EXCHANGE BARRIERS AND EXACT RECOVERY IN EXPERIMENTAL DESIGN

Anonymous authors
Paper under double-blind review

ABSTRACT

Batch active learning and optimal experimental design commonly select measurements through individual acquisition scores or greedy batch construction. We identify information and exchange barriers when a prediction target requires several measurements jointly. The closure objective's submodularity ratio is exactly zero, making ratio-based weak-submodularity guarantees vacuous. Any strategy accessing the instance only through order-j information values and purchased responses cannot distinguish two targets whose minimal closing sets are disjoint and have size $j+1$, regardless of its number of oracle calls. With known features, even greedy selection using exact risk can terminate at strict local optima. For paired designs with cycle-structured nuisance, escaping a complete-pair starting batch requires $\Theta(\min\{b, 1+\sqrt{t}\})$ replacements, where t scales nuisance strength and b is the budget. The same structure that obstructs local improvement enables global optimization: for paired designs with graph-structured nuisance, we recover hidden measurement pairs from the candidate Gram matrix and compute an exactly optimal batch in polynomial time. These results include full-row-rank designs with positive reconstruction residual. Witness search is polynomial on regular matroids; closing the uniform grand mean is NP-complete on complete three-factor interaction grids. On a real perovskite task, an automatically discovered three-measurement exchange reduces design risk by 13.4% beyond complete one/two-measurement descent. Removing cross-candidate nuisance correlations while preserving singleton risk and information scores reverses this gain, isolating joint geometry as the mechanism that pointwise acquisition scores leave unresolved.

1 INTRODUCTION

Modern machine learning increasingly runs on acquisition budgets. Choosing which examples to label, which reactions to run, and which measurements to buy connects batch active learning with optimal experimental design. Common approaches rank individual candidates or build batches greedily. They include uncertainty and leverage scores, prediction-oriented information gain, and greedy C-optimal design. The central question is whether such decisions can recover value that only emerges from acquiring several measurements together.

Consider measurements that each carry a weak target signal mixed with nuisance components. A particular set cancels those components and reconstructs the target; its reconstruction coefficients form a *witness*. Stronger individual distractors can attract greedy selection. At a fixed budget of four, replacing one, two or three distractors by members of the reconstructing set increases risk; replacing all four reduces it to 9.6% (Figure 1). Substitution contrasts across unmeasured factor combinations motivate this reconstruction problem in materials and reaction design.

Our results identify the barriers and a constructive solution. Joint reconstruction forces the closure objective's submodularity ratio to zero (Theorem 2.2). When feature rows are withheld, order-j information cannot identify a closing set of size $j+1$ (Theorem 3.1). With known features, the minimum improving exchange in paired cycle designs grows with nuisance strength (Theorem 4.6). Yet hidden-pair recovery yields an exactly optimal batch in polynomial time for paired graph designs (Theorem 4.3).

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.

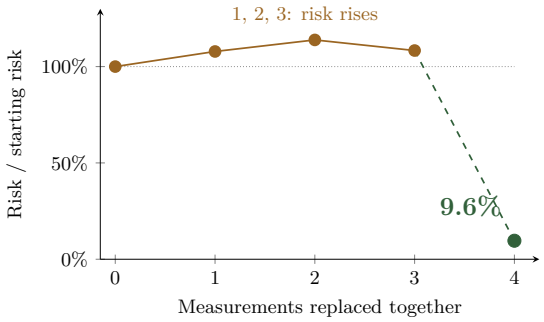

Figure 1: **Every partial replacement increases risk; only the complete set pays off.** Replacing one, two or three measurements raises risk; replacing four reduces it to 9.6%. The budget stays at four throughout. $\lambda = 0.01$, $\varepsilon = 0.2$, $\delta = 0.4$ (Theorem 3.3).

The exchange-size law is

$$r_* = \Theta_{a,\delta}\big(\min\{b, 1 + \sqrt{t}\}\big).$$

It holds from every complete-pair start: b is the budget, t scales the cycle Laplacian, a is signal strength, and $0 \le \delta < a$ is diagonal nuisance. Matching bounds exclude smaller exchanges and construct an improving balanced interval (Figure 2).

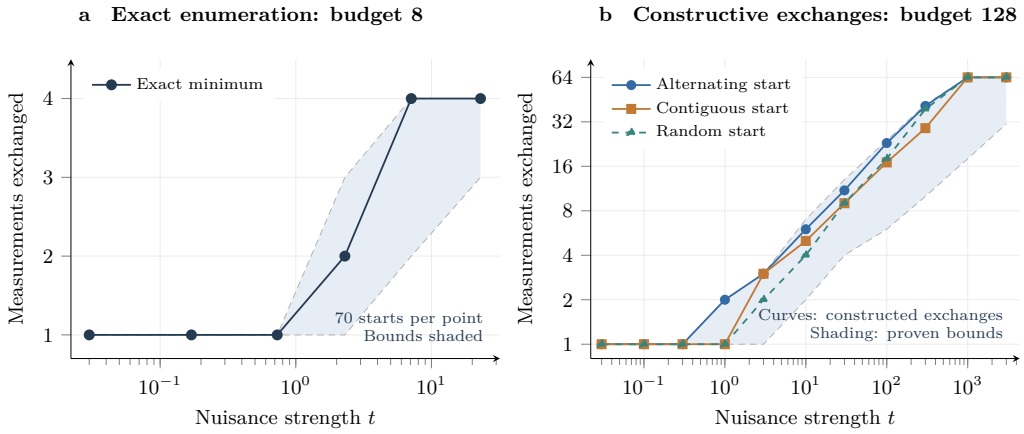

Figure 2: **Exchange size grows with nuisance strength.** Cycle designs with $a = 1$ and $\delta = 0.2$. (a) Budget 8: exhaustive minima agree across all 70 complete-pair starts at each tested t. (b) Budget 128: improving balanced-interval constructions from three fixed starts give upper bounds on the minimum exchange size. Shading joins proven bounds; lines connect tested scales.

Recovery covers connected weighted graphs with nuisance Gram $C = \delta I + tL$, including coinciding signal and nuisance eigenvalues. For $\delta > 0$, all candidate rows are independent, so the result extends to designs where every affordable batch has positive reconstruction residual. Witness search also has contrasting computational regimes. It is polynomial on regular matroids, while closing the uniform grand mean is NP-complete on complete three-factor interaction grids (Theorem 5.2).

Position within active learning and data selection. Optimal experimental design provides the prediction objective (Pukelsheim, 2006; Chaloner & Verdinelli, 1995; Meyer et al., 1996). EPIG targets predictive information (Bickford Smith et al., 2023); BatchBALD accounts for batch dependence through joint parameter information and greedy construction (Kirsch et al., 2019). We separate the information a criterion uses from the exchanges its optimizer can make. This distinction matters for batch annotation and data acquisition, including pipelines that curate training data for large models (Tirumala et al., 2023).

In HOIP, removing nuisance correlations reverses a three-measurement exchange's gain while preserving singleton scores, connecting the observed improvement to joint measurement geometry (Section 6).

2 Query Closure and the Exact Zero Ratio

Let $\varphi : \mathcal{X} \to \mathbb{R}^p$ be a feature map on a finite set of measurements, and let Φ_D stack the rows $\varphi_x^\top$ for $x \in D$. The observed set is T, the purchasable pool is $\mathcal{C} \subseteq \mathcal{X} \setminus T$, and a purchase S must satisfy $c(S) = \sum_{x \in S} c_x \le B$. A deployment target $\theta_q = \varphi_q^\top \beta$ cannot itself be purchased. In the real-data studies, these targets are substitution contrasts.

Observations satisfy $y_x = \varphi_x^\top \beta + \varepsilon_x$, with $\beta \sim N(0, \rho^2 I)$ and independent noise $\varepsilon_x \sim N(0, \sigma^2)$. Write $\lambda = \sigma^2/\rho^2$, $G_D = \Phi_D^\top \Phi_D$, and

$$R(S) = \sigma^2 r_\lambda^{T \cup S}(q) = \sigma^2 \varphi_q^\top (G_{T \cup S} + \lambda I)^{-1} \varphi_q.$$

For several targets we sum their risks. Let $P^{(D)}$ and $P_\perp^{(D)}$ project onto $\mathrm{row}(\Phi_D)$ and its orthogonal complement. The squared reconstruction residual $A(S) = \|P_\perp^{(T \cup S)} \varphi_q\|^2$ contributes $\rho^2 A(S)$ to risk (Theorem A.2).

Definition 2.1 (Witness and query closure)**.** A witness for q on D is a coefficient vector w satisfying $\Phi_D^\top w = \varphi_q$. The query is *closed* when such a vector exists, and *open* otherwise. Write $\mathrm{cl}_q(D) = \mathbf{1}\{\varphi_q \in \mathrm{row}(\Phi_D)\}$ and $\mathrm{cl}_Q(D) = \sum_{q \in Q} \mathrm{cl}_q(D)$. A minimal closing purchase loses closure when any purchased row is removed.

Closure depends on the measurement rows. The same weighted combination of their mean responses reconstructs the target, and a minimal purchase specifies exactly which measurements must be acquired together.

For $F(S) = \mathrm{cl}_q(T \cup S)$, use the submodularity ratio of Das & Kempe (2018, Definition 2):

$$\gamma_{U,k}(F) = \min_{\substack{L \subseteq U,\ |S| \le k \\ S \cap L = \varnothing}} \frac{\sum_{x \in S} [F(L \cup \{x\}) - F(L)]}{F(L \cup S) - F(L)}, \qquad 0/0 := 1.$$

Theorem 2.2 (Exact zero submodularity ratio)**.** *Suppose q is open under T and U contains a minimal closing purchase C of size $m \ge 2$. Then $\gamma_{U,k}(F) = 0$ for every $k \ge 2$.*

Proof. Choose distinct $c_1, c_2 \in C$, let $L = C \setminus \{c_1, c_2\}$, and set $S = \{c_1, c_2\}$. Minimality gives $F(L) = F(L \cup \{c_1\}) = F(L \cup \{c_2\}) = 0$, whereas $F(L \cup S) = F(C) = 1$. The displayed ratio is $(0 + 0)/1 = 0$; monotonicity of closure supplies the matching lower bound. □

This identity holds at every noise level because closure is an algebraic property. For predictive mutual information, the corresponding construction gives $0 \le \gamma_{U,k}(\psi_\lambda) \le c_0 / \log(1/\lambda) \to 0$ (Corollary A.12). The next section resolves which closing sets remain hidden when the permitted interaction order is fixed.

3 Oracle Lower Bounds and Exchange Barriers

We separate access to information from the moves used in optimisation. An interaction-order bound limits which sets an oracle can evaluate; an exchange-size bound limits how an algorithm can change its current batch. Table 1 locates familiar acquisition rules along these two dimensions.

3.1 An interaction-order oracle lower bound

An order-j value oracle returns $\psi_{\mathcal{I}}(A) = I(y_q; y_A \mid y_T)$ for any set $|A| \le j$. A strategy in Π_j accesses the instance through these values, candidate identities and costs, and responses to purchased measurements. It may make arbitrarily many oracle calls, randomise, and adapt its purchases; the feature rows are not part of this oracle's input (Definition B.1).

Theorem 3.1 (Indistinguishable closing sets)**.** *For every $j \ge 1$, there are two instances on $2(j+1)$ candidates whose order-j oracle values and entire purchasable response laws coincide, but whose minimal closing sets are disjoint and have size $j+1$. Every strategy in Π_j buying at most $j+1$ measurements closes the query with probability at most $1/2$ on a uniformly chosen instance. Failed purchases have risk ratio $\Omega(1/\lambda)$ relative to a closing purchase as $\lambda \to 0$.*

Table 1: **Information used to choose a batch.** Interaction order concerns the objective a rule evaluates; exchange size concerns the move it can make. The two barriers in this section address these separately.

method	information order	batch construction
uncertainty; ridge leverage (El Alaoui & Mahoney, 2015)	pointwise scores	rank individual candidates
EPIG (Bickford Smith et al., 2023)	single-query MI, Π_1	rank individual candidates
second-order predictive-MI truncation	pairwise MI, Π_2	optimise the truncated objective
CoreSet; BADGE (Sener & Savarese, 2018; Ash et al., 2020)	embedding geometry	distances or gradient diversity
BatchBALD (Kirsch et al., 2019)	joint parameter MI	greedy batch expansion
greedy predictive MI (exact)	growing batch; $> j$	exact objective, one cell per step
Witness exchange	**complete reconstruction**	**joint proposal, exact-risk acceptance**

Proof. Put $k = j + 1$. Join vertices a, b by two internally disjoint paths P, P' of k edges each, and give an edge $e = (u, v)$ the row $\varphi_e = e_u - s_e e_v$, with $s_e \in \{-1, 1\}$. Choose path-sign products $+1$ on P and -1 on P'. The $2k$ purchasable rows are identical in both instances. Their query rows are $q_1 = e_a - e_b$ and $q_2 = e_a + e_b$, respectively; T is empty. The signed sum along P reconstructs q_1, and the signed sum along P' reconstructs q_2.

All accessible information values coincide. Add the query as an edge between a and b. Every cycle then has at least $k + 1 = j + 2$ edges. For $|A| \le j$, both A and $A \cup \{q\}$ are forests. On a forest, a sign change of vertex coordinates, followed by row sign changes, transforms either signing into the other. Thus the observation Gram matrices are related by conjugation with a diagonal ± 1 matrix, preserving every covariance determinant. Writing $\Sigma_D = \rho^2 \Phi_D \Phi_D^\top + \sigma^2 I$,

$$\psi(A) = \frac{1}{2} \log \frac{\det \Sigma_A \det \Sigma_q}{\det \Sigma_{A \cup \{q\}}}$$

is identical in both instances. The purchased responses have the same joint law because their feature rows are identical.

The transcript cannot select both closing paths. Couple the two runs with the same internal randomness and the same full candidate-response vector. Every answer, purchase, and observed response agrees, so the final set S agrees. Within $|S| \le k$, a query can close only along a complete a–b path: any dangling edge has an uncancelled internal endpoint. The path P closes only q_1, and P' only q_2. Their disjoint union costs $2k$, so S closes at most one target. Averaging over the two instances proves the $1/2$ bound.

The risk separation follows from the remaining target direction. On a failed purchase, $\eta = \|P_{\mathrm{row}(\Phi_S)^\perp} q\|^2 > 0$, giving risk at least $\rho^2 \eta$. The complete path has coefficients of magnitude one, so its risk is at most $\sigma^2 k$. Their ratio is at least $\eta/(k\lambda)$. There are finitely many feasible failed purchases at fixed j, hence their minimum positive η supplies the stated uniform order. The bound is attained in success probability by purchasing either complete path. □

Proposition 3.2 (Matching order upper bound)**.** *For one Gaussian deployment target and cardinality budget k, an order-k value oracle finds a globally minimum-risk batch with at most $\sum_{r=0}^{k} \binom{n}{r}$ calls.*

Proof. Evaluate every feasible set and maximize predictive mutual information. Its conditional prior entropy is fixed, and its posterior entropy is strictly increasing in the scalar target variance. Independent query observation noise adds the same constant to that variance. □

Thus order $k - 1$ can fail at budget k, whereas order k suffices, using $n^{O(k)}$ calls. This is polynomial for fixed k. Across multiple queries, size-two witnesses can still give two-at-a-time greedy approximation ratio $2/(B - 1)$; Appendix A.21 gives the construction and relates shared candidates to supermodular width (Chen et al., 2019).

3.2 EXACT RISK STILL ADMITS LOCAL TRAPS

Even an algorithm given the full design can stop at a poor batch if it only accepts small improving exchanges. Let u be a unit target orthogonal to nuisance directions. Measurements $s_i = \varepsilon u + v_i$ share a target component, with the v_i forming a regular simplex and summing to zero. Distractors $d_k = \delta u + w_k$ use mutually orthogonal nuisance directions. Their complete risk comparison yields the following result.

Theorem 3.3 (Exchange barrier)**.** *For every $m \geq 2$ and $\lambda > 0$, take a budget of m, m simplex measurements and m distractors with $\delta^2/\varepsilon^2 > (1+\lambda)(m-1)/(1+\lambda(m-1))$. Greedy C-optimal terminates at an all-distractor batch that is a strict local minimum under every exchange of at most $m-1$ measurements. The m-measurement exchange purchases the unique minimal closing set and lowers r_λ by $A/\lambda\,(1+o(1))$ as $\lambda \to 0$, where $A > 0$ is the all-distractor squared reconstruction residual.*

Risk comparison. For a simplex measurements in a fixed batch of size B, risk is decreasing in

$$\frac{a\varepsilon^2(1+\lambda)}{(m-a)/(m-1)+\lambda} + (B-a)\delta^2.$$

The stated condition makes every $a < m$ replacement reduce this quantity. At $a = m$, nuisance cancellation restores the target direction. Appendix A.7 derives the expression and the low-noise limit; Theorem A.4 gives the exact greedy activation threshold.

4 RECOVERY, GLOBAL OPTIMA, AND EXCHANGE SIZE

With unit prior and observation-noise variances, consider $2b$ measurements and an even budget $b \geq 4$:

$$x_j^{\pm} = \sqrt{a}\,u_j \pm v_j, \qquad \varphi_q = b^{-1/2}\sum_{j=1}^{b} u_j, \qquad R(S) = \varphi_q^\top (I + X_S^\top X_S)^{-1}\varphi_q.$$

Here $a > 0$, the u_j are orthonormal and orthogonal to every v_i, and $C = VV^\top \succeq 0$ is an arbitrary nuisance Gram matrix. The algorithm receives the candidate rows and φ_q; pair labels, candidate order, the latent signal basis, and row orientations are hidden.

Lemma 4.1 (Exact risk and monotonicity)**.** *Let F, P, Z index fully purchased, singly purchased, and unpurchased pairs, with $f = |F|$, $p = |P|$, $2f + p = b$, and singleton signs $s \in \{-1,1\}^P$. Put $c_0 = 2a/(1+2a)$ and*

$$D = C_{PP} - C_{PF}(C_{FF} + \tfrac{1}{2}I)^{-1}C_{FP}.$$

Then

$$R(S) = 1 - \frac{1}{b}\left[c_0 f + a\, s^\top ((1+a)I + D)^{-1} s\right]. \tag{1}$$

Moreover $x^\top D x = \min_y\{z^\top C_{F\cup P, F\cup P} z + \frac{1}{2}\|y\|^2\}$ for $z = (y, x)$, so D increases with C in Loewner order; $C \succeq \delta I$ implies $D \succeq \delta I$.

Pair sums and differences measure signal and nuisance separately. Conditioning on the differences gives D; conditioning on the singles gives (1). This Schur-complement representation also connects the graph case to Kron reduction (Dörfler & Bullo, 2013).

Theorem 4.2 (Recovery from the candidate Gram matrix)**.** *Orient rows to have positive query inner product, all equal to $\sqrt{a/b}$, and set $G = XX^\top$, $M = G - 2aI$. Every original pair i, j satisfies $M_{i,:} + M_{j,:} = 0$. Any perfect matching under this identity gives a valid paired representation, $u'_k = (x_i + x_j)/(2\sqrt{a})$, $v'_k = (x_i - x_j)/2$, with the same a and φ_q, in polynomial time. If $\eta = \min_k |\lambda_k(C) - a| > 0$, the original pairs are unique: every false pair has row-sum residual at least 2η. Nearest-residual matching tolerates an oriented shifted-Gram perturbation of operator norm less than $\eta/\sqrt{2}$.*

Proof. An original pair satisfies $G(e_i + e_j) = 2a(e_i + e_j)$. Opposite nonzero rows of M occur with equal multiplicities, and zero rows occur in even number, so matching them always finishes. The

identity makes the recovered u'_k orthonormal and orthogonal to every v'_ℓ. Their query projections are $1/\sqrt{b}$, giving the stated target. In the pair-sum/difference basis, $M = \mathrm{diag}(0, 2(C - aI))$; a false pair has a difference component of norm one. The margin and perturbation bound follow (Appendix A.18). □

Theorem 4.3 (Exact global optimum). *Let $\delta = \lambda_{\min}(C)$. If $\delta < a$ and $Cs = \delta s$ for a sign vector s, one row per pair with signs s attains*

$$R^* = \frac{1+\delta}{1+a+\delta}.$$

If $\delta \geq a$, any $b/2$ complete pairs attain $R^* = (1+a)/(1+2a)$. At $\delta = a$, a sign eigenvector also gives an optimal one-per-pair batch. Under the sign-eigenvector hypothesis, if $\delta < a$ is simple, Theorem 4.2 and an eigendecomposition recover the optimum in polynomial time, including when a is another nuisance eigenvalue.

Proof. The bound $D \succeq \delta I$ gives

$$R(S) \geq 1 - \frac{1}{b}\left[\frac{2af}{1+2a} + \frac{a(b-2f)}{1+a+\delta}\right].$$

This is affine in f, minimized at $f = 0$ for $\delta < a$ and $f = b/2$ for $\delta > a$. The stated batches attain these endpoints. Under that hypothesis, for simple $\delta < a$ the 2δ-eigenvector of G is unique, has equal-magnitude entries, and has opposite entries within every recovered pair. Its signs give the optimum (Appendix A.18). □

Every connected weighted graph with Laplacian L supplies the hypothesis: $C = \delta I + tL$, $t > 0$, has simple minimum eigenvector $\mathbf{1}/\sqrt{b}$. Arbitrary nuisance-row sign changes preserve it. For $\delta > 0$, the eigenvalues of G are $2a$ and $2\lambda_k(C)$, all positive. Thus all $2b$ rows are independent, and the unique representation of φ_q uses every row: *every budget-b batch has positive reconstruction residual, while the recovered batch is globally optimal.*

4.1 How many measurements must change?

Write $\mu = \min_\ell \lambda_{\min}(C_{-\ell,-\ell})$. It measures the smallest nuisance energy of a unit combination omitting a pair; in graph designs these deleted matrices are grounded Laplacians (Pirani & Sundaram, 2016).

Proposition 4.4 (Half-budget exchange). *If $\mu > a$, greedy risk minimization selects $b/2$ complete pairs, and fewer than $b/2$ replacements cannot improve their risk. If $\theta = s^\top Cs/b < a$ for a sign vector, its one-per-pair batch has risk at most $(1+\theta)/(1+a+\theta)$, so the minimum improving exchange size is $b/2$.*

Under exact cancellation $Cs = 0$, deleting coordinate j from s gives $\mu \leq \|v_j\|^2/(b-1)$. Hence $\mu > a$ requires nuisance energy above $(b-1)a$ in every pair, or $a < 1/b$ for unit-norm candidates. Joint cancellation makes these individually nuisance-dominated measurements useful together. Appendix A.18 proves the proposition.

Proposition 4.5 (Exact exchange size in a symmetric family). *For $C = t(bI - \mathbf{1}\mathbf{1}^\top)$ and any complete-pair starting batch, put*

$$d_f = \frac{2tf(1+tb)}{1+2t(b-f)}, \qquad r_* = \min\{1 \leq r \leq b/2 : d_{b/2-r} < a\}.$$

Then r_* is the minimum improving exchange size. Greedy selects complete pairs when $d_{b/2-1} > a$, and $\mu = t$.

The knob t controls the exchange size throughout this family: $b = 8$, $a = 1$, and $t = 0.1, 0.25, 0.5, 1$ give $r_* = 1, 2, 3, 4$. In particular, the three-measurement threshold at $\mu = 0.5 < a$ extends beyond the sufficient condition above.

Theorem 4.6 (Exchange size on cycles)**.** *Let $C = \delta I + tL$ for the unweighted cycle on b vertices, $t > 0$, $0 \le \delta < a$, and any starting set of $b/2$ complete pairs. Set $\epsilon = a - \delta$, $d = \delta + \frac{1}{2}$, and $K = \epsilon$ for $\epsilon \le d$, otherwise $K = (\epsilon + d)^2/(4d)$. For each integer $1 \le r < b/2$, no exchange of at most r measurements improves risk if $t \ge 4Kr^2$. An improving r-exchange can be constructed in polynomial time if*

$$2r > \frac{1 + 2a}{a - \delta} \left(\sqrt{1 + \frac{4t}{1 + a + \delta}} - 1 \right). \tag{2}$$

An exchange of $b/2$ always attains the global optimum. Uniformly over starting pairs, b, and t,

$$r_* = \Theta_{a,\delta}\big(\min\{b, 1 + \sqrt{t}\}\big).$$

Proof sketch. An r-exchange leaves $p \le 2r$ singles and at least one unpurchased pair. Extend a unit singleton vector x by values y on full pairs and zeros on unpurchased pairs. Total variation on the cycle gives $\|z\|_\infty^2 \le \|z\|_2 (\sum_{ij} (z_i - z_j)^2)^{1/2}$. With $v = \|y\|^2$, the variational formula yields

$$x^\top D x \ge \delta + \min_{v \ge 0} \left[\frac{t}{p^2(1 + v)} + dv \right] \ge a \quad \text{when } t \ge 4Kr^2,$$

which rules out improvement for every sign choice. For the upper bound, slide a cyclic interval of $2r$ vertices: one contains exactly r starting pairs. Retain one aligned row from each and buy the other r singles. The conditional quadratic form is bounded below by $\mathbf{1}^\top (cI + tL_{2r}^D)^{-1} \mathbf{1}$, $c = 1 + a + \delta$. Solving its second-order difference equation gives the sufficient condition (2). Full details, including the closed form and general spectral comparison bounds, are in Appendix A.18. □

For cycle weights in $[w_{\min}, w_{\max}]$, replace t by $t w_{\min}$ in the obstruction and by $t w_{\max}$ in the construction. Fixed positive lower and finite upper weight bounds preserve the scaling. When $\eta > 0$, recovered pairs and nuisance signs identify the cycle's edges and traversal order.

5 Witness Search, Matroids, and Sparse Recovery

The recovery theorem uses the feature representation directly. Its combinatorial object belongs to a classical optimisation problem.

Circuits through a specified target. Contract the span of the observed rows T, and form the real vector matroid on the projected candidates and the projected target row. A minimal closing purchase S gives a circuit $S \cup \{q\}$ in this matroid: the rows in S are independent, and adding q creates one dependency. Conversely, any circuit containing q supplies a witness. Thus a cheapest closing purchase is a minimum-weight circuit through q, or equivalently

$$\min_w \sum_{x \in \mathcal{C}} c_x \mathbf{1}\{w_x \ne 0\} \quad \text{subject to} \quad \Phi_{T \cup \mathcal{C}}^\top w = \varphi_q.$$

This is sparse linear-system optimisation (Garey & Johnson, 1979; Natarajan, 1995); closure of a specified purchase is a rank test.

The regular-matroid class and factorial designs. In a graphic matroid, a circuit through a specified edge is a cycle through that edge, found by a shortest path between its endpoints after removing it. The cographic case is a cut problem. Seymour's regular-matroid decomposition combines graphic and cographic matroids with R_{10} (Seymour, 1980). Minimum-weight circuit search is polynomial on this class (Kapadia, 2014). Three-factor interaction grids cross this structural boundary.

Proposition 5.1 (Three-factor interaction designs are not regular)**.** *The real row matroid of the $3 \times 3 \times 3$ factorial design with all main and pair effects is not binary, and hence is neither regular nor graphic. The conclusion holds for every larger such grid with at least three levels per factor.*

Proof. The eight-row sets C, D listed in Appendix B.1 have $|C \cap D| = 3$. Exact rank computation gives $\operatorname{rank}(C) = 7$ and $\operatorname{rank}(C \setminus \{c\}) = 7$ for every $c \in C$; thus C is a circuit. The full design has rank 19; $E \setminus D$ has rank 18, restored to 19 by each element of D, so D is a cocircuit. In a binary matroid every circuit–cocircuit intersection has even size (Oxley, 2011, Chapter 9). The odd intersection excludes binarity. Larger factorial grids contain this matroid as a restriction, which proves the final assertion. □

Purchase complexity on the factorial grid. The grid also supports a direct hardness result for a target fixed by its geometry.

Theorem 5.2 (Closing the grand mean is NP-complete)**.** *Let $E = [n]^3$ have an intercept, all one-hot main effects and all one-hot pair effects, with uniform grand-mean target*

$$\bar{q}_n = \frac{1}{n^3} \sum_{x \in E} \varphi_x = \left(1, \frac{1}{n}\mathbf{1}_{3n}, \frac{1}{n^2}\mathbf{1}_{3n^2}\right).$$

Given observed T, unit-cost purchases from $E \setminus T$, and budget B, deciding whether some $|S| \le B$ satisfies $\bar{q}_n \in \mathrm{row}(\Phi_{T \cup S})$ is NP-complete. This holds for valid partial Latin squares T and $B = n^2 - |T|$, with arbitrary real witness coefficients.

Proof. A Latin completion supplies n^2 triples whose coefficients $1/n^2$ reconstruct $\bar{q}_n$. Conversely, rescale a feasible witness by n^2. Its n^2 nonzero AB coordinates require at least n^2 rows. The budget permits exactly that many, forcing one row per AB position with coefficient one and retaining every observed entry. The AC and BC equations then impose the Latin constraints. Partial Latin-square completion is NP-complete (Colbourn, 1984); a rational rank test gives membership in NP. Appendix A.19 gives the full reduction. □

The target matters: a single held-out cell on the complete grid has a seven-cell $2 \times 2 \times 2$ witness, searchable in polynomial time at fixed factor count. The theorem addresses the uniform grand mean.

From minimum support to finite-noise design. Optimal experimental design estimates specified linear functionals (Pukelsheim, 2006; Chaloner & Verdinelli, 1995; Mutný & Krause, 2022); follow-up designs resolve confounding by acquiring additional runs (Meyer et al., 1996). Here the exact finite-noise objective is

$$\min_{c(S) \le B} R(S) = \min_{\sum_{x \in C} c_x \mathbf{1}\{w_x \neq 0\} \le B} \left[\rho^2 \|\Phi_{T \cup C}^\top w - \varphi_q\|^2 + \sigma^2 \|w\|^2\right]$$

(Theorem A.3). The support chooses the measurements, and the coefficients trade reconstruction residual against observation noise. The exact-zero result identifies the regime in which the positive-ratio framework used in transductive active learning (Hübotter et al., 2024; Das & Kempe, 2018) does not supply a positive guarantee for closure.

Witness exchange uses reconstruction coefficients to propose a joint purchase and accepts it by exact risk. For categorical designs, alternating paths and contrasts supplied by measured rows generate the proposal. For general features, the equivalent quadratic coefficient objective is sparsified by reweighting, following established sparse recovery and experimental-design tools (Candès et al., 2008; Alexanderian et al., 2014). The resulting supports are evaluated and refined by single-measurement exchange. Appendices A.17 and B.3 give the algorithms; the paired recovery procedure of Theorem 4.3 provides a polynomial exact solution in its stated design family.

6 Empirical Mechanisms and Search Cost

Numerical design risks are reported in observation-noise units, R/σ^2. Exhaustive checks of 42 recovery configurations and 384,388 batches, together with budget-128 interval constructions, confirm the theoretical predictions; configurations and executable checks are in Appendix C.14.

6.1 A necessary joint exchange in a real design

On the original HOIP seed 7 task, with budget 12 and $\lambda = 0.0087$, automatic support search proposes a three-measurement exchange that reduces design risk from 396.754 to 343.758, or 13.4%. The starting batch is the best archived search output followed by complete one/two-measurement descent. Full-pool checks exclude every improving one- or two-measurement exchange, so three is both necessary and achieved. The proposal buys the Br, F and I measurements of the AcA–Pb composition family together. All partial replacements drawn from this proposal raise risk (Figure 3a).

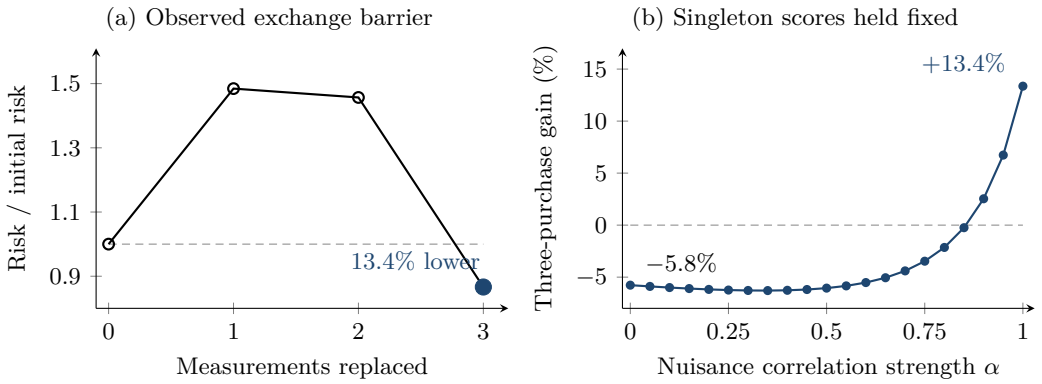

Figure 3: **A real exchange barrier depends on joint nuisance geometry.** (a) All 20 three-purchase choices in the HOIP initial/proposed union, with nine shared measurements fixed. One- and two-measurement partial exchanges raise risk by at least 48.4% and 45.7%. (b) Changing only cross-candidate nuisance correlations reverses the complete proposal's value. Every singleton risk and information score stays fixed. Each gain uses the incumbent risk at the same α.

6.2 Isolating the joint geometry

Condition on the observations and nine shared purchases, then split the whitened candidate rows into query-space components U and orthogonal nuisance components V. Replace the nuisance Gram $VV^\top$ by $\alpha VV^\top + (1-\alpha)\,\mathrm{diag}(VV^\top)$. This intervention preserves each candidate's norm, query covariance, and singleton predictive information gain. It changes only the nuisance correlations between measurements.

At $\alpha = 1$, the measured three-purchase gain is 13.4%; at $\alpha = 0$, the same proposal increases risk by 5.8% (Figure 3b). The gain therefore depends on joint nuisance correlations. Summed singleton information scores favour the incumbent. The original budget-24 comparison repeats this reversal across all ten HOIP tasks: Circuit has lower design risk than EPIG at $\alpha = 1$, and higher risk at $\alpha = 0$, with both selected batches fixed. Thus cross-candidate geometry determines the batch preference while singleton information scores remain unchanged (Appendix C.15).

6.3 The computational consequence

Across twenty HOIP tasks from common greedy starts, witness exchange captures 66.8% of complete one/two-measurement descent's mean task-normalised risk reduction with 1500 risk evaluations, or 0.063% of its median count. Its median final risk is within 3.3% of complete descent. At $\lambda = 0.0087$, median held-out RMSE falls from 0.114 to 0.081, a 29% reduction. Table 4 and Appendix C.13 provide the paired records and evaluation accounting.

Fixed-batch interventions on learned reaction features also identify finite-noise geometry as a determinant of exchange value: projection preserves exact reconstruction relations, yet reverses 8 of 37 distinct accepted multi-measurement improvements (Appendix C.16). The full representation study and continuous neural-feature experiments are in Appendices C.12 and C.14.

7 Conclusion

In paired designs with graph-structured nuisance, hidden structure yields a globally optimal batch even when improvement requires changing many measurements together. On cycles, the necessary exchange size grows on the square-root scale, with a matching constructive bound. These results cover full-row-rank designs with positive reconstruction residual at every affordable batch. HOIP interventions connect real joint improvements to nuisance correlations that singleton scores leave unresolved.

REFERENCES

Derek T. Ahneman, Jesús G. Estrada, Shishi Lin, Spencer D. Dreher, and Abigail G. Doyle. Predicting reaction performance in C–N cross-coupling using machine learning. *Science*, 360(6385): 186–190, 2018. doi: 10.1126/science.aar5169.

Alen Alexanderian, Noemi Petra, Georg Stadler, and Omar Ghattas. A-optimal design of experiments for infinite-dimensional Bayesian linear inverse problems with regularized ℓ_0-sparsification. *arXiv preprint arXiv:1308.4084*, 2014. URL https://arxiv.org/abs/1308.4084v2.

Jordan T. Ash, Chicheng Zhang, Akshay Krishnamurthy, John Langford, and Alekh Agarwal. Deep batch active learning by diverse, uncertain gradient lower bounds. In *International Conference on Learning Representations (ICLR)*, 2020.

Freddie Bickford Smith, Andreas Kirsch, Sebastian Farquhar, Yarin Gal, Adam Foster, and Tom Rainforth. Prediction-oriented Bayesian active learning. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, volume 206 of *Proceedings of Machine Learning Research*, pp. 7331–7348, 2023.

Emmanuel J. Candès, Michael B. Wakin, and Stephen P. Boyd. Enhancing sparsity by reweighted ℓ_1 minimization. *Journal of Fourier Analysis and Applications*, 14(5):877–905, 2008. URL https://web.stanford.edu/ boyd/papers/rwl1.html.

Ivano E. Castelli, David D. Landis, Kristian S. Thygesen, Søren Dahl, Ib Chorkendorff, Thomas F. Jaramillo, and Karsten W. Jacobsen. New cubic perovskites for one- and two-photon water splitting using the computational materials repository. *Energy & Environmental Science*, 5(10):9034–9043, 2012. doi: 10.1039/C2EE22341D.

Kathryn Chaloner and Isabella Verdinelli. Bayesian experimental design: a review. *Statistical Science*, 10(3):273–304, 1995.

Wei Chen, Shang-Hua Teng, and Hanrui Zhang. Capturing complementarity in set functions by going beyond submodularity/subadditivity. In *Innovations in Theoretical Computer Science (ITCS)*, 2019. doi: 10.4230/LIPIcs.ITCS.2019.24.

Charles J. Colbourn. The complexity of completing partial Latin squares. *Discrete Applied Mathematics*, 8(1):25–30, 1984. doi: 10.1016/0166-218X(84)90075-1.

Abhimanyu Das and David Kempe. Approximate submodularity and its applications: Subset selection, sparse approximation and dictionary selection. *Journal of Machine Learning Research*, 19 (3):1–34, 2018.

Andreea Deniţiu, Stefania Petra, Claudius Schnörr, and Christoph Schnörr. Phase transitions and cosparse tomographic recovery of compound solid bodies from few projections. *arXiv preprint arXiv:1311.0423*, 2013.

Florian Dörfler and Francesco Bullo. Kron reduction of graphs with applications to electrical networks. *IEEE Transactions on Circuits and Systems I: Regular Papers*, 60(1):150–163, 2013.

Ahmed El Alaoui and Michael W. Mahoney. Fast randomized kernel ridge regression with statistical guarantees. In *Advances in Neural Information Processing Systems (NIPS)*, volume 28, 2015.

M. R. Garey and D. S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman, 1979.

Jonas Hübotter, Bhavya Sukhija, Lenart Treven, Yarden As, and Andreas Krause. Transductive active learning: Theory and applications. In *Advances in Neural Information Processing Systems*, volume 37, pp. 124686–124755, 2024. doi: 10.52202/079017-3961.

Alon Itai and Michael Rodeh. Finding a minimum circuit in a graph. *SIAM Journal on Computing*, 7(4):413–423, 1978.

Rohan Kapadia. Finding shortest circuits in binary matroids. The Matroid Union blog, 2014. https://matroidunion.org/?p=1106.

Andreas Kirsch, Joost van Amersfoort, and Yarin Gal. BatchBALD: Efficient and diverse batch acquisition for deep Bayesian active learning. In *Advances in Neural Information Processing Systems*, volume 32, 2019.

R. D. Meyer, D. M. Steinberg, and G. E. P. Box. Follow-up designs to resolve confounding in multifactor experiments. *Technometrics*, 38(4):303–313, 1996. doi: 10.1080/00401706.1996.10484538.

Mojmír Mutný and Andreas Krause. Experimental design for linear functionals in reproducing kernel Hilbert spaces. In *Advances in Neural Information Processing Systems*, volume 35, pp. 20175–20188, 2022. doi: 10.52202/068431-1467.

B. K. Natarajan. Sparse approximate solutions to linear systems. *SIAM Journal on Computing*, 24 (2):227–234, 1995.

James Oxley. *Matroid Theory*. Oxford University Press, 2nd edition, 2011.

Stefania Petra and Christoph Schnörr. TomoPIV meets compressed sensing. *Pure Mathematics and Applications*, 20(1–2):49–76, 2009.

Mohammad Pirani and Shreyas Sundaram. On the smallest eigenvalue of grounded Laplacian matrices. *IEEE Transactions on Automatic Control*, 61(2):509–514, 2016.

Friedrich Pukelsheim. *Optimal Design of Experiments*, volume 50 of *Classics in Applied Mathematics*. SIAM, classic edition, 2006. doi: 10.1137/1.9780898719109.

Philippe Schwaller, Alain C. Vaucher, Teodoro Laino, and Jean-Louis Reymond. Prediction of chemical reaction yields using deep learning. *Machine Learning: Science and Technology*, 2(1): 015016, 2021. doi: 10.1088/2632-2153/abc81d.

Ozan Sener and Silvio Savarese. Active learning for convolutional neural networks: a core-set approach. In *International Conference on Learning Representations (ICLR)*, 2018.

P. D. Seymour. Decomposition of regular matroids. *Journal of Combinatorial Theory, Series B*, 28 (3):305–359, 1980.

Kushal Tirumala, Daniel Simig, Armen Aghajanyan, and Ari Morcos. D4: Improving LLM pretraining via document de-duplication and diversification. In *Advances in Neural Information Processing Systems*, volume 36, 2023. doi: 10.52202/075280-2348. URL `https://arxiv.org/abs/2308.12284`.

Namrata Vaswani and Wei Lu. Modified-CS: Modifying compressive sensing for problems with partially known support. *IEEE Transactions on Signal Processing*, 58(9):4595–4607, 2010.

AI USE STATEMENT

A Large language model was used to assist with audit and language polishing.

Responsibility. The authors have reviewed all AI-assisted content and take full responsibility for the entire contents of this paper, including any errors.

REPRODUCIBILITY STATEMENT

Every numerical claim in this paper is produced by a script in the accompanying code and archived as raw output. The exact-zero, oracle, recovery, exchange-scaling, and purchase-complexity results are established in Sections 2, 3.1, 4 and 5; exact rank certificates and finite-instance numerical checks accompany them. Supporting derivations are in Appendices A, A.18 and A.19. The recovery and scaling checks archive every tested configuration, starting batch, and figure source value, together with the Latin-reduction and threshold scripts. The acquisition experiments of Section 6 run from a repository-relative data path with fixed seeds; the verification index is provided in `reproducibility/verification.md`. Figure data is generated from the archived experiment output.

A PROOFS

Throughout, Φ_D stacks $\{\varphi_x^\top\}_{x\in D}$ for $D \subseteq \mathcal{X}$, $G_D = \Phi_D^\top \Phi_D$, and $P^{(D)}, P_\perp^{(D)}$ project onto $\mathrm{row}(\Phi_D)$ and its orthogonal complement. We write $\lambda = \sigma^2/\rho^2$ and $r_\lambda(q) = \varphi_q^\top (G_T + \lambda I)^{-1}\varphi_q$. For candidate set S and query set Q, write $\nu_T(S,Q) = r(T\cup S) + r(T\cup Q) - r(T) - r(T\cup S\cup Q)$, where $r(\cdot)$ is the row rank. This counts the query directions that S newly brings into range. We write $\kappa = \rho/\sigma = 1/\sqrt{\lambda}$ for the signal-to-noise ratio. Numerical checks and their scripts are indexed in `reproducibility/verification.md` in the accompanying code.

A.1 CLOSURE IS DETERMINED BY THE DESIGN

Theorem A.1 (No feedback channel)**.** *Every closure-layer object ($\nu_T(S,Q)$, the closure predicate $\mathbf{1}\{\varphi_q \in \mathrm{row}\,\Phi_D\}$, the witness families, the circuit gains) is a deterministic functional of (D,Q) and does not depend on y. Adaptivity can change which cells are bought, never what a purchase closes.*

A.2 RECONSTRUCTION AND DESIGN RISK

Set $r_\lambda(q) = \varphi_q^\top (G_T + \lambda I)^{-1}\varphi_q$. The risk decomposition connects query reconstruction to the acquisition objective. Proofs are in Appendix A.

Theorem A.2 (Exact two-layer decomposition)**.** *Let $\lambda = \sigma^2/\rho^2$. Then the Bayes risk (posterior variance) of θ_q, the minimax affine risk $\min_a \sup_{\|\beta\|\le\rho} \mathbb{E}[(a^\top y_T - \theta_q)^2]$, and $\sigma^2 r_\lambda(q)$ all coincide, and*

$$\sigma^2 r_\lambda(q) \;=\; \underbrace{\rho^2 \|P_\perp^{(T)}\varphi_q\|^2}_{\textit{closure layer (bias)}} \;+\; \underbrace{\sigma^2 \textstyle\sum_{s_i>0} \frac{a_i^2}{s_i^2+\lambda}}_{\textit{precision layer (variance)}} .$$

Here s_i are the nonzero singular values of Φ_T, and a_i are the query coordinates along the corresponding right singular vectors. The theorem fixes the exchange rate at $\lambda = \sigma^2/\rho^2$. The first term vanishes exactly when q closes and remains unchanged as observation noise decreases.

Theorem A.3 (Acquisition is sparse witness selection)**.** *Let C index the purchasable cells with costs c_x, and let $\Phi = \Phi_{T\cup C}$ stack the measured and purchasable rows. Then*

$$\min_{S\subseteq C:\, c(S)\le B} \sigma^2 r_\lambda^{T\cup S}(q) \;=\; \min_w \left\{\rho^2\|\Phi^\top w - \varphi_q\|^2 + \sigma^2\|w\|^2\right\} \quad \textit{s.t.} \quad \sum_{x\in C} c_x\, \mathbf{1}\{w_x \neq 0\} \;\le\; B,$$

with the same minimising support. The identity holds for the Bayes and the minimax-affine risks, for unequal costs, and with $\|w\|^2$ replaced by $w^\top \Sigma w$ under heteroscedastic noise; the measured rows T enter unpenalised in the constraint. For a set Q of queries the same statement holds with w a matrix W, Frobenius norms, and a group-row budget $\sum_x c_x \mathbf{1}\{\|W_{x,:}\|_2 > 0\} \le B$.

Three strategies now optimise the same objective over different supports. Closure search restricts it to the face $\|\Phi^\top w - \varphi_q\| = 0$, greedy C-optimal explores it locally, and exact optimisation searches all budget-feasible supports.

Unrepresented query component. The squared part of the query left outside the selected measurements' span is $A(S) = \|P_\perp^{(T\cup S)}\varphi_q\|^2$. It is zero exactly when the query can be reconstructed. Appendix D.1 develops the resulting tradeoff between closure and precision.

A.3 THE TWO-LAYER DECOMPOSITION

Proof of Theorem A.2. (i) Bayes risk. With $\beta \sim N(0,\rho^2 I)$ and $y_T = \Phi_T\beta + \varepsilon$, the posterior covariance is $\Sigma_{\mathrm{post}} = (\sigma^{-2}G_T + \rho^{-2}I)^{-1} = \sigma^2(G_T + \lambda I)^{-1}$, so $\mathrm{Var}(\theta_q \mid y_T) = \sigma^2 r_\lambda(q)$.

(ii) Minimax affine risk. For an affine rule $a^\top y_T$,

$$\sup_{\|\beta\|\le\rho} \mathbb{E}\big[(a^\top y_T - \theta_q)^2\big] = \rho^2\|\Phi_T^\top a - \varphi_q\|^2 + \sigma^2\|a\|^2,$$

the supremum being attained at β aligned with $\Phi_T^\top a - \varphi_q$. Minimising this strictly convex quadratic gives the normal equation $(G_T + \lambda I)\,\tilde{a} = \varphi_q$ with $a^* = \Phi_T \tilde{a}$, i.e. $a^* = \Phi_T (G_T + \lambda I)^{-1}\varphi_q$ and $\Phi_T^\top a^* - \varphi_q = -\lambda (G_T + \lambda I)^{-1}\varphi_q$. Substituting,

$$\rho^2\lambda^2 \varphi_q^\top (G_T + \lambda I)^{-2}\varphi_q + \sigma^2 \varphi_q^\top (G_T + \lambda I)^{-1} G_T (G_T + \lambda I)^{-1}\varphi_q = \sigma^2 r_\lambda(q),$$

where the last step uses $\rho^2\lambda^2 = \sigma^2\lambda$ and $\lambda I + G_T = (G_T + \lambda I)$.

(iii) Layer split. Let $G_T = \sum_{s_i>0} s_i^2 v_i v_i^\top$ be the spectral decomposition on $\mathrm{row}(\Phi_T)$ and $a_i = \langle \varphi_q, v_i\rangle$. Then $(G_T + \lambda I)^{-1} = \sum_{s_i>0} \frac{v_i v_i^\top}{s_i^2+\lambda} + \frac{P_\perp^{(T)}}{\lambda}$, so

$$\sigma^2 r_\lambda(q) = \sigma^2 \sum_{s_i>0} \frac{a_i^2}{s_i^2+\lambda} + \frac{\sigma^2}{\lambda}\|P_\perp^{(T)}\varphi_q\|^2 = \sigma^2 \sum_{s_i>0} \frac{a_i^2}{s_i^2+\lambda} + \rho^2 \|P_\perp^{(T)}\varphi_q\|^2. \qquad \square$$

A.4 Acquisition as sparse witness selection

Proof of Theorem A.3. Fix a measured set D. By Theorem A.5(c) and $\lambda = \sigma^2/\rho^2$,

$$\sigma^2 r_\lambda^D(q) = \sigma^2 \min_w \Big[\|w\|^2 + \tfrac{1}{\lambda}\|\Phi_D^\top w - \varphi_q\|^2\Big] = \min_w \Big[\sigma^2\|w\|^2 + \rho^2\|\Phi_D^\top w - \varphi_q\|^2\Big],$$

since $\sigma^2/\lambda = \rho^2$. Take $D = T \cup S$ and index w by $T \cup C$, constraining $w_x = 0$ for $x \in C \setminus S$; then $\Phi_D^\top w = \Phi^\top w$ and $\|w\|$ is unchanged, so

$$\min_{S:\, c(S)\le B} \sigma^2 r_\lambda^{T\cup S}(q) = \min_{S:\, c(S)\le B}\ \min_{\mathrm{supp}(w)\subseteq T\cup S} \Big[\rho^2\|\Phi^\top w - \varphi_q\|^2 + \sigma^2\|w\|^2\Big].$$

The two sides of the claimed identity now bound each other. Given a feasible pair (S, w), the vector w satisfies $\sum_{x\in C} c_x \mathbf{1}\{w_x \neq 0\} \le c(S) \le B$, so it is feasible for the right-hand program. Conversely, given w feasible there, put $S = \mathrm{supp}(w) \cap C$; then $c(S) \le B$ and w is supported in $T \cup S$. The two programs therefore have the same value, and S read off a minimiser w is a minimising support.

The measured rows enter $\|w\|^2$ but not the budget, which is the unpenalised-T clause. The statement for the minimax affine risk is the same computation, the two risks being equal by Theorem A.2. Unequal costs are already carried by c_x. For noise covariance Σ the inner problem is $\min_w[w^\top \Sigma w + \rho^2\|\Phi^\top w - \varphi_q\|^2]$, obtained by repeating the derivation of Theorem A.5(c) with Σ in place of $\sigma^2 I$; the support argument above is unchanged. For a query set Q, summing the per-query risks gives $\min_W[\rho^2\|\Phi^\top W - \Phi_Q^\top\|_F^2 + \sigma^2\|W\|_F^2]$, and a cell purchased for one query is available to all, so the budget couples the columns of W through its rows: the constraint is $\sum_x c_x \mathbf{1}\{\|W_{x,:}\|_2 > 0\} \le B$. $\square$

A.5 The frontier

Proof of Proposition D.1. Write $F_t(w) = t\,b(w) + v(w)$ with $b(w) = \|\Phi^\top w - \varphi_q\|^2$ and $v(w) = \|w\|^2$. For $t > 0$, F_t is strictly convex with Hessian $2(t\,\Phi\Phi^\top + I) \succ 0$, so the minimiser is unique and solves $(t\,\Phi\Phi^\top + I)w = t\,\Phi\varphi_q$, giving $w_S(t) = t(t\,\Phi\Phi^\top + I)^{-1}\Phi\varphi_q$, a rational and hence analytic function of t on $(0,\infty)$.

For each fixed w the map $t \mapsto F_t(w)$ is affine with slope $b(w) \ge 0$. Thus $R_S(t) = \sigma^2 \min_w F_t(w)$ is a pointwise infimum of affine functions with non-negative slopes: concave, non-decreasing, and continuous on the open half-line. Danskin's theorem applies because the minimiser is unique, giving $R_S'(t) = \sigma^2 b(w_S(t))$. Concavity makes R_S' non-increasing, so $b(w_S(t))$ decreases. For the variance, the Legendre relation $v(w_S(t)) = R_S(t)/\sigma^2 - t\,b(w_S(t))$ gives $\frac{d}{dt}v(w_S(t)) = -t\,R_S''(t)/\sigma^2 \ge 0$, so v increases.

If S closes q there is $w^\star$ with $\Phi^\top w^\star = \varphi_q$, so $R_S(t) \le \sigma^2\|w^\star\|^2$ for all t: R_S is concave, increasing and bounded, hence $R_S'(t) \to 0$, i.e. $b(w_S(t)) \to 0$. Feasible points converge to the affine set $\{\Phi^\top w = \varphi_q\}$ and v is coercive there, so $w_S(t)$ converges to its minimum-norm point and $v(w_S(t)) \to V(S)$. $\square$

Proof of Theorem D.2. Let P project onto $\mathrm{row}(\Phi_{T\cup S})$. Since $\Phi^\top w \in \mathrm{row}(\Phi)$ for every w, $\|\Phi^\top w - \varphi_q\|^2 = \|\Phi^\top w - P\varphi_q\|^2 + A(S)$ with $A(S) = \|P_\perp^{(T\cup S)}\varphi_q\|^2$, and the second term is constant in w. Hence

$$r_\lambda(S) = \frac{A(S)}{\lambda} + \min_w \Big[\|w\|^2 + \tfrac{1}{\lambda}\|\Phi^\top w - P\varphi_q\|^2\Big].$$

The remaining minimum is the same object for a query that the support does close, so by Proposition D.1 it converges to $V(S) = \min\{\|w\|^2 : \Phi^\top w = P\varphi_q\}$ as $\lambda \to 0$. This gives $r_\lambda(S) = A(S)/\lambda + V(S) + o(1)$.

Substituting the expansion for S_g and $S_\star$ and letting $\lambda \to 0$: if $A(S_\star) > 0$ both numerator and denominator are $\Theta(1/\lambda)$ and the ratio tends to $A(S_g)/A(S_\star)$; if $A(S_\star) = 0 < A(S_g)$ the numerator is $\Theta(1/\lambda)$ against a finite denominator $V(S_\star)$, so the ratio diverges at rate $1/\lambda$ with constant $A(S_g)/V(S_\star)$; if both squared residual normes vanish the $1/\lambda$ terms are absent and the ratio tends to $V(S_g)/V(S_\star)$. □

Proof of Theorem D.3. Let $\mathcal{S}$ be the budget-feasible supports, a finite set, and write $A_{\min} = \min_{S\in\mathcal{S}} A(S)$ and $\mathcal{S}_0 = \{S : A(S) = A_{\min}\}$. By Theorem D.2, $r_\lambda(S) = A(S)/\lambda + V(S) + o(1)$ uniformly over the finite set. For $S \notin \mathcal{S}_0$ and any $S_0 \in \mathcal{S}_0$,

$$r_\lambda(S) - r_\lambda(S_0) = \frac{A(S) - A_{\min}}{\lambda} + \big(V(S) - V(S_0)\big) + o(1) \;\longrightarrow\; +\infty,$$

so for all λ below some threshold the minimiser lies in $\mathcal{S}_0$. Restricted to $\mathcal{S}_0$ the $1/\lambda$ terms are equal and cancel, leaving $r_\lambda(S) = A_{\min}/\lambda + V(S) + o(1)$, whose minimiser is the element of $\mathcal{S}_0$ with smallest V. That is the lexicographic program. If some feasible support closes q then $A_{\min} = 0$, so every limiting minimiser closes q.

The finiteness of $\mathcal{S}$ also gives the staircase claim of Proposition D.1: the global risk is the pointwise minimum of finitely many concave curves, so the attaining support is constant except at their crossings. □

A.6 The activation barrier

Theorem A.4 (Activation barrier)**.** *Greedy C-optimal declines every prefix of the circuit, and therefore never completes it, exactly when*

$$\delta^2(\lambda + \varepsilon^2 + 1) > \varepsilon^2(\lambda + \delta^2 + 1) \iff (\delta^2 - \varepsilon^2)(\lambda + 1) > 0 \iff \delta^2 > \varepsilon^2.$$

The threshold does not depend on λ or on the circuit size m; the cost of the failure does, growing without bound as $\lambda \to 0$ and with m.

Proof sketch. After j distractors, the risk is $(1+\lambda)/[\lambda(1+\lambda+j\delta^2)]$. One more distractor adds δ^2 to the final denominator; either bundle member adds ε^2. Thus the same inequality makes greedy choose a distractor at every step. The complete pair instead cancels v and has risk $1/(2\varepsilon^2+\lambda)$, which stays bounded as noise decreases. Simplex vectors extend this cancellation to m members.

Proof of Theorem A.4. Reduce to the state after the measured rows, so that the query direction is u and the design is empty. The circuit is the pair $s_1 = \varepsilon u - v,\ s_2 = \varepsilon u + v$ with $v \perp u$: neither alone carries u, while $s_1 + s_2 = 2\varepsilon u$ does. The distractors are $D_k = \delta u + w_k$ with $\{w_k\}$ orthonormal and orthogonal to u, v, so each carries u by itself.

Direct computation of $r_\lambda(u \mid \cdot)$ on these designs gives closed forms, verified to 10^{-8} by `verify_trap_proof.py`:

$$r_j = \frac{1+\lambda}{\lambda(1+\lambda+j\delta^2)}, \qquad r_{j,s_1} = \frac{1+\lambda}{\lambda(1+\lambda+j\delta^2+\varepsilon^2)}, \qquad r_{\{s_1,s_2\}} = \frac{1}{2\varepsilon^2+\lambda},$$

where r_j is the risk after j distractors. Greedy, at a state with j distractors bought, compares the risk after one more distractor against the risk after s_1; since the risk is decreasing in the denominator,

$$r_{j+1} < r_{j,s_1} \iff 1+\lambda+(j+1)\delta^2 > 1+\lambda+j\delta^2+\varepsilon^2 \iff \delta^2 > \varepsilon^2,$$

and the same comparison holds for s_2 by symmetry. The condition contains neither λ nor j: it is the *same* inequality at every prefix, which is why greedy declines the circuit at each step rather than at some steps, and never completes it while distractors remain. Written in the equivalent form obtained by clearing denominators against the self-variances, the condition is $\delta^2(\lambda+\varepsilon^2+1) > \varepsilon^2(\lambda+\delta^2+1)$, whose cross terms cancel to leave $(\delta^2-\varepsilon^2)(\lambda+1) > 0$.

The cost is unbounded. With budget B greedy holds B distractors and risk $r_B = \Theta(1/\lambda)$, while the optimum holds the circuit and risk $r_{\{s_1,s_2\}} \to 1/(2\varepsilon^2)$, a finite limit: the ratio is $\Theta(1/\lambda)$, which is the diverging case of Theorem D.2 with $A(S_g) > 0 = A(S_\star)$. The generalisation to a circuit of m cells replaces $2\varepsilon^2$ by $m\varepsilon^2$ in $r_{\{s_1,\dots,s_m\}}$ and leaves the selection condition unchanged, so the threshold does not move with m while the ratio it costs grows with it. □

A.7 The exchange barrier

Proof of Theorem 3.3. Take the simplex-family instance of Appendix A.11 with budget $B \geq m$ and at least B distractors, and the stated ratio $\delta^2/\varepsilon^2 > (1+\lambda)(m-1)/(1+\lambda(m-1))$, which exceeds 1 for every $m \geq 2$, so $\delta > \varepsilon$.

Greedy. By Lemma A.11 a distractor step always gains δ^2 while a circuit step at $a = 0$ gains ε^2; greedy therefore starts on distractors, and since its circuit gain at any reached a would require first taking a circuit step at $a = 0$, it stays at $a = 0$ and fills the budget with distractors.

Local minimum. An exchange of $j \leq m-1$ cells at the all-distractor batch replaces j distractors by j circuit members (distractor-for-distractor swaps leave the risk unchanged, and no circuit member is present to remove). By Lemma A.11 the risk is a strictly decreasing function of $\mathcal{R}(a,b) = a\varepsilon^2(1+\lambda)/W(a) + b\delta^2$ with $W(a) = \frac{m-a}{m-1} + \lambda$, so the exchange strictly increases the risk iff $\varepsilon^2(1+\lambda)/W(j) < \delta^2$. Since W is decreasing in j, the binding case is $j = m-1$, where $(1+\lambda)/W(m-1) = (1+\lambda)(m-1)/(1+\lambda(m-1))$: exactly the stated ratio condition, uniformly over $j \leq m-1$.

The m-exchange. Replacing m distractors by the full bundle gives $a = m$, $W(m) = \lambda$, and by Lemma A.10 $t_{m,B-m} \to 1$ as $\lambda \to 0$, so the risk $r = (1 - t^{(\lambda)})/\lambda$ stays bounded; the all-distractor batch has $t_{0,B} = B\delta^2/(1+\lambda+B\delta^2)$ bounded away from 1, so its risk is $A/\lambda \cdot (1+o(1))$ with squared residual norm $A = 1/(1+B\delta^2)$. The difference is $A/\lambda \cdot (1+o(1))$. Lemma A.9 makes $\{s_1,\dots,s_m\}$ the unique minimal closing set, so the improving m-exchange is the witness bundle and no other. □

A.8 Leverage as the ℓ_2 relaxation

Theorem A.5 (Leverage is the ℓ_2 relaxation of witness search)**.** *(a) θ_q is unbiasedly estimable iff $\varphi_q \in \mathrm{row}(\Phi_T)$ iff a witness exists iff the matroid circuit condition of Definition 2.1 holds. (b)* $\mathrm{Var}(\hat{\theta}^{\mathrm{BLUE}}) = \sigma^2 \min\{\|w\|^2 : \Phi_T^\top w = \varphi_q\} = \sigma^2 \varphi_q^\top G_T^+ \varphi_q$. *(c)* $r_\lambda(q) = \min_w \left[\|w\|^2 + \frac{1}{\lambda}\|\Phi_T^\top w - \varphi_q\|^2\right]$.

By (b) the price of answerability is the cheapest chain of measurements that answers; by (c) ridge leverage is, verbatim, the penalty relaxation of the witness program, informative *after* closure and silent before it.

Proof of Theorem A.5. (a) An affine rule $w^\top y_T$ is unbiased for θ_q for all β iff $\Phi_T^\top w = \varphi_q$, which is solvable iff $\varphi_q \in \mathrm{row}(\Phi_T)$. The matroid form is Definition 2.1: writing M for the vector matroid of Φ, φ_q lies in the span of $\{\varphi_x\}_{x\in T}$ iff q is not a coloop of $\{q\} \cup T$, i.e. iff $\{q\} \cup T$ contains a circuit through q with all other elements in T.

(b) Among unbiased rules, $\mathrm{Var}(w^\top y_T) = \sigma^2\|w\|^2$, so the BLUE solves $\min\{\|w\|^2 : \Phi_T^\top w = \varphi_q\}$, whose minimiser is the pseudoinverse solution $w = \Phi_T G_T^+ \varphi_q$ with value $\varphi_q^\top G_T^+ \varphi_q$.

(c) The map $w \mapsto \|w\|^2 + \frac{1}{\lambda}\|\Phi_T^\top w - \varphi_q\|^2$ is strictly convex with stationarity condition $\lambda w + \Phi_T(\Phi_T^\top w - \varphi_q) = 0$, i.e. $(\Phi_T\Phi_T^\top + \lambda I)w = \Phi_T\varphi_q$, so $w^\star = \Phi_T(G_T + \lambda I)^{-1}\varphi_q$ by the push-through identity $(\Phi_T\Phi_T^\top + \lambda I)^{-1}\Phi_T = \Phi_T(G_T + \lambda I)^{-1}$. Then $\Phi_T^\top w^\star - \varphi_q = -\lambda(G_T + \lambda I)^{-1}\varphi_q$,

and the objective evaluates to

$$\varphi_q^\top (G_T+\lambda I)^{-1} G_T (G_T+\lambda I)^{-1} \varphi_q + \lambda\, \varphi_q^\top (G_T+\lambda I)^{-2} \varphi_q = \varphi_q^\top (G_T+\lambda I)^{-1} \varphi_q = r_\lambda(q). \quad \square$$

Part (c) is the identity the paper turns on: the $\lambda \to 0$ limit recovers the hard constraint of (b), whose feasibility is exactly the combinatorial condition of (a). On a graphic design, (b) is the effective resistance between the two endpoints of the contrast.

Lemma A.6 (No ℓ_2-silent circuit)**.** *If $\nu_T(\{s_1, s_2\}, \{q\}) = 1$ then at least one of s_1, s_2 strictly reduces the ℓ_2 residual of q.*

Proof of Lemma A.6. Write $z = P_\perp^{(T)} \varphi_q$ and $z_i = P_\perp^{(T)} \varphi_{s_i}$. Since q is open, $z \neq 0$. "s_i strictly reduces the ℓ_2 residual of q" means $\langle z, z_i \rangle \neq 0$; suppose instead $\langle z, z_1 \rangle = \langle z, z_2 \rangle = 0$. From $\nu_T(\{s_1, s_2\}, \{q\}) = 1$ we have $\varphi_q \in \mathrm{row}(\Phi_{T \cup \{s_1, s_2\}}) = \mathrm{row}(\Phi_T) + \mathrm{span}\{\varphi_{s_1}, \varphi_{s_2}\}$; applying $P_\perp^{(T)}$ gives $z \in \mathrm{span}\{z_1, z_2\}$. Together with $z \perp z_1, z_2$ this forces $\langle z, z \rangle = 0$, i.e. $z = 0$, a contradiction. $\square$

The consequence used in Section 3 is negative and constrains our own constructions: a separating instance cannot be built by driving single candidate scores to zero, so it must be built by making them *equal*, which is what the decoys in Lemma A.11 do.

A.9 Mutual information and measurement precision

Fix a query block Φ_Q and a measured design D. With $G_D = \Phi_D^\top \Phi_D = \sum_{s_i > 0} s_i^2 v_i v_i^\top$, write

$$\Phi_Q (G_D + \lambda I)^{-1} \Phi_Q^\top = A_D(\lambda) + \lambda^{-1} B_D, \qquad B_D = \Phi_Q P_\perp^{(D)} \Phi_Q^\top, \tag{3}$$

where $A_D(\lambda) = \Phi_Q \sum_{s_i > 0} v_i v_i^\top / (s_i^2 + \lambda) \Phi_Q^\top$. The rank d_D of B_D equals $r(D \cup Q) - r(D)$. In an orthonormal basis for $\mathrm{range}(B_D) \oplus \ker(B_D)$, write $B_D = \mathrm{diag}(B_{D,1}, 0)$ and $M_D(\lambda) = I + A_D(\lambda)$ in matching blocks. Define

$$\begin{aligned} c_D &= \log\det B_{D,1} + \log\det M_{D,22}(0), \\ H_D &= \|\Phi_Q G_D^+\|_F^2, \qquad J_D = \mathrm{tr}\Big(B_{D,1}^{-1} M_{D,11}(0) \Big). \end{aligned} \tag{4}$$

Empty determinants are one; $J_D = 0$ when $d_D = 0$. These constants include both the query geometry and the nonzero measurement singular values.

Theorem A.7 (Mutual-information expansion with an explicit remainder)**.** *For fixed T, S, Q and every $\lambda > 0$,*

$$I(y_Q; y_S \mid y_T) = \frac{\nu_T(S, Q)}{2} \log \frac{1}{\lambda} + \frac{c_T - c_{T \cup S}}{2} + E(\lambda),$$

with the finite-noise bound

$$|E(\lambda)| \leq \frac{\lambda}{2} \big(H_T + J_T + H_{T \cup S} + J_{T \cup S} \big). \tag{5}$$

The coefficient records the number of query directions reconstructed by the purchase. The remainder quantifies when this leading term approximates the finite-noise information for the given measurement scales.

Lemma A.8 (Exact block identity)**.** *Let $B \succeq 0$ have rank d and $M = I + A$ with $A \succeq 0$. In a basis adapted to $\mathrm{range}(B) \oplus \ker(B)$, write $B = \mathrm{diag}(B_1, 0)$, $B_1 \succ 0$, and split M into blocks. For every $\lambda > 0$,*

$$\log\det(I + A + \lambda^{-1} B) = d \log \frac{1}{\lambda} + \log\det B_1 + \log\det M_{22} + \log\det(I_d + \lambda B_1^{-1} \dot{M}),$$

where $\dot{M} = M_{11} - M_{12} M_{22}^{-1} M_{21}$. Moreover, $M_{22} \succeq I$, $0 \prec \dot{M} \preceq M_{11}$, and

$$0 \leq \log\det(I_d + \lambda B_1^{-1} \dot{M}) \leq \lambda\, \mathrm{tr}(B_1^{-1} \dot{M}).$$

Proof. The Schur determinant formula factors the determinant as $\det M_{22} \cdot \lambda^{-d} \det B_1 \cdot \det(I_d + \lambda B_1^{-1}\dot{M})$. Positive definiteness of M gives $\dot{M} \succ 0$; the subtracted term is positive semidefinite, giving $\dot{M} \preceq M_{11}$. The matrix $B_1^{-1}\dot{M}$ is similar to the positive definite matrix $B_1^{-1/2}\dot{M}B_1^{-1/2}$, so $\log(1+x) \le x$ applied to its eigenvalues gives the bound. □

Proof of Theorem A.7. The linear–Gaussian observation model gives

$$I(y_Q; y_S \mid y_T) = \tfrac{1}{2}\big[f_T(\lambda) - f_{T\cup S}(\lambda)\big], \quad f_D(\lambda) = \log\det\big(I + \Phi_Q(G_D + \lambda I)^{-1}\Phi_Q^\top\big).$$

Apply Lemma A.8 to (3). This yields $f_D = d_D \log(1/\lambda) + c_D + e_D$. The spectral formula for A_D implies

$$0 \preceq A_D(0) - A_D(\lambda) \preceq \lambda\Phi_Q(G_D^+)^2\Phi_Q^\top, \qquad \operatorname{tr}[A_D(0) - A_D(\lambda)] \le \lambda H_D. \tag{6}$$

Along the segment joining $M_{D,22}(\lambda)$ and $M_{D,22}(0)$, the inverse norm is at most one. Integrating the derivative of $\log\det$ therefore bounds the change in this block's log determinant between $-\lambda H_D$ and zero. The remaining term in Lemma A.8 lies between zero and λJ_D, since $\dot{M}_D(\lambda) \preceq M_{D,11}(\lambda) \preceq M_{D,11}(0)$. Consequently $|e_D| \le \lambda(H_D + J_D)$. Finally,

$$d_T - d_{T\cup S} = [r(T\cup Q) - r(T)] - [r(T\cup S\cup Q) - r(T\cup S)] = \nu_T(S, Q),$$

and subtracting the two expansions proves the result. □

Repeated measurements. Suppose every row of S is observed m times with independent noise. The measurement Gram matrix becomes $G_T + mG_S$, while the row space and hence $B_{T\cup S}$ remain fixed for every $m \ge 1$. The finite-noise effect follows directly from Gaussian conditioning. For a scalar query q, set

$$P_T = (G_T + \lambda I)^{-1}, \qquad K = \Phi_S P_T \Phi_S^\top, \qquad z = \Phi_S P_T q.$$

Write the positive eigenvalues and orthonormal eigenvectors of K as h_i, u_i. The risk in observation-noise units is exactly

$$r_m = q^\top P_T q - \sum_{h_i>0} \frac{m(u_i^\top z)^2}{1 + mh_i}. \tag{7}$$

Indeed, Woodbury gives $(G_T + mG_S + \lambda I)^{-1} = P_T - P_T\Phi_S^\top(m^{-1}I + K)^{-1}\Phi_S P_T$; z is orthogonal to $\ker K$. For a direction with $u_i^\top z \neq 0$, its risk reduction is the fraction $mh_i/(1+mh_i)$ of its infinite-repeat reduction. Half of that reduction is attained exactly when $mh_i \ge 1$. Thus repetition acts through the conditional measurement eigenvalues and their query projections.

For example, let $T = \{e_1\}$, $q = e_2$, and $S = \{\varepsilon e_2 - e_3, \varepsilon e_2 + e_3\}$. Then $B_T = 1$ and $B_{T\cup S} = 0$ for every $\varepsilon > 0$, whereas the query-aligned eigenvalue of K is $2\varepsilon^2/\lambda$. Writing $h = 2\varepsilon^2$, the exact formulas are

$$r_m = \frac{1}{\lambda + mh}, \qquad I(y_q; y_{S^{(m)}} \mid y_T) = \frac{1}{2}\log\frac{1 + 1/\lambda}{1 + 1/(\lambda + mh)}.$$

The latent-query variance is halved when $mh \ge \lambda$. This threshold includes the measurement signal scale ε^2.

A.10 Adaptivity is vacuous at the closure layer

Proof of Theorem A.1. Every closure-layer object listed is a function of $\operatorname{row}(\Phi_D)$ alone: $r(\cdot)$ and hence $\nu_T(S, Q)$ are ranks; the closure predicate is membership of φ_q in a row space; the witness family is the solution set of $\Phi_D^\top w = \varphi_q$; the gains are the positive eigenvalues of $\Phi_Q P_\perp^{(D)} \Phi_Q^\top$. None involves y.

Since none involves y, the conditional law of every closure-layer object given the responses equals its unconditional law: responses carry no information about them, which is the statement. Two consequences. A strategy whose decisions read *only* closure-layer objects is deterministic (induct on t: $D_0 = T$ fixed, so its objects and hence s_1 are; if $s_1, \dots, s_{t-1}$ are deterministic so is s_t), so adaptivity adds nothing within that class. And for a general adaptive strategy, any bound whose quantities are closure-layer objects evaluates identically on every response path, so it holds for adaptive and randomised strategies with no further argument. □

Consequently closure-layer acquisition is offline combinatorial planning, and the lower bounds below need no Yao-style argument: they are already statements about every adaptive and randomised strategy. Adaptivity retains its usual role at the precision layer, where σ estimation, model selection and repeat decisions do read y.

A.11 THE SIMPLEX CONSTRUCTION

Fix $m \geq 2$ and let $v_1, \ldots, v_m$ be regular-simplex vertices in $u^\perp$: $\|v_i\| = 1$, $\langle v_i, v_j \rangle = -\frac{1}{m-1}$ for $i \neq j$, $\sum_i v_i = 0$. Put

$$s_i = \varepsilon u + v_i \ (i \in [m]), \qquad d_k = \delta u + w_k \ (k \in [K]),$$

with $\{w_k\}$ orthonormal and orthogonal to u and to every v_i. Then $\sum_i s_i = m\varepsilon u$, so $[m]$ closes u. Take T orthogonal to all of these, so that conditioning on y_T is vacuous and we may set $T = \varnothing$ (this is the only role T plays; nothing below depends on it).

Lemma A.9 (No proper subset closes)**.** *For $A \subsetneq [m]$, $u \notin \operatorname{span}\{s_i\}_{i \in A}$.*

Proof. $\{v_i\}_{i \in A}$ is linearly independent for $|A| < m$ (the simplex Gram $\frac{m}{m-1} I - \frac{1}{m-1} J$ restricted to A is positive definite). If $u = \sum_{i \in A} c_i s_i = \varepsilon(\sum_i c_i) u + \sum_i c_i v_i$, matching the u-component gives $\varepsilon \sum_i c_i = 1$, so $\sum_i c_i \neq 0$, while matching the $u^\perp$ component gives $\sum_{i \in A} c_i v_i = 0$, so $c \equiv 0$ by independence, contradicting $\sum_i c_i \neq 0$. □

For a purchased set S write $c = \Phi_S u$ and $K_S = \Phi_S \Phi_S^\top$, and define the *effective increment*

$$t_S^{(\lambda)} := c^\top (K_S + \lambda I)^{-1} c, \qquad t_S := t_S^{(0)}.$$

By the kernel form of r_λ, $r^S := u^\top (G_S + \lambda I)^{-1} u = \big(1 - t_S^{(\lambda)}\big)/\lambda$, so r^S is a strictly decreasing function of $t_S^{(\lambda)}$; $t_S = 1$ iff S closes u.

Lemma A.10 (Exact increment, mixed sets)**.** *For a set with a circuit members and b distractors, the effective increment toward u is*

$$t_{a,b} = \frac{N}{N + W(1+\lambda)}, \qquad W = \frac{m-a}{m-1} + \lambda, \qquad N = a\varepsilon^2(1+\lambda) + b\delta^2 W,$$

and $\nu = 0$ whenever $a < m$. Setting $b = 0$ or $a = 0$ recovers the pure cases; the cross term $ab\varepsilon^2\delta^2$ cancels identically.

Proof of Lemma A.10. Let A hold a circuit members and b distractors, $c = \Phi_A u = (\varepsilon \mathbf{1}_a, \delta \mathbf{1}_b)^\top$. Every block of $K_A = \Phi_A \Phi_A^\top$ is a multiple of I plus a multiple of J: the circuit block is $\alpha I_a + \beta J_a$ with $\alpha = \frac{m}{m-1}$, $\beta = \varepsilon^2 - \frac{1}{m-1}$; the distractor block is $I_b + \delta^2 J_b$; the cross block is $\varepsilon\delta J_{a \times b}$. Hence $V = \operatorname{span}\{(\mathbf{1}_a, 0), (0, \mathbf{1}_b)\}$ is invariant under $K_A + \lambda I$, and since $c \in V$ the increment $t_{a,b} = c^\top (K_A + \lambda I)^{-1} c$ reduces to a 2×2 solve. In the basis $(\mathbf{1}_a, 0), (0, \mathbf{1}_b)$ the operator is

$$M = \begin{pmatrix} P & \varepsilon\delta b \\ \varepsilon\delta a & S \end{pmatrix}, \qquad P = W + a\varepsilon^2, \quad S = (1+\lambda) + b\delta^2, \quad W := \tfrac{m-a}{m-1} + \lambda,$$

using $\alpha + \lambda + a\beta = W + a\varepsilon^2$. Then $M^{-1}(\varepsilon, \delta)^\top = \frac{1}{\det M}(\varepsilon(S - b\delta^2),\ \delta(P - a\varepsilon^2))^\top = \frac{1}{\det M}(\varepsilon(1+\lambda),\ \delta W)^\top$, and since $\langle \mathbf{1}_a, \mathbf{1}_a \rangle = a$, $\langle \mathbf{1}_b, \mathbf{1}_b \rangle = b$,

$$t_{a,b} = \frac{a\varepsilon^2(1+\lambda) + b\delta^2 W}{\det M} = \frac{N}{\det M}.$$

Finally $\det M = PS - ab\varepsilon^2\delta^2 = W(1+\lambda) + Wb\delta^2 + a\varepsilon^2(1+\lambda) + ab\varepsilon^2\delta^2 - ab\varepsilon^2\delta^2 = N + W(1+\lambda)$, the cross term cancelling identically. Setting $b = 0$ gives $t_{a,0} = \varepsilon^2 a(m-1)/[(m-a) + \lambda(m-1) + a\varepsilon^2(m-1)]$ and $a = 0$ gives $t_{0,b} = \delta^2 b/(1 + \lambda + b\delta^2)$, the two pure cases. For $a < m$, $W > 0$ so $t_{a,b} < 1$ and $\nu = 0$. □

Lemma A.11 (Exact selection condition)**.** *$t_{a,b}$ is strictly increasing in $\mathcal{R}(a,b) = a\varepsilon^2(1+\lambda)/W + b\delta^2$. A distractor step therefore gains exactly δ^2 and a circuit step exactly $\varepsilon^2(1+\lambda)\frac{(m-1)(m+c)}{(m-a-1+c)(m-a+c)}$ with $c = \lambda(m-1)$,* neither depending on *b. At $a = 0$ the distractor wins iff $\delta > \varepsilon$, and the threshold rises with a, diverging as the last circuit member closes the query.*

Proof of Lemma A.11. $t_{a,b} = \big(1 + W(1+\lambda)/N\big)^{-1}$ is strictly increasing in $N/W = \mathcal{R}(a,b)$. Appending a distractor sends $\mathcal{R} \mapsto \mathcal{R} + \delta^2$ regardless of a and b. Appending a circuit member changes only the first term, by $\varepsilon^2(1+\lambda)\big[\frac{a+1}{W(a+1)} - \frac{a}{W(a)}\big]$. Writing $W(a) = \frac{m-a+c}{m-1}$ with $c = \lambda(m-1)$,

$$\frac{a+1}{W(a+1)} - \frac{a}{W(a)} = (m-1)\,\frac{(a+1)(m-a+c) - a(m-a-1+c)}{(m-a-1+c)(m-a+c)} = \frac{(m-1)(m+c)}{(m-a-1+c)(m-a+c)},$$

the numerator collapsing to $m + c$. Neither gain involves b. At $a = 0$ the circuit gain is $\varepsilon^2(1+\lambda)\frac{(m-1)(m+c)}{(m-1+c)(m+c)} = \varepsilon^2$ using $c = \lambda(m-1)$, so the distractor wins iff $\delta^2 > \varepsilon^2$; the denominator decreases in a, so the threshold rises, and at $a = m-1$ it diverges because that purchase closes the query. □

A.12 The submodularity ratio

Proof of Theorem 2.2. Recall $\gamma_{U,k}(F) = \min\big\{\sum_{x\in S}[F(L\cup\{x\}) - F(L)]/[F(L\cup S) - F(L)]\big\}$ over $L \subseteq U$ and S with $|S| \le k$, $S \cap L = \varnothing$, with $0/0 := 1$.

Take $F(A) = \mathrm{cl}_q(T \cup A)$, let C be a minimal closing bundle for q of size $m \ge 2$ contained in the pool, and put $L = C \setminus \{c_1, c_2\}$, $S = \{c_1, c_2\}$, which is admissible for every $k \ge 2$. Both L and $L \cup \{c_i\}$ are proper subsets of the circuit C, so by minimality of circuits neither closes q and the numerator is $0 + 0 = 0$ exactly. The denominator is $F(C) - F(L) = 1 - 0 = 1 > 0$. Hence $\gamma_{U,k} \le 0$, and $\gamma_{U,k} \ge 0$ by monotonicity of cl_q, so $\gamma_{U,k} = 0$.

At $k = 1$ the numerator and denominator coincide term by term, so $\gamma_{U,1}(F) = 1$ for *every* set function F; the $k = 1$ value certifies nothing. Finally γ does not stratify: the argument used a size-2 set S regardless of m, so $\gamma_{U,2} = 0$ for every circuit size, and γ cannot distinguish a pairwise rule from a j-wise one. □

The closure objective records whether a reconstruction exists. The next corollary gives the corresponding ratio for the finite-noise mutual-information objective.

Corollary A.12 (Ratio of the mutual-information objective)**.** *With $\psi_\lambda(A) = I(y_q; y_A \mid y_T)$ and the instance of Appendix A.11, there is a constant c_0 depending on the instance but not on λ such that for every $k \ge 2$*

$$0 \;\le\; \gamma_{U,k}(\psi_\lambda) \;\le\; \frac{c_0}{\log(1/\lambda)} \;\xrightarrow[\lambda\to 0]{}\; 0.$$

Proof. Take L and S as above. Every set involved in the numerator has $\nu = 0$, so by Theorem A.7 its mutual information is bounded by a λ-free constant; monotonicity of ψ_λ in A makes the numerator non-negative and at most $2\bar{C}$ for such a constant $\bar{C}$. The denominator is $\psi_\lambda(C) - \psi_\lambda(L) \ge \frac{1}{2}\log\frac{1}{\lambda} - C_0$ for a finite constant C_0, because $\nu_T(C, \{q\}) = 1$ and the remainder is $O(\lambda)$. Take λ small enough that the denominator is at least $\frac{1}{4}\log(1/\lambda)$ and set $c_0 = 8\bar{C}$. Non-negativity is monotonicity again. □

So the weak-submodularity fallback degrades as $1/\log(1/\lambda)$ in the noise level and vanishes in the algebraic limit: it is not merely loose in the open region, it carries no rate there.

A.13 Exact predictive information and greedy selection

Theorem A.13 (Greedy trap, simplex form)**.** *Let $\delta > \varepsilon > 0$ and $K \ge B \ge 2$. For every $\lambda > 0$, $\pi_\psi^{\mathrm{greedy}}$ (Definition B.2) purchases only distractors for B steps, and its objective is bounded by $\frac{1}{2}\log(1 + B\delta^2)$, a constant independent of λ, while the two-element planner attains $\frac{1}{2}\log\frac{1}{\lambda} + C \to \infty$. The ratio diverges as $\sigma/\rho \to 0$. At $B = 6$, $\lambda = 10^{-8}$ greedy buys six distractors for 0.3365 against the planner's 6.56.*

Proof. For the pair construction of Theorem A.4, predictive information is $\psi_\lambda(S) = \frac{1}{2}\log[(1 + 1/\lambda)/(1 + r^S)]$. It is strictly decreasing in the posterior risk r^S, so exact-MI greedy makes the same

choices as C-optimal greedy: B distractors. Substituting the risks derived in the activation-barrier proof gives

$$\psi_\lambda(D_B) = \frac{1}{2}\log\frac{(1+\lambda)(1+\lambda+B\delta^2)}{(1+\lambda)^2+\lambda B\delta^2} < \frac{1}{2}\log(1+B\delta^2),$$

whereas the complete pair gives

$$\psi_\lambda(\{s_1,s_2\}) = \frac{1}{2}\log\frac{(1+\lambda)(2\varepsilon^2+\lambda)}{\lambda(1+2\varepsilon^2+\lambda)} = \frac{1}{2}\log(1/\lambda)+O(1).$$

Their ratio therefore diverges as $\lambda \to 0$. □

A.14 Verification of the oracle construction

The complete argument is in Section 3.1. `verify_oracle_construction.py` constructs the identical candidate rows and the two targets directly. It checks every accessible set for $j = 1, \ldots, 6$ at four prior/noise settings: from 20 to 25,904 set/parameter pairs, every oracle gap is numerically zero. It also enumerates every budget-size purchase, from 6 to 3432 sets, verifying that exactly one path closes each target and that the two paths are disjoint. The equivalent signing construction in `verify_twin_all_orders.py` additionally checks that order $j+1$ separates the instances.

A.15 Structured shift is necessary

The shift is *structured* when the prior scale of β on the queries' open directions equals its scale on the directions closed by T (blockwise in the block-resolved form), the isotropic prior above being the simplest case. Under this assumption, the scale measured on closed directions transfers to open ones. Theorem A.14 establishes its role in identifying allocations; Appendices C.8 and D.2 give the measured family variation and batch selection over parameter intervals.

Two properties are used repeatedly: monotonicity, $D \subseteq D' \Rightarrow \mathrm{cl}_Q(D) \le \mathrm{cl}_Q(D')$, and the dependence of cl_Q on the design alone. Theorem A.1 establishes that observed responses leave the closure of a fixed design unchanged.

Theorem A.14 (No pre-measurement decision without structured shift)**.** *There are pairs of instances, identical in design and identical in the distribution of y_T, whose risk-optimal allocations are opposite, structure-first in one, precision-first in the other, each by a constant factor. Consequently any allocation rule computed from (y_T, Φ), deterministic or randomised, selects the suboptimal arm on one member of the pair with probability at least $1/2$.*

Proof. Let $u = P_\perp^{(T)}\varphi_q / \|P_\perp^{(T)}\varphi_q\|$ and let the two priors agree on $\mathrm{row}(\Phi_T)$, with scale ρ on u under π_1 and 0 under π_2. Since $\Phi_T\beta = \Phi_T(P^{(T)}\beta)$, the law of y_T is the same Gaussian under both. Give the budget two layer-disjoint options: a single-direction closing purchase S_{struct} along u with gain g, and a precision purchase S_{stat} with gain $P > 0$ on supported directions. By the spectral split (Theorem A.2), the risk drops are

$$\Delta_1(\mathrm{struct}) = \rho^2\beta\,\frac{g\kappa^2}{g\kappa^2+1}, \quad \Delta(\mathrm{stat}) = \sigma^2 P \qquad (\beta = \|P_\perp^{(T)}\varphi_q\|^2),$$

under π_1, while under π_2 the structural drop is 0. Choosing ρ with $g\kappa^2 \ge 1$ and $\rho^2\beta \ge 4\sigma^2 P$ makes structure-first optimal by a factor ≥ 2 under π_1 and precision-first optimal outright under π_2. A rule that is a function of (y_T, Φ) has the same output distribution on both instances; whichever arm it selects with probability $\ge 1/2$ is the suboptimal arm of one member. (Numerical check: `verify_structured_shift.py`, law-of-y_T equality exact, margins as stated.) □

A.16 Numerical construction for the recovery theorem

The complete-pair barrier is Proposition 4.4; its proof and the recovery proofs appear in Appendix A.18. The path and cycle experiments use $C = \gamma^2[\rho L + (1-\rho)D]$, with D the graph degree matrix and $\gamma^2 = 2a/\min_\ell \lambda_{\min}(L_{-\ell,-\ell})$. Choosing $\rho = 1 - \theta/(\gamma^2\,\mathrm{tr}(D)/b)$ gives $\mathbf{1}^\top C\mathbf{1}/b = \theta$. The fixed grid and enumeration checks are recorded in Appendix C.14.

A.17 COEFFICIENT-BASED PROPOSALS

Let $R(S) = \operatorname{tr}[Q(H + X_S^\top X_S)^{-1}Q^\top]$, $H \succ 0$, $K = I + XH^{-1}X^\top$, $Z = XH^{-1}Q^\top$, and $W_* = K^{-1}Z$. Eliminating the measured coefficients in Theorem A.3, or completing the square using the Woodbury identity, gives the reconstruction cost

$$J(W) = R_{\text{all}} + \operatorname{tr}[(W - W_*)^\top K(W - W_*)], \qquad \min_{\operatorname{rowsupp}(W)\subseteq S} J(W) = R(S).$$

Thus the row support is exactly the purchased batch. For $F(W) = J(W) + \tau \sum_i \log(1 + \|W_i\|^2/\epsilon)$, the tangent upper bound of the concave logarithm at the current row energies is a quadratic with minimiser $W_{\text{new}} = (K + D)^{-1}Z$, where $D_{ii} = \tau/(\epsilon + \|W_i\|^2)$. It touches F at W, so

$$F(W) - F(W_{\text{new}}) \geq \operatorname{tr}[(W_{\text{new}} - W)^\top (K + D)(W_{\text{new}} - W)] \geq 0.$$

This is the fixed-penalty descent property used to generate supports. Integer batches are compared using R itself. The full schedule and single-measurement refinement are specified in Appendix C.14.

A.18 GENERAL NUISANCE GRAM MATRICES

Use the paired design of Section 4, and write $q = \varphi_q$, $C = VV^\top$, and $c_0 = 2a/(1 + 2a)$. For a batch let F, P, Z index its full, single, and unpurchased pairs. Thus $f = |F|$, $p = |P|$, $2f + p = b$, and $|Z| = f$. The vector $s \in \{-1, 1\}^P$ records the singleton signs. Empty matrix blocks have their usual vacuous meaning.

Proof of Lemma 4.1. For each full pair, the normalized sum and difference observations have independent unit noise and feature vectors $\sqrt{2a}\, u_j$ and $\sqrt{2}\, v_j$. The sums reduce target risk by c_0/b per pair. Conditioning on the differences leaves singleton nuisance covariance

$$D = V_P(I + 2V_F^\top V_F)^{-1}V_P^\top = C_{PP} - C_{PF}(C_{FF} + \tfrac{1}{2}I)^{-1}C_{FP}.$$

The signal coordinates of the singletons have independent unit priors and remain independent of this nuisance. Their observation covariance is $(1 + a)I + E_s D E_s$, where $E_s = \operatorname{diag}(s)$. Gaussian conditioning therefore gives their risk reduction as $a\mathbf{1}^\top[(1+a)I + E_s D E_s]^{-1}\mathbf{1}/b = as^\top[(1+a)I + D]^{-1}s/b$, proving Equation (1). The same formula holds for intermediate batches with $2f + p \leq b$. Completing the square gives, for every $x \in \mathbb{R}^P$,

$$x^\top D x = \min_{y \in \mathbb{R}^F} \left\{ \begin{pmatrix} y \\ x \end{pmatrix}^\top C_{F\cup P, F\cup P} \begin{pmatrix} y \\ x \end{pmatrix} + \tfrac{1}{2}\|y\|^2 \right\}. \tag{8}$$

The matrix $C_{FF} + \frac{1}{2}I$ is positive definite even if C is singular. This identity shows that D increases in Loewner order with C, and hence the risk of each fixed batch increases with C. In particular, $C \succeq \delta I$ implies $D \succeq \delta I$. □

This elimination is a Schur-complement calculation; its graph-energy form is related to Kron reduction (Dörfler & Bullo, 2013).

Pairing and its spectral margin.

Proof of Theorem 4.2. Orient each candidate x_i so that $x_i^\top q > 0$. All these inner products then equal $\sqrt{a/b}$, so $a = b(x_i^\top q)^2$ is observable. Set $G = XX^\top$ and $M = G - 2aI_{2b}$. For an original pair i, j, its sum is $2\sqrt{a}\, u_k$, and

$$G(e_i + e_j) = 2a(e_i + e_j), \qquad M_{i,:} + M_{j,:} = 0.$$

Group the nonzero rows of M by their value. Each value and its negative have equal multiplicity, because an original pairing exists. The zero rows occur an even number of times. Consequently, matching rows in opposite classes, and pairing the zero rows arbitrarily, always produces a perfect matching in polynomial time.

For any such matching define $h_k = (e_i + e_j)/\sqrt{2}$, $d_k = (e_i - e_j)/\sqrt{2}$, and

$$u'_k = \frac{x_i + x_j}{2\sqrt{a}}, \qquad v'_k = \frac{x_i - x_j}{2}.$$

The vectors h_k, d_k form an orthonormal basis of candidate space, and $Gh_k = 2ah_k$. It follows that $\langle u'_k, u'_\ell \rangle = \mathbf{1}\{k = \ell\}$ and $\langle u'_k, v'_\ell \rangle = 0$. Moreover, $\langle u'_k, q \rangle = 1/\sqrt{b}$. Since $\|q\| = 1$, these projections exhaust its norm, giving $q = b^{-1/2} \sum_k u'_k$. Every recovered matching therefore yields a valid paired representation.

In the original pair-sum and pair-difference basis,

$$G \cong \mathrm{diag}(2aI_b, 2C), \qquad M \cong \mathrm{diag}(0, 2(C - aI_b)). \tag{9}$$

Let $\eta = \min_k |\lambda_k(C) - a| > 0$. A false pair has a pair-difference component of norm one, so its residual $\|M(e_i + e_j)\|$ is at least 2η; a true pair has zero residual. If the oriented shifted matrix changes to $\widetilde{M}$ with $\|\widetilde{M} - M\|_{\mathrm{op}} \le e$, a true residual is at most $\sqrt{2}e$ and every false residual is at least $2\eta - \sqrt{2}e$. Thus $e < \eta/\sqrt{2}$ makes every true partner the unique nearest residual partner. Error in the estimated shift $2aI$ is included in e. □

The global optimum.

Proof of Theorem 4.3. Set $\delta = \lambda_{\min}(C)$. The bound $D \succeq \delta I$ and Equation (1) imply

$$R(S) \ge 1 - \frac{1}{b}\left[\frac{2af}{1+2a} + \frac{a(b-2f)}{1+a+\delta}\right]. \tag{10}$$

This affine function of f is minimized at $f = 0$ when $\delta < a$ and at $f = b/2$ when $\delta > a$. If $Cs = \delta s$ for a sign vector, the all-single batch with signs s attains the former endpoint, $(1+\delta)/(1+a+\delta)$. Every batch of $b/2$ full pairs attains the latter endpoint, $(1+a)/(1+2a)$. At $\delta = a$, full-pair batches always attain the bound. An all-single batch attains it exactly when its sign vector belongs to the minimum eigenspace: equality in $s^\top[(1+a)I + C]^{-1}s \le b/(1+2a)$ has precisely this condition.

For the algorithm, compute a valid matching using Theorem 4.2 and form its nuisance Gram matrix C'. Equation (9) holds for every valid matching, so the spectrum of C' equals that of C, with multiplicity. If $\delta \ge a$, buy any $b/2$ recovered full pairs. If $\delta < a$ is simple and its eigenvector is a sign vector divided by $\sqrt{b}$, the unique unit eigenvector of G at 2δ has all $2b$ entries of magnitude $1/\sqrt{2b}$. For every recovered pair its two entries sum to zero, because its eigenvalue differs from $2a$. Its coordinates in the recovered pair-difference basis thus have magnitude $1/\sqrt{b}$. They provide a sign eigenvector of C' and the optimal all-single batch. Matching, matrix construction, and eigendecomposition require polynomially many arithmetic operations. This argument includes matrices having a as a nonminimal nuisance eigenvalue. □

For every connected positively weighted graph, $C = \delta I + tL$, $t > 0$, has simple minimum eigenvector $\mathbf{1}/\sqrt{b}$; diagonal sign conjugation preserves the condition. If $\delta > 0$, Equation (9) also shows that the $2b$ candidate rows are independent. Their unique representation of q is

$$q = \frac{1}{2\sqrt{ab}} \sum_{j=1}^{b} (x_j^+ + x_j^-),$$

which uses every row. Thus the finite-noise optimum above also applies when every budget-b batch leaves the query open.

A sufficient half-budget barrier.

Proof of Proposition 4.4. Let $\mu = \min_\ell \lambda_{\min}(C_{-\ell,-\ell}) > a$. Every proper principal submatrix of C is at least μI. If $F \cup P$ is proper, Equation (8) gives $D \succeq \mu I$ and singleton gain at most $ap/(1+a+\mu)$. For S_0 consisting of $b/2$ full pairs, write $R_0 = (1+a)/(1+2a)$. A mixed batch with $0 < f < b/2$ then satisfies

$$R(S) - R_0 \ge \frac{b/2 - f}{b}\left[\frac{2a}{1+2a} - \frac{2a}{1+a+\mu}\right] > 0.$$

Fewer than $b/2$ replacements retain a full pair, so $f > 0$ and $|F \cup P| = b - f < b$. Such exchanges either give this strict inequality or produce another full-pair batch with risk R_0.

After f full pairs and one singleton, completing that pair gives total gain $(f+1)c_0/b$. Opening a second pair gives total gain at most $[fc_0 + 2a/(1+a+\mu)]/b < (f+1)c_0/b$. Here the relevant union has at most $f+2 \le b/2+1 < b$ pairs. Induction therefore shows that single-measurement greedy completes each pair it opens and finishes at a full-pair batch.

For every size-b batch, $q^\top(I + X_S^\top X_S)q = 1+a$. Cauchy–Schwarz gives $R(S) \ge 1/(1+a)$. For an all-single batch with signs s, put $\theta = s^\top Cs/b$. The normalized sum of its observations has signal $\sqrt{a}\, q^\top \beta$ and independent noise variance $1+\theta$, including the nuisance prior. Using that scalar alone gives

$$R(S_s) \le \frac{1+\theta}{1+a+\theta}, \qquad \frac{R(S_s)}{R^*} \le \frac{(1+a)(1+\theta)}{1+a+\theta}.$$

When $\theta < a$, this upper bound is below R_0 and S_s is exactly $b/2$ replacements from S_0. The minimum improving exchange size is therefore $b/2$.

For completeness, the nearest-sum pairing rule also works under this condition. A true pair has squared sum norm $4a$. For distinct i, j, choose $\ell \notin \{i, j\}$ and apply the Rayleigh bound in $C_{-\ell,-\ell}$ to obtain $C_{ii} + C_{jj} - 2|C_{ij}| \ge 2\mu$. Every false partner has cost at least $2a + 2\mu > 4a$. If $Cs = 0$ for a sign vector, the kernel is one-dimensional: two independent null vectors would yield a nonzero null vector with a zero coordinate, contradicting a positive definite deleted principal matrix. The null eigenvector therefore recovers the signs and attains $R^* = 1/(1+a)$. □

For a connected graph, deleted principal Laplacians are grounded Laplacians and are positive definite (Pirani & Sundaram, 2016). Thus $C = tL + \delta I$, with $0 \le \delta < a$, satisfies $\delta < a < \mu$ for sufficiently large t. A path provides such examples with nuisance interaction graph of treewidth one. On any tree, a traversal minimizes $s^\top Cs$ by assigning $s_i s_j = -\operatorname{sign}(C_{ij})$ on each nonzero edge.

The symmetric exchange formula and spectral comparisons.

Proof of Proposition 4.5. The calculation also covers $C = \delta I + t(bI - \mathbf{1}\mathbf{1}^\top)$, $0 \le \delta < a$, $t \ge 0$. Set

$$A = \delta + tb, \qquad \beta_f = \frac{t(1+2\delta+2tb)}{1+2\delta+2t(b-f)}, \qquad \alpha = 1 + a + A.$$

The rank-one inverse of $C_{FF} + \frac{1}{2}I$ gives $D = AI_p - \beta_f \mathbf{1}\mathbf{1}^\top$. Therefore

$$R(S) = 1 - \frac{1}{b}\left[c_0 f + a\left\{\frac{p}{\alpha} + \frac{\beta_f(\mathbf{1}^\top s)^2}{\alpha(\alpha - \beta_f p)}\right\}\right]. \tag{11}$$

Since $\beta_f \ge 0$, aligned singleton signs minimize the risk. With $p = b - 2f$, their residual eigenvalue is

$$d_f(\delta, t) = A - \beta_f(b-2f) = \delta + \frac{2tf(1+2\delta+tb)}{1+2\delta+2t(b-f)}. \tag{12}$$

For $p > 0$ their risk difference from R_0 is

$$R_{\min}(f) - R_0 = \frac{ap[d_f(\delta,t) - a]}{b(1+2a)[1+a+d_f(\delta,t)]}. \tag{13}$$

The quantity d_f is nondecreasing in f, strictly so for $t > 0$. An r-exchange breaks at most r original full pairs and hence has $f \ge b/2 - r$. This rules out improvement if $d_{b/2-r} \ge a$. Conversely, break r original full pairs, retain aligned singletons, and buy r aligned singletons from previously unpurchased pairs. This attains $f = b/2 - r$ at exchange size r. Since $d_0 = \delta < a$,

$$r(\delta, t) = \min\{r \in \{1, \dots, b/2\} : d_{b/2-r}(\delta, t) < a\} \tag{14}$$

is the exact minimum improving exchange size. At $\delta = 0$ this is the formula in Proposition 4.5.

For its greedy condition, after f full pairs the best two-singleton gain uses aligned signs and residual eigenvalue $e_f = tb - 2t(1+2tb)/[1+2t(b-f)]$. Completing a singleton gives larger gain exactly when $e_f > a$. The sequence e_f decreases for $0 \le f \le b/2 - 1$, with final value $d_{b/2-1}(0, t)$. The stated strict condition therefore implies pair completion at every step. Finally, a deleted principal submatrix has eigenvalues t and tb, so $\mu = t$. □

In this symmetric family the maximal barrier $r_* = b/2$ holds exactly when $d_1(0,t) \geq a$, or equivalently

$$t \geq \frac{a(b-1) - 1 + \sqrt{[1-a(b-1)]^2 + 2ab}}{2b}.$$

At equality the relevant exchanges tie. For example, $b = 8, a = 1, t = 1/2$ gives $r_* = 3$ although $\mu = 1/2 < a$.

Now suppose $Cs = \delta s$, $s \in \{-1,1\}^b$, $\delta < a$, and write the remaining eigenvalues as $\lambda_2 \leq \cdots \leq \lambda_b$. Switch nuisance signs so that $s = \mathbf{1}$. Spectral decomposition gives

$$\delta I + (\lambda_2 - \delta)(I - \mathbf{1}\mathbf{1}^\top/b) \preceq C \preceq \delta I + (\lambda_b - \delta)(I - \mathbf{1}\mathbf{1}^\top/b).$$

All three matrices give the same risk R_0 at a full-pair batch. Monotonicity from Equation (8) and the exact thresholds in Equation (14) therefore imply, for every such start,

$$r\left(\delta, \frac{\lambda_2 - \delta}{b}\right) \leq r_* \leq r\left(\delta, \frac{\lambda_b - \delta}{b}\right). \tag{15}$$

Indeed, the lower comparison rules out each smaller exchange, and an aligned exchange improving the upper comparison also improves C. The bounds coincide when the nonminimal eigenvalues coincide.

Cycles: obstruction and construction.

Proof of Theorem 4.6. Let $C = \delta I + tL$ for the unweighted cycle, $0 \leq \delta < a$, and let H be any $b/2$ starting full pairs. Define

$$\epsilon = a - \delta, \qquad d = \delta + \tfrac{1}{2}, \qquad c = 1 + a + \delta, \qquad K = \begin{cases} \epsilon, & \epsilon \leq d, \\ (\epsilon + d)^2/(4d), & \epsilon > d. \end{cases}$$

Consider an exchange of at most $r < b/2$ measurements. It retains $f \geq b/2 - r > 0$ full pairs and has $p \leq 2r$ singletons. If $p = 0$ it ties R_0. Otherwise $Z \neq \varnothing$. For $\|x\|^2 = 1$, extend x to a cycle vector z by $z_P = x$, $z_F = y$, $z_Z = 0$. Equation (8) becomes

$$x^\top D x = \delta + \min_y \left[t \sum_{ij \in E} (z_i - z_j)^2 + d\|y\|^2 \right]. \tag{16}$$

Because z has a zero coordinate, the total variation of z_i^2 around the cycle is at least $2\|z\|_\infty^2$. Cauchy–Schwarz yields

$$2\|z\|_\infty^2 \leq \sum_{ij \in E} |z_i^2 - z_j^2| \leq 2\|z\|_2 \left[\sum_{ij \in E} (z_i - z_j)^2 \right]^{1/2}.$$

Using $1 \leq p\|z\|_\infty^2$ and $\|z\|_2^2 = 1 + \|y\|^2$ in Equation (16) gives

$$x^\top D x \geq \delta + \min_{v \geq 0} \left[\frac{t}{p^2(1+v)} + dv \right].$$

For $k = t/p^2$, this minimum is k when $k \leq d$ and $2\sqrt{kd} - d$ otherwise. It is at least ϵ exactly when $k \geq K$. Thus $t \geq 4Kr^2$ implies $D \succeq aI$ for every feasible exchange under consideration. Equation (1) then gives $R(S) \geq 1 - [c_0 f + ap/(1+2a)]/b = R_0$, proving the obstruction for all choices of singleton signs.

For the construction, fix $1 \leq r < b/2$. Each cycle vertex belongs to exactly $2r$ of the b cyclic intervals of length $2r$. Hence the average value of $|P \cap H|$ over these intervals is r. Sliding an interval changes this integer by at most one, so some interval P has $|P \cap H| = r$. Break those r full pairs, retain aligned singletons, and add aligned singletons at the other r vertices of P. Retain every pair in $H \setminus P$. This is a feasible r-exchange.

Conditioning reduces the nuisance, so $D \preceq C_{PP} = \delta I + tL_p^D$, where $p = 2r$ and L_p^D has diagonal entries 2 and adjacent off-diagonal entries -1. Therefore the actual singleton quadratic form satisfies

$$\mathbf{1}^\top [(1+a)I + D]^{-1} \mathbf{1} \geq Q_p := \mathbf{1}^\top (cI + tL_p^D)^{-1} \mathbf{1}. \tag{17}$$

Let

$$z = \frac{2t}{c + 2t + \sqrt{c^2 + 4ct}} \in (0, 1).$$

The solution of $(cI + tL_p^D)w = \mathbf{1}$ is

$$w_i = \frac{1}{c}\left[1 - \frac{z^i + z^{p+1-i}}{1 + z^{p+1}}\right], \qquad 1 \le i \le p.$$

It satisfies the two zero boundary conditions and the recurrence $(c + 2t)w_i - tw_{i-1} - tw_{i+1} = 1$. Summing it gives

$$Q_p = \frac{1}{c}\left[p - \frac{2z(1 - z^p)}{(1 - z)(1 + z^{p+1})}\right] \ge \frac{1}{c}\left[p - \left(\sqrt{1 + 4t/c} - 1\right)\right], \tag{18}$$

using $2z/(1 - z) = \sqrt{1 + 4t/c} - 1$. The constructed batch improves exactly when its actual quadratic form in Equation (17) exceeds $p/(1 + 2a)$. Consequently, $Q_p > p/(1 + 2a)$ is a sufficient condition, and Equation (18) guarantees it whenever

$$2r > \frac{1 + 2a}{a - \delta}\left(\sqrt{1 + \frac{4t}{1 + a + \delta}} - 1\right). \tag{19}$$

Finding a balanced interval requires only a scan of the cycle. At $r = b/2$, retaining and adding aligned singletons reaches the all-single optimum from Theorem 4.3.

More explicitly, the two arguments give

$$\min\left\{\frac{b}{2}, \left\lfloor\sqrt{\frac{t}{4K}}\right\rfloor + 1\right\} \le r_* \le \min\left\{\frac{b}{2}, \left\lfloor\frac{1 + 2a}{2(a - \delta)}\left(\sqrt{1 + \frac{4t}{1 + a + \delta}} - 1\right)\right\rfloor + 1\right\}.$$

For fixed a, δ, these imply $r_* = \Theta_{a,\delta}(\min\{b, 1 + \sqrt{t}\})$, uniformly over the starting set H. □

If cycle weights satisfy $0 < w_{\min} \le w_e \le w_{\max} < \infty$, use $tw_{\min}$ in the obstruction and $tw_{\max}$ in the construction, by the quadratic-form comparison $w_{\min}L \preceq L_w \preceq w_{\max}L$. The same order of growth holds with constants depending on $a, \delta, w_{\min}, w_{\max}$ when both weight bounds are independent of b and t. If $\eta > 0$, the pairing theorem recovers the original cycle pairs; the minimum eigenvector fixes their nuisance signs, and the negative off-diagonal entries then give the edges and cyclic order. This also makes the bounded-exchange construction available from hidden feature coordinates. The global algorithm of Theorem 4.3 applies to every valid recovered pairing, including the resonant cases.

A.19 Grand-mean witness purchase on a factorial grid

Proof of Theorem 5.2. Reduce from partial Latin-square completion, which is NP-complete (Colbourn, 1984). An instance is an explicitly represented $n \times n$ partial array whose prescribed symbols lie in $[n]$ and do not repeat within a row or column. Associate symbol k in position (i, j) with the triple (i, j, k) and let T contain these triples. Construct the complete grid $E = [n]^3$, with feature row φ_{ijk} containing the intercept, all three one-hot main effects, and all three one-hot pair effects. Set

$$\bar{q}_n = \frac{1}{n^3}\sum_{x \in E}\varphi_x = \left(1, \frac{1}{n}\mathbf{1}_{3n}, \frac{1}{n^2}\mathbf{1}_{3n^2}\right), \qquad B = n^2 - |T|.$$

All cells in $E \setminus T$ are purchasable at unit cost. The grid, target, and budget have size polynomial in the array representation.

If L is a Latin completion, its n^2 triples contain every pair of factor levels exactly once. Its main-effect counts are all n, so $n^{-2}\sum_{x \in L}\varphi_x = \bar{q}_n$. Buying $L \setminus T$ therefore closes the target within budget.

Conversely, suppose $|S| \le B$ and real coefficients w satisfy $\Phi_{T \cup S}^\top w = \bar{q}_n$. Rescale them by n^2. Every AB coordinate of the rescaled target equals one, and each feature row contributes to exactly one such coordinate. Hence a representation must use at least n^2 distinct rows with nonzero

coefficients. But $|T \cup S| \le n^2$, so it uses exactly one row at each AB position, uses every row in $T \cup S$, and assigns each rescaled coefficient the value one. The AC and BC equations then require each symbol to occur once in every row and column. These triples form a Latin square extending every entry of T. This support argument applies to arbitrary real witness coefficients.

Finally, a proposed purchase set is a polynomial-size certificate. Its closure can be checked by comparing the rational ranks of $\Phi_{T\cup S}$ and that matrix with $\bar{q}_n^\top$ appended. Rational elimination has polynomial bit complexity, proving membership in NP and completing the reduction. □

The verification script `verify_latin_reduction.py` compares floating-point span solves with Latin completion on 67,930 budget-sized purchase sets, covering all 35 valid order-two partial squares and two order-three examples. It also checks integer certificates for the accepted sets. The theorem concerns the grand mean: a single held-out cell on the full three-factor grid has a seven-cell witness from an alternating $2 \times 2 \times 2$ cube.

A.20 Uniform coefficient recovery on the factorial grid

The pair-effect operator is the three-direction orthogonal projection operator in discrete tomography: both record all pairwise marginal sums. Main-effect and intercept coordinates are sums of these coordinates and leave its kernel unchanged. The sharp uniform three/four threshold is established by Petra & Schnörr (2009, Proposition 3.2 and Section 4.6); see also Deniţiu et al. (2013, Table 2). The following block proof also applies to rectangular grids.

Theorem A.15 (Uniform basis-pursuit threshold)**.** *Let $A = \Phi^\top$ for the complete $n_1 \times n_2 \times n_3$ pair-effect grid, with every $n_i \ge 2$. Every $h \in \ker A$ satisfies*

$$\|h\|_1 \ge 8\|h\|_\infty. \tag{20}$$

Every vector w with at most three nonzero entries is the unique minimizer of $\|z\|_1$ subject to $Az = Aw$. Distinct four-sparse vectors can have the same image under A.

Proof. View h as an array. The equations $Ah = 0$ say that every line sum is zero. Fix (i_0, j_0, k_0), and in each factor partition the levels into the distinguished level and its complement. These partitions give eight disjoint blocks. Their sums form a $2 \times 2 \times 2$ array with zero line sums, so they equal alternating signs times $h_{i_0 j_0 k_0}$. Applying the triangle inequality within each block gives $\|h\|_1 \ge 8|h_{i_0 j_0 k_0}|$; choose a maximum-magnitude cell to prove Equation (20).

For $|S| \le s < 8$, the bound yields

$$\|h_S\|_1 \le \frac{s}{8}\|h\|_1, \qquad \|h_S\|_1 \le \frac{s}{8-s}\|h_{S^c}\|_1. \tag{21}$$

For $s \le 3$ and nonzero $h \in \ker A$, the ratio is strictly below one. If w is supported on S, then

$$\|w + h\|_1 \ge \|w\|_1 - \|h_S\|_1 + \|h_{S^c}\|_1 > \|w\|_1.$$

This proves uniqueness. Finally, an alternating ± 1 array on any two levels of each factor is a kernel vector. Its positive and negative parts have four nonzero entries each and the same image. □

Column deletion preserves Equation (20) by extending kernel vectors by zero. For r factors with all interactions through order $r-1$, the same partition into 2^r blocks gives $\|h\|_1 \ge 2^r\|h\|_\infty$, uniform recovery for $s < 2^{r-1}$, and an alternating-cube ambiguity at $s = 2^{r-1}$. With positive penalties c_i, weighted basis pursuit recovers every s-sparse vector whenever $s \max_i c_i < (8-s)\min_i c_i$, since Equation (21) then gives $\sum_{i\in S} c_i|h_i| < \sum_{i\notin S} c_i|h_i|$.

When observed coefficients on T are free and paid indices are $\mathcal{C}$, the relevant condition for a paid support $S \subseteq \mathcal{C}$ is the partial null-space inequality

$$\|h_S\|_1 < \|h_{\mathcal{C}\setminus S}\|_1 \quad \text{for all } h \in \ker \Phi_{T\cup\mathcal{C}}^\top \text{ with } h_{\mathcal{C}} \ne 0. \tag{22}$$

It gives uniqueness of the paid coefficients by the same triangle inequality; differences supported entirely on T do not change them. The sufficient count condition $2|S| + |T| < 8$ follows from Equation (20), because

$$\|h_{\mathcal{C}\setminus S}\|_1 - \|h_S\|_1 = \|h\|_1 - \|h_T\|_1 - 2\|h_S\|_1 \ge \left(1 - \frac{|T| + 2|S|}{8}\right)\|h\|_1.$$

Penalizing only coefficients outside a known set is the Modified-CS formulation (Vaswani & Lu, 2010). Purchase weights require the corresponding weighted partial inequality; coefficient recovery and minimum support cost are different optimization criteria.

A.21 Witness size and batched greedy selection

Proposition A.16. *There are linear designs in which every query has a unique minimal closing set of size two, but greedy purchase of two measurements at a time has approximation ratio tending to zero for total closure count.*

Proof. Take even budget $B \geq 4$ and $2B$ candidate rows given by distinct coordinate vectors. Split them into blocks U, V of size B each. For every unordered pair $\{u, v\} \subset U$, include query $e_u + e_v$. Partition V into $B/2$ disjoint pairs and, for each pair, include both $e_u + e_v$ and $e_u - e_v$. Each query has norm $\sqrt{2}$ and closes exactly when its two coordinates have been purchased.

Initially, and after any number of purchases of complete pairs from V, each remaining such pair closes two queries. Two candidates from U close one query, and any other new pair closes none. Thus two-at-a-time greedy spends its entire budget on V and closes B queries. A size-B selection containing k candidates from U closes at most $\binom{k}{2} + B - k$ queries. This convex function of $k \in [0, B]$ is maximized at an endpoint, and for $B \geq 4$ its maximum is $\binom{B}{2}$, achieved by selecting U. The greedy ratio is $B/\binom{B}{2} = 2/(B-1) \to 0$. The same choices maximize marginal closure per purchased candidate among batches of size at most two. □

Chen et al. (2019, Theorem 3.2) give batched greedy a $1 - \exp[-1/(d+1)]$ guarantee when d is the supermodular width of the complete objective. This parameter can grow when queries share candidates. For a single nonconstant Boolean closure function, its width is the largest minimal closing-set size minus one: deleting one element from a minimal closing set gives a supermodular set; conversely, a supermodular set T and its distinguished element must occur together in a minimal closing set, since deleting any element of T removes that marginal gain. In the construction above, $U \setminus \{v\}$ is a supermodular set for the total closure count: the marginal of $v \in U$ equals the number of its selected neighbors. Its width is therefore at least $B - 1$, despite every individual query having a size-two witness.

Proof of Proposition 3.2. For a scalar Gaussian target with fixed prior predictive variance v_0, the predictive information of a batch is $\frac{1}{2}\log[v_0/v(S)]$. Here $v(S)$ is its posterior target variance, plus the same independent query-noise variance for every batch if the target is a future noisy observation. Thus maximizing information minimizes risk. An order-k value oracle can evaluate every feasible batch using at most $\sum_{r=0}^{k}\binom{n}{r}$ calls; selecting a maximizer proves the fixed-budget upper bound. □

A.22 Perturbation of strict finite-noise barriers

Proposition A.17 (Stability of a strict risk barrier)**.** *Fix $\lambda > 0$, a finite feature matrix X, target q, and a finite family of feasible batches. Let $r_X(S) = q^\top(\lambda I + X_S^\top X_S)^{-1}q$. Suppose every allowed local neighbor S of S_0 satisfies $r_X(S) - r_X(S_0) \geq \Delta > 0$, and some feasible S_* satisfies $r_X(S_0) - r_X(S_*) \geq G > 0$. Put $K = \|X\|_{\mathrm{op}}$. Every perturbation with $\|E\|_F \leq \varepsilon$ preserves these strict comparisons if*

$$2\|q\|^2\lambda^{-2}(2K\varepsilon + \varepsilon^2) < \min\{\Delta, G\}. \tag{23}$$

Proof. Uniformly over all batches,

$$\|(X_S + E_S)^\top(X_S + E_S) - X_S^\top X_S\|_{\mathrm{op}} \leq 2K\varepsilon + \varepsilon^2.$$

Both regularized Gram matrices have inverse norm at most $1/\lambda$. The resolvent identity $A^{-1} - B^{-1} = A^{-1}(B - A)B^{-1}$ therefore gives

$$|r_{X+E}(S) - r_X(S)| \leq L(\varepsilon) := \|q\|^2\lambda^{-2}(2K\varepsilon + \varepsilon^2).$$

Each difference of two risks changes by at most $2L(\varepsilon)$. Condition (23) preserves the signs of all the stated comparisons simultaneously. □

For independent $E_{ij} \sim N(0, \tau^2)$ and N matrix entries, the chi-square bound $\Pr[\chi^2_N > N + 2\sqrt{Nx} + 2x] \le e^{-x}$ gives, with probability at least $1 - \zeta$,

$$\|E\|_F \le \tau \left[\sqrt{N} + \sqrt{2\log(1/\zeta)}\right].$$

Substituting this radius into Equation (23) yields a Gaussian survival guarantee. In particular, strict finite-noise barriers occupy open neighborhoods of feature matrices. The strict gap assumption is material: complete-pair batches in the paired family can tie, whereas this proposition controls strictly separated comparisons.

`verify_perturbation.py` checks the four-measurement simplex construction at $\lambda = 0.1$, signal parameters $\varepsilon^2 = 0.2$ and $\delta^2 = 0.6$. The minimum partial-exchange gap is 0.2700 and the full-exchange gain is 2.0317. All 100 sampled perturbations at the certified Frobenius radius 1.688×10^{-4} preserve the comparisons among the $\binom{8}{4} = 70$ batches. The corresponding independent Gaussian entry standard deviation 1.465×10^{-5} gives the stated survival guarantee with probability at least 0.99.

B Strategy information and witness search

Instances and strategies. An *acquisition instance* is $\mathcal{I} = (\Phi_T, \mathcal{C}, Q, c, B)$ with candidate pool $\mathcal{C} \subseteq \mathcal{X} \setminus T$, deployment query set Q, costs $c : \mathcal{C} \to \mathbb{R}_{>0}$ and budget B; no cell whose row constitutes some φ_q, $q \in Q$, lies in $\mathcal{C}$: a query is a prediction about something we cannot measure. A *strategy* π maps an instance to a purchase sequence, adaptive and randomised strategies both admitted. Reported quantities: model risk $\sigma^2 r_\lambda$, held-out error on the true responses, the closure count, and the budget to first closure.

Lower bounds restrict access to oracle *values* while allowing arbitrary subset inspection: $\pi \in \Pi_j$ reads an order-j value oracle on $\psi_{\mathcal{I}}(A) = I(y_Q; y_A \mid y_T)$ (Definition B.1, Appendix B), and Section 3 treats it and untruncated greedy under one ceiling.

Definition B.1 (The class Π_j)**.** $\pi \in \Pi_j$ if π interacts with the instance only through an *order-j value oracle*: it may query the acquisition objective $\psi_{\mathcal{I}}(A) = I(y_Q; y_A \mid y_T)$ at any A with $|A| \le j$, adaptively, with randomisation, and may post-process the answers by an arbitrary function; it may not query $\psi_{\mathcal{I}}$ at any set of size $> j$.

Theorem 3.1 covers this entire class, including arbitrary post-processing. The exact-joint greedy rule below can evaluate larger sets; its barrier concerns the sequence of purchases (Theorem A.13).

Definition B.2 (Greedy on the exact joint objective)**.** $\pi^{\text{greedy}}_{\psi}$ selects $s_t \in \arg\max_c [\psi(S_{t-1} \cup \{c\}) - \psi(S_{t-1})]$.

For $j < B$ it evaluates ψ on sets larger than j and so lies outside Π_j. Table 1 separates this predictive-information objective from BatchBALD's parameter-information objective. Section 3 treats restricted information and greedy search separately (Theorems 3.1 and A.13).

Proposition B.3 (Geometry carries no closure information)**.** *For the row-distance and D^2-sampling rules instantiated here, pairwise row geometry does not determine closure. These geometry-based rules use a different input from the value oracle of Definition B.1. On the* spread *family (the unit-simplex circuit of Section 3 with orthonormal distractors) the circuit members are the mutually farthest triple and k-center collects all m of them within its first purchases; on the* clumped *family, $s_i = \varepsilon u + a\, v_i$ with $a \le 0.2$, k-center buys no circuit member while any distractor remains, and D^2-sampling closed the circuit in 0 of 200 runs (`verify_geometric_rules.py`). Whether a geometric rule completes a circuit is decided by an alignment the closure event does not constrain.*

Proof, clumped side. The first pick maximises the norm: $\delta^2 + 1 > \varepsilon^2 + a^2$, a distractor. While only distractors are selected, a fresh distractor's min-distance^2 to the selection is 2 (orthonormal w_k), while a circuit member's is $(\varepsilon - \delta)^2 + a^2 + 1 < 2$ for $a \le 0.2$, so max-min picks the distractor; induction runs until the distractors are exhausted. For D^2-sampling the same gap bounds each round's circuit-pick probability, and collecting all m compounds it; the numerical certificate covers the stated family. The spread side is the archived run. □

B.1 THE THREE-FACTOR DESIGN IS NOT GRAPHIC

Estimability is decided by real rank. The relevant matroid is therefore $M_{\mathbb{R}}$, the vector matroid of Φ over $\mathbb{R}$.

Certificate for Proposition 5.1. In the $3\times3\times3$ design with all pair effects, $C = \{(0,1,0),(0,1,2),(0,2,0),(0,2,2),(1,1,0),(1,1,2),(1,2,0),(1,2,2)\}$ is a circuit of $M_{\mathbb{R}}$ and $D = \{(0,0,1),(0,2,2),(1,0,2),(1,1,0),(1,1,2),(1,2,1),(2,1,1),(2,2,0)\}$ is a cocircuit, with $|C \cap D| = 3$. In a binary matroid every circuit meets every cocircuit in an even number of elements (Oxley, 2011, Ch. 9), so $M_{\mathbb{R}}$ is not binary, hence neither regular nor graphic. The $3\times3\times3$ cells sit inside any $a\times b\times c$ design with $a, b, c \geq 3$ as a restriction, and non-binariness passes to the larger matroid, so the conclusion holds for every such larger grid.

The script `verify_nonbinary_exact.py` checks both memberships from their definitions using rational arithmetic: the full rank is 19, the circuit rank and all eight deletion ranks are 7, and the complement rank 18 becomes 19 after restoring any element of D. The archived certificate includes the explicit rows and all ranks.

B.2 COMPLEXITY AND THE ℓ_1 ENGINE

Witness search over $\mathbb{R}$ is a sparse-recovery problem by identity: the circuits of $M_{\mathbb{R}}$ are the minimal supports of the real kernel of $\Phi^\top$, so a minimum witness through q is a minimum-support w with $\Phi^\top w = \varphi_q$, and the purchase-weighted version weights that support by cost. That is MINIMUM WEIGHT SOLUTION TO LINEAR EQUATIONS, NP-complete over the rationals (Garey & Johnson, 1979, MP5), and its approximate form is NP-hard as well (Natarajan, 1995). Deciding whether a given purchase closes q stays polynomial, since it is a rank test, and producing *some* closing purchase is a linear solve; the cost is in finding the cheapest one.

The tractable frontier is Seymour's. On graphic designs the minimum witness is a shortest cycle through an edge (Itai & Rodeh, 1978) (our two-factor layer), on cographic designs a minimum cut, and the decomposition of regular matroids (Seymour, 1980) composes the two (with the sporadic R_{10}) into polynomial girth for the whole regular class (exposition in Kapadia, 2014). Proposition 5.1 places the three-factor estimability matroid outside that class: it is not binary, so it is not regular and not graphic, and no shortest-path-type search is complete there. The ceiling of Appendix C.3 is the measured shape of that fact, and the ℓ_1 engine below is what recovers the queries the path search misses.

The ℓ_1 engine. Between the exact search (ℓ_0, hard as above) and ridge leverage (ℓ_2, polynomial but silent before closure, Theorem A.5c) sits the remaining convex point of the family: weighted basis pursuit, $\min \sum_{x\in\mathcal{C}} c_x|w_x|$ subject to $\Phi_{T\cup\mathcal{C}}^\top w = \varphi_q$ with T-entries free. Its support is a witness set by construction, so the engine is sound, and it is one LP per query (9 ms median on Castelli). On the family instances it closes all queries recovered by the chained search, plus 11 of 72 queries that search misses, at cost no worse in 94% of the shared cases (`exp_l1_witness_probe.py`). Substituted into CIRCUIT's closure layer under the ceiling protocol it attains the pool-feasibility maximum from budget 80 on, $0.980/0.990/0.990$ against the chained $0.840/0.930/0.930$, where 0.990 is the rank-test bound on what the pool can close at all (`exp_l1_circuit.py`). The standard exactness condition applies verbatim: when the design's kernel satisfies the null-space property at the bundle size, ℓ_1 recovers the minimum-cost witness; establishing it for structured factorial designs is open.

B.3 WITNESS EXCHANGE IMPLEMENTATION

Witness exchange starts from a greedy batch and repeats three actions. First, use the reconstruction equation to find measurements that work together. Second, make room for them by removing incumbent measurements with the smallest cost in risk. Third, evaluate the proposed batch and accept it only if its exact risk improves (Algorithm 1). Theorem 3.3 establishes why proposing a complete bundle can cross barriers to smaller exchanges.

The bundle generator. The generator finds measurements that reconstruct the target together. In the end-to-end baseline named CIRCUIT, it is followed by bundle selection and ridge-leverage conditioning to spend the full purchase budget. Witness exchange uses the same generator to propose changes to an existing greedy batch.

For a query's factor pair, the generator searches the graph of co-occurring factor levels (the concurrence graph). An alternating path specifies measurements whose signed contributions reconstruct a contrast. Where one such path is insufficient, the generator combines paths using a **bridge**: a contrast supplied by already measured cells. In Figure 4, the purple measured column supplies the correction that lets the four green cells reconstruct the target. The resulting weights specify which cells enter together; the exact-risk comparison then decides whether to buy them. Appendices B.2 and C.3 give the generators and Appendix C.7 the selection procedure for overlapping witnesses.

Algorithm 1 Witness exchange: structure proposes, exact risk disposes

Require: measured T, pool $\mathcal{C}$, queries Q, costs c, budget B
1: $S \leftarrow$ greedy C-optimal batch under budget B
2: **while** some $q \in Q$ has residual under $T \cup S$ and proposals remain **do**
3: $U \leftarrow$ minimum-cost witness bundle closing q given $T \cup S$ (shortest alternating path, chained through a bridge when q's factor pair admits no witness; design-only, no risk calls)
4: $D \leftarrow$ the $|U|$ cells of S with smallest removal marginals
5: **if** $R\big((S \setminus D) \cup U\big) < R(S)$ **then** $S \leftarrow (S \setminus D) \cup U$
6: **end while**
7: polish S by single-cell exchange with the remaining evaluation budget
Ensure: batch S with $c(S) \le B$

Proposals from continuous features. For an arbitrary candidate matrix X, query matrix Q, and measured precision $H = G_T + \lambda I$, put $K = I + XH^{-1}X^\top$ and $Z = XH^{-1}Q^\top$. The reconstruction cost of Theorem A.3, after eliminating the measured coefficients, is a quadratic in the candidate coefficients W. Penalising its nonzero rows by $\tau \sum_i \log(1 + \|W_i\|^2/\epsilon)$ gives the update

$$W_{\text{new}} = \left(K + \operatorname{diag}_i \frac{\tau}{\epsilon + \|W_i\|^2} \right)^{-1} Z.$$

This update decreases the penalised objective at fixed τ, ϵ. We rank rows by their coefficient norms, propose the top-budget support, and retain the best exact-risk batch after single-measurement refinement. This generator uses the feature and query matrices directly; the recovery procedure in Theorem 4.3 supplies the exact solution for its paired design family. Appendix A.17 gives the derivation and fixed search schedule.

C EXPERIMENTAL PROTOCOLS AND RESULTS

C.1 DOMAIN SETTINGS

Synthetic: $6\times6\times3$ design, $|Q| = 12$. Castelli (Castelli et al., 2012): the $14\times14\times7$ subgrid, response `icsd_min_abs_charge_residual`, three-way structure with marginal weights $0.242/0.242/0.205$, 10% initial coverage, $|Q| = 20$, pool 400. Resistor network: 7×7 grid, 49 nodes and 84 edges, rows are incidence vectors $e_u - e_v$.

C.2 WORKED EXAMPLE FOR FIGURE 4

The illustration uses a $4 \times 2 \times 3$ grid with feature vector

$$\varphi_{abx} = \big(1, e_a, e_b, e_x, e_a \otimes e_b, e_a \otimes e_x, e_b \otimes e_x\big),$$

where e_a, e_b, e_x are one-hot vectors of lengths 4, 2, and 3. The coefficients have prior $\beta \sim N(0, I)$ and observations have independent noise variance $\sigma^2 = 0.01$. Write $f^a_{ij} = \mathbb{E}[y_{ab_i x_j} \mid \beta]$ and take $\theta_q = f^\star_{11} - f^\star_{22}$, where $\star$ denotes the fourth family. The measured set is

$$T = \{(A_1, b_i, x_j) : i = 1, 2,\ j = 1, 2, 3\} \cup \{(A_2, b_1, x_1), (A_2, b_1, x_3), (A_2, b_2, x_2), (A_3, b_1, x_2), (A_3, b_2, x_2)\}.$$

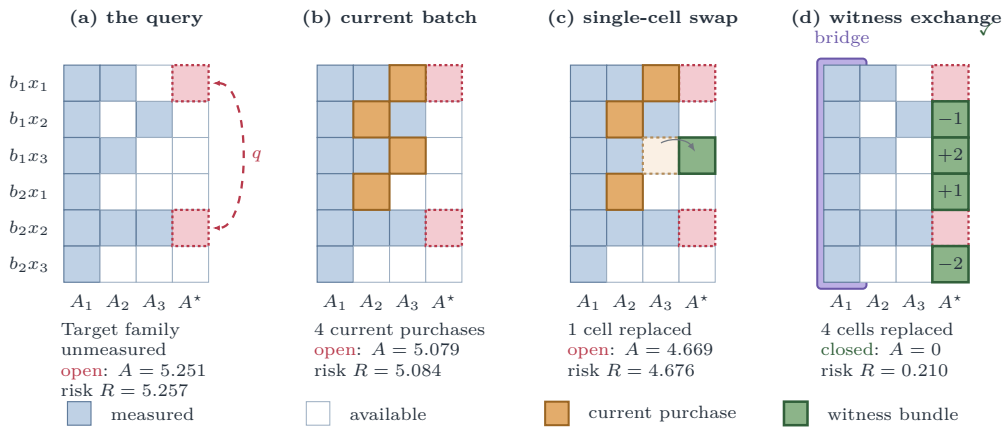

Figure 4: **Reconstructing a query by exchanging a bundle.** (a) The target is $f(A^\star, b_1x_1) - f(A^\star, b_2x_2)$; its two dashed cells cannot be purchased. (b–d) At budget four, one replacement reduces risk but leaves the query open; replacing the batch with the four green cells closes it. Their displayed coefficients, together with a contrast in the purple measured column (the bridge), reconstruct the target. A is the squared unrepresented query component and R the posterior variance. This illustrative pairwise model uses $\rho^2 = 1$ and $\sigma^2 = 0.01$; Appendix C.2 gives the reconstruction and exhaustive verification.

The 11 remaining cells other than the two query cells are purchasable, at unit cost and budget four. The illustrated current batch is $S_0 = \{(A_2, b_1, x_2), (A_2, b_2, x_1), (A_3, b_1, x_1), (A_3, b_1, x_3)\}$. Panel (c) replaces (A_3, b_1, x_3) with $(A^\star, b_1, x_3)$. Panel (d) replaces S_0 with $U = \{(A^\star, b_1, x_2), (A^\star, b_1, x_3), (A^\star, b_2, x_1), (A^\star, b_2, x_3)\}$. The mean-response reconstruction is

$$\theta_q = 2f^\star_{13} - 2f^\star_{23} + f^\star_{21} - f^\star_{12} + \Delta_1,$$
$$\Delta_1 = f^1_{11} + f^1_{12} - 2f^1_{13} - f^1_{21} - f^1_{22} + 2f^1_{23}.$$

The contrast Δ_1 uses the measured family A_1 and cancels its family-specific main and pairwise effects. The remaining shared $B \times X$ contrast supplies exactly the correction required by the four new cells. Applying these coefficients to noisy observations yields an unbiased estimator of θ_q.

For each panel, $A = \|P_\perp^{T\cup S}\varphi_q\|^2$ and $R = 0.01\varphi_q^\top(G_{T\cup S} + 0.01I)^{-1}\varphi_q$. Their respective values are $(5.251, 5.257)$ before purchase, $(5.079, 5.084)$ for S_0, $(4.669, 4.676)$ after the single-cell replacement, and $(0, 0.210)$ for U. The script `figures/teaching_example/compute_example.py` checks the reconstruction coordinate by coordinate and enumerates all purchase sets of size at most four. The set U is the unique closing purchase at this budget and attains the lowest risk among all 330 four-cell batches. Exact greedy also selects U in this illustration; Theorem 3.3 supplies the separate construction establishing strict local-search barriers.

C.3 The two knobs behind the ceiling

Bundle construction and selection. Each open query's bundle is a shortest alternating path on the concurrence graph, chained through a bridge where the query's own factor pair admits no witness, forced by structure, since the design matroid is not graphic (Appendix B.1). Deciding closure is a rank test and producing some closing purchase a linear solve; finding the *cheapest* one is NP-hard (Appendix B.2), and weighted basis pursuit closes everything the pool can close ($0.930 \to 0.990$). Where several bundles compete under the budget, variable elimination maximizes the number of completed witnesses in a fixed menu. Its cost depends exponentially on the width of the witness-sharing graph; the measured elimination widths are 4 to 6 on these HOIP pools (Proposition C.2 and Appendix C.7). Exact-risk evaluation then accounts for the reconstruction error and observation noise.

Search and candidate-pool effects. Appendix B.2 predicts a ceiling for any path-type search on this design: the matroid is not regular, so shortest-path witnesses cannot be complete. The single-colour search saturates at 0.670 on a pool of 400, budget 80 through 320 unchanged. Two mechanisms could produce that, and varying them independently separates them: chained cross-colour witnesses raise the ceiling to 0.930, while *enlarging the pool to* 1200 *raises the single-colour arm on its own to* 1.000, so a restricted pool and an incomplete search each contribute. The ℓ_1 witness engine closes the search half completely: at pool 400 it reaches $0.980/0.990/0.990$ at budgets $80/160/320$, where 0.990 is the rank-test bound on what that pool can close at all (Appendix B.2; `exp_l1_circuit.py`).

C.4 The network domain

The design matroid is graphic, so witness search reduces to a shortest path and is *complete*: this domain carries no ceiling of the kind Appendix C.3 diagnoses on the chemistry grid. *Closure* (7×7 grid, 12 seeds, ±1.96 se): CIRCUIT attains closure fraction 0.342 ± 0.056 at budget 4 against EPIG's 0.275 ± 0.055, 0.758 ± 0.066 vs. 0.483 ± 0.107 at budget 12, and 1.000 ± 0.000 vs. 0.925 ± 0.060 at budget 30. *Deployment error* (`exp_network_state.py`, 10 seeds per task, budget 12): node potentials solved from random current injections, instrument reads $y_e = v_u - v_w + \varepsilon$, queries drawn from node pairs at Manhattan distance ≤ 3 (near) or $\geq L-1$ (far). The transfer bias is the root mean squared residual norm of the query directions under the instrumented design and the price is $\sigma\sqrt{\|w\|^2}$ at the median witness path length, both design-computable.

C.5 Representation and query closure

Proposition C.1 (Representation scoping)**.** *Let a factor be carried by a feature map ϕ and let L be the set of its measured levels. If the held-out level $\ell^\star$ satisfies $\phi(\ell^\star) = \sum_{\ell\in L} c_\ell\,\phi(\ell)$ and the same relation holds blockwise in every interaction the design carries, then $\varphi_q \in \mathrm{row}(\Phi_T)$ for every contrast q at $\ell^\star$: the query is closed with no purchase at all. One-hot encoding never satisfies this, since $\phi(\ell^\star)$ is orthogonal to every measured level; a k-dimensional continuous encoding satisfies it generically once k falls below $|L|$.*

C.6 Oracles and statistical standard

Table 2: **Held-out RMSE and closure fraction at budget** 24 (HOIP volume, never-measured A-site family, 10 shared seeds). CIRCUIT and greedy C-optimal both complete the query bundles at this budget; point-scoring rules attain closure at most 0.01 and incur three to six times the error. A zero-budget imputation baseline attains 0.769 ± 0.096. Ratios are paired per-seed, baseline over CIRCUIT.

method	RMSE	closure	paired ratio
CIRCUIT	0.051 ± 0.010	1.00	–
greedy C-optimal	0.064 ± 0.018	1.00	1.105 ($p = 0.275$)
EPIG	0.183 ± 0.039	0.00	3.625 ($p = 0.002$)
ridge leverage	0.205 ± 0.066	0.01	4.094
CoreSet	0.291 ± 0.029	0.00	5.609

Six oracles are defined by their feasible sets and available information. *Full-pool exact optimum*: the minimiser over all budgeted subsets of the complete candidate pool, obtained by branch and bound with a certified gap. *Restricted-pool exact optimum*: the minimiser over a pool built from the union of the arms' own selections plus a seeded random completion. This oracle is *biased in favour of the arms that generated the pool*, because their selected cells enter the pool construction. *Certified near-optimum*: an incumbent carrying an explicit optimality gap, quoted with the gap. *Design oracle*: chooses between arms by the expected risk of the task, using the design only. *Hindsight oracle*: chooses by the realised held-out error, providing a retrospective upper bound on the performance of pre-measurement rules. *Robust-support oracle*: the minimiser of worst-case regret over a stated κ interval.

The statistical unit is the task, meaning a (dataset, response, depth, budget) cell. Ratios are formed per task and then aggregated; intervals come from a hierarchical bootstrap that resamples tasks and then seeds within tasks, giving tasks equal weight. Paired tests are used wherever the arms share seeds, which they do throughout. Event counts accompany the diagnostic estimates.

C.7 Elimination on the witness-sharing graph

Proposition C.2 (Exact completion count for a fixed witness menu)**.** *Let $W_1, \dots, W_M$ be specified witness supports on n purchasable cells, with unit costs and integer budget $B \geq 1$. Join cells that occur in a common W_ℓ. Given an elimination order of width τ, variable elimination computes*

$$\max_{|S| \leq B} \sum_{\ell=1}^{M} \mathbf{1}\{W_\ell \subseteq S\}$$

in $O((n+M)B^2 2^{\tau+1})$ time.

Proof. Give each cell a binary purchase variable and a unit-cost factor. Each witness contributes the indicator that all variables in its scope are one. Every scope is a clique and hence lies in an elimination bag. Store the best score for each bag assignment and budget. Combine scores by budget convolution and eliminate a cell by maximizing over its two assignments. Each table has at most $2^{\tau+1}(B+1)$ entries; a merge costs $O(2^{\tau+1}B^2)$. There are $O(n+M)$ factors and merges. Convolving the final component tables enforces the common budget. □

With one listed witness per query, this score counts queries certified by those witnesses. Multiple witnesses for one query require a single OR factor over their completion indicators to count that query once; the width is then measured on those factor scopes. The first stage of Theorem D.3 instead minimizes $A(S) = \sum_{q \in Q} \|P_\perp^{(T \cup S)} \varphi_q\|^2$, which uses the magnitudes of the remaining residuals. The completion-count program provides an exact optimizer for its stated discrete objective.

On real HOIP instances the min-fill width of that graph is 4 to 6, so the largest bag has 5 to 7 cells: elimination returns the menu optimum in 5–8 ms. The five pools contain between $50{,}388$ and $86{,}493{,}225$ subsets of size 12. Agreement with exhaustive enumeration is $60/60$ on random instances small enough to enumerate.

The optimisation uses a fixed collection of witness columns. At $B = 6$, its predicted closure counts agree with the returned supports on $5/5$ instances. Adding an alternative witness closes one additional query on $2/5$ instances, identifying column generation as the complementary step to exact selection.

C.8 Family variation and parameter intervals

The structured-shift model (Appendix A.15) uses the measured families to set a reference scale for the held-out family. On the fully covered HOIP grid, we compare the metal-by-halide interaction rate of each family with that of the others, in fitted-residual units (`exp_structured_shift_test.py`). Across nine responses and sixteen families (144 pairs), the ratio has median 0.74, spans 0.11–4.20, and falls within a factor of two of one on 69% of pairs; per-response medians range from 0.64 to 0.92. This measured variation sets the parameter intervals used to evaluate batch stability in Appendices C.10 and D.2.

C.9 Certification against the full pool

On this grid $\hat{\kappa}$ is 2.31 from the within-cell replicate spread and 10.7 from the full-coefficient estimator, so the operating range is $\lambda \in [0.0087, 0.187]$.

Ratios to the certified incumbent over the complete 200-candidate pool; the two rows at $\lambda = 0.187$ are the same tasks searched harder, and the gap column bounds how far the incumbent may still sit above the true optimum.

λ	search	greedy med.	greedy max	$> 2\times$	closure-seeking	gap
0.187	400 nodes	1.0045	1.5	0%	1.187	28.2%
0.187	6200 nodes	1.0097	1.5	0%	1.187	22.3%
0.0087	400 nodes	1.0776	9.6	30%	1.345	89.5%

Each ratio is measured against the certified incumbent, so it is a lower estimate of the excess against the true optimum and the gap column bounds the understatement. The worst task costs between 1.5 and 9.6 times the incumbent depending on which $\hat{\kappa}$ is believed, and held-out error moves the same way (0.1202 against greedy's 0.1310 at $\lambda = 0.187$).

Increasing the search from 400 to 6200 nodes reduces the reported gap from 28.2% to 22.3% at $\lambda = 0.187$; the table retains both bounds. Below the operating range, the ratio reaches 748.6 at $\lambda = 10^{-4}$, following the $1/\lambda$ scaling of the mechanism.

C.10 Batch stability and interval widening

Relative regret of 15329% corresponds to 154.29 times the oracle model risk: the unclosed query retains an $A(S)\kappa^2$ term while a closing design has bounded risk (Theorem D.2).

Table 3: Worst-case relative model-risk regret over nine values from $\hat{\kappa}/3$ to $3\hat{\kappa}$, with reference $\hat{\kappa} = 10.7$. HOIP volume, unseen A-family, budget 10, enumeration pool 18. n counts the retained budget-feasible supports for each policy.

policy	n	worst-case regret (median)	held-out RMSE
robust support	10	0.0%	0.048
greedy C-optimal	10	0.1%	0.047
closure-seeking search	10	0.0%	0.042
boundary arm selection	8	15329%	0.072
static 50/50 split	10	15727%	0.075

For each support, relative regret is $\max_{\kappa \in \mathcal{K}}[R_\kappa(S)/R^\star_\kappa - 1]$, where $\mathcal{K}$ is the nine-point grid and $R^\star_\kappa$ is the minimum over all $\binom{18}{10}$ supports. The reference value 10.7 comes from the full-coefficient estimate on HOIP volume. For example, on seed 0 the boundary-selected support has model risk 925.157 in noise-variance units at $\kappa = 32.1$, against 6.145 for the enumerated optimum. Its residual query mass is 0.894; the optimum closes the queries. Thus the $A(S)\kappa^2$ contribution explains the large ratio. Boundary-selected supports fall outside the enumeration pool on two seeds; the other eight yield the reported median. The complete numerator, denominator and support records are included as `regret_details.json`.

Worst-case objective regret as the κ interval of Appendix D.2 widens; the answer of Table 3 is unchanged across a factor of 81 in κ.

interval	robust support	boundary arm selection	greedy C-optimal
$\hat{\kappa}/3 \ldots 3\hat{\kappa}$	0.0%	14955%	0.7%
$\hat{\kappa}/4.2 \ldots 4.2\hat{\kappa}$	0.0%	29277%	0.9%
$\hat{\kappa}/9 \ldots 9\hat{\kappa}$	0.5%	134304%	1.5%

C.11 Matched exchange comparison

Figure 5 shows how a complete exchange changes the batch. Random multi-cell exchange with the same evaluation budget accepts *zero* moves across all twenty tasks. Witness proposals identify cells that jointly reconstruct a query. The thresholds of Theorem 3.3 are visible in the mechanism case: on the task where the witness arm reaches 100% of the certified gap (its final risk 70.9 against the certified lower bound 70.6, the certified optimum itself) single-cell search recovers 0.2%; on another task single-cell search finds nothing at either ridge scale while k=2 witness exchanges close 27 and 39%. Each accepted exchange spends tens of evaluations (removal marginals plus one proposal) against the controls' 1500.

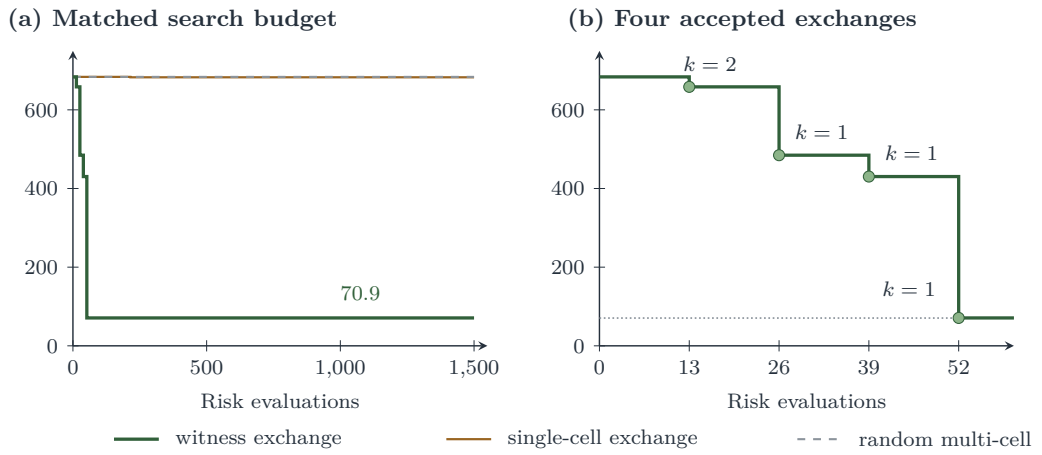

Figure 5: **Four accepted exchanges turn reconstruction into risk reduction.** HOIP seed 9, $\lambda = 0.0087$, purchase budget 12. (a) All search arms use 1500 exact-risk evaluations from the same greedy batch. (b) The first 52 evaluations contain the four accepted witness moves: one $k = 2$ exchange followed by three $k = 1$ exchanges. The batch risk falls from 683.8 to 70.9; the certified lower bound is 70.6. Curves are step functions from the archived evaluation counts and risks. Aggregate results over all twenty tasks are in Table 4.

When the exchange pays. The benefit concentrates on tasks where the greedy batch leaves valuable query bundles incomplete. In full-pool searches, the largest greedy-to- incumbent risk ratio reaches 9.6 at $\lambda = 0.0087$ (Appendix C.9). Exact-risk acceptance selects improving proposals from that batch.

The budget-24 HOIP comparison evaluates the value of completing these bundles (Table 2). Both CIRCUIT and greedy C-optimal attain full closure; the point-scoring baselines attain closure at most 0.01 and incur three to six times the error.

Four accepted moves. On the 100% task the exchange executes four moves: a k=2 exchange buying two cells of one candidate family, then three k=1 consolidations, taking the squared residual norm $5.38 \to 0$, closure $0.58 \to 1.00$, risk $683.8 \to 70.9$, and held-out error $0.099 \to 0.053$.

Table 4: **Controlled comparison at matched evaluation budgets.** Median over ten tasks per ridge scale; gap closed is the share of the distance from the greedy value to the certified family-menu lower bound. All arms start from the same greedy batch with 1500 risk evaluations. Witness exchange closes the largest gap share and gives the lowest median held-out RMSE at both scales.

λ	metric	greedy	1-swap	random-k	witness
0.0087	certified gap closed	–	0.1%	0.0%	**16.8%**
0.0087	held-out RMSE	0.114	0.117	0.114	**0.081**
0.187	certified gap closed	–	0.0%	0.0%	**10.5%**
0.187	held-out RMSE	0.099	0.110	0.099	**0.086**

C.12 REACTION-YIELD PREDICTION WITH FROZEN NEURAL FEATURES

Data and tasks. We use the processed Buchwald–Hartwig benchmark of Ahneman et al. (2018); Schwaller et al. (2021): sheet `FullCV_01` of the authors' input workbook contains 3955 unique recipes spanning 15 aryl halides, 3 bases, 4 ligands and 22 additives. Each additive defines a background. Nineteen backgrounds contain all 180 recipes; the other three contain 178, 178 and 179. All reported yields are retained and the five missing recipes are unavailable for purchase. The ten `FullCV` sheets are permutations of the same observations. We use one sheet. The supplement records the source revision, workbook hash, factor identities and sheet-identity check.

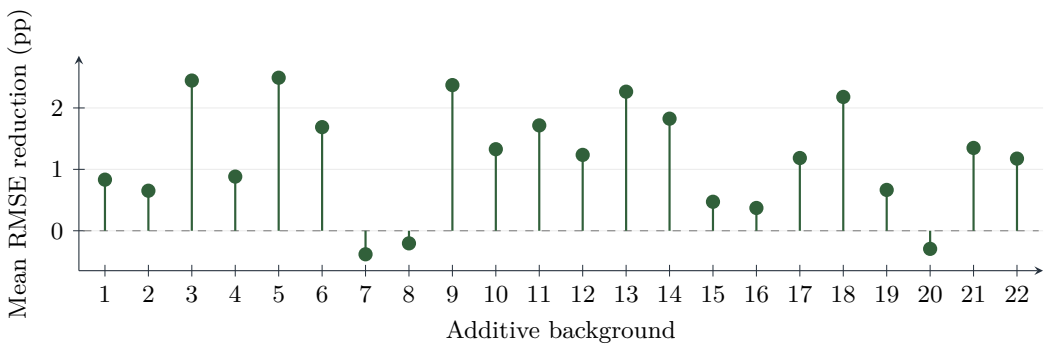

Figure 6: **Neural-feature exchanges improve reaction-yield prediction across additive backgrounds.** Full frozen ChemBERTa features, $\lambda = 0.0087$. Each point is the mean of ten paired task differences in yield-contrast RMSE: single-cell search minus witness exchange. Positive values favour witness exchange; both arms use 1500 risk evaluations and 12 purchases. All 22 backgrounds are shown in fixed identifier order.

For each background, seeds 500–509 select an unmeasured aryl-halide family and initial observations with probability 0.10 outside that family. The task predicts up to 12 ligand contrasts within the held-out family; query endpoints are excluded from the purchase pool. The candidate cap is 200, and the purchase budget is 12. All 220 generated tasks meet the protocol's eligibility criteria. Task averages summarise repeated splits of the same measured table; the 22 backgrounds are contexts within one dataset.

Representation and reconstruction. We apply frozen ChemBERTa to the aryl-halide SMILES, using masked mean pooling of the final hidden layer. The checkpoint revision and cached embeddings are included in the supplement. Centring the 15 feature rows gives rank 14. The full representation retains every nonzero singular direction, and the eight-dimensional representation retains the first eight, accounting for 99.794% of the centred squared singular values. Each representation is normalised to unit mean squared row norm. Bases and ligands remain categorical; the design includes all pairwise interactions.

Let x_c be an original categorical design row and $\widetilde{x}_c = x_c M$ its feature representation for the fixed factorwise map M. Any original reconstruction identity transfers directly:

$$x_q = \sum_{c \in W} a_c x_c \quad \Longrightarrow \quad \widetilde{x}_q = \sum_{c \in W} a_c \widetilde{x}_c.$$

The categorical identities therefore generate valid proposals in the neural feature space. Proposal acceptance evaluates the exact Gaussian risk using the transformed rows. Across all settings, the maximum squared relative reconstruction residual is 6.6×10^{-30}, and primal and dual risk calculations agree to 4.1×10^{-11} absolute error.

Development and fixed evaluation settings. The preceding HOIP development experiment tested categorical features and frozen ChemBERTa representations of dimensions 4, 8 and full rank 15, at both acquisition scales. It motivated transferring the eight-dimensional representation with $\lambda = 0.0087$ as the primary reaction setting. Before running reaction acquisitions, the protocol also specified full features, categorical features and $\lambda = 0.187$ as comparisons. Table 5 reports all six settings. The full-feature comparison at $\lambda = 0.0087$ is the result presented in Appendix C.12.

What feature compression changes. The eight-dimensional primary setting has mean paired RMSE reduction -0.068 percentage points. Compression preserves the reconstruction identity above: it changes the represented directions and their risk weights, while every original witness remains valid. At $\lambda = 0.0087$, both representations generate proposals on $218/220$ tasks, but full features accept multi-cell moves on 92 tasks versus 49 for eight components. Their mean relative unrepresented query masses at the greedy batch are 0.0111 and 0.0103, respectively. Thus high re-

tained feature energy does not imply identical exchange decisions. The observed contrast motivates evaluating proposals in the chosen representation's exact risk.

A projection audit checks all 880 neural task settings. The full design maps linearly to the eight-dimensional design with maximum entrywise error 1.6×10^{-14}; projection preserves closure on every fixed batch tested. Witness residuals remain below 6.6×10^{-30}. The code and task-level geometry are in `reaction_exchange/representation_diagnostic/`.

All arms share the greedy C-optimal starting batch within each task and representation. Single-cell search, random two-to-four-cell exchange and witness exchange each receive 1500 risk evaluations. The witness arm first evaluates reconstruction-guided proposals, then uses remaining evaluations for single-cell search. Its structural checkpoint is archived separately. Acquisition uses design rows and queries. Prediction fits a Gaussian linear readout with ridge 0.001 to the initial and purchased observations; response centring and scaling use the initial observations. Reported error is RMSE of measured yield contrasts in yield percentage points.

Table 5: **Complete reaction-yield comparison.** Each row contains 220 tasks. G, S, R and W denote greedy, single-cell, random multi-cell and witness exchange. The four RMSE columns are medians in yield percentage points. $\overline{\Delta}$ is the mean paired RMSE of S minus W; W/T/L counts tasks with lower/equal/higher error for W. Backgrounds counts positive mean paired reductions among the 22 backgrounds. Median paired reduction is zero in every setting.

Features	λ	G	S	R	W	$\overline{\Delta}$	W/T/L	Backgrounds
8 dims	0.0087	19.22	18.48	19.13	18.96	−0.068	29/152/39	6/22
8 dims	0.187	20.01	20.71	19.66	19.60	1.817	21/179/20	20/22
Full	0.0087	21.60	19.46	19.00	18.28	1.193	59/127/34	19/22
Full	0.187	19.96	20.07	19.71	20.04	0.118	16/179/25	15/22
Categorical	0.0087	16.25	16.51	16.32	17.05	−0.335	54/105/61	11/22
Categorical	0.187	16.19	15.75	16.00	16.19	−0.365	61/93/66	11/22

Complete records and reproduction. The supplement contains all 1320 task settings, selected cells, trajectories, structural checkpoints, per-task errors and per-background summaries. It also includes the HOIP development grid, both response variables, frozen embeddings and protocols. Identical acquisition problems across backgrounds reuse search outputs with the same random stream, then evaluate each background's own measured yields. This gives 240 distinct design settings and 1,080,000 physically executed search evaluations, corresponding to 5,940,000 evaluations under the per-task budget accounting. The command `python run.py reactions` reproduces the full reaction grid from the included workbook and embeddings.

C.13 Search efficiency and convergence

All controls use the twenty original HOIP tasks, purchase budget 12, and the same greedy C-optimal batch. Complete one/two-cell descent scores every one-cell and two-cell replacement in the full candidate pool, takes the best improving batch, and repeats until a full sweep finds no improvement. Its count includes the final sweep. Simulated annealing uses symmetric one-to-four-cell proposals, initial temperatures 0.01, 0.1, and 1 times starting risk, and checkpoints at 1500 and 15000 evaluations. Temperatures cool by a factor of 1000 per 1500 evaluations and restart at that cadence; all settings are reported. The ℓ_1 control replaces the proposal generator with weighted basis pursuit and retains the same drop ranking, exact-risk acceptance and remaining-budget single-cell search.

Efficiency and convergence measure different parts of the search. At 1500 evaluations, witness exchange attains mean normalised risks 0.729 and 0.931 at the two noise scales (Figure 7). Complete descent attains 0.628 and 0.864 after median counts 2,384,073 and 2,399,535. It matches witness exchange on ten tasks and finds a lower-risk batch on ten. The pooled median paired risk ratio is 1.033; the mean is 1.142. Witness exchange recovers 66.8% of complete descent's mean task-normalised risk reduction with 0.063% of its median evaluation count. At 15000 evaluations, the 0.1-temperature annealing arm attains mean normalised risk 0.671 at $\lambda = 0.0087$; all other endpoints are shown in the figure. This comparison measures candidate-risk evaluations. Complete descent uses

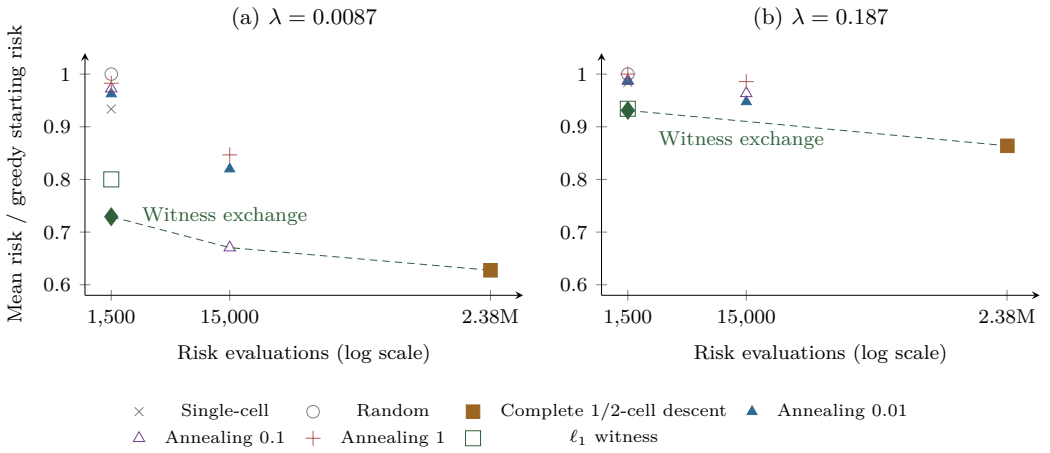

Figure 7: **Reconstruction identifies improving batches with fewer risk evaluations.** Each panel averages ten task risks after dividing by their common greedy starting risk. Horizontal positions are median evaluation counts; all annealing temperatures and both budgets are shown. Dashed segments join the observed Pareto endpoints. Counts include vectorised candidate scores in complete descent.

vectorised closed-form updates; implementation timings and the common-oracle check are recorded separately.

The full protocol, all 240 task/arm rows, accepted moves, final neighbourhood sweeps and timing records are supplied in `experiments/search_controls/`. The original eighty control rows agree with the preceding exchange archive, and primal risk checks agree to 2.5×10^{-11}. Figure data are regenerated by `figures/prepare_search_efficiency.py`.

C.14 Automatic support discovery and hidden-pair recovery

Inputs and schedule. The coefficient-based generator receives X, Q, H, an incumbent batch and a cardinality budget b. Risks in this study use $\sum_q r_\lambda^{T \cup S}(q)$. Four starts use W_*, the fitted incumbent coefficients plus $0.05W_*$, and two rowwise log-normal perturbations of W_* with log standard deviation 0.2. The penalty multipliers are $0.01, 0.1, 1, 10$ times $(R_{\text{incumbent}} - R_{\text{all}})/b$; ϵ follows $1, 0.1, 0.01, 0.001$ times the median nonzero row energy of W_*. Each phase takes at most 60 updates, stopping at relative coefficient change 10^{-7}. We record the top-b support every five updates and at each phase end, rank distinct supports by exact risk, and refine the best eight by complete single-measurement descent. The incumbent remains among the available outputs.

Original HOIP tasks. All twenty original tasks retain their full candidate pools, observations, queries and purchase budget 12: ten seeds at each of $\lambda = 0.187$ and 0.0087. For each task we take the best archived feasible batch across the existing searches and complete one/two-measurement descent. The automatic generator directly improves two of these local optima before single-measurement refinement (Table 6). The other tasks retain the incumbent. Each improved starting batch was independently rechecked by conditioned Cholesky evaluation of its full one/two-measurement neighbourhood, with relative tolerance 10^{-8}.

Table 6: **Direct proposals beyond completed local search.** Both are original HOIP seed 7 tasks, budget 12; model risk is evaluated on the unchanged full design. The replacement count describes the returned proposal. At $\lambda = 0.0087$, excluding all one/two-measurement improvements also certifies that three is the minimum improving exchange size.

λ	Initial risk	Proposal risk	Reduction	Replacements
0.187	63.878918	63.261014	0.967%	5
0.0087	396.754335	343.757586	13.358%	3

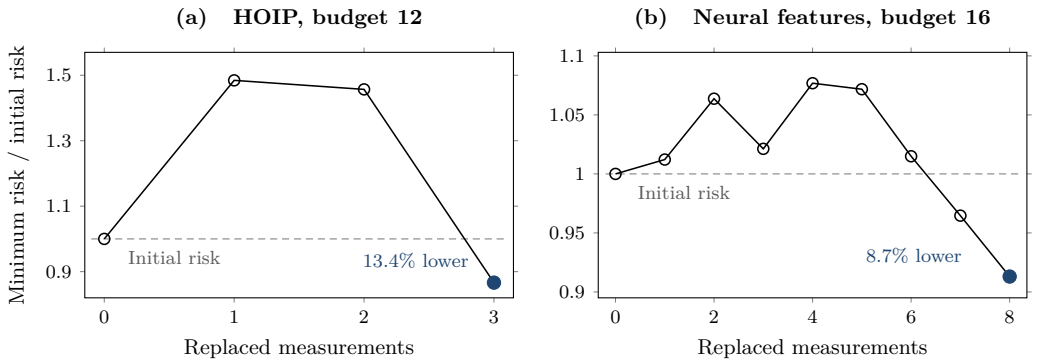

Figure 8: **Automatically discovered batches cross measured exchange barriers.** Minimum model risk at each replacement count, divided by the initial risk. Both initial batches have passed complete one/two-measurement descent. Enumeration retains measurements shared by the initial and final batches: (a) 20 combinations, HOIP seed 7, $\lambda = 0.0087$; (b) 12,870 combinations, Diabetes seed 9200, $\lambda = 10^{-4}$, one query, empirical neural-gradient features. The first improvement within each proposal union requires three and seven replacements.

Continuous neural features. The fixed grid contains Friedman, Diabetes, California Housing and Digits; three seeds (9200–9202), three ridge scales $(10^{-4}, 0.01, 0.1)$ and query counts $1, 4, 64$, for 108 configurations. Each has 128 candidates and budget 16. Empirical network gradients are projected by CountSketch into 128 feature dimensions; classification uses the trace gradient-kernel objective. Each incumbent is the best output of the seven archived optimisation arms and joint rounding, followed by complete one/two-measurement descent. Proposal generation and single-point refinement improve model risk in 15 configurations, by up to 8.697%. The raw supports and subsequent refinement are stored separately in the per-task records. The displayed Diabetes example is the largest reduction: $33.948535 \to 30.996178$, with eight replacements. It shares eight measurements with the initial batch. Enumerating all $\binom{16}{8} = 12{,}870$ combinations of the remaining rows gives the curve in Figure 8; the seven-replacement threshold refers to this union with shared rows fixed.

The complete-pair barrier configurations. We test paths and cycles, $b \in \{4, 6, 8, 16, 32, 64, 128\}$, $a \in \{10, 100\}$, and residual $\theta \in \{0, 0.1\}$: 56 fixed configurations. Candidate order is permuted, coordinates are densely rotated, and individual row signs are independently flipped. All hidden pairs are correctly recovered. The 24 configurations with $b \leq 8$ also enumerate every budget-feasible batch. The remaining 32 use the analytical lower bound and constructive upper bound of Proposition 4.4 and its proof. All zero-residual outputs attain the global optimum; at residual 0.1 the certified factor is at most 1.098912.

Gram recovery and exchange scaling. The fixed 42-case suite comprises 24 unweighted cycles with $b \in \{6, 8\}$, $\delta \in \{0, 0.2\}$, and $t \in \{0.03, 0.17, 0.73, 2.3, 7.1, 23\}$; nine weighted cycles at $b = 8$ with $\delta \in \{0.2, 1, 1.4\}$ and $t \in \{0.17, 2.3, 23\}$; six dense signed Gram matrices; and three additional cases covering spectral coincidences and ambiguous original pair labels. All use $a = 1$. The weighted cycle edges are $(0.7, 1.1, 0.9, 1.4, 0.8, 1.2, 1.0, 1.3)$. The dense matrices have nuisance eigenvalues δ and $\delta + \mathrm{geomspace}(0.3, 10, 7)$, with random orthogonal eigenvectors perpendicular to $\mathbf{1}$, followed by independent nuisance sign changes. The seed is 120926. All 42 outputs match exhaustive budget-b optimization.

The general spectral exchange bounds are checked at every complete-pair start with $\delta < a$: 1,646 starts in total, including 1,296 cycle starts. Direct feature-space inverses and the conditional risk formula differ by at most 1.23×10^{-15}. The pair-identification perturbation bound passes 32 checks; the interval inverse formula passes 20 checks. For budget 128, $a = 1$, $\delta = 0.2$, the nuisance strengths are $0.03, 0.1, 0.3, 1, 3, 10, 30, 100, 300, 1000, 3000$. Starts are alternating pairs, the first half of the pairs, and a random half with seed 120927. The interval routine returns 33 improving exchanges; its complete 616 candidate-risk evaluations are archived.

Reproduction. The supplement includes the solver, all 128 design-search inputs, complete per-task records and the enumerated profiles for Figure 8. Replay commands for the HOIP improvement and all 56 recovery configurations are given in its `continuous_design/README.md`. The general-Gram and cycle checks are in `research/continuous_design/general_gram/`; `verify.py` regenerates the 42-configuration record, and `scaling.py` regenerates the budget-128 constructions. The Latin reduction, symmetric exchange formula, and perturbation checks are in `research/continuous_design/theory_extensions/`. `reproducibility/verification.md` lists the replay commands and the archived outputs for each result.

C.15 HOIP JOINT-GEOMETRY INTERVENTION

We reuse the original seed 7, $\lambda = 0.0087$ task, its complete one/two-exchange local optimum, and the automatic three-measurement proposal. The nine shared purchases form S_0. With original observations T, set $H = \lambda I + \Phi_{T\cup S_0}^\top \Phi_{T\cup S_0}$ and whiten candidate and query rows as $\widetilde{X} = XH^{-1/2}$ and $\widetilde{Q} = QH^{-1/2}$. Let P project onto $\mathrm{row}(\widetilde{Q})$, and set $U = \widetilde{X}P$, $V = \widetilde{X}(I-P)$, $B = \widetilde{X}\widetilde{Q}^\top$. For the three remaining purchases S, the intervened risk is

$$R_\alpha(S) = \|\widetilde{Q}\|_F^2 - \mathrm{tr}\left[B_S^\top (I + U_S U_S^\top + N_{\alpha,SS})^{-1} B_S\right], \qquad N_\alpha = \alpha VV^\top + (1-\alpha)\,\mathrm{diag}(VV^\top).$$

Every N_α is positive semidefinite. An explicit feature realization concatenates U, $\sqrt{\alpha}V$, and $\sqrt{1-\alpha}\,\mathrm{diag}(\|v_i\|)$, with queries $(\widetilde{Q}, 0, 0)$. Candidate variances and query covariances stay fixed, so every singleton risk reduction and joint predictive-information score stays fixed. The intervention changes the observation covariances between distinct candidates. It tests the role of those covariances in batch value.

Table 7: **The actual measurements in the HOIP exchange.** Composition labels follow the data. Information scores condition on $T \cup S_0$; ranks are among the 153 remaining candidates.

Purchase	Composition	Information gain	Rank
Removed	HyA–Ge–Cl	0.114576	20
Removed	HyA–Ge–F	0.157694	10
Removed	HyA–Pb–Br	0.407132	1
Added	AcA–Pb–Br	0.309202	2
Added	AcA–Pb–F	0.151947	12
Added	AcA–Pb–I	0.104843	26

At each $\alpha \in \{0, 0.05, \ldots, 1\}$ we enumerate all 20 choices of three rows from the six exchanged measurements. At $\alpha = 0$, the initial and proposed risks are 806.703 and 853.260; at $\alpha = 1$ they are 396.754 and 343.758. Both original batches close nine of the twelve queries. The total unclosed reconstruction residual decreases from 2.89971 to 2.42654, connecting the improvement to reconstruction quality beyond closure count. Direct feature-space evaluation agrees with the covariance calculation to 8×10^{-12}.

The original EPIG comparison. We replay the ten tasks of Table 2: budget at most 24, $\lambda = 10^{-4}$, and task seeds 500–509. Condition only on the original observations and hold both selected batches fixed while applying the same covariance intervention. CIRCUIT selects 12–16 purchases and EPIG selects 24. EPIG has the larger summed singleton information score in every task. With the observed nuisance correlations, CIRCUIT has lower design risk in all ten tasks; with independent nuisance it has higher risk in all ten. The latter median risk ratio is 1.301. Candidate scores, actual purchase counts, and both intervened risks are recorded for every task. This links the baseline comparison to joint measurement geometry through an intervention that preserves the entire singleton score vector.

C.16 REPRESENTATION CHANGES EXCHANGE VALUES AT FIXED PURCHASES

The reaction-feature projection in Appendix C.12 retains 99.794% of singular-value energy and preserves existing exact reconstruction relations. We isolate its finite-noise effect by fixing the observations, queries, and the pre/post batches of every archived accepted multi-measurement move,

then changing only the feature map used to evaluate those purchases. The complete representation has 163 recorded moves across the two noise scales, representing 37 distinct search moves because several response backgrounds share the same design. Of these, 44 records (8 distinct moves) change from improving to non-improving under the eight-dimensional representation. All occur at $\lambda = 0.0087$.

A second intervention retains the full-feature normalization and multiplies only the discarded coordinates by $\beta \in \{0, 0.01, 0.1, 0.25, 0.5, 1\}$. At $\beta = 0$, the same 44 records change sign, isolating the effect of discarded directions from overall rescaling. In the reverse comparison, 73 of the 74 moves accepted in eight dimensions still improve under full features.

D Risk computation and parameter dependence

D.1 The risk as noise varies

Fix a support S and write $t = \kappa^2 = \rho^2/\sigma^2$. Theorem A.3 makes the inner problem a ridge fit, so each support carries a one-dimensional path

$$R_S(t) = \sigma^2 \min_w \{t\, b(w) + v(w)\}, \qquad b(w) = \|\Phi^\top w - \varphi_q\|^2, \quad v(w) = \|w\|^2.$$

Two quantities travel with it. The *squared residual norm* $A(S) = \|P_\perp^{(T\cup S)}\varphi_q\|^2$ is the part of the query no witness on S can represent, and the *witness variance* $V(S) = \min\{\|w\|^2 : \Phi^\top w = P^{(T\cup S)}\varphi_q\}$ is the price of representing the part that can be. The exact risk is the lower envelope of the lines $t \mapsto t\, b(w) + v(w)$ over *all* witnesses w (already in the scalar case $\Phi = [1]$, $\varphi_q = 1$ a closed support has $R_S(t) = \sigma^2 t/(1+t)$, strictly below its own asymptote $tA + V = \sigma^2$) and $(A(S), V(S))$ is the line that envelope approaches as $t \to \infty$. Supports are compared by these asymptotes: the frontier of achievable (A, V) points is the low-noise organising object, exact in the limit of Theorem D.2 and an upper envelope proxy at finite t.

Proposition D.1 (Path structure)**.** *For $t > 0$ the minimiser $w_S(t)$ is unique and analytic, R_S is continuous, concave and increasing, $b(w_S(t))$ decreases and $v(w_S(t))$ increases, and the envelope identity $R_S'(t) = \sigma^2 b(w_S(t))$ holds. If S closes q then $b(w_S(t)) \to 0$ and $v(w_S(t)) \to V(S)$.*

The global risk, a pointwise minimum of concave increasing functions, *moves continuously while the optimal support moves as a staircase*.

The trichotomy below is what makes failure predictable from the design matrix alone: the limiting ratio between any two supports depends only on their squared residual normes.

Theorem D.2 (Low-noise trichotomy)**.** *As $\lambda = 1/t \to 0$, $r_\lambda(S) = A(S)/\lambda + V(S) + o(1)$. Comparing a greedy support S_g with an optimal $S_\star$, the risk ratio has exactly three limits:*

$$\frac{r_\lambda(S_g)}{r_\lambda(S_\star)} \longrightarrow \begin{cases} A(S_g)/A(S_\star), & A(S_\star) > 0, \\ \infty \text{ at rate } 1/\lambda, & A(S_\star) = 0 < A(S_g), \\ V(S_g)/V(S_\star), & A(S_g) = A(S_\star) = 0. \end{cases}$$

Which case a task falls in is a property of the design, and Section 6 verifies all three. On these grids λ lies between 0.0087 and 0.187; at $\lambda = 0.187$ the two-term expansion misses the measured ratio by a median 5.7%; the trichotomy says which term wins eventually, and Appendix C.9 reports the finite-noise comparisons.

The low-noise limit selects closure first; the exact risk compares closure and precision at finite κ.

Theorem D.3 (Closure limit)**.** *As $\kappa \to \infty$ the optimal budgeted support obeys the lexicographic program: first minimise $A(S)$ over budget-feasible supports; then, among the minimisers, minimise $V(S)$. In particular, if closure is reachable within budget the limiting support closes the query.*

At finite κ the optimum trades squared residual norm against witness variance, and the frontier describes that trade.

D.2 BATCH SELECTION OVER A PARAMETER INTERVAL

One support can serve the whole admissible range of the exchange rate that prices closure. A computable gap Γ quantifies the remaining decision uncertainty and bounds how much any probe could add.

The measured family-to-family variation defines an admissible range for the exchange rate (Appendix C.8). This makes batch selection an optimisation over that range.

We select a batch over the admissible parameter interval. In practice κ is estimated with an interval at least as wide as the measured shift span (Appendix C.8), and for $[\kappa_L, \kappa_U]$ the experiments choose the *robust support* by minimising $\max_\kappa [R_\kappa(S)/R^\star_\kappa - 1]$ over nine geometrically spaced parameter values, re-fitting the witness once the data are in. Across the tested intervals, the robust support attains median worst-case regret of 0.0–0.5% (Table 3 and Appendix C.10).

Definition D.4 (Decision-closure gap)**.** Let $\mathcal{U}_T$ be the set of shift laws consistent with the measured design, and for $\theta \in \mathcal{U}_T$ let R_θ be the risk it induces. Write

$$\Gamma(T) \;=\; \min_{S:\, c(S)\le B} \; \sup_{\theta \in \mathcal{U}_T} \Big[R_\theta(S) - \min_{S'} R_\theta(S') \Big].$$

$\Gamma(T) = 0$ says a single support is optimal for every admissible θ, fully determining the batch decision. Small $\Gamma(T)$ says the same to a tolerance, and no probe can be worth more than Γ; large $\Gamma(T)$ leaves the responses something to contribute only if some purchasable cell distinguishes the disagreeing shifts (buy that cell first) and where no such cell exists the minimax support is the final answer, the case these designs realise, most of the shift space being unreachable from this pool.

The decision is a design object, Theorem A.1 one level up: for fixed θ the minimiser $S^\star(\theta)$ is a functional of $(\Phi_T, \mathcal{C}, Q, c, B)$ and of θ, never of y, so $\Gamma(T)$ is computable before any purchase once $\mathcal{U}_T$ is fixed. Responses can refine the set of admissible rates.

On these designs the decision is nearly closed. Taking $\mathcal{U}_T$ to be the measured shift span applied to each operating ridge scale and solving for the optimal support exactly at every point of it: at $\lambda = 0.0087$ the admissible set visits 2 supports, 86% of them sharing the midpoint's, with $\Gamma(T) = 0.55\%$; at $\lambda = 0.187$, 3 supports, 52%, and 3.94%. A set of shifts spanning a factor of forty in the open-block rate is served by one support to within four percent.

The certifying engine. Each quantity the procedure needs is one solve of the same program: Theorem A.3 leaves a group-cardinality-constrained sparse ridge problem with an exact mixed-integer second-order cone form and no big-M (perspective constraints $\|W_{x,:}\|^2 \le u_x z_x$, budget on binary z), agreeing with enumeration on 8 of 8 instances small enough to enumerate. The reported numerical values are certified by exact enumeration on restricted pools. Solver configurations and statuses are provided in `reproducibility/bound_computations.tex`.

D.3 FAMILY ELIMINATION AND CERTIFIED LOWER BOUNDS

This section builds the measuring instruments the experiments lean on. An exact reparameterisation collapses the model to a small shared block, which makes the full risk computable at real size; a Frank–Wolfe scheme over per-family menus then yields both a certified lower bound and a feasible batch, so improvement can be reported as a fraction of a certified gap.

Theorem A.3 says what to minimise and Section 5 searches one face of it. This section solves the objective itself, by first writing it in coordinates that make a purchase inside one family a small object.

An exact reparameterisation. A full one-hot pairwise model on an $n_A \times n_B \times n_C$ grid splits into a part every family shares and a part local to each level of A: $f(a,b,x) = g(b,x)^\top \gamma + h(b,x)^\top \eta_a$, with g the intercept, B, X and $B\times X$ terms and h the within-family intercept, B and X. On the $16\times 3\times 4$ grid that is $20 + 16\cdot 8 = 148$ coefficients, which is p.

The proposition below is what makes the full risk computable at real size: 148 coefficients collapse to 18 exactly, so a purchase inside one family becomes a small matrix added to a shared block.

Proposition D.5 (Family elimination). *On the $n_B \times n_C$ cells g has rank $n_B n_C$ and h has rank $n_B + n_C - 1$; the remaining directions are never observed and never queried. Let θ collect the $n_B n_C$ shared directions together with the local directions of the h families the queries touch, and let η_a be the local directions of the others. Then, with $\Sigma_0^{-1} = \lambda I$ and $\Omega_a = \lambda^{-1} I$ matching the prior of Theorem A.2,*

$$J_a(D_a) = A_{D_a}^\top \big(I + \lambda^{-1} B_{D_a} B_{D_a}^\top\big)^{-1} A_{D_a}, \qquad R(S) = \operatorname{tr}\Big[Q^\top \Big(\lambda I + \sum_a J_a(T_a \cup S_a)\Big)^{-1} Q\Big],$$

exactly, with shared dimension $n_B n_C + (n_B + n_C - 1)h$. For the grid here and one query family that is $12 + 6 = 18$ against 148.

The reduced and full computations are compared at $\lambda = 0.187, 0.0087$ and 10^{-4} and agree to a worst relative error of 3.3×10^{-13}. It compresses the evaluation of the risk; the purchase combinatorics is addressed next.

Family menus and a budget oracle. Enumerate, inside each family, every budget-feasible local purchase set, and store its cost with its matrix $J_{a,o}$. A batch is one menu item per family with total cost at most B, and the objective $f(J) = \operatorname{tr}(Q^\top J^{-1} Q)$ is convex in the summed matrix with $\nabla f(\bar{J}) = -\bar{J}^{-1} Q^\top Q \bar{J}^{-1}$. Linearising at any positive definite $\bar{J}$ turns the inner minimisation into a scalar score per menu item, so a budget dynamic program solves it exactly, returning both a valid lower bound and a genuine integer batch. Iterating gives a Frank–Wolfe scheme whose every iterate is feasible.

The strengthening over relaxing cells one at a time is that nuisances are eliminated *after* the local purchase set is fixed, so no fractional purchase inside a family can cancel a nuisance that no real purchase cancels. Mixing across families is still possible, and Lemma D.6 is what governs it.

On the real tasks. The family-menu formulation gives a stronger lower bound than the per-cell relaxation on all eight HOIP tasks, spanning four instances at both operating ridge scales. Task-level gaps, feasible batches, and convergence measurements are recorded in `reproducibility/family_menu_tasks.tex`. Lemma D.6 explains the noise dependence of the relaxation gap.

The relaxation gap is characterised in Lemma D.6 and Corollary D.7. The bound comparisons, configurations and solver records are supplied in `reproducibility/bound_computations.tex` and `reproducibility/family_menu_tasks.tex`.

D.4 Menu-hull mixing

Lemma D.6 (The divergence a mixture creates). *Let $G_s \succeq 0$ be finitely many information matrices with weights $\alpha_s > 0$ summing to one, $\bar{G} = \sum_s \alpha_s G_s$, and write $N_s = \ker G_s$ and $\bar{N} = \ker \bar{G} = \bigcap_s N_s$. With $r_\lambda(G) = q^\top (G + \lambda I)^{-1} q$,*

$$\sum_s \alpha_s r_\lambda(G_s) - r_\lambda(\bar{G}) \;=\; \frac{1}{\lambda} \sum_s \alpha_s \big\| (P_{N_s} - P_{\bar{N}}) q \big\|^2 \;+\; O(1),$$

and the coefficient is non-negative, vanishing exactly when every $P_{N_s} q = P_{\bar{N}} q$.

Corollary D.7. *If every budget-feasible integer plan leaves $P_{N_s} q \neq 0$ and some feasible convex combination closes q, the ratio between the best integer risk and that fractional value diverges as $1/\lambda$.*