisions are drawn from it. Every component has mean $\mathbf{p}$ and

Table 5: Target Λ: worst-case stability at which each block family's designs are mixed with independent draws and the common uniform. Excess: e_G of the least-disagreement mixture as a percentage of its value under independent draws, mean over 4 Adult, 8 ACS Income or 8 ACS Coverage batches. Paired wins: batches in which the graph family has the smaller excess.

		Excess (%)		
Dataset	Target Λ	Graph	Random	Paired wins
Adult	2	88.0	99.8	4/4
	4	73.0	99.4	4/4
	8	57.4	98.6	4/4
	16	45.1	97.2	4/4
ACS Income	2	84.2	99.5	8/8
	4	65.3	98.7	8/8
	8	48.7	97.2	8/8
	16	38.3	94.5	8/8
ACS Coverage	2	84.5	99.2	8/8
	4	65.8	97.9	8/8
	8	45.8	95.4	8/8
	16	31.2	90.5	8/8

Table 6: t: target loss deviation at $\delta = 0.05$. Cap: selected cap, the same in every batch. r_δ: certified radius of (16). d_G: disagreement of the selected design; $\Delta d_G = d_G(P_{\mathrm{ind}}) - d_G(P_{\mathrm{selected}})$. Means over 4 Adult or 8 ACS batches.

Dataset	t (pp)	Cap	r_δ (pp)	d_G (%)	Δd_G (pp)
Adult	3	4	2.39	12.96	1.00
	5	8	3.53	12.34	1.62
ACS Income	3	4	2.39	24.90	0.33
	5	8	3.75	24.72	0.51
ACS Coverage	3	4	2.39	12.41	0.39
	5	8	3.74	12.15	0.65

entrywise nonnegative covariance, so

$$\Lambda\left(\sum_j \lambda_j P_j\right) = \sum_j \lambda_j \Lambda(P_j), \qquad d_G\left(\sum_j \lambda_j P_j\right) = \sum_j \lambda_j d_G(P_j),$$

and the least d_G at a target $\Lambda \in \{2, 4, 8, 16\}$ is a linear program whose optimum mixes at most two designs. The common uniform has Λ between 62 and 547, so every target is reached. Table 5 reports the excess of the best mixture at each target. The graph family has smaller excess in every batch at every target; at $\Lambda = 16$ it keeps 13–46% of the independent excess across batches, against 75–97% for random blocks. A mixture has the stated Λ but not the block-size tail bound.

E.6 Selecting a cap from a loss-deviation target

This subsection supports the loss example in the paragraph on choosing a design in Section 5.2. For target loss deviations of 3 and 5 percentage points at $\delta = 0.05$, we evaluate the cap limit in (16) on each batch and select the largest available $\kappa \in \{1, 2, 4, 8, 16\}$ below it; cap one is independence and the other caps use the graph blocks of Section 5.2. Every batch admits a cap at both targets. Table 6 reports, for each dataset and target, the selected cap, its certified radius and the disagreement it attains. The guarantee holds for labels fixed before the draw.

E.7 Selecting a cap from a label-free group target

This subsection supports the group-rate example in the paragraph on choosing a design. For a group statistic the uniform bound $\mathrm{Var}(T_a) \le s_{\mathbf{p}}^2 \Lambda(P) \|a\|_\infty^2$ is loose, so the coefficient-specific scale S_a in (15) admits a larger cap. For the action rate of each group and their difference on Adult, we compare

Table 7: Adult, $\delta = 0.05$. t: target deviation of the statistic. Cap: largest cap in $\{1, 2, 4, 8, 16, 32, 64\}$ certified by the rule of Section 4 with the coefficient-specific scale S_a or the full-batch scale $s_{\mathbf{p}}^2 \|a\|_\infty^2$; a/b means that the four batches select caps a and b, and – that the rule certifies no cap. Excess: graph-block excess relative to independent draws at the selected cap, mean over batches. Exceedance: largest share of 2,000 draws exceeding the target under either selected design, against the allowed 5%.

Statistic	t (pp)	Cap, S_a	Cap, full	Excess, S_a	Excess, full	Exceedance (%)
rate in S_0	3	1	–	1.000	–	0.15
rate in S_0	5	2	1/2	0.889	0.916	0.00
rate in S_0	10	8	4	0.590	0.746	0.05
rate in S_1	3	4	2	0.746	0.889	0.10
rate in S_1	5	8	4	0.590	0.746	0.00
rate in S_1	10	32	16	0.347	0.453	0.05
S_1 minus S_0	3	1	–	1.000	–	0.30
S_1 minus S_0	5	2	1/2	0.889	0.916	0.05
S_1 minus S_0	10	8	4	0.590	0.746	0.05

Table 8: Means over 4 Adult or 16 ACS batches of the excess divided by its value under independent draws and of the attained Λ, at each block cap. W: merge along edges in decreasing weight; H: merge the pair with the largest saved excess H_{AB}; R: the pair with the largest H_{AB}/C_{AB}; S: per batch and cap, the smallest exact excess of W, H and R.

		$e_G/e_G(P_{\mathrm{ind}})$				attained Λ			
Data	Cap	W	H	R	S	W	H	R	S
Adult	2	0.889	0.889	0.889	0.889	1.88	1.89	1.88	1.89
	4	0.763	0.746	0.753	0.746	3.61	3.64	3.43	3.64
	8	0.634	0.590	0.598	0.590	6.97	7.21	6.51	7.21
	16	0.528	0.455	0.460	0.453	13.60	14.29	12.58	14.01
ACS	2	0.876	0.860	0.876	0.860	1.75	1.85	1.75	1.85
	4	0.738	0.684	0.723	0.684	3.12	3.51	3.02	3.51
	8	0.606	0.501	0.553	0.501	5.70	6.84	5.39	6.80
	16	0.499	0.351	0.401	0.351	10.43	13.71	9.87	13.71

the cap selected with S_a against the same rule with the full-batch scale $s_{\mathbf{p}}^2 \|a\|_\infty^2$, for targets of 3, 5 and 10 percentage points, $\delta = 0.05$ and caps $\{1, 2, 4, 8, 16, 32, 64\}$, where cap one is independence. The selection uses only the marginals and the group labels. In Table 7 the coefficient-specific rule selects twice the cap at the 10-point targets and for the S_1 rate at every target, and it certifies independence where the full-batch rule certifies no cap. No selected design exceeds its target in more than 0.3% of draws, against the allowed 5%.

E.8 Choosing the block partition

This subsection supports the choice among greedy partitions in Section 4. Every merge order starts from singleton randomized vertices and merges two blocks joined by an edge while the merged block has at most κ vertices. W merges along edges in decreasing weight, H picks the pair with the largest saved excess H_{AB}, and R the pair with the largest H_{AB}/C_{AB}; ties follow the weight order, and R coincides with W at cap 2. The selection S, used for every graph block in Section 5.2, keeps the order with the smallest exact excess (12) for each batch and cap. Table 8 compares the orders at the same cap. S lowers excess relative to W in all 80 batch–cap pairs; the reduction is smallest at cap 2 (0.001 of the independent excess on Adult and 0.016 on ACS) and grows with the cap, to $0.53 \to 0.45$ on Adult and $0.50 \to 0.35$ on ACS at cap 16. S also raises the attained Λ: at the same attained Λ, S is below W in 55 of 79 comparable pairs and above in 24, by at most 0.034 of the independent excess, so part of the gain comes from using more of the cap.

Table 9: Adult. k: nearest neighbours; weight: $1/(1 + \text{distance})$ (inverse) or 1 (unit). Edges: mean number per batch. Excess at cap 16: graph-block excess relative to independent draws, range over batches. Gap: random minus graph-block excess at the same attained Λ, mean of batch means and range over single comparisons. Graph better: comparisons won by graph blocks.

k	weight	edges	excess at cap 16	mean gap	gap range	graph better
5	inverse	7854	0.305–0.309	0.453	0.183–0.675	16/16
5	unit	7854	0.340–0.349	0.419	0.164–0.638	16/16
10	inverse	15542	0.438–0.466	0.320	0.108–0.538	16/16
10	unit	15542	0.486–0.494	0.288	0.092–0.486	16/16
20	inverse	30904	0.591–0.602	0.213	0.062–0.385	16/16
20	unit	30904	0.617–0.630	0.191	0.050–0.359	16/16
40	inverse	61138	0.709–0.720	0.138	0.036–0.265	16/16
40	unit	61138	0.731–0.740	0.121	0.027–0.244	16/16

E.9 Sensitivity to the graph

This robustness check is not used in the main text. We repeat the Adult comparison on eight graphs, with $k \in \{5, 10, 20, 40\}$ nearest neighbours and inverse-distance or unit weights, keeping the data, the postprocessor and hence $\mathbf{p}$ fixed. As in Appendix E.4, random blocks extend to caps 32 and 64, and each graph-block point is compared with the random-block excess interpolated at the same attained Λ. In Table 9 graph blocks win all 128 comparisons, and the gain shrinks as the graph becomes denser.

F The exact-capacity lottery of Section 5.3

This section gives the data, designs and checks of Section 5.3; the design itself is in Appendix C.3.

F.1 Data, scores and graph

Of the 75 surveyed villages of Banerjee et al. (2013), the 72 with at least 120 surveyed people have 16,681 people; the data are the public replication files of Jackson et al. (2012). The score is a logistic regression for current non-employment on age, education, native status and caste, fitted on four of five folds of villages and applied to the fifth (held-out area under the ROC curve 0.67 to 0.71); it serves only to set $\mathbf{p}$. Within a village $p_i = \min(1, \lambda s_i)$ with $\sum_i p_i = K$, and no p_i reaches one. The graph joins every pair of surveyed members of a household and every surveyed pair named as close relatives. The help-seeking edges of Section 5.3 are pairs in which one person names the other as someone they ask for help with a decision and which are not in the graph. We study six villages at size ranks $1, 14, 26, 39, 51, 64$ of 72 (124 to 304 people; K = 12–30 at $B = 0.1$ and 37–91 at $B = 0.3$).

F.2 Designs, selection and checks

This subsection gives the candidates behind Table 1, their checks and the results not shown in that table. The graph-aware candidates use block caps 4 and 8 and $\beta \in \{0.5, 1, 2, 5\}$, with the partition built by merging along edges in decreasing $w_{ij}\gamma_{ij}$ and, as a control, two random partitions of the same caps. For each certified cap $\kappa \in \{2, 4, 8\}$ we keep the candidate with the smallest d_G among those with $\bar{\Lambda} \le \kappa$. At $\kappa = 8$ the selected designs sit at the top of the β grid, so the gains there may be understated. The pivotal method processes people in survey order and is evaluated from 20,000 draws. Without the quota, conditional Poisson sampling exceeds a five-point gender deviation with probability 0.005 to 0.10 at $B = 0.1$ and 0.08 to 0.21 at $B = 0.3$, and the pivotal method with 0.006 to 0.11 and 0.11 to 0.26; the quota designs never do. Each selected design was sampled 10,000 times: every draw selected exactly K people with an admissible gender count, the largest standardized marginal deviation was 3.3, and the sampled d_G was within 2.0 standard errors of the exact value. Figure 4 shows every candidate: for each village and budget, its certified $\bar{\Lambda}$ against its d_G, one curve per partition as β grows.

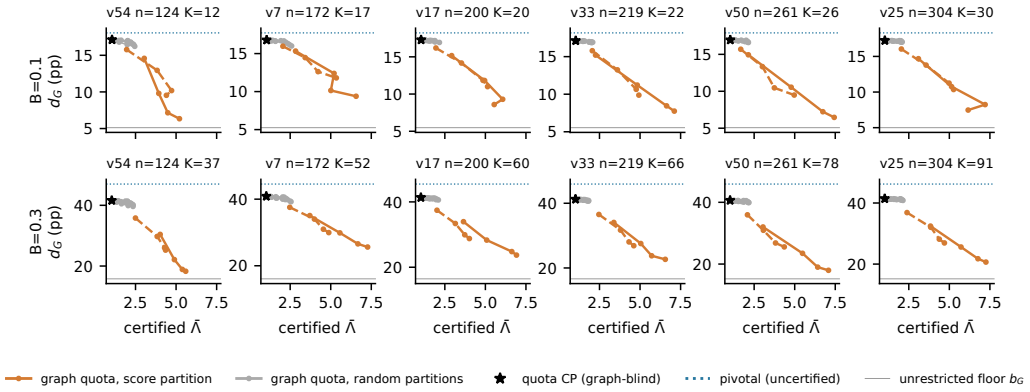


Figure 4: One panel per village (columns) and budget (rows; n and K in the title). Horizontal: certified $\bar{\Lambda}$; vertical: d_G in percentage points. Orange: graph-aware quota design on the score partition (dashed: block cap 4; solid: 8), one point per $\beta \in \{0.5, 1, 2, 5\}$; grey: the same on two random partitions; star: quota CP; dotted: pivotal method (uncertified); thin line: unrestricted floor b_G.

Results not in Table 1. On the two random partitions the same interaction lowers disagreement relative to quota CP by only 0.15 to 1.74 percentage points. The pivotal method has higher disagreement than quota CP in all twelve village–budget pairs, by 0.93 to 5.39 points, and without a quota its gender deviation is not controlled (above). On help-seeking edges outside the graph, the selected graph-aware designs change disagreement relative to quota CP by -0.09 to 3.52 points, so the benefit stays on the relations the design uses.