JOINT MEASUREMENT DESIGN FOR PREDICTION UNDER OBSERVATION-POLICY SHIFT

Anonymous authors
Paper under double-blind review

ABSTRACT

A change in measurement policy can alter predictive risk even when the underlying data distribution remains fixed. We study budgeted passive follow-up for prediction under known observation policies, starting from nonignorably missing records. In a finite-state model, exact risk identification requires joint observation of every nonzero interaction in the target-policy loss. Even a simple logistic observation rule can therefore require full-record audits. We formulate a convex program that jointly chooses partial measurements and a loss reconstruction. It controls bias pointwise and mean-squared error under a declared reference envelope. Observable risk contrasts can reduce the information needed to compare predictors. For approximation selection, we distinguish the smallest error bound from the smallest actual error. We construct an unbiased candidate-MSE criterion from a separately charged validation sample and prove a finite-menu selection guarantee. Bias-matched logistic experiments separate the effects of approximation, reconstruction, and allocation. Paid-pilot experiments on repeated semi-synthetic image splits show larger prediction gains from reconstruction than from learning the allocation, whose gains are smaller and inconsistent. Validation and reference fitting can also cost more than they save. The results characterize the information identified by partial measurements, the error left by approximation, and the design decisions supported by measurements an investigator acquires.

1 INTRODUCTION

A predictor can fail after a change in measurement policy even when the underlying process remains fixed. An event-triggered sensor may record a value when an unrecorded signal is unusual. A model trained on those records then faces a different prediction problem under a new trigger. If the historical trigger depends on hidden information, a more expressive predictor does not by itself identify the distribution of omitted values.

Consider three binary sensor values. A fixed, exogenous historical mask reveals at most two on each record, while a deployment trigger depends on their product. The laws $2^{-3}(1 + \theta x_1 x_2 x_3)$ have identical one- and two-coordinate marginals for every θ, but different trigger probabilities. Measuring each coordinate on different records leaves this ambiguity unresolved. A randomized audit of all three on the *same* record resolves it. The deployment loss determines the required joint measurement, which can include values the predictor never reads.

Exact identification does not settle the finite-budget design problem. A logistic trigger can have a nonzero interaction involving every coordinate even when most high-order coefficients are small. Exact support coverage then requires full-record audits. At a finite sample size, partial measurements may instead suffice for the required accuracy if their approximation error is controlled and their acquisition cost is lower.

We study this tradeoff for passive randomized follow-up of nonignorably missing records. Our central result is a convex program that chooses audit probabilities and a reconstruction of a specified target loss with controlled bias. The bias bound holds pointwise, while the variance bound uses a declared reference envelope. Exact reconstruction recovers the interaction-coverage characterization. The estimand also matters. Observable risk contrasts can select a predictor without identifying its absolute risk. Learning the full target observation/outcome table can require more information than either task.

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.

The experiments separate fixed-query estimation from prediction learning. We charge the measurements used to fit a pilot reference, run the reference optimizer, and compare four estimators on identical acquired data. The target families have overlapping, disjoint, or dense scopes, with heterogeneous costs, partial native records, and informative context. Every candidate feature is relevant to the policy family. Dense-logistic tests measure finite-budget approximation gains and the persistent bias that can favor exact full audits as the sample size grows.

Position relative to prior work. Double sampling addresses informative missingness by recovering complete data on a selected follow-up sample, allowing arbitrary first-stage coarsening and smooth functionals (Levis et al., 2024). Optimal two-phase sampling allocates follow-up using auxiliary information, costs, and efficient-estimation criteria (Gilbert et al., 2014). We instead allow *partial* second-stage recovery. The target loss determines which information may remain unidentified and how much approximation error can be tolerated. Loss-difference observability is established in partial monitoring (Bartók et al., 2014). We apply that principle to distinguish predictor selection from absolute-risk estimation. Appendix A compares the results and assumptions.

Missingness-shift robustness (Rockenschaub et al., 2024), missing-data recoverability (Mohan & Pearl, 2021), active-acquisition evaluation (von Kleist et al., 2025), and Learning-to-Measure (Kobayashi et al., 2025) use different assumptions about the available information. Our analysis builds on inverse-inclusion weighting, limited-attribute learning, and optimal design (Horvitz & Thompson, 1952; Hazan & Koren, 2012; Bullins et al., 2016; Elfving, 1952). Convex estimation of a linear functional through a fixed affine observation map is also established (Juditsky & Nemirovski, 2009). We optimize the channel mixture under acquisition costs together with a pointwise approximation allowance. The additional structure is this joint choice, rather than fixed-channel estimability or convexity itself.

2 A PASSIVE-AUDIT MODEL

Records are independent. A finite context H, recoverable features $X \in \{-1, 1\}^d$, an eventual finite outcome Y, and a historical mask $M \subseteq [d]$ have an otherwise unrestricted joint law. The pre-audit record is

$$O = (H, M, X_M). \tag{1}$$

An audit selects $A \subseteq \overline{M}$ with known probability $q(A \mid O)$ and reveals X_A, subject to

$$A \perp (X_{\overline{M}}, Y) \mid O, \qquad \mathbb{E}[c(A, O)] \le b. \tag{2}$$

The outcome is subsequently observed on every record, including unaudited records. It may enter an estimator but not the current audit decision. Audits recover existing values without changing the outcome. Recovery from a buffered sensor record fits this model; a diagnostic intervention that changes treatment does not. We consider one measurement batch per record and exclude arbitrary within-record adaptive sensing.

A specified target kernel $\pi_j(m' \mid h, x)$ changes the observation policy without changing the law of (H, X, Y). For predictor f_j and loss $0 \le \ell \le 1$, its complete-data loss and risk are

$$\begin{aligned} g_j(h, x, y) &= \sum_{m'} \pi_j(m' \mid h, x) \ell\big(f_j(h, m', x_{m'}), y\big), \\ R_j &= \mathbb{E}[g_j(H, X, Y)]. \end{aligned} \tag{3}$$

The target kernel is a known input to the estimand and may depend on the state. Source audits do not identify an unknown deployment mechanism. A predictor family gives a family of losses. Estimating the complete target observation/outcome table instead requires queries $k_j(z \mid h, x)\mathbf{1}\{Y = y\}$ for each target observation symbol and outcome. These tasks need not require the same measurements.

3 EXACT INFORMATION REQUIREMENTS

A query is identified when all complete-data laws inducing the same law of (O, Y, A, X_A) give it the same mean. Let $\rho_{oy} = \mathbb{P}(O = o, Y = y) > 0$. In cell $o = (h, m, x_m), y$, only $u = x_{\overline{m}}$ is unknown.

Table 1: Assumptions and the guarantees they support. A fitted reference affects precision, not the exact pointwise reconstruction identity.

Assumption	Consequence and boundary
Known target kernel, passive randomized audits, eventual outcomes	Exact scores are valid without a model for historical missingness.
Unrestricted full-support completion laws	Uncovered nonzero interactions imply nonidentification. Structural restrictions can weaken this necessity.
A declared reference or density envelope	Supports variance optimality or an MSE upper bound, not universal efficiency.
Independent, paid pilot; frozen main design	Reference fitting does not bias fixed-query scores. Pilot cost reduces the main acquisition allowance.

For $S \subseteq \overline{m}$, define $\chi_S(u) = \prod_{i \in S} u_i$ and $\chi_\emptyset = 1$. The unique Walsh expansion is

$$g_{j,o,y}(u) = \sum_S \alpha_{j,o,y,S} \chi_S(u), \qquad \alpha_{j,o,y,S} = 2^{-|\overline{m}|} \sum_u g_{j,o,y}(u) \chi_S(u). \tag{4}$$

This is an algebraic basis expansion, not an independence or uniformity assumption on the data. Write

$$p_S(o) = \sum_{a \supseteq S} q(a \mid o), \qquad p_\emptyset(o) = 1. \tag{5}$$

Theorem 1 (Exact partial-audit identification)**.** *Over unrestricted conditional full-support completion laws, $R_1, \ldots, R_J$ are identified if and only if, in every positive-mass cell,*

$$\alpha_{j,o,y,S} \neq 0 \implies p_S(o) > 0 \quad \textit{for every } j, S. \tag{6}$$

If a condition fails, every interior complete-data law has arbitrarily close alternatives with the same audited-data distribution and a different target risk.

An audit of a reveals a marginal whose row space is spanned by characters with $S \subseteq a$. An uncovered character lies in the null space of every supported marginalization channel. Perturbing the completion law in that direction leaves the observed law unchanged, proving necessity. Inclusion weighting proves sufficiency. Appendix K gives the full argument and an explicit MNAR pair with a positive fraction of native complete records.

When (6) holds, a fixed-query score is

$$W_j = \alpha_{j,O,Y,\emptyset} + \sum_{\substack{\emptyset \neq S \subseteq A \\ \alpha_{j,O,Y,S} \neq 0}} \frac{\alpha_{j,O,Y,S} \chi_S(X_S)}{p_S(O)}, \qquad \widehat{R}_j = n^{-1} \sum_{i=1}^n W_{ij}. \tag{7}$$

Proposition 2 (Design-based validity)**.** *For every complete record, $\mathbb{E}_A[W_j \mid H, X, Y, M] = g_j(H, X, Y)$. Thus $\widehat{R}_j$ is unbiased under any complete-data law satisfying the randomization, including laws without full support.*

The proof is the identity $\mathbb{E}[\mathbf{1}\{S \subseteq A\}/p_S(O) \mid O] = 1$. No historical observation propensity is estimated. Scores can be signed and need not lie in the loss range.

3.1 Selecting a predictor can require less information

Proposition 3 (Observable contrasts)**.** *For one target policy and a finite predictor class, fix f_0 and let $d_f = g_f - g_{f_0}$. All risk contrasts are identified exactly when every nonzero Walsh coefficient of every d_f is covered. This suffices to identify an optimal predictor, even if the absolute risks are not identified. Uniform contrast-estimation error at most e gives excess risk at most $2e$ for a contrast minimizer.*

For example, valid squared losses can be $g_\pm = 1/2 + \chi_{\{1,2,3\}}/8 \pm x_1/8$. Their difference is $x_1/4$. Auditing coordinate one identifies the comparison but neither absolute risk. Appendix B constructs the observation kernel and gives a decision-identification linear program. Contrast identification is sufficient, but a pointwise dominant predictor can be selected without it. Policy-specific unknown offsets need not cancel under a worst-policy objective.

4 DESIGNING MEASUREMENTS FOR BOUNDED ERROR

Exact identification can require full audits even for a simply parameterized policy. For $g(x) = \sigma(0.7 + \sum_i \beta_i x_i)$, with β_i equally spaced from 0.2 to 0.8, every Walsh coefficient is nonzero in enumerations at $d = 2, 4, 6, 8$. At $d = 8$ the full-interaction coefficient is approximately 4.54×10^{-4}. A mask-only predictor can induce this loss. Thus a simple nonlinear observation policy can require full-record audits for exact identification. Any relaxation must account for the omitted information rather than silently threshold it away.

4.1 A JOINT RECONSTRUCTION AND ACQUISITION PROGRAM

Fix a pre-audit stratum and enumerate its complete states $u \in [N]$, including the eventual outcome when needed. Audit a has a known channel C_a whose columns are output distributions. Its output includes native information and the eventual outcome. For a strictly positive reference law $p_0 > 0$, let $r_a = C_a p_0$ after deleting impossible output rows. Write $v_a(z) = q_a h_a(z)$ for a reconstruction contribution. Introduce an intercept ζ, a residual allowance ϵ, and an energy bound t. For a declared $\kappa \geq 1$, solve

$$\begin{aligned} \min_{q,v,\zeta,\epsilon,t} \quad & \epsilon^2 + \kappa t/n \\ \text{subject to} \quad & \left\| \zeta \mathbf{1} + \sum_a C_a^\top v_a - g \right\|_\infty \leq \epsilon, \\ & \sum_{a,z} r_a(z) \frac{v_a(z)^2}{q_a} \leq t, \\ & q \in \Delta(\mathcal{A}), \quad \sum_a q_a c_a \leq b, \quad \epsilon, t \geq 0. \end{aligned} \tag{8}$$

At $q_a = 0$, the perspective term is zero if $v_a = 0$ and infinite otherwise. This convention prevents unpurchased actions from carrying information.

Theorem 4 (Joint bias–variance design)**.** *Program* (8) *is convex. Every finite feasible point defines the score $W = \zeta + v_A(Z)/q_A$ on selected actions. Its absolute bias for $R_p = p^\top g$ is at most ϵ under every completion law. If $p \leq \kappa p_0$ coordinatewise, the mean of n independent scores satisfies*

$$\mathbb{E}_p[(\widehat{R} - R_p)^2] \leq \epsilon^2 + \kappa t/n. \tag{9}$$

With a zero-cost action, a finite menu, and strictly positive reference output probabilities, an optimum exists even when exact identification is impossible.

The proof uses the quadratic perspective (Boyd & Vandenberghe, 2004) and the pointwise identity $\mathbb{E}[W \mid u] = \zeta + \sum_a C_a^\top v_a$. The residual bounds bias, and $p \leq \kappa p_0$ bounds the centered second moment. The program optimizes an MSE upper bound within its reconstruction class; it does not characterize minimax risk over arbitrary estimators. Appendix C gives the proof, the global multi-stratum extension with audit probabilities tied across outcomes, and the multiple-loss version. Replacing reference energy by its maximum over a specified set of laws preserves convexity. Maximizing over all point masses gives a distribution-free bound, usually at greater conservatism.

For a uniform reference and observable character dictionary $\mathcal{T}$, the program becomes

$$\begin{aligned} \min_{q,\beta_0,\beta,\epsilon} \quad & \epsilon^2 + \frac{\kappa}{n} \sum_{S \in \mathcal{T}} \frac{\beta_S^2}{p_S(q)} \\ \text{subject to} \quad & \max_u \left| \beta_0 + \sum_{S \in \mathcal{T}} \beta_S \chi_S(u) - g(u) \right| \leq \epsilon, \end{aligned} \tag{10}$$

with the same simplex and budget constraints. Acquisition probabilities and reconstruction coefficients are optimized jointly, so the coefficients need not equal a truncated Walsh expansion. Appendix C.4 proves equivalence to the channel program for the given menu. The implemented degree-L version uses all actions of size at most L, explicit numerical coverage floors, and residual checks on the complete finite cube.

4.2 A USEFUL APPROXIMATION CLASS

Proposition 5 (Smooth-gate approximation). *Let $g(x) = \sigma(a + \beta^\top x)$ and $B = \|\beta\|_1 > 0$. For any $\rho > 1$ satisfying $B(\rho - \rho^{-1})/2 < \pi/2$, there is a Walsh polynomial of degree at most L with uniform error at most $2\rho^{-L}/(\rho - 1)$. If two mask-branch losses lie in $[0, 1]$ and depend on a core of at most s coordinates, their sigmoid-gated mixture has an approximation of degree at most $\min(d, s + L)$ with the same error.*

The proof combines analytic Chebyshev approximation with $x_i^2 = 1$ (Appendix D). It gives approximate low-order structure, not exact sparsity. Large B weakens the bound and can make low-degree approximation statistically inadequate. The finite program also enumerates completion constraints and is not polynomial-time in unrestricted dimension. A degree-L action menu has $\sum_{k=0}^{L} \binom{d}{k}$ members.

4.3 CHOOSING THE APPROXIMATION FROM AVAILABLE INFORMATION

For a finite menu of feasible reconstructions, minimizing $B_j = \epsilon_j^2 + \kappa t_j/n$ uses only the declared reference envelope. This selects the smallest bound, which need not correspond to the smallest true-law MSE. A user-specified bias tolerance can first restrict the menu to $\epsilon_j \le \epsilon_\star$.

A second choice uses a separately paid validation block. Freeze all candidates before recovering its $v \ge 2$ complete records. Write $h_j(u) = \mathbb{E}[W_j \mid u]$, $e_j = h_j - g$, $m_j = h_j - \zeta_j$, and $a_j(u) = \mathbb{E}[(W_j - \zeta_j)^2 \mid u]$. For $\mathcal{U}_v(z) = \sum_{i \ne k} z_i z_k/[v(v-1)]$, the observable criterion

$$\widehat{M}_j = \mathcal{U}_v(e_j) + \frac{1}{n}\{\overline{a}_j - \mathcal{U}_v(m_j)\} \tag{11}$$

is unbiased for candidate j's MSE on a fresh n-record analysis sample. Its terms use acquired validation values and the known score, without a simulator's probability law. Proposition 7 gives simultaneous error radii r_j and $M_{\widehat{j}_+} \le \min_j(M_j + 2r_j)$ for the selector minimizing $\widehat{M}_j + r_j$. The unpenalized selector minimizes $\widehat{M}_j$. Validation cost is reserved before constructing candidates, and validation records are excluded from analysis. The procedure's cost and noise can outweigh its benefit.

5 LEARNING A REFERENCE WITHOUT ASSUMING IT IS CORRECT

The exact finite-channel reconstructor permits a singular information matrix. For a chosen reference, let $B_\perp$ have orthonormal columns spanning $\mathbf{1}^\perp$, put $D_a = \mathrm{diag}(C_a p_0)$, and define

$$I(q) = \sum_a q_a B_\perp^\top C_a^\top D_a^{-1} C_a B_\perp, \qquad b_g = B_\perp^\top g.$$

A query is estimable precisely when $b_g \in \mathrm{range}\, I(q)$. The minimum-reference-variance score uses $I(q)^\dagger b_g$ and has variance $b_g^\top I(q)^\dagger b_g$. Its reconstruction identity holds under every true law, including a misspecified reference. Appendix M gives the expression and proof. The numerical routine rejects directions exceeding its declared range tolerance and reports the reconstructed query's pointwise residual as a numerical bias allowance. This tolerance-qualified computation is distinct from symbolic identification (Appendix G.3).

An independent, fully audited pilot fits the reference before the main design is frozen. If the pilot costs C_0 and n records remain, their conditional expected measurement cap is $(B_{\mathrm{tot}} - C_0)/n$. Same-pilot comparisons receive the same remaining allowance; the separate zero-pilot comparison retains the entire budget. Conditional expected total spending is therefore at most B_{tot}. The reference-fitting pilot is excluded from the main estimate and is separate from the optional validation block. It learns a completion reference for allocation and reconstruction, not an unknown deployment policy or hidden interaction support.

Reference error affects efficiency. If $(1 - \eta)p_0 \le p \le (1 + \eta)p_0$ with $\eta < 1$, a reference-optimal design and reconstruction have true variance at most $(1+\eta)/(1-\eta)$ times the true-reference optimum over the same feasible budget class. Appendix E proves this comparison, its multiple-query extension,

and conditional pilot validity. A small pilot need not satisfy a useful multiplicative bound. We therefore compare fitted-reference and complete-law designs and account for the cost of fitting, rather than assuming a precise reference is available for free.

Algorithm 1 Budgeted measurement design and reconstruction

1: Specify the target kernels, finite query family, feasible channels, and costs. Choose risks, fixed-policy contrasts, or table queries according to the intended task.
2: Reserve any reference-fitting pilot and independent validation block within B_{tot}. A zero-pilot design uses a declared reference and retains the whole allowance.
3: Subtract realized construction-pilot cost and the predetermined validation reservation. Fit a positive reference using only the construction pilot, then set the remaining per-record cap.
4: Construct feasible candidates using exact reconstruction or (8). Select by a prespecified bias tolerance, the bound, or independent paid validation. Freeze the selected design.
5: **for** each main record **do**
6: Draw the audit from $q(\cdot \mid O)$ before inspecting hidden values or the outcome. Store the selected action and its probability, then acquire its requested values.
7: After the outcome is available, evaluate the prescribed score using only the channel output.
8: **end for**
9: Average scores for fixed-query estimation, or fit the stated finite-table/class learner. Report approximation allowances separately from sampling uncertainty and evaluate fitted predictors independently.

6 Experiments and failure regimes

The logistic study evaluates biased *fixed-query estimation*. The paid-pilot study evaluates exact *table reconstruction and prediction*. Parity and correctly specified model controls examine information requirements and structural assumptions. The fixed-query MSE theorem does not give a finite-sample Brier-risk guarantee for a fitted predictor.

Dense-policy estimation. We enumerate logistic targets at $d = 4, 6, 8$, sample sizes $n = 250, 1000, 4000$, and cost caps $b = 0.5, 1, 2$. The nonuniform evaluation law is bounded by 1.8 times the uniform reference. Each case compares exact full audits, bias-aware full audits, fixed uniform subsets with optimized reconstruction, and joint degree-1, degree-2, and degree-3 designs. All optimizations use the same target, reference, envelope, sample size, and cap. The evaluation law is used only to calculate bias and variance afterward.

Table 2: Bias-matched eight-coordinate logistic comparison at $n = 250$ and cost cap 0.5. Every row spends 0.5 in expectation. MSE and its envelope are in units of 10^{-3}. The chosen envelope and actual MSE need not order designs alike.

Design	Joint size	Residual	MSE	Bound
Exact full	8	0.00000	4.3969	8.1030
Bias-aware full	8	0.02364	3.7643	7.4826
Uniform two-subsets	2	0.09108	3.4689	14.0271
Joint degree two	2	0.09109	3.0347	13.1906

At $d = 8, n = 250, b = 0.5$, joint degree two reduces actual MSE by 19.4% relative to bias-aware full auditing and 12.5% relative to fixed two-subset acquisition (Table 2). The degree is prespecified; choosing by the bound would select bias-aware full auditing instead. At $n = 4000, b = 2$, degree one has 15.2 times exact-full MSE. Its singleton menu spends only one unit despite the cap of two. Appendix S reports the full bias-matched grid, expected costs, and numerical objective gaps.

Charged approximation selection. For eight-coordinate targets, we reserve 20% of the total allowance for complete validation records. We construct the same eight candidates with the remaining budget and apply (11) before drawing the analysis sample. Across 2,000 repetitions per setting, the raw validation selector at $N_{\mathrm{tot}} = 250, b = 0.5$ has mean conditional MSE 0.004125 versus

0.003764 for bound selection without validation. The reservation buys only three complete validation records. Neither raw nor confidence-penalized validation improves mean MSE over the no-validation bound selector in the six tested resource settings. Thus these feasible, fully costed procedures do not establish a general benefit from validation. Appendix S.2 separates conditional MSE, fresh-sample squared error, and selection uncertainty.

Paid-pilot prediction. The learning study has four binary candidate features, all relevant to the target-policy family, with costs $(0.5, 1, 1.5, 2)$. Always-observed context is informative. A nonignorable native mask reveals either no candidate feature or the first one. The three target families have overlapping, disjoint, or dense scopes. Each policy reveals one core feature or an empty observation, giving three observation symbols. We compare full/union audits, uniform scopes, uniform-reference optimization, pilot-reference optimization, and an inaccessible complete-law reference design. Full and union strategies coincide here, so neither pays for known irrelevant features.

Budgets are $0.4, 0.8, 1.6$ cost units per candidate record, *including pilot acquisition*. Up to 20% is reserved for fully recovered pilot records. Each conditional reference is a smoothed 32-state feature/outcome histogram. The optimizer minimizes the sum of target-table query variances, a surrogate rather than an exact Brier-risk objective. All methods share the candidate bundles, context, native records, pilot, and remaining budget. Forty repetitions generate 4,000-record synthetic samples or stratified 1,200/597 train/test partitions of the public digit images (Alpaydin & Kaynak, 1998; scikit-learn developers, 2026). The image experiments are semi-synthetic, with imposed censoring and target policies. Their uncertainty includes split variation within the fixed image collection, not new-writer or deployment variation.

For each acquisition design, four estimators use identical acquired data: scope inverse-inclusion weighting, a pilot-augmented scope estimator, character reconstruction, and learned-reference channel reconstruction. The last three use query information available without the complete scope. All joint tables undergo the same simplex projection and vanishing pseudocount before fitting a Brier predictor. Primary comparisons hold learned-reference reconstruction fixed across acquisition policies. Table 3 reports Brier risks, and Table 4 reports paired differences. Appendix Q gives the same-audit estimator contrasts and budget curves.

Table 3: Brier risk at total expected cost 0.8 per candidate record, including the pilot. Parentheses give standard errors in units of 10^{-5} across 40 repetitions. Every acquisition method uses learned-reference reconstruction. Complete-law allocation is an inaccessible reference comparator.

Data	Family	Full/union	Uniform scopes	Uniform ref.	Pilot ref.	Complete law
Synthetic	Overlap	0.20849 (4)	0.20836 (4)	0.20838 (4)	0.20839 (4)	0.20840 (4)
	Disjoint	0.21853 (6)	0.21858 (6)	0.21846 (4)	0.21849 (5)	0.21850 (5)
	Dense	0.21438 (4)	0.21438 (4)	0.21430 (4)	0.21427 (4)	0.21429 (3)
Digits	Overlap	0.22723 (64)	0.22681 (61)	0.22669 (65)	0.22669 (62)	0.22662 (63)
	Disjoint	0.22440 (61)	0.22451 (68)	0.22447 (64)	0.22441 (64)	0.22461 (62)
	Dense	0.22513 (69)	0.22513 (69)	0.22474 (67)	0.22479 (68)	0.22477 (67)

Reconstruction accounts for the largest Brier gain. In dense digits at budget 0.8, the same pilot-reference acquisitions give 0.23550 with scope weighting and 0.22479 with learned-reference reconstruction. Holding reconstruction fixed, uniform-reference and pilot-reference allocation give 0.22474 and 0.22479. Table 11 reports the estimator comparison. The trace-variance criterion and nonlinear Brier objective are reported separately because learning the allocation reference does not consistently improve prediction.

Is the reference-fitting pilot worth its cost? A separate sweep reserves $0, 10, 20,$ or 40% of the budget before observing pilot values. The zero-pilot uniform-reference strategy uses all records and the entire measurement allowance. Positive fractions fit the same smoothed reference and use the remaining records for analysis. The sweep covers both datasets, all families and budgets, with 40 repetitions. Table 20 gives costs, reference error, conditional table variance, and Brier risk. At digits budget 0.8, the family-averaged Brier risk is 0.22501 without a pilot and 0.22529 at the 20% pilot reservation. The grid is a prespecified resource sensitivity analysis, not a pilot fraction tuned on held-out outcomes.

Table 4: Pilot-reference minus uniform-reference allocation on digits at total budget 0.8, with the same pilot and reconstruction. Brier differences and paired standard errors are in units of 10^{-5} over 40 repetitions. Trace ratios compare the attained conditional table-variance criterion, not prediction loss.

Family	Brier difference	Paired SE	Trace ratio
Overlap	-0.68	9.95	0.991
Disjoint	-6.75	10.34	0.985
Dense	$+4.59$	15.99	0.991

Misspecification and controls. We perturb retained support, the specified target kernel, and the completion reference separately. Omitting a nonzero interaction creates bias even with correct audit probabilities. An incorrect target policy changes the risk being estimated even with complete measurements. An incorrect reference preserves exact reconstruction but can increase variance. Appendix R reports these sensitivities and numerical design costs. The parity control gives matching budget rates and a d/k variance ratio when a known k-coordinate parity is the only unknown direction. This advantage disappears at $k = d$ (Theorem 10). In the correctly specified hidden-state control, a joint model achieves Brier risk 0.172714 without audits, versus 0.173318 for the scope-coverage learner. Under these valid structural assumptions, additional auditing is unnecessary (Table 25).

7 Discussion

Measurement requirements depend on the task: absolute-risk estimation, predictor comparison, or target-table learning. Exact identification requires observability of the relevant loss directions. Finite-budget estimation also permits unresolved directions if their approximation error is controlled. Joint reconstruction and acquisition formalizes this tradeoff. A paid pilot can improve the reference used for precision without making that reference an identifying assumption. Bound selection, a prescribed bias tolerance, and paid validation support different decisions. A favorable population result at one degree does not ensure that an investigator with limited data will select it.

The results assume finite known target kernels, passive recovery, and eventual observed outcomes. They do not learn an unspecified deployment mechanism or justify measurements that change the process. Dense policies, poor reference fits, and large completion spaces can remove the practical savings. The experiments test controlled statistical mechanisms and finite-state implementation, rather than performance of a deployed sensing system. Informative-sampling correction and joint dynamical observation models remain complementary approaches (Vanderschueren et al., 2023; Iakovlev & Lähdesmäki, 2025; Yalavarthi et al., 2026).

AI USE STATEMENT

Generative AI was used for grammatical checks and proofreading. It was also used for reference/literature search and verification. The authors assume full responsibility for the content.

ETHICS STATEMENT

The experiments use synthetic records and a public handwritten-digit benchmark. No new human-subject data were collected. Passive additional measurement is an explicit modeling assumption, not a recommendation to collect sensitive data without consent. In applications, acquisition costs should include privacy, burden, and downstream intervention effects where relevant. A measurement procedure that changes a person's treatment or outcome falls outside the model. The semi-synthetic censoring rules are controlled constructions and should not be interpreted as evidence about a real institution's observation practices.

REPRODUCIBILITY STATEMENT

Proofs, experimental methods, and additional controls are in the appendix. Code, data, and complete results are provided in the supplement.

REFERENCES

E. Alpaydin and C. Kaynak. Optical recognition of handwritten digits. UCI Machine Learning Repository, 1998. URL https://doi.org/10.24432/C50P49. Dataset; CC BY 4.0.

Gábor Bartók, Dean P. Foster, Dávid Pál, Alexander Rakhlin, and Csaba Szepesvári. Partial monitoring—classification, regret bounds, and algorithms. *Mathematics of Operations Research*, 39(4):967–997, 2014. doi: 10.1287/moor.2014.0663.

Stephen Boyd and Lieven Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004. URL https://web.stanford.edu/~boyd/cvxbook/.

Brian Bullins, Elad Hazan, and Tomer Koren. The limits of learning with missing data. In *Advances in Neural Information Processing Systems*, volume 29, 2016. URL https://papers.nips.cc/paper_files/paper/2016/hash/955a1584af63a546588caae4d23840b3-Abstract.html.

G. Elfving. Optimum allocation in linear regression theory. *The Annals of Mathematical Statistics*, 23(2):255–262, 1952. doi: 10.1214/aoms/1177729442.

Peter B. Gilbert, Xuesong Yu, and Andrea Rotnitzky. Optimal auxiliary-covariate-based two-phase sampling design for semiparametric efficient estimation of a mean or mean difference, with application to clinical trials. *Statistics in Medicine*, 33(6):901–917, 2014. doi: 10.1002/sim.6006.

Elad Hazan and Tomer Koren. Linear regression with limited observation. In *Proceedings of the 29th International Conference on Machine Learning*, 2012. URL https://arxiv.org/abs/1206.4678.

Wassily Hoeffding. A class of statistics with asymptotically normal distribution. *The Annals of Mathematical Statistics*, 19(3):293–325, 1948. doi: 10.1214/aoms/1177730196.

D. G. Horvitz and D. J. Thompson. A generalization of sampling without replacement from a finite universe. *Journal of the American Statistical Association*, 47(260):663–685, 1952. doi: 10.1080/01621459.1952.10483446.

Valerii Iakovlev and Harri Lähdesmäki. Learning spatiotemporal dynamical systems from point process observations. In *International Conference on Learning Representations*, 2025. URL https://proceedings.iclr.cc/paper_files/paper/2025/hash/b35df4ede4411f6831b94452bb98e43b-Abstract-Conference.html.

Anatoli B. Juditsky and Arkadi S. Nemirovski. Nonparametric estimation by convex programming. *The Annals of Statistics*, 37(5A):2278–2300, 2009. doi: 10.1214/08-AOS654. URL https://arxiv.org/abs/0908.3108.

Yuta Kobayashi, Zilin Jing, Jiayu Yao, Hongseok Namkoong, and Shalmali Joshi. Learning-To-Measure: In-context active feature acquisition. *arXiv preprint arXiv:2510.12624*, 2025. URL https://arxiv.org/abs/2510.12624.

Alexander W. Levis, Rajarshi Mukherjee, Rui Wang, Heidi Fischer, and Sebastien Haneuse. Double sampling for informatively missing data in electronic health record-based comparative effectiveness research. *Statistics in Medicine*, 43(30):6086–6098, 2024. doi: 10.1002/sim.10298.

Karthika Mohan and Judea Pearl. Graphical models for processing missing data. *Journal of the American Statistical Association*, 116(534):1023–1037, 2021. doi: 10.1080/01621459.2021.1874961.

Ryan O'Donnell. *Analysis of Boolean Functions*. Cambridge University Press, 2014. doi: 10.1017/CBO9781139814782.

Patrick Rockenschaub, Zhicong Xian, Alireza Zamanian, Marta Piperno, Octavia-Andreea Ciora, Elisabeth Pachl, and Narges Ahmidi. Robust prediction under missingness shifts. *arXiv preprint arXiv:2406.16484*, 2024. URL https://arxiv.org/abs/2406.16484.

scikit-learn developers. load_digits: Optical recognition of handwritten digits dataset. scikit-learn documentation, 2026. URL https://scikit-learn.org/1.8/modules/generated/sklearn.datasets.load_digits.html. Accessed September 19, 2026.

Lloyd N. Trefethen. *Approximation Theory and Approximation Practice*. Society for Industrial and Applied Mathematics, 2013. doi: 10.1137/1.9781611972405.

Toon Vanderschueren, Alicia Curth, Wouter Verbeke, and Mihaela van der Schaar. Accounting for informative sampling when learning to forecast treatment outcomes over time. In *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pp. 34855–34874, 2023. URL https://proceedings.mlr.press/v202/vanderschueren23a.html.

Henrik von Kleist, Alireza Zamanian, Ilya Shpitser, and Narges Ahmidi. Evaluation of active feature acquisition methods for time-varying feature settings. *Journal of Machine Learning Research*, 26 (60):1–84, 2025. URL https://www.jmlr.org/papers/v26/23-1635.html.

Vijaya Krishna Yalavarthi, Randolf Scholz, Christian Klötergens, Kiran Madhusudhanan, Stefan Born, and Lars Schmidt-Thieme. Reliable probabilistic forecasting of irregular time series through marginalization-consistent flows. In *International Conference on Learning Representations*, 2026. URL https://proceedings.iclr.cc/paper_files/paper/2026/hash/899c6a43b9976e1077522fe5a39cafa3-Abstract-Conference.html.

A RESULT-LEVEL RELATION TO EARLIER WORK

We study how to choose a partial measurement experiment for a deployment loss. Linear estimability and randomized follow-up for informative missingness are established foundations. Tables 5 and 6 distinguish those foundations from the results developed here.

Table 5: Comparison with double sampling and optimal two-phase sampling. "Not an estimand" means that the earlier paper does not formulate a change of deployment observation kernel, not that it assumes that kernel is unknown.

Dimension	Earlier sampling results	Partial-audit formulation here
Source missingness	Levis et al. (2024), including Appendix A, allow arbitrary first-stage coarsening, with assumptions on follow-up selection and recovery. Gilbert et al. (2014) study designed second-phase sampling conditional on auxiliary variables.	Historical (H, X, Y, M) is unrestricted. Known, positive audit probabilities identify the particular directions they reveal.
Follow-up	Levis et al. recover the complete missing information for selected follow-up cases. Gilbert et al. sample an expensive endpoint given inexpensive auxiliary information.	Different actions reveal different coordinate bundles. No action need recover every hidden coordinate when the requested query is estimable.
Target mechanism	A deployment observation kernel is not the estimand of those sampling results.	Each kernel π_j is supplied. It can induce interactions absent from a predictor and is not identified automatically by the audit.
Identification	Levis et al.'s generalized coarsening result identifies the complete-data law, hence smooth complete-data functionals, under their recovery assumptions.	Theorem 1 identifies a specified loss even when the completion law remains nonidentified. Its necessity is over an unrestricted interior finite model.
Statistical target	Smooth functionals and efficient estimation in Levis et al. Means or mean differences in Gilbert et al.	Fixed predictive risks, contrasts, or target-table probabilities are distinct query families. Proposition 3 applies to a fixed-policy ranking.
Design criterion	Follow-up selection and efficient influence functions already motivate cost-aware sampling and variance reduction. Bias–variance tradeoffs also appear in model-assisted missing-data estimation.	Equation (8) jointly varies the partial observation experiment and its approximate reconstruction, with a uniform residual and a declared MSE envelope.
Computation	Sampling probabilities and nuisance fits exploit their specified statistical models and auxiliary covariates.	The exact finite channel and the implemented approximation enumerate completion states. Union menus can be exponential. The support bound does not remove this dependence.

What changes at the theorem level. A full-recovery indicator distinguishes no recovery from recovery of the whole missing vector. A menu of partial channels changes the estimable subspace. Theorem 1 characterizes that subspace for coordinate measurements and an observation-policy loss, without requiring recovery of the complete law. The result specializes finite linear estimability; the proof is included to make this connection explicit, not to introduce row-space estimability as a new principle.

Theorem 4 treats a specified loss outside an affordable channel span. It jointly chooses an experiment and an approximate loss, bounds the pointwise reconstruction error, and uses a reference envelope to control precision. The full-recovery sampling results discussed here do not state this partial-channel residual-constrained program or the smooth-gate acquisition consequence of Proposition 5. Our finite-state analysis, in turn, does not subsume their semiparametric efficiency theory, flexible nuisance learning, or clinical applications. The program uses classical perspective convexity and gives an MSE upper bound, rather than a general minimax theorem.

Convex estimation through a fixed channel. Juditsky & Nemirovski (2009), Sections 2–3, construct estimators of linear functionals under known affine observation maps and prove constant-factor minimax guarantees for their statistical families. Their discrete-distribution example is directly relevant. With q fixed, our audit law is $p \mapsto (q_a C_a p)_a$, and the target is $g^\top p$. This is an established fixed-channel estimation setting. We additionally choose q under measurement costs, while the residual program controls the part of g left unresolved by the partial channels. Theorem 4 bounds MSE under a declared envelope and does not inherit their minimax guarantee. The distinction is the optimized observation experiment and its guarantee, rather than convexity alone.

Table 6: Relation to learning, observability, and approximation. These are distinctions in the formulated problem and guarantee, not claims of universal superiority.

Work	Established result or objective	Relationship to this paper
Partial monitoring, Bartók et al. (2014)	Feedback observability of loss differences distinguishes regret regimes in finite games.	Observable contrasts are explicitly credited to this framework. Here the action purchases a measurement, the query is a deployment loss, and acquisition cost is distinct from predictive loss. No new regret classification is claimed.
Limited attributes, Hazan & Koren (2012); Bullins et al. (2016)	Learning with restricted attribute observations, including information-theoretic limitations.	Limited joint access is not new. The observation kernel here changes the complete-data loss, and the historical coarsening can be unrestricted MNAR. The parity separation assumes the relevant interaction is specified.
Acquisition evaluation, von Kleist et al. (2025)	Identification and evaluation of active acquisition policies in time-varying feature settings, with explicit retrospective-data assumptions.	Our added measurements supply information absent from the historical law. The finite one-batch model does not replace their sequential evaluation theory.
Learning-to-Measure, Kobayashi et al. (2025)	In-context active acquisition learned from incomplete retrospective observations under identifying assumptions.	We learn a reference from paid prospective audits, not their in-context acquisition model. The MI control is not a reproduction of that method.
Missingness-shift robustness, Rockenschaub et al. (2024)	Conditions and predictors for robustness to changes in missingness.	We do not claim that a changed missingness mechanism is a new problem. We ask which new partial measurements restore a specified loss functional.
Boolean analysis and analytic approximation, O'Donnell (2014); Trefethen (2013)	Character expansions and exponential polynomial approximation of analytic functions.	These give a coordinate interpretation and a smooth-gate approximation class. Neither the basis nor the analytic approximation theorem is new.

Limits of the novelty claim. The method assumes the deployment policy and query structure are specified. It does not discover a hidden sparse support with sublinear feature access or provide a computationally efficient procedure for arbitrary high-dimensional neural losses. Its contribution combines an estimand-specific partial follow-up experiment with controlled reconstruction error. The implementation tests the value of fitting a reference rather than assuming it. In the target-aware comparisons, reference optimization gives modest predictive gains once competing estimators exploit the same acquired information.

B Decision identification and observable contrasts

Partial monitoring distinguishes loss estimation from action comparison (Bartók et al., 2014). We apply that distinction to a fixed deployment kernel and a separate randomized measurement experiment. A measurement action need not be a prediction action, and acquisition cost is separate from predictive loss.

Proof of Proposition 3. For a fixed target policy and reference predictor f_0, set $d_f = g_f - g_{f_0}$. Linearity of the Walsh transform gives the coefficient of d_f as $\alpha_{f,o,y,S} - \alpha_{f_0,o,y,S}$. Applying

Theorem 1 to this finite family identifies all contrasts exactly when each nonzero difference coefficient is covered. The proof of that theorem does not require nonnegative queries, so a loss difference in $[-1, 1]$ is admissible.

Since $R(f) = R(f_0) + \mathbb{E}d_f$, minimizing the identified contrasts minimizes the original risk, even when $R(f_0)$ is unknown. For an estimator with $|\widehat{D}_f - \mathbb{E}d_f| \le e$ simultaneously, let f^* minimize R and let $\hat{f}$ minimize $\widehat{D}$. Then

$$R(\hat{f}) - R(f^*) = \mathbb{E}d_{\hat{f}} - \mathbb{E}d_{f^*} \le \widehat{D}_{\hat{f}} - \widehat{D}_{f^*} + 2e \le 2e.$$

The same reasoning applies to a specified weighted average of policy risks, by first forming its aggregate complete-data loss. It does not apply separately to each policy before taking a maximum. □

A valid squared-loss construction. Take $d \ge 3$, $Y = 0$, $K = \{1, 2, 3\}$, and four distinct mask labels m_0, m_1, m_2, m_3. For example, use $\emptyset, \{1\}, \{2\}, \{1, 2\}$. Assign probabilities

$$\begin{aligned} \pi(m_0 \mid x) &= \tfrac{1}{4} + \tfrac{1}{8}\chi_K(x), \quad & \pi(m_1 \mid x) &= \tfrac{1}{4} + \tfrac{1}{8}x_1, \\ \pi(m_2 \mid x) &= \tfrac{1}{4} - \tfrac{1}{8}x_1, & \pi(m_3 \mid x) &= \tfrac{1}{4} - \tfrac{1}{8}\chi_K(x). \end{aligned} \tag{12}$$

They lie in $[1/8, 3/8]$ and sum to one. Predictor f_+ returns one at m_0, m_1 and zero elsewhere, while f_- returns one at m_0, m_2 and zero elsewhere. Both ignore feature values. Their squared target losses are

$$g_\pm(x) = \tfrac{1}{2} + \tfrac{1}{8}\chi_K(x) \pm \tfrac{1}{8}x_1, \qquad g_+(x) - g_-(x) = \tfrac{1}{4}x_1.$$

In the native-empty stratum, audits of coordinate one identify their risk difference but not either absolute risk. This is not a feature-sparsity assumption on the complete law. For example, laws proportional to $1 + \theta x_1 + \psi\chi_K(x)$, with $|\theta| + |\psi| < 1$, admit both an identifiable sign of the risk difference and a freely varying common risk component.

Contrast identification is sufficient, not necessary, for a decision. If $g_f \le g_k$ pointwise for all k, f is optimal without identifying any nonconstant contrast. More generally, let $\mathcal{P}(t) = \{p \in \Delta_N : K_q p = t\}$ be the completion-law fiber of a finite observed experiment. A given f is optimal under every law in that fiber if and only if

$$\max_{p \in \mathcal{P}(t)} (g_f - g_k)^\top p \le 0 \quad \text{for every } k. \tag{13}$$

Each maximum is a linear program. Necessity follows by evaluating the inequality at every compatible law; sufficiency follows by reading the same inequalities as risk comparisons. For the interior model with an interior compatible point, using the closed fiber does not alter the supremum, since mixing any boundary point with the interior point approaches it from within the fiber. Thus observable contrasts characterize contrast estimation, not every possible decision-identification mechanism.

Why a worst-policy objective is different. Consider two predictors with policy-risk rows $(0.1, 0.4)$ and $(0.3, 0.2)$. The policy-specific contrasts are unchanged when a common offset is added to each column. Offsets $(0.3, 0)$ make the first predictor preferable under the maximum risk, giving maxima 0.4 and 0.6. Offsets $(0, 0.3)$ instead give maxima 0.7 and 0.5. All risks remain in $[0, 1]$. Therefore separate within-policy contrast identification is not, by itself, a valid reduction of worst-policy selection. Absolute offsets, a shared cross-policy offset, or another identified comparison of the maximum is additionally needed.

C Joint approximation and measurement design

C.1 Finite channels and quadratic perspectives

Fix a pre-audit stratum. Enumerate its full states by $u \in [N]$. The state includes the later outcome whenever it is not constant. Each action a has a known column-stochastic observation channel C_a, and outputs can include the outcome and native measurements at zero additional cost. The action is chosen before any of these later outcomes are seen. Rows of C_a that are identically zero are removed. Let $p_0 > 0$ be a reference law, and put $r_a = C_a p_0 > 0$.

For $q \in \Delta(\mathcal{A})$, vectors v_a indexed by action outputs, and an intercept ζ, define

$$\tilde{g} = \zeta \mathbf{1} + \sum_a C_a^\top v_a, \qquad \mathcal{E}_0(q, v) = \sum_{a,z} r_a(z)\, \psi(v_a(z), q_a), \tag{14}$$

where $\psi(v, q) = v^2/q$ for $q > 0$, $\psi(0, 0) = 0$, and $\psi(v, 0) = +\infty$ for $v \neq 0$. The last convention is essential. An action selected with probability zero cannot carry a nonzero reconstruction weight.

Proof of Theorem 4. For positive q, the epigraph of ψ is given by $qt \geq v^2$, $q, t \geq 0$, equivalently by the positive semidefiniteness of $\left(\begin{smallmatrix} q & v \\ v & t \end{smallmatrix}\right)$. The closed extension at zero has the same epigraph. It is a convex cone, so ψ is convex. Equivalently, convexity follows from the weighted Cauchy–Schwarz inequality. The energy in (14) is therefore jointly convex in (q, v). The reconstruction residual constraints are affine inequalities in (ζ, v, ϵ), and ϵ^2 is convex on $\epsilon \geq 0$. All simplex, budget, and epigraph constraints are convex.

For any finite-energy feasible point, $v_a = 0$ when $q_a = 0$. On a selected action define

$$W = \zeta + \frac{v_A(Z)}{q_A}.$$

For each full state u, conditional randomization and the channel definition give

$$\mathbb{E}[W \mid u] = \zeta + \sum_{a: q_a > 0} q_a \sum_z C_a(z \mid u) \frac{v_a(z)}{q_a} = \tilde{g}(u).$$

Thus $|\mathbb{E}_p W - \mathbb{E}_p g| \leq \|\tilde{g} - g\|_\infty \leq \epsilon$ under every p, without a reference-law assumption. If $p(u) \leq \kappa p_0(u)$ for all u, positivity of the channel implies $C_a p \leq \kappa C_a p_0$ componentwise. Hence

$$\begin{aligned} \mathrm{Var}_p(W) &\leq \mathbb{E}_p[(W - \zeta)^2] \\ &= \sum_{a: q_a > 0} \sum_z (C_a p)_z \frac{v_a(z)^2}{q_a} \leq \kappa \mathcal{E}_0(q, v) \leq \kappa t. \end{aligned}$$

For n independent records and independent action draws, the average score has variance $\mathrm{Var}_p(W)/n$. Its mean squared error is variance plus squared bias, proving

$$\mathbb{E}_p[(\widehat{R} - R_p)^2] \leq \epsilon^2 + \kappa t/n.$$

This is an upper bound for the returned estimator, not an assertion of minimax optimality over all nonlinear estimators or exact optimality for every p in the envelope.

Finally, suppose a zero-cost action is available. The constant reconstruction $\zeta = 0$, $v = 0$, $\epsilon = \|g\|_\infty$, $t = 0$ is feasible at any nonnegative budget. A finite objective sublevel bounds ϵ and t. Since $q_a \leq 1$ and every retained $r_a(z) > 0$, it also bounds every $v_a(z)$ by the energy constraint. The residual constraints then bound ζ. The simplex bounds q, and the closed perspective epigraph makes the sublevel closed. In finite dimensions it is compact, so the minimum is attained. This also explains why allowing approximation remains feasible when the exact target is not identified by the menu. □

C.2 Several targets and distributional uncertainty

For a finite query family, give each query its own $\zeta_j, v_j, \epsilon_j, t_j$ but use one q. Minimize T subject to

$$\epsilon_j^2 + \kappa t_j/n \leq T \quad (j \in [J])$$

and the corresponding reconstruction and energy constraints. This remains convex and bounds the largest of the individual mean squared errors. It does not bound the expectation of the maximum squared error without an additional argument.

The density-ratio assumption can instead be replaced by a specified convex set $\mathcal{U}$ of completion laws. Impose

$$\sum_{a,z} (C_a p)_z \psi(v_a(z), q_a) \leq t \quad \text{for every } p \in \mathcal{U},$$

and minimize $\epsilon^2 + t/n$. The preceding proof holds for every $p \in \mathcal{U}$. If $\mathcal{U}$ is the convex hull of a finite list of laws, checking its vertices suffices because the energy is linear in p. If $\mathcal{U} = \Delta_N$, the constraints become

$$\sum_{a,z} C_a(z \mid u)\psi(v_a(z), q_a) \le t \quad (u \in [N]). \tag{15}$$

This supplies an all-laws bound, typically more conservative than a reference envelope. Merely fitting p_0 does not establish a useful value of κ or membership in a smaller uncertainty set.

C.3 Multiple pre-audit strata and timing

For the global independent-record experiment, write a full state as (o, u), where u includes the eventual outcome. Use $q_a(o)$ for every completion and outcome in that pre-audit stratum. Define a global reference $P_0(o, u) = \rho^0_o p_{0,o}(u)$, use a single intercept ζ, and replace the energy by

$$\sum_o \rho^0_o \sum_{a,z} (C_{a,o} p_{0,o})_z \psi(v_{a,o}(z), q_a(o)).$$

Impose the pointwise residual constraint separately at every (o, u). A global density envelope $P \le \kappa P_0$ gives the same proof and MSE bound. One can enforce $\sum_a q_a(o)c(a, o) \le b$ for each o, which implies the expected cost constraint under any distribution of O. Alternatively, when the relevant stratum weights are specified, a weighted average budget can be used. A budget constraint based on estimated weights alone is not a guarantee under arbitrary different true stratum frequencies.

The current outcome enters the channel output, not the action label. This uses the same $q_a(o)$ across outcome completions without revealing the outcome to acquisition. The conditional-cell reference formulas for exact reconstruction express the same timing constraint.

C.4 The implemented Walsh program

Let the reference be uniform on a binary completion cube and let $\mathcal{T}$ contain all nonconstant characters observable under a chosen action menu. Every possible reconstructed query belongs to the span of the empty character and $\mathcal{T}$. Write it as $\tilde{g} = \beta_0 + \sum_{S \in \mathcal{T}} \beta_S \chi_S$. The minimum-reference-variance score for this query is the exact inclusion score with α replaced by β. Its reference mean is β_0 and its centered energy is

$$\sum_{S \in \mathcal{T}} \frac{\beta_S^2}{p_S(q)}.$$

Consequently, the uniform-reference version of the joint program is exactly (10), not a heuristic coefficient threshold. To verify equivalence, start with any feasible general reconstruction. Its conditional mean $\tilde{g}$ has a Walsh representation on the observable span. Replacing its score by the minimum-reference-variance reconstruction and choosing intercept β_0 leaves $\tilde{g}$ and its residual unchanged while weakly decreasing the energy. Conversely, every feasible Walsh point yields a feasible channel point through $v_a = q_a(h_a - \beta_0)$.

For the degree-L experiments, the menu consists of the empty action and every subset of size at most L, so $\mathcal{T}$ is precisely the nonempty subsets of size at most L. There are $\sum_{k=0}^{L} \binom{d}{k}$ actions and 2^d completion constraints. The numerical routine imposes the explicit coverage floor $p_S \ge 10^{-7}$ to evaluate smooth derivatives. It therefore solves a slightly restricted convex program. This floor is disclosed in the configuration and is not an assertion that an unsupported coefficient is exactly zero. The reported residual is recomputed directly on every cube state after optimization, independently of the solver's epigraph variable.

For $g = \alpha_0 + a\chi_K$ and a fixed coverage $p > 0$, consider the one-character score with coefficient β. Under $\kappa = 1$ the objective is

$$(a - \beta)^2 + \frac{\beta^2}{np}.$$

It is minimized at $\beta = anp/(1 + np)$, with value $a^2/(1 + np)$. This calculation distinguishes finite-sample error control from exact support recovery. The program shrinks a poorly measured interaction and accounts for the resulting bias rather than declaring the interaction absent.

C.5 Concentration with explicit approximation error

For any fixed feasible score, suppose $|W - \zeta| \leq B$ on all possible action outputs. Hoeffding's inequality and the preceding bias bound give, with probability at least $1 - \delta$,

$$|\widehat{R} - R| \leq \epsilon + B\sqrt{\frac{2\log(2/\delta)}{n}}. \tag{16}$$

The finite maximum B can be computed from the chosen score. It can be large when a design has rare actions, so (16) may be conservative. For a finite class of $J|\mathcal{F}|$ queries, replace $\log(2/\delta)$ by $\log(2J|\mathcal{F}|/\delta)$ and take the largest ϵ and B. A union bound then gives simultaneous approximate-risk control, and the empirical-risk minimization argument adds at most twice this bound to the oracle risk. For single-policy contrasts, the same argument can be applied directly to the contrast queries. For an independently chosen pilot design, all these arguments apply conditionally on the pilot.

D Smooth policies can be approximately low order

Exact Walsh sparsity is not guaranteed by a small number of logistic parameters. Proposition 5 instead concerns approximation in the uniform norm. Its analytic approximation ingredient is standard (Trefethen, 2013); the acquisition consequence follows by reducing powers of binary variables.

Proof of Proposition 5. If $B = \|\beta\|_1 = 0$, the gate is constant and the claim is immediate. Otherwise define $F(t) = \sigma(a + Bt)$ on $[-1, 1]$. Let E_ρ be the Bernstein ellipse parameterized by $t = (w + w^{-1})/2$ with $|w| = \rho$. The largest absolute imaginary part of $a + Bt$ on and inside this ellipse is $B(\rho - \rho^{-1})/2 < \pi/2$. The sigmoid has no pole there. Moreover, writing $z = s + it$ gives

$$|1 + e^{-z}|^2 = 1 + e^{-2s} + 2e^{-s}\cos t \geq 1 \quad \text{when } |t| \leq \pi/2.$$

Thus $|\sigma(z)| \leq 1$ throughout that ellipse.

For completeness, put $G(w) = F((w + w^{-1})/2)$. It is analytic in the annulus $\rho^{-1} < |w| < \rho$ and invariant under $w \mapsto w^{-1}$. Its Laurent coefficients therefore have matching positive and negative indices. Cauchy's formula on a circle approaching radius ρ bounds the coefficient of w^k by ρ^{-k}. Since $T_k((w + w^{-1})/2) = (w^k + w^{-k})/2$, the Chebyshev coefficient of degree $k \geq 1$ is bounded by $2\rho^{-k}$. Truncation at degree L gives a polynomial P_L satisfying

$$\sup_{|t|\leq 1} |F(t) - P_L(t)| \leq \sum_{k=L+1}^{\infty} 2\rho^{-k} = \frac{2\rho^{-L}}{\rho - 1}.$$

For every cube point, $t = \beta^\top x/B$ lies in $[-1, 1]$. The composite polynomial $P_L(\beta^\top x/B)$ has ordinary total degree at most L. Reducing powers with $x_i^2 = 1$ cannot increase the degree, so its Walsh expansion has no term of order exceeding L.

For two mask branches with losses $L_0(x), L_1(x) \in [0, 1]$ depending only on a core of at most s coordinates, write the exact loss as

$$g = L_0 + \sigma(a + \beta^\top x)(L_1 - L_0).$$

Replacing the sigmoid by the polynomial changes g by at most the same uniform error, since $|L_1 - L_0| \leq 1$. Each branch loss is a function on an s-coordinate cube and therefore has Walsh degree at most s. The product has degree at most $s + L$. Intersecting with the ambient degree bound gives $\min\{d, s + L\}$. This proves the stated extension without assuming a polynomial predictor on its own observed coordinates. □

The bound guarantees existence of a polynomial approximation. It does not bound the coefficient tail of a truncated Walsh expansion, which is a different approximation. The numerical program optimizes retained coefficients and checks the residual directly, without using the conservative analytic constant. As B grows, the permitted ρ approaches one and the guarantee deteriorates. The limitation is statistical as well as computational. The logistic experiment reports both truncation-tail diagnostics and optimized residuals.

E PILOT AUDITS, REFERENCE ERROR, AND POLICY ERROR

E.1 CONDITIONAL VALIDITY AFTER A PAID PILOT

Let $\mathcal{D}_0$ consist of pilot records fully recovered by randomized passive audits. A reference law, finite query family, and feasible action distribution can be arbitrary measurable functions of this pilot. All of them are frozen before an independent analysis sample is inspected. Conditional on $\mathcal{D}_0$, the reference and design are fixed, and every exact reconstructor satisfies

$$\mathbb{E}[W \mid \mathcal{D}_0, H, X, Y, M] = g(H, X, Y).$$

This identity follows pointwise and does not require a consistent reference estimator. It establishes conditional and unconditional unbiasedness for fixed queries, but it does not make the training risk of an adaptively chosen predictor unbiased. Simultaneous finite-class bounds or held-out prediction evaluation are still needed.

Let there be N_{tot} candidate records, of which n_0 are reserved for the pilot. Record the realized pilot measurement cost C_0, and let the total expected acquisition allowance be B_{tot}. When $C_0 \le B_{\mathrm{tot}}$, impose on each of the remaining $n = N_{\mathrm{tot}} - n_0$ records the conditional expected cost cap

$$b_{\mathrm{main}} = \frac{B_{\mathrm{tot}} - C_0}{n}. \tag{17}$$

Then $\mathbb{E}[C_{\mathrm{total}} \mid \mathcal{D}_0, \text{pre-audit records}] \le C_0 + n b_{\mathrm{main}} = B_{\mathrm{tot}}$. This is an expected-cost guarantee, not a deterministic bound on realized randomized spending. To ensure $C_0 \le B_{\mathrm{tot}}$ beforehand, reserve at most $\lfloor \xi B_{\mathrm{tot}}/c_{\max} \rfloor$ pilot records, where $c_{\max}$ bounds full recovery cost and $0 < \xi < 1$. The experiment uses $\xi = 0.2$, and does not reuse pilot rows for the main table estimate. Baselines receive and pay for the same pilot.

E.2 A CONDITIONAL ORACLE COMPARISON

For one finite experiment, let $\mathcal{H}(q)$ be the pointwise-unbiased reconstructors of a fixed query, and define $V_p(q,h) = \mathrm{Var}_{p,q}(h)$. Consider the same feasible design class under p and p_0; in particular, per-stratum budget caps cannot change with an unobserved completion law.

Proposition 6 (Reference distortion)**.** *Suppose $(1-\eta)p_0 \le p \le (1+\eta)p_0$ coordinatewise for $0 \le \eta < 1$. If (q_0, h_0) minimizes reference variance and (q_*, h_*) minimizes true variance over the same feasible pointwise-unbiased pairs, then*

$$V_p(q_0, h_0) \le \frac{1+\eta}{1-\eta} V_p(q_*, h_*). \tag{18}$$

The statement also holds for nonnegative weighted sums or maxima of the query variances. If the reference solution is within additive τ of the reference optimum, the right-hand side acquires the term $(1+\eta)\tau$.

Proof. For any real constant c and fixed (q, h), nonnegativity of the squared deviation and the density bounds imply

$$(1-\eta)\mathbb{E}_{p_0,q}(h-c)^2 \le \mathbb{E}_{p,q}(h-c)^2 \le (1+\eta)\mathbb{E}_{p_0,q}(h-c)^2.$$

Minimizing over c gives $(1-\eta)V_{p_0}(q,h) \le V_p(q,h) \le (1+\eta)V_{p_0}(q,h)$. Therefore

$$\begin{aligned} V_p(q_0,h_0) &\le (1+\eta)V_{p_0}(q_0,h_0) \\ &\le (1+\eta)\{V_{p_0}(q_*,h_*) + \tau\} \\ &\le \frac{1+\eta}{1-\eta} V_p(q_*,h_*) + (1+\eta)\tau. \end{aligned}$$

Nonnegative sums and maxima preserve both comparison inequalities. The proof also covers a numerically near-optimal solution if a valid suboptimality bound τ is available; a floating-point gap is a diagnostic rather than an interval-arithmetic bound. □

The proposition does not guarantee that a small pilot attains small η. Multiplicative error can be large in a sparsely populated completion space despite small average error. We report pilot size, reference ℓ_1 error, and oracle-design comparisons separately; a fit diagnostic alone gives no population precision guarantee.

E.3 A DESIGN-VALID AUGMENTATION BASELINE

Let $S(o)$ be a sufficient missing-feature scope for a query and $p_S(o) > 0$. Let $b(O, Y)$ be any function estimated on an independent pilot and measurable from the native record and eventual outcome. Then

$$W^{\mathrm{aug}} = b(O, Y) + \frac{\mathbf{1}\{S(O) \subseteq A\}}{p_{S(O)}(O)}[g(H, X, Y) - b(O, Y)] \tag{19}$$

is pointwise unbiased. Indeed, conditional on the complete record, the bracket and the baseline are fixed and the mean inclusion weight is one. This is the familiar augmented two-phase principle (Levis et al., 2024; Gilbert et al., 2014), instantiated here for a sufficient partial scope. It remains valid if the imputation model is wrong. Unlike the channel reconstructor, it need not use every lower-order part of a partially acquired scope.

In the paid-pilot study, b is the fitted conditional query mean given the eventual outcome and any native feature value. The four estimators receive identical realized measurements for a given acquisition policy. Signed reconstructed tables are projected onto the probability simplex before fitting a Brier predictor. Projection stabilizes finite-sample estimates but does not make the predictor unbiased. If the unprojected table converges, continuity of Euclidean projection and vanishing pseudocounts give consistency at each positive-mass conditional prediction. This is the argument of Proposition 16, with projection before normalization.

E.4 MISSPECIFIED TARGET POLICIES

Let π be the actual target policy and $\tilde{\pi}$ the specified policy used by the estimator. For a fixed predictor and a loss in $[0, 1]$, define

$$\delta_\pi = \sup_{h,x} \mathrm{TV}\{\pi(\cdot \mid h, x), \tilde{\pi}(\cdot \mid h, x)\}.$$

The variational characterization of total variation gives $|g_\pi - g_{\tilde{\pi}}| \le \delta_\pi$ pointwise. Thus an approximation with residual ϵ relative to $g_{\tilde{\pi}}$ has total worst-case bias at most $\epsilon + \delta_\pi$ for the actual policy. To verify the TV step directly, split the signed difference of the two probability vectors into its positive and negative parts. Each has total mass TV, and integrating a function between zero and one changes the expectation by at most that mass. Averaging over (H, X, Y) proves the risk bound. Additional audits remove neither an incorrect target kernel nor its induced estimand error.

F CHOOSING AN APPROXIMATION BEFORE THE ANALYSIS SAMPLE

The approximation degree must be chosen without the evaluation law. This appendix gives two rules. One minimizes an explicit error envelope without validation observations. The other estimates actual mean-squared error from an independently acquired validation block. The latter uses standard distinct-observation products (Hoeffding, 1948). The contribution here is a measurable selection procedure with its acquisition cost charged, rather than a new unbiased-estimation principle.

F.1 ENVELOPE SELECTION AND A USER-SPECIFIED BIAS TOLERANCE

Condition on any independent construction pilot. Let a finite candidate menu contain feasible pairs (q_j, W_j) for the same scalar target g, with residual bounds ϵ_j, reference energies t_j, and a common future sample size n. Assume $p \le \kappa p_0$ for the reference used in each stated bound, and put $B_j = \epsilon_j^2 + \kappa t_j / n$. Choosing $j_B \in \arg\min_j B_j$ uses only known targets, references, costs, and the declared envelope. Theorem 4 gives

$$\mathbb{E}[(\widehat{R}_{j_B} - p^\top g)^2] \le \min_j B_j.$$

This is a comparison of upper bounds. It does not assert $\mathrm{MSE}_{j_B} \le \mathrm{MSE}_j$ for each actual completion law. An investigator who requires a bias tolerance $\epsilon_\star$ can instead restrict the candidate set to $\epsilon_j \le \epsilon_\star$ and minimize B_j there. An empty restricted set means that the requested tolerance is unsupported by that menu and budget. The tolerance and envelope must be specified independently of the hidden evaluation law.

For any selection index measurable from independent preliminary data, the fresh analysis sample still has conditional MSE at most the selected B_j, when its envelope assumption holds. Averaging over the preliminary data gives an unconditional bound by the expectation of the selected envelope, not necessarily its minimum. A fitted reference does not establish the density envelope on its own.

F.2 An observable estimate of candidate MSE

Let U denote the complete finite record in one stratum. A candidate action distribution q_j and score W_j are fixed before a validation sample is acquired. Channel randomness, including randomized outputs when present, is averaged in

$$\begin{aligned} h_j(u) &= \mathbb{E}[W_j \mid U = u], & m_j(u) &= h_j(u) - \zeta_j, \\ a_j(u) &= \mathbb{E}[(W_j - \zeta_j)^2 \mid U = u], & e_j(u) &= h_j(u) - g(u). \end{aligned} \tag{20}$$

Here ζ_j is any fixed centering constant, conveniently the reconstruction intercept. These are known functions of a *fully recovered* validation record. Evaluating the finite channel or target kernel at that record does not require the unknown distribution p.

For $v \geq 2$ independent complete validation records, define the cross-product average

$$\mathcal{U}_v(z) = \frac{(\sum_{i=1}^v z_i)^2 - \sum_{i=1}^v z_i^2}{v(v-1)} = \frac{1}{v(v-1)} \sum_{i \neq k} z_i z_k.$$

The criterion implemented by the validation selector is

$$\widehat{M}_j = \mathcal{U}_v(e_j) + \frac{1}{n} \left\{ \frac{1}{v} \sum_{i=1}^v a_j(U_i) - \mathcal{U}_v(m_j) \right\}. \tag{21}$$

It can be negative. Clipping it is not part of the unbiased construction, and the implementation retains negative values.

Proposition 7 (Paid validation of a fixed candidate menu)**.** *Let the J candidate scores be fixed independently of the validation records and of a subsequent n-record analysis sample. For every completion law p, $\widehat{M}_j$ in* (21) *is unbiased for the fresh-sample MSE M_j of candidate j. Suppose $\|e_j\|_\infty \leq \epsilon_j$, $\|m_j\|_\infty \leq L_j$, and $0 \leq a_j \leq A_j$. Set*

$$D_j = 2\epsilon_j^2 + \frac{A_j + 2L_j^2}{n}, \qquad r_j = D_j \sqrt{\frac{\log(2J/\delta)}{2\lfloor v/2 \rfloor}}. \tag{22}$$

With probability at least $1 - \delta$, simultaneously $|\widehat{M}_j - M_j| \leq r_j$. Consequently, the raw selector $\widehat{j} \in \arg\min_j \widehat{M}_j$ satisfies

$$M_{\widehat{j}} \leq \min_j M_j + 2 \max_j r_j.$$

The upper-confidence selector $\widehat{j}_+ \in \arg\min_j (\widehat{M}_j + r_j)$ satisfies the sharper candidate-dependent comparison

$$M_{\widehat{j}_+} \leq \min_j (M_j + 2r_j).$$

These statements concern fresh-sample conditional MSE, not the realized squared error of the selected estimate. They do not require a reference density envelope.

Proof. Independence of two validation records gives $\mathbb{E}\mathcal{U}_v(e_j) = (\mathbb{E}e_j)^2$ and $\mathbb{E}\mathcal{U}_v(m_j) = (\mathbb{E}m_j)^2$. Also $\mathbb{E}a_j = \mathbb{E}(W_j - \zeta_j)^2$. Therefore

$$\mathbb{E}\widehat{M}_j = (\mathbb{E}h_j - \mathbb{E}g)^2 + \frac{\mathbb{E}(W_j - \zeta_j)^2 - (\mathbb{E}W_j - \zeta_j)^2}{n} = (\mathbb{E}W_j - \mathbb{E}g)^2 + \frac{\mathrm{Var}(W_j)}{n} = M_j.$$

Centering by ζ_j changes neither the identity nor the variance and avoids subtracting unnecessary constant terms numerically.

For the deviation bound, write (21) as the order-two U-statistic with symmetric kernel

$$K_j(u, u') = e_j(u)e_j(u') + \frac{a_j(u) + a_j(u')}{2n} - \frac{m_j(u)m_j(u')}{n}.$$

The kernel lies in

$$\left[-\epsilon_j^2 - L_j^2/n,\ \epsilon_j^2 + (A_j + L_j^2)/n\right],$$

a deterministic interval of width D_j. Put $s = \lfloor v/2 \rfloor$. For each permutation of the validation indices, average K_j over its first s disjoint pairs. Averaging these pair averages over all permutations gives the U-statistic, including when one record is unused in each odd-sized permutation. For a fixed permutation the s pair evaluations are independent.

We include the bounded-variable exponential bound. For a variable Z in an interval of width D, let $\psi(t) = \log \mathbb{E}e^{tZ}$. The tilted law is supported on the same interval, so $\psi''(t) = \mathrm{Var}_t(Z) \le D^2/4$; this last inequality follows by centering at the midpoint and then minimizing the squared deviation over constants. Twice integrating, with $\psi(0) = 0$, yields $\log \mathbb{E}e^{t(Z-\mathbb{E}Z)} \le t^2 D^2/8$. Applying this to the independent pair average bounds its centered log moment-generating function by $t^2 D_j^2/(8s)$. Jensen's inequality over permutations gives the same bound for $\widehat{M}_j - M_j$. Exponential Markov and optimization over t give

$$\mathbb{P}(|\widehat{M}_j - M_j| > r) \le 2\exp(-2sr^2/D_j^2).$$

The case $D_j = 0$ is constant and immediate. A union bound proves (22). On that simultaneous event, comparison with an MSE minimizer $j_\star$ gives

$$M_{\widehat{j}} \le \widehat{M}_{\widehat{j}} + r_{\widehat{j}} \le \widehat{M}_{j_\star} + r_{\widehat{j}} \le M_{j_\star} + r_{j_\star} + r_{\widehat{j}}.$$

For the upper-confidence selector, comparison with any j instead gives $M_{\widehat{j}_+} \le \widehat{M}_{\widehat{j}_+} + r_{\widehat{j}_+} \le \widehat{M}_j + r_j \le M_j + 2r_j$. □

Budget and information accounting. Reserve v validation records before observing their missing values. If full recovery costs at most c_{full} per record, reserve vc_{full} and construct candidates using the remaining allowance $B_{\mathrm{tot}} - vc_{\mathrm{full}}$ and $n = N_{\mathrm{tot}} - v$. Every candidate must obey this same remaining expected cap. Validation rows choose a candidate but are not reused in its analysis estimate. Thus

$$\mathbb{E}C_{\mathrm{total}} \le vc_{\mathrm{full}} + nb_{\mathrm{main}} \le B_{\mathrm{tot}}.$$

When actual validation costs vary, the candidate budget here is based on the *reservation*, not a data-dependent refund. Otherwise the candidates could depend on the validation data, invalidating the proposition as stated. A distinct construction pilot may be used, with its cost also subtracted and all assertions conditioned on that pilot. References or candidates must not be refitted on the validation block used in (21). The logistic experiment has constant full-recovery cost d, so reserved and realized validation costs coincide.

Choosing the degree by minimum simulated population MSE would require an unavailable law. Here, candidate query values and score moments use only acquired validation rows, and the population law is used afterward to evaluate the selected design. A small validation block can still select poorly, and its uniform radius can exceed the actual MSE by a large margin. The experiments report both outcomes, which are compatible with unbiasedness.

G Bias-matched controls and numerical optimality diagnostics

G.1 Full/no-audit reconstruction

For uniform reference, let the full-audit probability be $r = \min\{b/d, 1\}$. The full/no-audit control uses $W = c$ without an audit and $W = c + (\widetilde{g}(U) - c)/r$ with a full audit. The reference-energy program reduces to

$$\min_{c,\epsilon \ge 0,\widetilde{g}} \epsilon^2 + \frac{\kappa}{nr}\mathbb{E}_0(\widetilde{g} - c)^2 \quad \text{subject to } |\widetilde{g}(u) - g(u)| \le \epsilon.$$

For fixed c, ϵ, the optimum in each coordinate is the point of $[g(u) - \epsilon, g(u) + \epsilon]$ nearest c. Hence

$$\widetilde{g}(u) = c + \mathrm{sign}(g(u) - c)(|g(u) - c| - \epsilon)_+.$$

The remaining convex objective is

$$F(c, \epsilon) = \epsilon^2 + \frac{\kappa}{nr}\mathbb{E}_0(|g - c| - \epsilon)_+^2.$$

The squared hinge has a continuous derivative, so direct differentiation gives the two derivatives used in the implementation. A minimizing c lies between $\min g$ and $\max g$; moving an outside value toward that interval cannot increase any hinge term. It also makes the uniform mean of $\widetilde{g} - c$ zero by the first-order condition. Thus its energy equals the minimum-reference-variance expression for the same reconstructed query. Increasing r weakly decreases this objective, so using the largest feasible full-audit probability is valid. The control has the same permission to approximate as the partial-audit program.

For the fixed degree-L control, an audit occurs with probability $\min\{b/L, 1\}$ and selects a uniformly random L-subset. The remaining probability selects the empty set. The retained dictionary contains all characters of degree at most L. We optimize the intercept, coefficients, and uniform residual at these fixed coverages. This controls for improved reconstruction while leaving the acquisition distribution unchanged.

G.2 A DUAL LOWER BOUND FOR THE APPROXIMATION OBJECTIVE

Write $\Phi_{us} = \chi_s(u)$, $P_{sa} = \mathbf{1}\{s \subseteq a\}$, $\lambda = \kappa/n$, and $\mathcal{Q} = \{q \geq 0 \colon \mathbf{1}^\top q = 1, c^\top q \leq b\}$. For any $w \in \mathbb{R}^N$ with $\mathbf{1}^\top w = 0$, put $s = \Phi^\top w$ and

$$E(w) = \max_{q \in \mathcal{Q}} \sum_s \frac{(Pq)_s s_s^2}{4\lambda}.$$

Then the following is a lower bound on the optimum of the closed-perspective program, even without an imposed coverage floor,

$$\mathcal{L}(w) = -w^\top g - \frac{\|w\|_1^2}{4} - E(w). \tag{23}$$

To prove this directly, any feasible residual $r = c_0\mathbf{1} + \Phi\beta - g$ satisfies $w^\top r \leq \|w\|_1 \epsilon$. Consequently

$$\begin{aligned}\epsilon^2 + \lambda \sum_s \frac{\beta_s^2}{p_s} &\geq \epsilon^2 - \|w\|_1 \epsilon + w^\top(\Phi\beta - g) + \lambda \sum_s \frac{\beta_s^2}{p_s} \\ &\geq -\frac{\|w\|_1^2}{4} - w^\top g - \sum_s \frac{p_s s_s^2}{4\lambda} \geq \mathcal{L}(w).\end{aligned}$$

The middle step completes the square in ϵ and each β_s. At $p_s = 0$, feasibility forces $\beta_s = 0$, and that term satisfies the same inequality by the closed convention. This is weak duality and needs no claim of strict feasibility or strong duality.

For a candidate w, replacing it by tw, $t \geq 0$, gives $tA - t^2 D$, where $A = -w^\top g$ and $D = \|w\|_1^2/4 + E(w)$. If $D > 0$, its maximizing scale is $t = A_+/(2D)$. If $w = 0$, the bound is zero. We construct w from the residual-constraint multipliers, subtracts its mean to enforce the constant constraint, and recomputes this lower bound. The linear maximum defining E occurs at a vertex with at most two nonzero acquisition probabilities, so the simplex-with-cost oracle evaluates it by enumerating affordable actions and affordable/expensive pairs. No evaluation distribution is used.

A second diagnostic linearizes the convex primal objective and minimizes that affine function over the feasible polyhedron. To keep this linear oracle bounded, use the constant reconstruction to obtain a feasible objective $U = (\max g - \min g)^2/4$. Every optimum satisfies

$$\epsilon \leq \sqrt{U}, \quad |c_0 - \mathbb{E}_0 g| \leq \sqrt{U}, \quad |\beta_s - \alpha_s| \leq \sqrt{U}.$$

The latter two follow by averaging the residual against the constant and each character. These boxes contain an optimum and may be added to the linear oracle without invalidating its lower bound. The reported diagnostic takes the stronger available lower bound. A dual or supporting-plane gap is a floating-point check, not an outward-rounded numerical certificate. The theorem's MSE bound depends on the returned residual and energy, not on the solver's optimality status.

G.3 TOLERANCE-QUALIFIED FINITE-CHANNEL SCORES

The exact range theorem is a real-arithmetic statement. The public numerical routine separates eigenvalue resolution, relative range acceptance, and the residual of the score actually returned. The

relative rank tolerance is 10^{-12}, and the relative range tolerance is also 10^{-12}. The eigenvalue cutoff is the rank tolerance times $\max\{1, \lambda_{\max}(I)\}$. A target's discarded tangent norm is compared with the range tolerance times its tangent norm, with the smallest positive normal float added only to handle the zero target.

After constructing a score, the routine computes every coordinate of

$$r_j(u) = \sum_a q_a \mathbb{E}[W_j \mid u, a] - g_j(u).$$

The quantity $\max_u |r_j(u)|$ is the numerical bias allowance for the returned score. The supplement provides the residuals and tolerances. This allowance is computed in floating point and is not a rigorous enclosure of rounding error.

In the two-coordinate regression case, the menu is $\{\emptyset, \{1\}\}$, $q = (1/2, 1/2)$, the reference is uniform, and $g = 1/2 + x_1/5 + 10^{-10}x_2$. The default range contract rejects the unsupported component. Relaxing the relative range tolerance to 10^{-8} permits a score only under that declared tolerance and reports its approximately 10^{-10} pointwise residual. Under the positive law $p(x) = (1 + x_2/2)/4$, this component contributes approximately -5×10^{-11} bias. No finite tolerance is presented as a proof that every smaller unsupported coefficient is zero.

H Additional exact-design and learning statements

H.1 Exact reference variance and a convex allocation

Within the stated model, unbiasedness is distribution-free, whereas variance optimality depends on a reference. Specify a law P_0 that retains cell masses ρ_{oy} and makes missing coordinates uniform conditional on (O, Y). Put

$$V_j^0(q) = \mathrm{Var}_\rho(\alpha_{j,O,Y,\emptyset}) + \sum_{o,y} \rho_{oy} \sum_{\emptyset \neq S \subseteq \overline{m}} \frac{\alpha^2_{j,o,y,S}}{p_S(o)}. \tag{24}$$

Zero coefficients contribute zero, including at zero coverage.

Theorem 8 (Reference-optimal reconstruction and allocation)**.** *Under P_0, $\mathrm{Var}(W_j) = V_j^0(q)$. Among all single-record scores whose audit expectation equals g_j at every complete record, the score in* (7) *has minimum reference variance. The multi-target design problem*

$$\min_q \max_{j \in [J]} V_j^0(q) \quad \textit{subject to} \quad q(\cdot \mid o) \in \Delta(\mathcal{A}_o), \quad \sum_o \rho_o \sum_a q(a \mid o) c(a, o) \leq b \tag{25}$$

is convex, with value $+\infty$ when a required interaction is unsupported.

Orthogonality makes the reference second moment diagonal in the characters. More generally, any pointwise-unbiased action-specific score must allocate coefficient α_S across the actions containing S. Cauchy–Schwarz bounds its contribution below by α_S^2/p_S, attained by (7). In (25), each p_S is affine in q and the inverse is convex. The cell masses and variance model are specified reference inputs or pilot estimates. The action still cannot depend on the current outcome.

A nonuniform reference distribution admits an analogous minimum-variance linear reconstruction through a finite observation-channel information matrix. Appendix M derives the score, its all-laws unbiasedness, and its convex design criterion. Neither reference choice asserts variance optimality when that reference is misspecified.

H.2 Reducing the acquisition menu

Let $\mathcal{S}_o$ contain the nonempty interactions required by any target and outcome in cell o, and let $r_o = |\mathcal{S}_o|$. Under monotone costs, an audit a can be replaced by

$$a^\circ = \bigcup \{S \in \mathcal{S}_o : S \subseteq a\} \tag{26}$$

without changing required coverages or increasing cost, provided these replacement bundles are feasible.

Proposition 9 (Menu reduction and sparse designs)**.** *Assume $c(a, o)$ is nondecreasing under set inclusion. If all unions of required interactions are feasible, an optimum of* (25) *can be sought over these unions. Any feasible design in cell o has another design supported on at most $r_o + 2$ actions with exactly the same required coverages and expected cost.*

The first assertion follows because a° contains exactly the required interactions in a. The second reduces a finite convex combination in $\mathbb{R}^{r_o+1}$. Appendix N gives the construction and a supporting-hyperplane optimality diagnostic. This does not yield a polynomial-time algorithm in ambient dimension: the interaction expansion and union menu can be exponential. The specified loss must provide the useful sparsity; a neural representation alone cannot recover unidentified directions.

I Sharp budgets and finite-class guarantees

I.1 A matched lower bound

For $K \subseteq [d]$, $|K| = k$, a mask-only predictor can induce the squared target loss

$$g(x) = \frac{1 + \gamma \chi_K(x)}{2}, \qquad 0 < \gamma \le 1. \tag{27}$$

For example, take $Y = 0$ and let the predictor output one or zero according to which of two masks the target policy chooses, with the first mask having the displayed probability.

Consider (27) and the family

$$P_\theta(x) = 2^{-d}(1 + \theta \chi_K(x)), \quad |\theta| \le \tfrac{1}{2}, \qquad R_\theta = \tfrac{1}{2} + \tfrac{\gamma}{2}\theta. \tag{28}$$

Here no features are observed natively, K is known, and each acquired coordinate costs one unit. An audit may depend on all past audited records, but its current bundle is selected before the current feature values are revealed.

Theorem 10 (Measurement-budget separation)**.** *For $n \ge 1$ and $0 < b \le k$, let $\mathfrak{R}_{n,b}$ be the minimax mean squared error over* (28)*, allowing all such audit strategies with expected total cost at most nb under every P_θ. Then*

$$\frac{\gamma^2}{256} \min\left\{1, \frac{k}{nb}\right\} \le \mathfrak{R}_{n,b} \le \frac{\gamma^2}{4} \min\left\{1, \frac{k}{nb}\right\}. \tag{29}$$

If only empty and full-record audits are allowed, the same bounds hold with k replaced by d. At $\theta = 0$, the unbiased estimators based on auditing K with probability b/k and $[d]$ with probability b/d have variances $\gamma^2 k/(4nb)$ and $\gamma^2 d/(4nb)$, respectively.

Only audits containing K distinguish the alternatives. Their expected count is at most nb/k, so transcript divergence is bounded by that count times the divergence of one parity observation. Two-point testing gives the lower bound, including history-adaptive audit probabilities. Appendix O proves the constants and matching upper bound. The variance comparison is against full-record audits, not every estimator, and K is specified rather than discovered from hidden data.

I.2 From risk identification to prediction learning

For a finite predictor class $\mathcal{F}$, apply (7) to every (j, f), using one audit design covering their union of required interactions. Define

$$B_q = \max_{j,f,o,y} \left(|\alpha_{j,f,o,y,\emptyset}| + \sum_{\substack{S \neq \emptyset \\ \alpha_{j,f,o,y,S} \neq 0}} \frac{|\alpha_{j,f,o,y,S}|}{p_S(o)} \right). \tag{30}$$

With probability at least $1 - \delta$, simultaneously over all pairs,

$$|\widehat{R}_j(f) - R_j(f)| \le \varepsilon_n := B_q \sqrt{\frac{2 \log(2J|\mathcal{F}|/\delta)}{n}}. \tag{31}$$

Consequently, $\hat{f} \in \arg\min_{f \in \mathcal{F}} \max_j \widehat{R}_j(f)$ satisfies

$$\max_j R_j(\hat{f}) \le \min_{f \in \mathcal{F}} \max_j R_j(f) + 2\varepsilon_n. \tag{32}$$

These conservative finite-sample bounds follow directly from bounded scores, not from a claim that the reference variance bounds all distributions. Their proof is in Appendix P. For finite observation and outcome spaces, estimating the target joint table instead identifies its Bayes predictor. The experiments use this route and also train small neural predictors on the same weighted tables. We make no uniform-generalization claim for an unrestricted neural class.

J Notation, timing, and model scope

Table 7 collects the notation. Main probability results assume independent records or episodes; coordinates within a record may be arbitrarily dependent. Audit expectations condition on all pre-existing variables, including hidden features and the later recorded outcome. Randomization ensures that this conditional audit distribution remains $q(\cdot \mid O)$.

Table 7: Notation used in the identification and design results.

Symbol	Meaning
H, X, Y	Context, recoverable features, subsequently recorded outcome.
M, O	Historical feature mask and record (H, M, X_M).
A, q, c, b	Audit bundle, randomization distribution, acquisition cost, expected budget.
π_j, f_j, g_j, R_j	Target observation kernel, predictor, complete-data target loss, its risk.
ρ_{oy}	Probability of the observed-record/outcome cell (o, y).
$u, \overline{m}$	Missing completion and its coordinate set in a fixed cell.
$\chi_S, \alpha_{j,o,y,S}$	Binary character and corresponding target-loss coefficient.
$p_S(o)$	Probability of jointly acquiring every coordinate in S.
$\mathcal{S}_o, \mathcal{A}_o$	Required interaction family and feasible audit menu.
C_a	Finite observation channel for audit a.
W_j, V_j^0	Pointwise-unbiased score and its variance under the uniform reference.
$P_0, p_0, I(q)$	Reference law, a conditional reference mass vector, tangent information matrix.

The pre-audit information does not contain Y. After an audit has been selected, Y is recorded even if no additional feature is acquired. This distinction permits outcome-dependent coefficients in a score without outcome-dependent randomization. The model allows an association between missingness and the outcome that is not explained by the recorded features. It does not interpret future outcomes as physically causing prior sensing decisions. A latent common cause is sufficient, and the semi-synthetic image experiment is explicitly a retrospective censoring construction.

The observation kernel at deployment is part of the estimand. If an unknown target kernel and a complete-data law can vary jointly without changing the source audit distribution, their target risks need not agree. Source-side audits alone therefore do not solve an unspecified target-policy problem. Likewise, if a requested value is physically unavailable or if obtaining it changes the outcome, the channel and the potential-outcome model must be changed before applying the results.

The full-support requirement is needed for the local necessity statement, not for the unbiased estimator. The estimator remains valid under an arbitrary distribution on the finite completion space whenever its pointwise identity holds on that space. If the support is known to be smaller, the finite-channel criterion below applies after restricting its columns to that support. Treating an unknown support restriction as known can incorrectly remove an unidentified direction.

K IDENTIFICATION FROM PARTIAL OBSERVATION CHANNELS

K.1 A FINITE-CHANNEL ESTIMABILITY LEMMA

Fix one cell (o, y) and enumerate its N possible completions. Let $p \in \mathbb{R}^N$ be their conditional mass vector. An audit a has an output alphabet $\mathcal{Z}_a$ and a known channel $C_a \in \mathbb{R}^{|\mathcal{Z}_a| \times N}$, where $(C_a)_{zu}$ is the probability of output z given completion u. Each column sums to one. Coordinate observation is the special case

$$(C_a)_{zu} = \mathbf{1}\{u_a = z\}. \tag{33}$$

Let $q_a = q(a \mid o)$. The joint output distribution, including the action label, is $q_a C_a p$. Define the stacked matrix K_q by stacking $q_a C_a$ over actions with $q_a > 0$. Action labels must be retained. Pooling the outputs of different actions can lose information and corresponds to a different channel.

Lemma 11 (Finite linear estimability)**.** *For a vector $g \in \mathbb{R}^N$, the functional $g^\top p$ is identified from $K_q p$ over the strictly positive probability simplex if and only if $g \in \mathrm{range}(K_q^\top)$. If g is outside this row space, every strictly positive p has arbitrarily close strictly positive alternatives with the same channel output and different values of $g^\top p$.*

Proof. If $g = K_q^\top h$, then $g^\top p = h^\top K_q p$, which is a function of the observed output distribution. Conversely, in finite-dimensional Euclidean space,

$$\mathrm{range}(K_q^\top) = (\ker K_q)^\perp.$$

If g is not in this range, there is $v \in \ker K_q$ with $g^\top v \neq 0$. Since each C_a is column stochastic and $\sum_a q_a = 1$, the columns of K_q also sum to one. Consequently,

$$\mathbf{1}^\top v = \mathbf{1}^\top K_q v = 0.$$

For strictly positive p, choose $\eta > 0$ so small that $p_u \pm \eta v_u > 0$ for all u. Both $p_+ = p + \eta v$ and $p_- = p - \eta v$ sum to one, and

$$K_q p_+ = K_q p_- = K_q p, \qquad g^\top p_+ - g^\top p_- = 2\eta g^\top v \neq 0.$$

The size of η can be made arbitrarily small. This proves both the necessity and the local statement. □

K.2 BINARY INTERACTIONS AND COORDINATE MARGINALS

For a binary completion cube of dimension D, the characters satisfy

$$2^{-D} \sum_u \chi_S(u)\chi_T(u) = \mathbf{1}\{S = T\}. \tag{34}$$

Indeed, their product is $\chi_{S \triangle T}$. If $S \neq T$, summing over any coordinate in $S \triangle T$ cancels the two signs. There are 2^D characters, so orthogonality also proves they form a basis of all functions on the cube and establishes (4).

For a coordinate audit a, every row of C_a is an indicator of the event $u_a = z$, hence is a function of u_a. Conversely, every function of u_a is a linear combination of these indicators. The row space is therefore the full space of functions of u_a. Its character basis is

$$\mathrm{row}(C_a) = \mathrm{span}\{\chi_S : S \subseteq a\}. \tag{35}$$

Scaling C_a by a positive q_a does not change this space. Stacking the channels takes the sum of their row spaces. By the linear independence of the character basis,

$$\mathrm{row}(K_q) = \mathrm{span}\{\chi_S : p_S(o) > 0\}. \tag{36}$$

Proof of Theorem 1. The audited-data law determines the cell masses ρ_{oy} because (O, Y) is observed, and it determines every $q_a C_a p_{oy}$ within each positive-mass cell. If all nonzero coefficients of each restricted loss correspond to characters in (36), Lemma 11 identifies each conditional mean. Their weighted sum with weights ρ_{oy} is R_j.

If the condition fails in cell (o, y) for target j, the restricted loss is not in that cell's row space. Lemma 11 supplies two arbitrarily close completion distributions with the same output distribution

but different conditional target means. Keep ρ_{oy}, all other completion distributions, and the audit kernel fixed. These choices define valid joint laws of (H, X, Y, M): assign mass $\rho_{oy} p_{oy}(u)$ to the complete record formed by o, y, u. The unrestricted model imposes no cross-cell constraint that would forbid this assignment. The audited-data laws agree, while R_j differs because $\rho_{oy} > 0$.

For the explicit character construction, let S be an uncovered interaction with $\alpha_{j,o,y,S} \neq 0$. Set $v_u = 2^{-D}\chi_S(u)$. Since the empty character is always covered, S is nonempty and $\sum_u v_u = 0$. Every supported action omits at least one coordinate in S, so summing v over the unobserved coordinates yields $C_a v = 0$. Finally,

$$g_{j,o,y}^{\top} v = \alpha_{j,o,y,S} \neq 0.$$

The positive perturbations $p \pm \eta v$ thus prove the stated conclusion directly. □

K.3 An indistinguishable pair with native complete records

The obstruction does not rely on having no complete historical records. Let $Y = 0$, let H be constant, and define, for $|\theta| < 1$,

$$\begin{aligned} \mathbb{P}_\theta(X = x, M = [d]) &= \tfrac{1}{2}\, 2^{-d}, \\ \mathbb{P}_\theta(X = x, M = \emptyset) &= \tfrac{1}{2}\, 2^{-d}(1 + \theta\chi_K(x)). \end{aligned} \tag{37}$$

These masses are positive and sum to one. Native complete records are uniform with total mass $1/2$, and the probability of an empty record is $1/2$, independently of θ. Thus the native observed law is identical for every θ. The marginal complete-data law is

$$\mathbb{P}_\theta(X = x) = 2^{-d}\big(1 + \tfrac{1}{2}\theta\chi_K(x)\big),$$

and the historical completeness probability is

$$\mathbb{P}_\theta(M = [d] \mid X = x) = \frac{1}{2 + \theta\chi_K(x)}.$$

Except at $\theta = 0$, the historical observation mechanism is nonignorable. Audits of an empty record that never contain K reveal only marginals independent of θ. Nevertheless, the target loss in (27) has risk $1/2 + \gamma\theta/4$. This is an explicit MNAR indistinguishable-pair construction with positive native completeness, separate from the no-native-measurement minimax family in Theorem 10.

L Unbiased estimation and exact reference optimality

Proof of Proposition 2. Fix (H, X, Y, M) and hence $O = o$. For each required interaction S, randomization gives

$$\mathbb{E}_A\left[\frac{\mathbf{1}\{S \subseteq A\}}{p_S(o)}\right] = 1.$$

All other quantities in that term of (7) are fixed. Summing the finitely many terms and applying (4) gives $\mathbb{E}_A W_j = g_j(H, X, Y)$. Taking expectation over the complete-data law proves $\mathbb{E}W_j = R_j$. Averaging independent scores preserves the expectation. The identity uses only q and never a model for $\mathbb{P}(M \mid H, X, Y)$. □

Proof of Theorem 8. First fix a cell and suppress (j, o, y) in the notation. Under its uniform reference completion law, nonconstant characters have mean zero. The conditional mean of the score is $\alpha_\emptyset$. For $S \neq T$, the product of their characters has zero reference mean, independently of which audit was drawn. For $S = T \neq \emptyset$, its square equals one, so

$$\mathbb{E}_0\left[\left(\frac{\mathbf{1}\{S \subseteq A\}\alpha_S\chi_S(U)}{p_S}\right)^2\right] = \frac{\alpha_S^2}{p_S}.$$

It follows that the conditional variance is $\sum_{S \neq \emptyset} \alpha_S^2 / p_S$. Applying the law of total variance over (O, Y) gives (24).

To prove optimality, consider any other action-specific score $h_a(u_a)$ whose conditional audit expectation equals $g(u)$ for every completion. An independently randomized score can be replaced by

its conditional mean given (a, u_a), without increasing variance, so deterministic functions suffice. Expand each such function in the characters it observes,

$$h_a(u_a) = \sum_{S \subseteq a} \beta_{a,S} \chi_S(u).$$

The pointwise reconstruction requirement and uniqueness of the character expansion imply

$$\sum_{a \supseteq S} q_a \beta_{a,S} = \alpha_S \quad \text{for every } S. \tag{38}$$

Its reference second moment, using orthogonality within each action, is

$$\mathbb{E}_0 h_A(U_A)^2 = \sum_S \sum_{a \supseteq S} q_a \beta_{a,S}^2.$$

For $p_S > 0$, Cauchy–Schwarz and (38) give

$$\alpha_S^2 \le \left(\sum_{a \supseteq S} q_a \right) \left(\sum_{a \supseteq S} q_a \beta_{a,S}^2 \right),$$

so the contribution for S is at least α_S^2 / p_S. If $p_S = 0$, the constraint requires $\alpha_S = 0$ and there is no contribution. Equality is attained by $\beta_{a,S} = \alpha_S / p_S$ on all positive-probability actions containing S. For the empty set, $p_\emptyset = 1$ and the minimizing constant coefficient is $\alpha_\emptyset$ in every action. These are exactly the coefficients of (7). All feasible scores have the same mean, so minimum second moment is equivalent to minimum variance. This argument applies separately in every cell and proves global reference optimality.

For convexity, $p_S(o)$ is a nonnegative affine function of q. The extended function $x \mapsto \alpha^2 / x$ is convex on $x > 0$ and equals $+\infty$ at zero when $\alpha \neq 0$. For $\alpha = 0$, define its contribution as zero everywhere. Thus each V_j^0 is convex, and so is their pointwise maximum. Simplex and expected-cost constraints are affine. This proves the last assertion. □

Theorem 8 optimizes over single-record scores that reconstruct the loss pointwise. It does not minimize finite-sample MSE over biased estimators or establish minimax design for every dependent completion law. For example, the constant estimator in the low-budget regime of Theorem 10 trades bias for lower variance.

M NONUNIFORM-REFERENCE RECONSTRUCTION

This section removes the uniform reference restriction while retaining a finite known channel. Fix a cell, a strictly positive reference mass vector $p_0 \in \mathbb{R}^N$, and channels C_a as in Appendix K.1. Remove channel-output rows that are identically zero. Then the output reference probabilities

$$r_a = C_a p_0, \qquad D_a = \operatorname{diag}(r_a)$$

are strictly positive. Let $B \in \mathbb{R}^{N \times (N-1)}$ have orthonormal columns spanning $\mathbf{1}^\perp$. When $N = 1$, the loss is constant and no acquisition is required; below assume $N \ge 2$. Define

$$I(q) = \sum_a q_a B^\top C_a^\top D_a^{-1} C_a B, \qquad b_g = B^\top g, \qquad \mu_g = p_0^\top g. \tag{39}$$

The notation b_g is a target vector and is unrelated to the acquisition budget b.

Theorem 12 (General reference reconstruction)**.** *A target g is identified if and only if $b_g \in \operatorname{range} I(q)$. Under this condition let $v_g = I(q)^\dagger b_g$, with $\dagger$ denoting the Moore–Penrose inverse, and define*

$$h_{g,a}(z) = \mu_g + \frac{(C_a B v_g)_z}{r_a(z)}. \tag{40}$$

Then

$$\sum_a q_a C_a^\top h_{g,a} = g, \qquad \operatorname{Var}_{p_0,q}(h_{g,A}(Z)) = b_g^\top I(q)^\dagger b_g. \tag{41}$$

The score is minimum variance under p_0 among all pointwise-unbiased reconstructions, and remains unbiased under every other completion distribution. The extended-valued function

$$F_g(q) = \begin{cases} b_g^\top I(q)^\dagger b_g, & b_g \in \operatorname{range} I(q), \\ +\infty, & \textit{otherwise} \end{cases} \tag{42}$$

is convex in q.

Proof. For any vector $v \in \mathbb{R}^{N-1}$,

$$v^\top I(q) v = \sum_a q_a \|D_a^{-1/2} C_a B v\|_2^2.$$

Consequently, $v \in \ker I(q)$ if and only if $C_a B v = 0$ for every supported action. Every vector $w \in \ker K_q$ satisfies $\mathbf{1}^\top w = 0$, so it can be written $w = Bv$. By Lemma 11, identification is equivalent to $g^\top B v = 0$ for every such v, or equivalently to $b_g \perp \ker I(q)$. Since $I(q)$ is symmetric, this is $b_g \in \operatorname{range} I(q)$.

Assume this condition and define h by (40). Put $L = \sum_a q_a C_a^\top h_a$. Because $C_a^\top \mathbf{1} = \mathbf{1}$, projection onto the tangent space gives

$$B^\top L = \sum_a q_a B^\top C_a^\top D_a^{-1} C_a B v_g = I(q) v_g = b_g = B^\top g.$$

It follows that $L - g$ is a constant vector. Its reference mean is zero, because

$$\begin{aligned} p_0^\top L &= \sum_a q_a r_a^\top \left(\mu_g \mathbf{1} + D_a^{-1} C_a B v_g\right) \\ &= \mu_g + \sum_a q_a \mathbf{1}^\top C_a B v_g = \mu_g + \mathbf{1}^\top B v_g = \mu_g. \end{aligned}$$

The constant must therefore be zero, proving pointwise reconstruction. Multiplying that identity by any probability vector proves all-laws unbiasedness.

The action-wise reference mean of h_a is μ_g, since $\mathbf{1}^\top C_a B v_g = 0$. Therefore its variance is

$$\sum_a q_a (C_a B v_g)^\top D_a^{-1} (C_a B v_g) = v_g^\top I(q) v_g = b_g^\top I(q)^\dagger b_g.$$

For any alternative pointwise reconstructor $\tilde{h}$, write $d_a = \tilde{h}_a - h_a$. Its pointwise difference obeys $\sum_a q_a C_a^\top d_a = 0$. The weighted cross product is

$$\sum_a q_a (h_a - \mu_g \mathbf{1})^\top D_a d_a = v_g^\top B^\top \sum_a q_a C_a^\top d_a = 0.$$

Both scores have reference mean μ_g, so expanding the square yields

$$\operatorname{Var}_0(\tilde{h}_A) = \operatorname{Var}_0(h_A) + \sum_a q_a d_a^\top D_a d_a \geq \operatorname{Var}_0(h_A).$$

This proves optimality without requiring $I(q)$ to be invertible.

Finally,

$$F_g(q) = \sup_{v \in \mathbb{R}^{N-1}} \{2 b_g^\top v - v^\top I(q) v\}. \tag{43}$$

When b_g is in the range, completing the square in the positive-eigenvalue subspace gives the value in (42). When it is outside the range, its nonzero projection onto the null space makes the supremum infinite. For fixed v, the expression in braces is affine in q, so the supremum is convex. □

A semidefinite epigraph for (42) is

$$\begin{pmatrix} I(q) & b_g \\ b_g^\top & t \end{pmatrix} \succeq 0. \tag{44}$$

To verify the equivalence without assuming invertibility, evaluate its quadratic form at (v, s) and minimize over the positive-eigenvalue subspace of $I(q)$. A null-space component of b_g would make

this form negative for some v when $s \neq 0$, so positive semidefiniteness first requires the range condition. Under that condition the minimum for fixed s is $s^2(t - b_g^\top I(q)^\dagger b_g)$, yielding precisely the epigraph inequality.

For several cells, the total reference variance adds $\mathrm{Var}_\rho(\mu_{g,O,Y})$ to the cell-weighted variances in (41). The same action probabilities must be used for all outcomes in the same pre-audit record. Multiple target losses are handled by taking the maximum of their total variances.

Reduction to the binary formula. For uniform $p_0 = \mathbf{1}/N$ on a binary cube, choose the nonconstant columns of B to be $\chi_S/\sqrt{N}$. Direct marginalization gives

$$I(q)_{ST} = N p_S \mathbf{1}\{S = T\}, \qquad (b_g)_S = \sqrt{N}\alpha_S.$$

Indeed, a character vanishes under a marginal that omits one of its coordinates; if $S \subseteq a$, then $(C_a B)_S(z) = \sqrt{N}\, 2^{-|a|}\chi_S(z)$ and $r_a(z) = 2^{-|a|}$. Substitution into (40) gives $\alpha_\emptyset + \sum_{S\subseteq a, S\neq\emptyset} \alpha_S \chi_S/p_S$, exactly (7).

The implementation uses an eigenvalue-based range check and Moore–Penrose reconstruction, including singular identified channels. Directions exceeding the declared range tolerance are rejected. Accepted scores expose their actual pointwise reconstruction residual (Appendix G.3). Its design optimizer uses an explicit small mixture over the supplied menu, not a requirement that full-record auditing be available. Appendix G.3 specifies the numerical rank and range tolerances.

N Design reduction and numerical diagnostics

Proof of Proposition 9. Fix a cell and let $\mathcal{S}$ be its required interaction family. For every action a, the replacement a° in (26) is a subset of a. If $S \in \mathcal{S}$ and $S \subseteq a$, it is one of the sets in that union, so $S \subseteq a^\circ$. Conversely, $S \subseteq a^\circ$ implies $S \subseteq a$ because $a^\circ \subseteq a$. Thus each required-coverage indicator is unchanged. Monotonicity gives $c(a^\circ, o) \leq c(a, o)$. Move the probability mass of a to a°, combining duplicate replacements. The simplex, all required coverages, and therefore the objective are preserved, while cost does not increase. Feasibility of the replacements is the stated menu assumption.

For support compression, enumerate $\mathcal{S} = \{S_1, \ldots, S_r\}$ and associate to each feasible action the vector

$$z_a = (\mathbf{1}\{S_1 \subseteq a\}, \ldots, \mathbf{1}\{S_r \subseteq a\}, c(a, o)) \in \mathbb{R}^{r+1}.$$

A design represents its coverage/cost vector as $z = \sum_a q_a z_a$. Suppose more than $r + 2$ actions have positive probability. Their augmented vectors $(1, z_a)$ are linearly dependent. Therefore there is a nonzero vector v supported on those actions satisfying

$$\sum_a v_a = 0, \qquad \sum_a v_a z_a = 0.$$

It has both positive and negative entries. Choose

$$\tau = \min_{a: v_a > 0} q_a/v_a > 0, \qquad q'_a = q_a - \tau v_a.$$

All probabilities remain nonnegative, at least one becomes zero, and both total probability and z are preserved. Repeat until at most $r + 2$ actions remain. This finite elimination proves the result constructively. Compressing each cell separately preserves the global cost and every term of (24). □

A sparse design need not have a small interaction dictionary. A generic loss can have $2^{|\overline{m}|} - 1$ nonempty coefficients, and its union menu can have 2^r actions. Efficient use therefore requires a tractable target family or an explicit factorization. Greedy feature-by-feature acquisition is not guaranteed to suffice: every proper subset can leave a k-way interaction unobserved.

N.1 Supporting-plane lower bounds

Write a one-cell objective as $F(q) = \max_j V_j(q)$ and let $\mathcal{Q}$ be the feasible simplex with its cost constraint. At a point $\bar{q}$ having positive required coverages,

$$V_j(q) \geq L_j(q) := V_j(\bar{q}) + \nabla V_j(\bar{q})^\top (q - \bar{q}).$$

For the variance in (24), if $D_{Sa} = \mathbf{1}\{S \subseteq a\}$, its derivative is

$$\frac{\partial V_j}{\partial q_a} = -\sum_{S \neq \emptyset} \frac{\alpha_{j,S}^2 D_{Sa}}{p_S^2}.$$

Therefore the linear program

$$L_* = \min_{q \in \mathcal{Q}, t} t \quad \text{subject to} \quad L_j(q) \leq t \quad \text{for all } j \tag{45}$$

satisfies $L_* \leq \min_{q \in \mathcal{Q}} F(q) \leq F(\bar{q})$. In exact arithmetic, $F(\bar{q}) - L_*$ is a valid optimality gap. The implementation solves both the convex problem and this linear program numerically. Its reported absolute discrepancy is not a proof of an exact nonnegative gap: roundoff and solver feasibility tolerances can reverse the inequality by a small amount. The numerical solution gives $F(\bar{q}) = 0.290498947299522$, a linear-program value 0.290498948638769, and an absolute discrepancy 1.34×10^{-9}. These values describe numerical accuracy rather than an exact-arithmetic certificate.

N.2 The overlapping-scope learning design

The six-coordinate scope controls use a tractable sufficient-scope design rather than the unknown true loss variance. Let $C = \{1, 2, 3\}$, let $G = \{4, 5, 6\}$, and let $S_j = C \cup \{3 + j\}$ for $j = 1, 2, 3$. The target observation table for policy j is a known function of X_{S_j} and Y. On a native-missing record, let an audit a cost $|a|$ and write $p_j = \mathbb{P}(S_j \subseteq A)$.

Proposition 13 (Optimal minimum scope coverage)**.** *For $0 < b \leq 6$, among all audit distributions with $\mathbb{E}|A| \leq b$,*

$$\max_q \min_{j=1,2,3} p_j = \frac{b}{6}.$$

The optimum is attained by auditing $U = C \cup G$ with probability $b/6$ and doing nothing otherwise.

Proof. Let $t(a) = |\{j : S_j \subseteq a\}|$. If $t(a) = 0$, then $t(a) \leq |a|/2$. If $t(a) \geq 1$, containing those scopes requires the three core coordinates and their $t(a)$ distinct gate coordinates. Thus $|a| \geq 3 + t(a) \geq 2t(a)$, since $t(a) \leq 3$. It follows that

$$\min_j p_j \leq \frac{1}{3} \sum_j p_j = \frac{1}{3} \mathbb{E}[t(A)] \leq \frac{1}{6} \mathbb{E}|A| \leq \frac{b}{6}.$$

The proposed union audit has cost $6(b/6) = b$ and contains every scope whenever it is selected, so $p_j = b/6$ for all j and the bound is attained. □

Uniform target-bundle acquisition selects each four-coordinate scope with probability $b/12$ and does nothing with the remaining probability. No one scope contains another, so its coverage of each is $b/12$. Full-record acquisition gives $p_j = b/d$. Independent acquisition of each coordinate with probability b/d gives $p_j = (b/d)^4$. These formulas concern the joint table-query scopes, not all possible lower-degree decompositions of their losses. The union design is coverage-optimal under this criterion, not necessarily variance-optimal for every table entry or dependent feature distribution.

O Measurement-budget lower and upper bounds

We first establish the observation structure of the parity family. For $S \not\supseteq K$, marginalizing P_θ over coordinates outside S sums out at least one coordinate of the nonconstant character χ_K. Hence

$$\mathbb{P}_\theta(X_S = z) = 2^{-|S|}, \qquad S \not\supseteq K, \tag{46}$$

independently of θ. If $S \supseteq K$, then

$$\mathbb{P}_\theta(X_S = z) = 2^{-|S|}(1 + \theta \chi_K(z)). \tag{47}$$

The parity $Z = \chi_K(X)$ has $\mathbb{P}_\theta(Z = 1) = (1 + \theta)/2$. Conditional on Z, all remaining distinctions in an observed covering bundle are independent of θ.

Lemma 14 (Divergence of an adaptive audit transcript)**.** *Let T_n contain all actions and their outputs from n independent records, where each current bundle is drawn using only the preceding transcript and independent randomization. For $0 < \theta \le 1/2$,*

$$\mathrm{KL}(\mathcal{L}_\theta(T_n), \mathcal{L}_{-\theta}(T_n)) = \mathbb{E}_\theta\left[\sum_{i=1}^{n} \mathbf{1}\{K \subseteq A_i\}\right] \theta \log \frac{1+\theta}{1-\theta}. \tag{48}$$

If expected total cost is at most nb, the divergence is at most $4\theta^2 nb/k$.

Proof. Conditional on a realized past transcript, the strategy uses the same action kernel under both parameter values. Its action-probability factors therefore cancel in the likelihood ratio. The current record is independent of this past. For a noncovering action, (46) makes the output law identical and its conditional divergence zero. For a covering action, (47) gives divergence

$$\begin{aligned}&\frac{1+\theta}{2}\log\frac{1+\theta}{1-\theta} + \frac{1-\theta}{2}\log\frac{1-\theta}{1+\theta}\\ &\qquad = \theta \log \frac{1+\theta}{1-\theta}.\end{aligned}$$

Factoring the joint transcript likelihood into its sequential conditional probabilities and taking its expected log gives the sum of these conditional divergences. This is (48), including randomized strategies. For $0 \le t \le 1/2$,

$$\log\frac{1+t}{1-t} = \int_0^t \frac{2}{1-s^2}\, ds \le 4t.$$

Every covering action has cost at least k, so the expected number of covering actions is at most nb/k. The stated upper bound follows. □

Proof of the lower bound in Theorem 10. For any admissible strategy, set

$$\theta_* = \min\left\{\frac{1}{2}, \sqrt{\frac{k}{16nb}}\right\}.$$

Lemma 14 gives transcript divergence at most $1/4$ between parameters θ_* and $-\theta_*$. Their total variation distance is at most $\sqrt{\mathrm{KL}/2} \le 1/2$. For completeness, this inequality follows by mapping both laws to a two-element partition attaining their total variation, applying the log-sum inequality to decrease relative entropy, and noting that the Bernoulli divergence in its first argument has second derivative $1/[p(1-p)] \ge 4$. Taylor expansion around its minimum gives Bernoulli divergence at least twice the squared difference of the probabilities.

The two target risks are separated by

$$\Delta = R_{\theta_*} - R_{-\theta_*} = \gamma\theta_*.$$

Turn any estimator $\widehat{R}$ into a test by choosing whichever of these two values is closer to $\widehat{R}$, breaking ties arbitrarily. On a test error, the squared estimation error for the true parameter is at least $\Delta^2/4$. The sum of testing error probabilities is at least one minus the total variation distance: for a decision event E, it equals $1 - \mathbb{P}_+(E) + \mathbb{P}_-(E) \ge 1 - \mathrm{TV}$. Therefore

$$\begin{aligned}\max_{\theta\in\{\theta_*,-\theta_*\}} \mathbb{E}_\theta(\widehat{R} - R_\theta)^2 &\ge \frac{\Delta^2}{8}(1-\mathrm{TV})\\ &\ge \frac{\gamma^2\theta_*^2}{16} \ge \frac{\gamma^2}{256}\min\left\{1, \frac{k}{nb}\right\}.\end{aligned}$$

The last inequality uses $\theta_*^2 = \min\{1/4, k/(16nb)\} \ge (1/16)\min\{1, k/(nb)\}$. Since the argument applies to every estimator and every admissible history-adaptive strategy, it proves the minimax lower bound. □

Proof of the upper bound and variance comparison. Audit K independently on each record with probability $p = b/k \le 1$. Let B_i be its selection indicator and $Z_i = \chi_K(X_i)$. The estimator

$$\widehat{R} = \frac{1}{2} + \frac{\gamma}{2n} \sum_{i=1}^{n} \frac{B_i Z_i}{p} \tag{49}$$

has expected total cost $nkp = nb$. Since $\mathbb{E}Z_i = \theta$, $Z_i^2 = 1$, and B_i is independent of Z_i,

$$\mathbb{E}(B_i Z_i/p) = \theta, \qquad \mathrm{Var}(B_i Z_i/p) = 1/p - \theta^2.$$

The estimator is unbiased and has mean squared error $\gamma^2(1/p - \theta^2)/(4n) \le \gamma^2 k/(4nb)$. The alternative constant estimator $1/2$ has worst-case squared error at most $\gamma^2/16$ over $|\theta| \le 1/2$. Choose between these two estimators using the known quantities n, b, k. Their smaller worst-case bound is at most

$$\frac{\gamma^2}{4} \min\left\{1, \frac{k}{nb}\right\},$$

which proves the upper bound.

For a strategy restricted to empty or full-record audits, only a full audit can reveal the parity and every such audit costs d. Replacing k by d in the divergence proof gives its lower bound. Auditing the full record with probability b/d gives its upper bound. At $\theta = 0$, the two unbiased estimator variances are exactly $\gamma^2 k/(4nb)$ and $\gamma^2 d/(4nb)$, with ratio d/k. □

The bound assumes that the same budget constraint holds under every member of the parameter family. Allowing a strategy to violate the budget on one of the testing alternatives would invalidate the divergence comparison. The result allows adaptation *between* records, but does not address a procedure that selects further coordinates of the same record after inspecting earlier acquisitions. It also presumes that K is known from the target specification. Discovering an unknown interaction is a different statistical problem.

P Prediction learning, concentration, and inference

P.1 Uniform risk estimates for a finite predictor class

Proposition 15 (Finite-class transfer learning)**.** *Let $\mathcal{F}$ be finite and let the fixed audit design cover every nonzero interaction of $g_{j,f}$ for $j \in [J]$ and $f \in \mathcal{F}$. Suppose B_q in* (30) *is finite. Then* (31) *and* (32) *hold. The conclusions also hold conditionally on independent pilot data used to select the design and the finite class, provided the resulting coverages are positive and B_q is computed for the selected design.*

Proof. Every character has absolute value one, and every inclusion indicator is at most one. Thus (30) gives $|W_{j,f}| \le B_q$ pointwise. Proposition 2 gives its mean $R_j(f)$. For independent variables in an interval of length $2B_q$, the exponential-moment bound for a centered bounded variable yields

$$\mathbb{P}\big(|\widehat{R}_j(f) - R_j(f)| > t\big) \le 2\exp\left(-\frac{nt^2}{2B_q^2}\right).$$

One way to obtain this inequality is to bound the exponential of a bounded variable by the chord between its interval endpoints, center it, and maximize the resulting two-point moment-generating function. Its logarithm has second derivative at most $(2B_q)^2/4 = B_q^2$, is zero at zero, and has zero first derivative there. Integrating twice gives log moment-generating function at most $\lambda^2 B_q^2/2$. Multiplying moment-generating functions over independent records, applying the exponential Markov inequality, and minimizing at $\lambda = t/B_q^2$ gives each one-sided bound above. Apply the same argument to the negative score and add the two tails.

A union bound over $J|\mathcal{F}|$ pairs and substitution of $t = \varepsilon_n$ gives (31). Let f^* minimize $\max_j R_j(f)$ over the finite class. On this event,

$$\begin{aligned}\max_j R_j(\hat{f}) &\le \max_j \widehat{R}_j(\hat{f}) + \varepsilon_n \\ &\le \max_j \widehat{R}_j(f^*) + \varepsilon_n \\ &\le \max_j R_j(f^*) + 2\varepsilon_n.\end{aligned}$$

The middle step is empirical optimality. This proves (32). Conditioning on independent pilot data makes the selected class and design fixed, and the analysis sample remains independent with the same pointwise randomization property. Applying the same argument conditionally proves the pilot version. □

The finite-class bound permits predictor selection on the training sample because it holds simultaneously over the class. Fixed-score unbiasedness alone does not make training-set evaluation of an adaptively fitted predictor unbiased. The experiments evaluate fitted predictors separately, using the exact synthetic population or held-out images.

P.2 Learning a target observation/outcome table

Let a target produce a finite observation symbol Z_j through a known kernel $k_j(z \mid h, x)$. This symbol includes any mask information available to the predictor. For each outcome value y and symbol z, define the complete-data query

$$g_{j,z,y}(h, x, y') = k_j(z \mid h, x)\mathbf{1}\{y' = y\}. \tag{50}$$

Its expectation is $t_{j,z,y} = \mathbb{P}_{\pi_j}(Z_j = z, Y = y)$. Applying Theorem 1 to the finite family of these queries exactly characterizes identification of the target joint table in the unrestricted model. It is sufficient for learning a Bayes predictor for any loss on that finite table. It can be stronger than identifying the risk of one already fixed predictor, which requires fewer queries.

Suppose a query can be computed from a known missing-feature scope $S_j(o)$ and the subsequently observed Y. A simple alternative to its full character decomposition is

$$W^{\text{scope}}_{j,z,y} = \frac{\mathbf{1}\{S_j(O) \subseteq A\}}{p_{S_j(O)}(O)} k_j(z \mid H, X)\mathbf{1}\{Y = y\}. \tag{51}$$

The kernel in this expression is evaluated only when the complete scope is visible. If the query is already computable from O, set its missing scope to empty and its weight to one. The same conditional inclusion argument proves unbiasedness. This score is nonnegative and often computationally simpler than a signed character reconstruction, but it need not have minimum reference variance.

For binary Y and squared loss, the target Bayes probability is

$$\eta_j(z) = \frac{t_{j,z,1}}{t_{j,z,0} + t_{j,z,1}}$$

on positive-mass target symbols. Let $C_{j,z,y} = \sum_{i=1}^n W^{\text{scope}}_{i,j,z,y}$ and use fixed pseudocounts $a > 0$,

$$\hat{\eta}_j(z) = \frac{C_{j,z,1} + a}{C_{j,z,0} + C_{j,z,1} + 2a}. \tag{52}$$

The experiments take $a = 1/2$. Because the scores are nonnegative, this estimate always lies in $[0, 1]$. The ratio and smoothing are not exactly unbiased even though the unnormalized table queries are.

Proposition 16 (Consistency of scope-weighted table learning)**.** *For independent records with fixed positive scope coverages,* (52) *converges almost surely to $\eta_j(z)$ on every positive-mass target symbol. Its target Brier risk converges to the Bayes risk.*

Proof. There are finitely many cells and each relevant coverage is positive, so each query score is bounded. The strong law gives $C_{j,z,y}/n \to t_{j,z,y}$ simultaneously over the finite table. Divide the numerator and denominator of (52) by n. The pseudocount terms tend to zero and each positive-mass

symbol has a strictly positive limiting denominator, proving ratio convergence. For any probability predictor f,

$$\mathbb{E}[(f(Z_j) - Y)^2 \mid Z_j] = \eta_j(Z_j)(1 - \eta_j(Z_j)) + (f(Z_j) - \eta_j(Z_j))^2.$$

The difference from Bayes risk is therefore a finite weighted sum of squared probability-estimation errors. Each tends to zero, and zero-mass symbols contribute nothing. This proves the risk statement. □

A neural model trained on the same nonnegative weighted table minimizes a well-defined empirical target loss. The table supplies the information; parameterization determines approximation and optimization error. Neither the table identity nor Proposition 15 gives a generalization bound for an unrestricted neural network.

P.3 Fixed-design confidence intervals

For one fixed target and audit design, the score is bounded on a finite space. If its variance v is positive, the classical independent-sample central limit theorem and the law of large numbers for W^2 give

$$\frac{\sqrt{n}(\widehat{R} - R)}{s_n} \implies N(0,1), \qquad s_n^2 = \frac{1}{n-1}\sum_{i=1}^{n}(W_i - \widehat{R})^2 \longrightarrow v.$$

Hence the usual interval $\widehat{R} \pm 1.96 s_n/\sqrt{n}$ has asymptotic 95% coverage. These claims require a fixed design with a nonzero variance, not a triangular sequence of increasingly rare audits without additional conditions. At very small joint inclusion probabilities, normal approximations can be poor even though all variables are bounded for each fixed design. The reported coverage experiments are finite-sample diagnostics, not assertions of exact or simultaneous coverage. Conservative finite-sample simultaneous intervals are supplied by (31) instead.

P.4 Explicit approximation error from truncation

If computational constraints lead to retaining only a supported set of characters $\mathcal{T}_{o,y}$, let $\tilde{g}$ be the resulting truncated expansion and estimate it by the corresponding unbiased reconstruction. Then

$$|\mathbb{E}\widehat{R}_{\mathcal{T}} - R| \le \sum_{o,y} \rho_{oy} \sum_{S \notin \mathcal{T}_{o,y}} |\alpha_{o,y,S}|. \tag{53}$$

To prove it, subtract the expansions pointwise, bound each character by one, and take expectation over cells. This is an approximation guarantee only when the omitted coefficient sum is controlled from the specified target loss. It is not identification of the original risk. In particular, an unsupported nonzero interaction cannot be declared irrelevant merely because it is difficult to observe.

Q Paid-pilot and bounded-error experiments

This section specifies the joint-approximation and learned-reference experiments. Appendix T gives the parity, fixed-scope, and hidden-state controls. Those controls use different budgets and datasets, so their results are reported separately.

Q.1 Dense logistic losses and an explicit approximation envelope

For $d \in \{2, 4, 6, 8\}$, the target query is

$$g(x) = \sigma\left(0.7 + \sum_{i=1}^{d} \beta_i x_i\right), \qquad (\beta_1, \ldots, \beta_d) = \mathrm{linspace}(0.2, 0.8, d).$$

It is a valid target-policy squared loss when $Y = 0$, a loss-one mask is selected with probability $g(x)$, and a mask-only predictor returns one or zero on the two masks. We evaluate the full Walsh transform on the complete cube. Table 8 reports the full-interaction coefficients and degree-two approximation

Table 8: Dense logistic interaction diagnostics. "Terms" counts coefficients exceeding 10^{-12} in magnitude, including the constant. Tail quantities concern ordinary degree-two truncation, not the optimized reconstruction.

d	Terms	Full coefficient	Tail ℓ_1	Max. residual
2	4	-0.00993049	0.000000	0.000000
4	16	$+0.00263797$	0.025741	0.025741
6	64	-0.00098708	0.138154	0.124735
8	256	$+0.00045424$	0.394150	0.277242

tails. Nonzero counts use a 10^{-12} reporting threshold. These numerical diagnostics apply to the enumerated examples; they do not symbolically certify every dimension or parameter choice. The reporting threshold is not used to decide identification.

For estimation, let

$$z(x) = \tanh\{0.6x_1 + 0.4x_2 + 0.5x_1x_2x_3\},$$
$$\tilde{z}(x) = \frac{z(x) - \bar{z}}{\max\{1, \max_u |z(u) - \bar{z}|\}}, \qquad p(x) = 2^{-d}\{1 + 0.8\tilde{z}(x)\},$$

where $\bar{z}$ is the uniform cube average. This strictly positive nonuniform law satisfies $p \le 1.8p_0$ for uniform p_0. The reference envelope is known by construction and is not estimated from evaluation outcomes. The grid has $d \in \{4, 6, 8\}$, $n \in \{250, 1000, 4000\}$, and $b \in \{0.5, 1, 2\}$. Coordinate costs are one. For each degree $L \in \{1, 2, 3\}$, solve (10) using all nonempty interactions and actions of order at most L, plus the empty action. We set $\kappa = 1.8$ and use SLSQP with analytic gradients, tolerance 10^{-13}, and at most 1,800 iterations per solve. Loose dual diagnostics trigger objective-rescaled restarts. The interaction-coverage floor is 10^{-7}. After optimization, every residual is recomputed on all 2^d states.

The exact full-audit comparator acquires $[d]$ with probability b/d and otherwise acquires nothing. It uses the centered score

$$\bar{g} + \frac{\mathbf{1}\{A = [d]\}}{b/d}\{g(X) - \bar{g}\}, \qquad \bar{g} = 2^{-d}\sum_x g(x).$$

This comparator is exactly unbiased. Its known reference center is not the true nonuniform risk. For each method, we enumerate the score's actual mean and variance under p and compute $\text{bias}^2 + \text{variance}/n$. These 108 calculations form the core degree sweep. Appendix S adds bias-matched full and fixed-subset reconstructions at every setting, giving 216 distinct method–setting evaluations. The calculations need neither asymptotic normality nor Monte Carlo standard errors. State enumeration uses floating-point computation, not interval arithmetic.

Table 9: Selected eight-coordinate logistic cases. MSE and the declared bound are in units of 10^{-3}. The two settings exhibit a limited-budget gain and a large-sample approximation failure. The complete grid is in the supplement.

n	b	Design	Bias	Residual	MSE	Bound
250	0.5	Exact full	+0.00000	0.00000	4.3969	8.1030
		Degree 1	-0.03118	0.12325	3.1578	19.1258
		Degree 2	-0.01865	0.09109	3.0347	13.1906
		Degree 3	-0.01030	0.04214	3.9137	8.6895
4000	2.0	Exact full	+0.00000	0.00000	0.0687	0.1266
		Degree 1	-0.03118	0.12325	1.0405	15.3137
		Degree 2	-0.01824	0.09108	0.3747	8.3716
		Degree 3	-0.00250	0.03148	0.0769	1.1190

Q.2 Candidate measurements, context, and native masks

The prediction experiment has four candidate binary features, with costs $c = (0.5, 1, 1.5, 2)$. Fully recovering an empty record costs five. An always-observed binary context is $H \in \{-1, 1\}$. Let $D \in \{0, 1\}$ indicate whether the first candidate is already available. The native mask is $M = \emptyset$ when $D = 0$ and $M = \{1\}$ when $D = 1$. In both datasets,

$$\mathbb{P}(D = 1 \mid X, Y, H) = \sigma\{-0.2 + 0.6H + 0.9X_3 - 0.8X_4 + 0.6(2Y - 1)\}. \tag{54}$$

This nonignorable censoring rule specifies a retrospective joint law. It does not presume that a prospective sensor observes a future outcome; a latent common cause can induce the association allowed by the formal model. All outcomes become available after the audit decision.

Allocation conditions on (H, D) and uses a common q across values of the already observed first feature. This permitted restriction reduces the number of fitted references. Reconstruction still uses the native feature when present and the eventual outcome. Each stratum has state (X_1, X_2, X_3, X_4, Y) with 32 possible values. The menu contains all subsets of missing candidates: 16 actions for $D = 0$ and eight for $D = 1$. Every reconstruction channel reveals Y, any native first feature, and the selected additional coordinates.

Three target families. A policy either reveals one core feature or returns an empty observation. Its symbol is $z \in \{0, 1, 2\}$ for empty, observed minus one, or observed plus one. Overlap policies use core X_1 and gates X_{j+1} for $j = 1, 2, 3$, with observation probabilities

$$r_j(x, h) = \sigma(0.4 + 0.3h + 1.1x_{j+1} + 0.4x_1).$$

Disjoint policies use core–gate pairs (X_1, X_2) and (X_3, X_4) with the same formula. Dense policies use cores X_1, X_2, X_3 and

$$r_j(x, h) = \sigma\{0.2(-1)^{j-1} + 0.3h + \beta_j^\top x\},$$

where β_j is the $(j - 1)$-position cyclic shift of $(0.2, 0.4, 0.6, 0.8)$. All four candidate features matter to each policy family. There are no extra known nuisance dimensions charged only to the full-audit comparator.

For each policy, the six queries $k_j(z \mid x, h)\mathbf{1}\{Y = y\}$, $z = 0, 1, 2$, $y = 0, 1$, specify the observation/outcome table. Reference design minimizes their unweighted sum of variances, rather than finite-sample Brier risk itself. A variance-optimal table design can therefore have higher realized predictive risk after projection, smoothing, and ratio estimation.

Q.3 Independent synthetic samples and repeated image partitions

Each synthetic repetition contains 4,000 records. Draw H uniformly and a latent sign S with $\mathbb{P}(S = 1 \mid H) = \sigma(0.8H)$. Conditional on S, the four features are independent noisy copies of S, with error probabilities $(0.12, 0.22, 0.32, 0.18)$. Draw

$$\mathbb{P}(Y = 1 \mid S, X, H) = \sigma(1.8S + 0.5X_1X_2 - 0.6H),$$

then use (54). Prediction is evaluated by exact summation over the synthetic population, with context weights one half. The complete-law allocation comparator receives its true conditional completion distributions.

The semi-synthetic image study uses the 1,797 examples in `load_digits`. Candidate features are thresholded intensities at zero-indexed pixel locations $(19, 27, 35, 28)$; context is the thresholded intensity at location 20. Intensities above eight map to one and all others to minus one. The outcome is one for digits five through nine. These locations are fixed before partitioning and are not selected from hidden training or test outcomes. Synthetic censoring and target kernels are imposed on the real images. The measurements are pixel queries within this controlled task, not a model of the cost of acquiring a natural image.

Repetition $r \in \{0, \ldots, 39\}$ uses a fresh stratified 1,200/597 training/test partition with random state $41000 + r$. The pilot is drawn from the 1,200 training images and charged to the budget, rather than provided as a separate fully observed sample. All 40 split-index pairs are included. Evaluation

averages the target kernel analytically over held-out images. Standard errors describe split, native-mask, and audit variation within this fixed image collection. They do not measure uncertainty for new writers or a genuine sensor-policy transition.

The image complete-law comparator uses inaccessible complete training records within each (H, D) stratum and adds 0.001 to every state count before normalization. It uses no test outcomes. This privileged allocation comparator is neither deployable from the paid pilot nor a Bayes-risk lower bound. Its finite-training reference need not equal the test distribution. The synthetic comparator instead uses the exact population law. The potentially ill-conditioned privileged histogram uses relative range tolerance 10^{-8} and records the returned score residual; feasible learned-reference scores use the public default 10^{-12}. In the primary table, complete-law allocation still uses the same *pilot-reference reconstruction* as the other acquisition methods. This holds reconstruction fixed when comparing allocations.

Q.4 Pilot cost, acquisition rules, and remaining allowance

At budget $b \in \{0.4, 0.8, 1.6\}$ per candidate record, the total expected allowance is $N_{\mathrm{tot}} b$. Reserve

$$n_0 = \max\left\{1, \left\lfloor \frac{0.2 N_{\mathrm{tot}} b}{5} \right\rfloor\right\}$$

pilot records using an independent random permutation. In this grid the floor is always at least one. Recover all their missing candidate values. The realized pilot cost is $C_0 = \sum_{i \in \mathrm{pilot}} (5 - 0.5 D_i)$, so it is at most 20% of the total allowance. Fit a separate 32-state conditional histogram for each (H, D), adding $8/32$ to every state count. This is a strictly positive reference even when a stratum has no pilot records.

The remaining $n = N_{\mathrm{tot}} - n_0$ records each receive expected audit-cost cap $(N_{\mathrm{tot}} b - C_0)/n$. The main table estimate excludes pilot records. Both reference and design are frozen before applying randomized measurements to the main records. Feasible comparators share the pilot, native observations, candidate menu, and pilot cost. The complete-law comparator additionally receives its privileged reference.

Full/union acquisition selects all missing candidates with probability scaled to the remaining cap. Uniform-scope acquisition selects a target scope uniformly, removes native coordinates, and scales the mixture to the cap. The two distributions coincide when all scopes are dense. Uniform-reference, pilot-reference, and complete-law allocations solve the reference-variance program over the entire menu. Feasible acquisition strategies do not use the current record's outcome. The privileged complete-training diagnostic is the stated exception: its fitted allocation can depend indirectly on main-record labels, but it does not use held-out test outcomes.

Reference optimization uses SLSQP with analytic gradients, tolerance 10^{-11}, and an initial limit of 180 iterations. A failed numerical check triggers a deterministic 650-iteration retry. The constraints impose a 10^{-4} mixture with a budget-feasible full-menu distribution. Every action therefore has a positive numerical floor, even when its additional allocation is zero. A supporting-plane linear program checks the objective gap. For the sum criterion, vertices are enumerated after subtracting lower bounds to avoid presolve failures near small exploration masses. Acceptance requires cost residual at most 10^{-7} and gap divided by $\max(1, \text{objective})$ below 5×10^{-5}. These are numerical checks, not exact optimality certificates.

Q.5 Four estimators on the same acquired records

Each realized acquisition design supplies four reconstructions of the same table queries. Scope HT uses $\mathbf{1}\{S_j \subseteq M \cup A\} g / p_{S_j \setminus M}$ and discards the query when the full sufficient scope is unavailable. Augmented HT uses the pilot conditional mean $b_j(H, D, X_M, Y)$ and score

$$b_j + \frac{\mathbf{1}\{S_j \setminus M \subseteq A\}}{p_{S_j \setminus M}} (g - b_j).$$

The baseline can use the eventual outcome because it is evaluated only after acquisition. Its design-based identity holds even under an incorrect imputation model. Character reconstruction uses the

Table 10: Pilot and total measurement accounting. Each row averages repetitions and policy families for the pilot-reference allocation. Realized spending is random and is not required to lie below its expected cap in every repetition.

Data	Budget/record	Pilot records	Pilot cost	Expected/record	Realized/record
Synthetic	0.4	64	304.62	0.4000	0.3990
	0.8	128	610.09	0.8000	0.8027
	1.6	256	1219.88	1.6000	1.5981
Digits	0.4	19	90.88	0.4000	0.3968
	0.8	38	181.85	0.8000	0.7989
	1.6	76	363.80	1.6000	1.6036

uniform-reference finite-channel score and can exploit lower-order observed directions. Learned-reference reconstruction applies the same range-aware formula with the pilot histogram. Score evaluation uses only native or acquired feature values and the eventual outcome.

Within an acquisition strategy, all four estimators use identical actions. Across strategies, common random numbers couple the action draws. Identical distributions, including dense full/union and uniform scopes, therefore produce identical acquisitions and identical risk estimates when reconstruction is fixed. Comparisons retain this pairing rather than treating methods as independent samples.

For each context and policy, divide the reconstructed count vector by the number of main records in that context, then project its six entries onto the probability simplex in Euclidean norm. Add $0.5/n_h$ to each outcome count when computing conditional probabilities, where n_h is the main-record count at that context. All methods use this learner. Unprojected query estimates are design-unbiased; projection and the probability ratio are not. With fixed positive coverages, the finite-dimensional law of large numbers and continuity at positive-mass symbols give consistency. This argument does not extend to an unrestricted neural-network class.

Table 11: Estimator comparison under the same pilot-reference acquisitions at total budget 0.8. Means are Brier risks; parenthesized standard errors are in units of 10^{-5}. These differences isolate reconstruction rather than acquisition.

Data	Family	Scope HT	Augmented HT	Character	Learned ref.
Synthetic	Overlap	0.20971 (13)	0.20858 (4)	0.20844 (4)	0.20839 (4)
	Disjoint	0.21974 (17)	0.21867 (7)	0.21855 (6)	0.21849 (5)
	Dense	0.21806 (33)	0.21565 (17)	0.21437 (5)	0.21427 (4)
Digits	Overlap	0.23129 (72)	0.22783 (61)	0.22680 (62)	0.22669 (62)
	Disjoint	0.22835 (74)	0.22528 (68)	0.22461 (60)	0.22441 (64)
	Dense	0.23550 (131)	0.22774 (85)	0.22513 (70)	0.22479 (68)

Q.6 Paired comparisons, full budget curves, and reference gaps

For each repetition, average the policy risks within a family. Those 40 independent experimental randomizations, not the number of policies or table cells, determine the reported standard error. For each comparator c, compute paired differences $R_r(\text{pilot}) - R_r(c)$, their mean, and their sample standard deviation divided by $\sqrt{40}$. No multiple-comparison correction is applied, so these are descriptive paired uncertainties, not a collection of simultaneous tests.

The study contains 14,400 prediction results and 14,400 stratum/method configurations, including heuristic strategies. Complete results are in the supplement. The complete-law reference variance evaluates the best reconstruction under that reference for the chosen allocation. It is a diagnostic of allocation, not the variance of the pilot score under an unknown population. Small ℓ_1 discrepancy does not establish the multiplicative envelope used in Proposition 6.

Table 12: Paired Brier differences, pilot-reference allocation minus each comparator, with learned-reference reconstruction fixed. Both means and standard errors are in units of 10^{-4}. Negative differences favor pilot allocation. The full budget grid is retained.

Data	Family	Budget	vs. full/union	vs. scopes	vs. uniform ref.
Synthetic	Overlap	0.4	−1.15 (0.73)	−0.29 (0.57)	+0.41 (0.38)
		0.8	−0.99 (0.40)	+0.34 (0.30)	+0.11 (0.18)
		1.6	−1.00 (0.24)	+0.05 (0.16)	+0.06 (0.10)
	Disjoint	0.4	−0.14 (0.88)	+1.03 (0.85)	+0.50 (0.59)
		0.8	−0.43 (0.44)	−0.88 (0.44)	+0.35 (0.28)
		1.6	−0.22 (0.21)	−0.59 (0.23)	−0.24 (0.17)
	Dense	0.4	−1.76 (0.89)	−1.76 (0.89)	+0.25 (0.60)
		0.8	−1.07 (0.35)	−1.07 (0.35)	−0.22 (0.26)
		1.6	−0.36 (0.28)	−0.36 (0.28)	+0.19 (0.19)
Digits	Overlap	0.4	−10.43 (3.37)	−1.86 (2.16)	−0.80 (1.49)
		0.8	−5.40 (2.61)	−1.18 (1.29)	−0.07 (1.00)
		1.6	+0.42 (1.13)	+1.77 (0.79)	+1.39 (0.50)
	Disjoint	0.4	−3.93 (4.48)	−4.38 (4.00)	+0.18 (2.43)
		0.8	+0.08 (1.70)	−1.03 (1.86)	−0.67 (1.03)
		1.6	−1.00 (1.24)	−0.44 (0.94)	+0.13 (0.79)
	Dense	0.4	−7.06 (2.83)	−7.06 (2.83)	−3.63 (1.94)
		0.8	−3.36 (2.18)	−3.36 (2.18)	+0.46 (1.60)
		1.6	−0.77 (0.93)	−0.77 (0.93)	+0.20 (0.69)

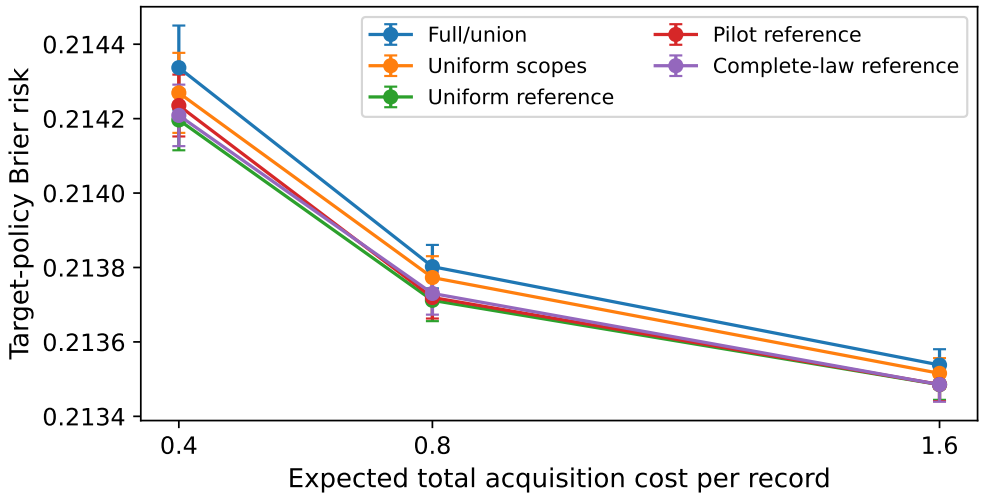

Figure 1: Synthetic target Brier risk versus total expected acquisition cost, including pilot recovery. Each repetition averages family-level risks before aggregation. Error bars show 1.96 standard errors over 40 repetitions. The small differences among target-aware allocations do not establish a general ordering.

We use seeds $731001 + r$ for sample/native-mask randomness and $470000 + r$ for pilot permutation. For acquisitions the seed is

$$951000 + r + 1000f + 100001\{b = 0.8\} + 200001\{b = 1.6\} + 300001\{\text{digits}\} + 500001\{h = 1\} + 100000D,$$

where $f = 0, 1, 2$ indexes overlap, disjoint, and dense families. The method index is deliberately absent to preserve the common-random-number comparison. Dataset split seeds are given above. Numerical libraries use one thread. The numerical environment and timing results are in the supplement.

Table 13: Complete-law trace-variance ratios relative to complete-law allocation on the same stratum and budget. Means average all repetitions, strata, and budgets within each family. Values above one represent allocation loss for this reference criterion, not a Brier-risk ratio.

Data	Family	Full/union	Uniform scopes	Uniform ref.	Pilot ref.
Synthetic	Overlap	1.218	1.034	1.004	1.003
	Disjoint	1.090	1.090	1.007	1.005
	Dense	1.150	1.150	1.015	1.008
Digits	Overlap	1.403	1.098	1.034	1.018
	Disjoint	1.193	1.193	1.046	1.024
	Dense	1.236	1.236	1.052	1.027

R MISSPECIFICATION AND COMPUTATIONAL BOUNDARIES

The finite-population calculations below distinguish three failure modes. All use four coordinates, $g(x) = \sigma(0.7 + \beta^\top x)$ with $\beta = (0.2, 0.4, 0.6, 0.8)$, and

$$p(x) \propto \exp\{0.5x_1 - 0.6x_2 + 0.8x_1x_2 + 0.35x_3x_4\}.$$

No samples are generated for these diagnostics. Expectations are sums over the 16-state population.

Incorrect support. The fixed actions are $(\emptyset, \{1, 2\}, \{3, 4\}, [4])$. Their probabilities are $(0.5, 0.2, 0.2, 0.1)$. Coefficients are retained only through a chosen degree. Correct audit probabilities preserve the retained query, not the omitted part. Table 14 reports the resulting bias, uniform residual, and omitted-coefficient ℓ_1 bound. The full-interaction coefficient is small but not zero. Dropping it leaves a nonzero bias; increasing the number of audited records would not remove that bias.

Table 14: Support misspecification under fixed, correctly randomized audits. The bias concerns the original untruncated target risk.

Degree	Bias	Max. residual	Tail ℓ_1	Variance
0	+0.024315	0.421305	0.472787	0.000000
1	+0.008311	0.082978	0.088253	0.199504
2	-0.001284	0.025741	0.025741	0.214441
3	-0.000442	0.002638	0.002638	0.209652
4	+0.000000	0.000000	0.000000	0.209513

Incorrect target kernel. Hold the nominal query fixed and let the actual target loss be $g_\delta(x) = \sigma(0.7+\beta^\top x+\delta x_4)$ for $\delta = 0, 0.1, 0.3, 0.6, 1$. Using the nominal query estimates the wrong risk even if every record is fully observed. Its exact bias is $p^\top(g - g_\delta)$, bounded by the uniform binary-kernel total variation $\max_x |g(x) - g_\delta(x)|$. Table 15 reports these quantities. Source audits cannot correct an unspecified or incorrect deployment kernel.

Table 15: Target-policy error is estimand error, not a reference-variance effect.

δ	Nominal risk	Actual risk	Bias	Uniform TV bound
0.0	0.611155	0.611155	+0.000000	0.000000
0.1	0.611155	0.607476	+0.003679	0.024855
0.3	0.611155	0.599616	+0.011539	0.074443
0.6	0.611155	0.587062	+0.024094	0.148885
1.0	0.611155	0.570145	+0.041010	0.242630

Incorrect completion reference. Optimize at expected cost 0.8 using $p_\lambda = (1 - \lambda)p + \lambda r$ for $\lambda = 0, 0.25, 0.5, 0.75, 1$. The endpoint r is either uniform or proportional to the exponential with all

four terms in the exponent above negated. Construct the exact reference reconstructor and evaluate it under the unchanged true p. Its pointwise identity survives every mixture, whereas variance need not remain optimal. The largest variance ratio here is 1.41 relative to matched-reference allocation and reconstruction. Table 16 also reports the multiplicative-envelope parameter. When it is at least one, the two-sided bound in Proposition 6 is inapplicable; a small-sample fit is not automatically certified by that proposition.

Table 16: Reference misspecification with exact scores. Variance ratios use a fixed true law. All computed biases have magnitude below 10^{-12}. A dash means that the two-sided multiplicative distortion bound is inapplicable because $\eta \geq 1$.

Endpoint	Mixture	ℓ_1 error	η	Var. ratio	Bound
Uniform	0.00	0.000	0.000	1.000	1.00
Uniform	0.25	0.144	0.778	1.007	8.02
Uniform	0.50	0.288	0.875	1.024	15.04
Uniform	0.75	0.432	0.913	1.069	22.07
Uniform	1.00	0.576	1.201	1.204	–
Reversed	0.00	0.000	0.000	1.000	1.00
Reversed	0.25	0.389	0.941	1.043	32.77
Reversed	0.50	0.778	0.969	1.205	64.55
Reversed	0.75	1.166	2.400	1.271	–
Reversed	1.00	1.555	15.990	1.414	–

Menu growth and measured runtime. For $r \in \{2, 4, 6, 8\}$, create three bounded loss queries with r pairwise interactions. Overlap uses pairs $\{1, i+1\}$, and disjoint uses $\{2i-1, 2i\}$. Positive coefficients are generated with seed $7613 + r$ and rescaled to sum to 0.3 per query, allowing an intercept 0.5 to keep each loss in $[0, 1]$. Enumerate all unions, use unit coordinate costs and a budget of one quarter of the largest action cost, and optimize the worst uniform-reference variance. Three repeated timings measure the same deterministic optimization problem. Support compression preserves the attained coverage/cost vector.

Table 17: Finite-menu computational scaling. Time is the median of three single-process measurements in seconds. Menu size is exactly 2^r in these constructions. Compressed support excludes probabilities below 10^{-8}.

Structure	Interactions	Dimension	Menu	Time (s)	Support
Disjoint	2	4	4	0.0017	3
Disjoint	4	8	16	0.0022	5
Disjoint	6	12	64	0.0083	7
Disjoint	8	16	256	0.1743	9
Overlap	2	3	4	0.0020	3
Overlap	4	5	16	0.0022	3
Overlap	6	7	64	0.0117	5
Overlap	8	9	256	0.3915	7

The results show exponential menu growth despite small support in the final distribution. They do not establish a polynomial-time structural algorithm. Complete-cube logistic designs reach a maximum dimension of eight. The smooth-gate theorem gives an approximation class but does not remove state enumeration from this implementation. Factorized constraints, separation oracles, or high-dimensional learned representations would require further analysis to preserve the residual bound.

S Operational selection and the cost of preliminary measurements

This study examines two decisions: choosing approximation and acquisition for a fixed loss, and allocating preliminary measurements to fit a reference for target-table reconstruction. Neither is tuned on the evaluation law or held-out image outcomes. The grids are fixed in advance. Complete-data laws are used only to generate records and evaluate the procedures.

S.1 Bias-matched acquisition controls

We use the logistic query and tilted finite population from Appendix Q. For $d \in \{4, 6, 8\}$, $n \in \{250, 1000, 4000\}$, and $b \in \{0.5, 1, 2\}$, every method receives the same query, uniform reference, envelope $\kappa = 1.8$, and acquisition cap. The candidate menu contains exact full auditing, bias-aware full auditing, and fixed or jointly optimized acquisition at degrees one, two, and three. Thus there are 216 method–configuration evaluations.

Exact full auditing uses the centered inverse-probability score. Bias-aware full auditing uses the same full/no-audit distribution and optimizes its reconstruction through the two-variable convex problem in Appendix G. The fixed degree-L control acquires a uniformly random L-subset with probability $\min\{b/L, 1\}$ and optimizes only reconstruction. The joint degree-L method optimizes both acquisition and reconstruction. These controls separate approximation, use of acquired information, and allocation. The population law is excluded from the optimizer's inputs.

The full control is solved by L-BFGS-B with analytic derivatives. The partial problems use SLSQP with analytic derivatives, tolerance 10^{-13}, and at most 1,800 iterations per solve. Joint designs retain the declared 10^{-7} interaction-coverage floor. When the numerical dual gap exceeds $10^{-4} \max\{B, 10^{-8}\}$, where B is the current objective, the same problem is restarted from its current solution with objective scales 10, 100, and 0.1. This changes conditioning, not the statistical objective. The smallest feasible returned objective and strongest available lower bound are retained. A run need not attain an exact optimizer for its recomputed residual/energy bound to hold.

Table 18 reports three representative settings, including the bias-dominated large-sample setting. Complete results for all 216 evaluations are in the supplement. The largest primal-minus-lower-bound diagnostic across the approximate controls is below 1.4×10^{-7}. It is a floating-point diagnostic, not an interval-arithmetic certificate. The declared envelope bounds exceed the corresponding population MSEs in every evaluated configuration.

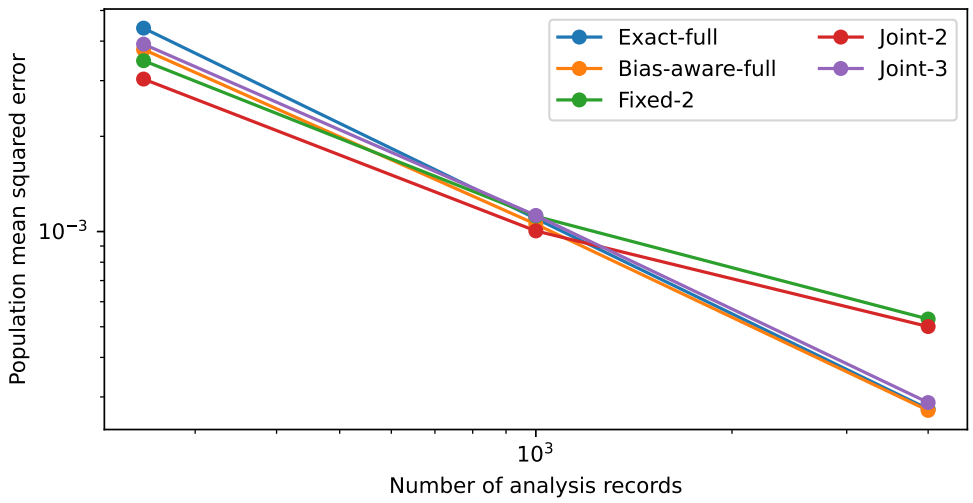

Figure 2: Fixed-query population MSE at $d = 8$ and $b = 0.5$. Each curve fixes the method before evaluating its MSE. A sweep over degrees does not itself provide an operational degree-selection rule.

Table 18: Bias-matched controls at $d = 8$. Both MSE and bound are multiplied by 10^3. Cost is expected acquired coordinates per analysis record, not the cap. ϵ is the recomputed pointwise residual. The gap compares the returned objective with a numerical lower bound on the same restricted problem. The exact-full row is not an approximation optimization.

n/b	Design	Cost	Bias	Residual	MSE	Bound	Gap
250/0.5	Exact-full	0.500	+0.00000	0.00000	4.3969	8.1030	–
	Bias-aware-full	0.500	-0.00380	0.02364	3.7643	7.4826	$0.0e+00$
	Fixed-1	0.500	-0.03118	0.12325	3.4955	19.7338	$2.4e-17$
	Joint-1	0.500	-0.03118	0.12325	3.1578	19.1258	$2.3e-08$
	Fixed-2	0.500	-0.01817	0.09108	3.4689	14.0271	$5.2e-10$
	Joint-2	0.500	-0.01865	0.09109	3.0347	13.1906	$1.1e-08$
	Fixed-3	0.500	-0.00753	0.04684	4.5381	10.3404	$2.6e-09$
	Joint-3	0.500	-0.01030	0.04214	3.9137	8.6895	$2.5e-09$
1000/0.5	Exact-full	0.500	+0.00000	0.00000	1.0992	2.0258	–
	Bias-aware-full	0.500	-0.00108	0.00639	1.0550	1.9838	$0.0e+00$
	Fixed-1	0.500	-0.03118	0.12325	1.6031	16.3265	$6.2e-17$
	Joint-1	0.500	-0.03118	0.12325	1.5186	16.1745	$2.1e-09$
	Fixed-2	0.500	-0.01825	0.09108	1.1180	9.7286	$9.5e-17$
	Joint-2	0.500	-0.01824	0.09108	1.0048	9.5194	$3.1e-08$
	Fixed-3	0.500	-0.00275	0.03156	1.3678	3.4602	$2.2e-10$
	Joint-3	0.500	-0.00282	0.03167	1.1245	3.0290	$1.6e-08$
4000/2	Exact-full	2.000	+0.00000	0.00000	0.0687	0.1266	–
	Bias-aware-full	2.000	-0.00007	0.00041	0.0685	0.1264	$8.4e-11$
	Fixed-1	1.000	-0.03118	0.12325	1.0511	15.3327	$1.7e-18$
	Joint-1	1.000	-0.03118	0.12325	1.0405	15.3137	$1.3e-07$
	Fixed-2	2.000	-0.01824	0.09108	0.3818	8.3847	$2.9e-16$
	Joint-2	2.000	-0.01824	0.09108	0.3747	8.3716	$2.8e-09$
	Fixed-3	2.000	-0.00249	0.03148	0.0916	1.1457	$3.8e-10$
	Joint-3	2.000	-0.00250	0.03148	0.0769	1.1190	$3.8e-09$

The bound and actual MSE can favor different designs. At $n = 250, b = 0.5$, joint degree-two acquisition improves actual MSE by 19.4% over bias-aware full auditing and by 12.5% over fixed two-subset acquisition, but its bound is larger than the bias-aware full bound. This gain holds for the evaluated law; minimizing the envelope would not select it. At $n = 4000, b = 2$, degree one can spend at most one coordinate per record. Its residual bias persists while full-audit variance decreases. Table 18 reports both the action restriction and actual expenditure.

S.2 Selecting a candidate using paid validation

The validation study fixes $d = 8$, total record counts $N \in \{250, 1000, 4000\}$, and total per-record caps $b \in \{0.5, 2\}$. Before drawing any validation values, it reserves

$$v = \left\lfloor \frac{0.2Nb}{8} \right\rfloor, \qquad n = N - v, \qquad b_{\mathrm{main}} = \frac{Nb - 8v}{N - v}.$$

The validation records are fully acquired at cost eight each and excluded from the analysis estimate. Each of the eight candidates above is constructed for the remaining n records and the same b_{main}. References remain uniform. No candidate or reference is fitted again using the validation block.

For each of 2,000 independent repetitions per resource setting, the raw selector minimizes (21) evaluated on the v paid complete records. The upper-confidence selector adds the radius in (22), using $\delta = 0.05$ and suprema over the known finite query/channel. Bound-after-validation pays for the same block but selects by the envelope. The prespecified degree-two comparator pays for that block without changing its degree. Bound-without-validation instead constructs all eight candidates using all N records and the entire allowance, then selects the smallest bound. The final comparison therefore accounts for both the sample and acquisition costs of validation.

Fresh main records and randomized audits provide an estimate for each selected candidate. Methods sharing validation also share main states and random uniforms. The no-validation comparator uses a separate main draw, so comparisons with it do not use common random numbers. After selection,

population MSE is calculated by complete enumeration. The study comprises 60,000 rule–repetition outcomes from 12,000 validation blocks shared across five rules, not 60,000 independent validation experiments. Complete results are in the supplement.

Table 19: Operational selection over 2,000 repetitions per resource setting. MSE, squared error, and their standard errors are multiplied by 10^3. Mean MSE averages the selected candidate's fresh-sample population MSE. Mean squared error is measured on the subsequently drawn main sample. Zero SE in the MSE column means the candidate is fixed, not that its realized estimator is deterministic. BV and B0 denote bound selection with and without paid validation, respectively. V selects the raw validation criterion, V+ adds the uniform radius, and J2 fixes degree two after reserving the same validation cost.

N/b	Choice	v	Mean MSE	SE	Mean sq. error	SE
250/0.5	B0	0	3.7643	0.0000	3.9873	0.1282
	BV	3	4.5047	0.0000	4.6010	0.1499
	V	3	4.1249	0.0135	4.2278	0.1334
	V+	3	4.6473	0.0000	4.7565	0.1437
	J2	3	3.6733	0.0000	3.6287	0.1163
250/2	B0	0	1.0543	0.0000	1.0802	0.0324
	BV	12	1.2927	0.0000	1.2761	0.0404
	V	12	1.3871	0.0078	1.4080	0.0444
	V+	12	1.3877	0.0002	1.4190	0.0463
	J2	12	1.1640	0.0000	1.2078	0.0347
1000/0.5	B0	0	1.0550	0.0000	1.0564	0.0346
	BV	12	1.2934	0.0000	1.3597	0.0441
	V	12	1.4102	0.0046	1.4664	0.0427
	V+	12	1.3890	0.0001	1.4548	0.0459
	J2	12	1.1646	0.0000	1.1890	0.0367
1000/2	B0	0	0.2718	0.0000	0.2713	0.0085
	BV	50	0.3389	0.0000	0.3473	0.0110
	V	50	0.4208	0.0031	0.4160	0.0122
	V+	50	0.3389	0.0000	0.3473	0.0110
	J2	50	0.5426	0.0000	0.5515	0.0140
4000/0.5	B0	0	0.2720	0.0000	0.2575	0.0083
	BV	50	0.3391	0.0000	0.3266	0.0099
	V	50	0.4378	0.0041	0.4264	0.0122
	V+	50	0.3391	0.0000	0.3266	0.0099
	J2	50	0.5428	0.0000	0.5505	0.0138
4000/2	B0	0	0.0685	0.0000	0.0616	0.0019
	BV	200	0.0856	0.0000	0.0807	0.0024
	V	200	0.0911	0.0007	0.0896	0.0028
	V+	200	0.0856	0.0000	0.0807	0.0024
	J2	200	0.3852	0.0000	0.3889	0.0062

Table 19 reports conditional MSE and realized squared error separately. At $N = 250, b = 0.5$, the reservation buys only three complete validation records. Raw validation gives mean conditional MSE 0.004125, versus 0.003764 for bound selection without validation. Neither validation rule improves mean MSE over the full-budget no-validation selector in any of the six resource settings. Upper-confidence radii are particularly conservative. These results are compatible with unbiased validation and its selection guarantee. They show that a small, charged sample can be ineffective for choosing among similar bias–variance tradeoffs. A prespecified residual tolerance or bound selection avoids that expense.

The criterion estimates a candidate's *future* MSE. It is neither a prediction loss nor an oracle choice based on the simulated law. Its unbiasedness also does not apply when the candidate is fitted on the same validation rows. Selection frequencies and confidence-event rates are in the supplement. Those indicators describe selection uncertainty, not confidence intervals for the selected risk estimate.

S.3 How much should the reference pilot cost?

The prediction study uses the same two datasets, 40 repetitions, target kernels, costs, native masks, and record splits as Appendix Q. The reservation shares are 0, 0.1, 0.2, and 0.4 of the total acquisition allowance. For each share τ and budget $b \in \{0.4, 0.8, 1.6\}$, the first

$$n_0 = \lfloor \tau N b/5 \rfloor$$

records in the prespecified random permutation form the reference pilot. Their actual recovery costs are charged. The remaining records receive the remaining expected acquisition allowance. The pilot histogram uses the same eight total pseudocounts as before. At zero pilot share, the reference is uniform, no records are removed for fitting, and the full acquisition budget is available for analysis.

Within each pilot-share comparison, full-union, uniform-reference, and fitted-reference allocations use the same learned-reference reconstruction. At zero share the reconstruction is uniform, and fitted- and uniform-reference allocations coincide. The three policy families have equal weight. Held-out image labels never select the pilot share. The grid measures resource sensitivity rather than estimating a universal optimum or tuning a policy.

We use the same reference optimizer and exploration floor as in the main prediction study. When the objective gap exceeds its tolerance, the optimizer restarts from the current distribution with tighter tolerance and scales 1, 0.1, and 10. Returned designs must satisfy the cost and objective-gap tolerances. Predictive risk does not enter this restart rule.

To separate reference quality from prediction, we also report the attained trace variance of the unprojected target-table estimate. Conditional on observed stratum counts, it is

$$T(q, W) = \sum_h \sum_m \frac{n_{hm}}{n_h^2} \sum_j \mathrm{Var}_{p_{hm}, q_{hm}}(W_j).$$

The reconstruction W is the score actually applied, not a new optimal score fitted to the evaluation law. On digits, p_{hm} is a positive-ridge complete-*training* histogram used only for this diagnostic. This is not a known image-population variance. On the synthetic task it is the specified population conditional law. The subsequent simplex projection and normalization are nonlinear, so this trace need not rank fitted Brier risks.

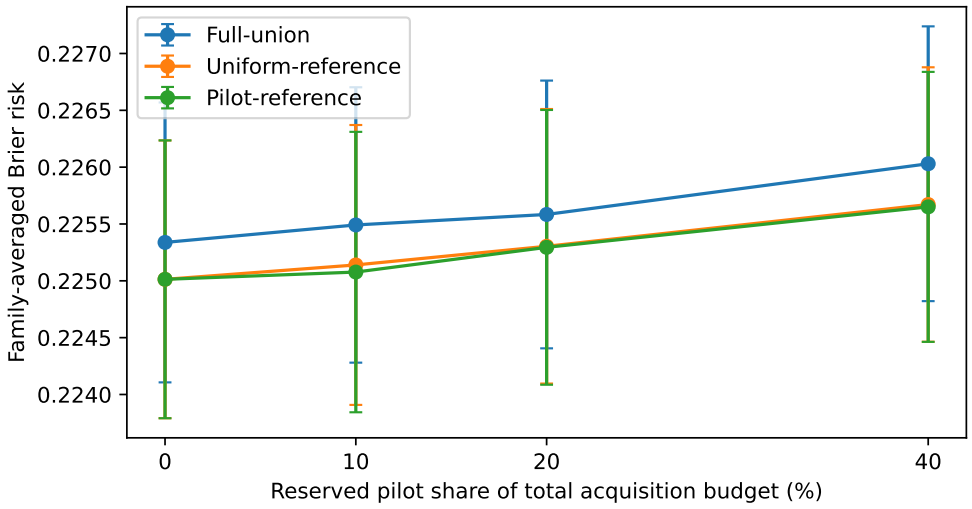

Figure 3: Reference-pilot cost and prediction. Family-averaged digits Brier risk at total budget $b = 0.8$ across four prespecified reservation shares. Error bars show descriptive means ± 1.96 standard errors over 40 repeated splits. Within a share, all methods use the same reconstruction reference. Across shares, both the reference and remaining main-sample resources change.

A larger pilot reduces the measured reference discrepancy but leaves fewer records and less acquisition budget for analysis. At digits budget 0.8, family-averaged Brier risk is 0.22501 at zero share and

Table 20: Reference-pilot resource tradeoff for fitted-reference allocation. Results first average equally over the three target-policy families within a repetition, then average over 40 repetitions. Trace is $10^3 T(q, W)$ and Brier SE is multiplied by 10^5. Reference discrepancy is the main-stratum-frequency-weighted ℓ_1 error, so it differs from an unweighted average of four conditional histograms. Pilot cost is the mean total cost, not cost per record. Zero pilot share uses the entire budget with a uniform reference.

Data	b	Pilot %	n_0	Pilot cost	Ref. ℓ_1	Trace	Brier	SE
Synthetic	0.4	0	0	0.0	0.897	6.639	0.21419	5.2
		10	32	152.4	0.710	6.707	0.21422	6.5
		20	64	304.6	0.634	7.152	0.21424	4.2
		40	128	610.1	0.530	8.756	0.21447	7.4
	0.8	0	0	0.0	0.897	3.599	0.21368	2.4
		10	64	304.6	0.634	3.569	0.21370	2.5
		20	128	610.1	0.530	3.779	0.21372	2.9
		40	256	1219.9	0.412	4.608	0.21391	3.1
	1.6	0	0	0.0	0.897	2.079	0.21346	1.9
		10	128	610.1	0.530	2.077	0.21346	1.5
		20	256	1219.9	0.412	2.218	0.21349	2.3
		40	512	2439.1	0.313	2.677	0.21358	2.6
Digits	0.4	0	0	0.0	1.010	24.185	0.22675	63.7
		10	9	43.0	0.929	25.668	0.22740	60.2
		20	19	90.9	0.850	27.654	0.22742	65.2
		40	38	181.8	0.731	33.711	0.22778	60.3
	0.8	0	0	0.0	1.010	12.973	0.22501	62.3
		10	19	90.9	0.850	13.420	0.22508	62.9
		20	38	181.8	0.731	14.141	0.22529	61.7
		40	76	363.8	0.599	17.031	0.22565	60.6
	1.6	0	0	0.0	1.010	7.369	0.22433	62.6
		10	38	181.8	0.731	7.433	0.22425	63.4
		20	76	363.8	0.599	7.847	0.22445	62.4
		40	153	732.9	0.455	9.359	0.22464	63.3

0.22529 at 20%. The table retains smaller differences and reversals at other budgets. Reference fit alone therefore does not justify the pilot expenditure. Paired comparisons with the zero-pilot allocation and the corresponding synthetic results are in the supplement.

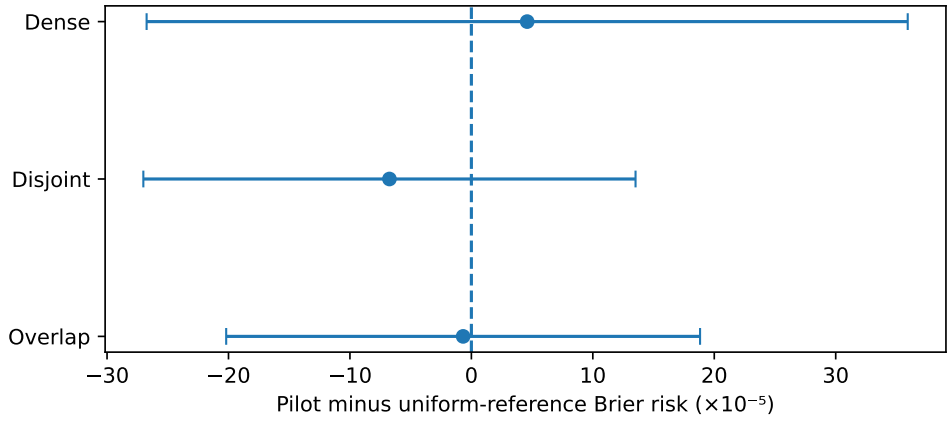

Figure 4: Allocation and reconstruction have different empirical roles. Pilot-minus-uniform-reference Brier differences on digits at $b = 0.8$, holding the 20% pilot and learned-reference reconstruction fixed. Bars are descriptive paired mean ± 1.96 SE, not simultaneous tests. The additional allocation gain is small and inconsistent across families.

The resource sweep contains 8,640 prediction rows and 34,560 stratum–method rows. Its 20% slice agrees with the corresponding same-pilot experiment to numerical precision. The largest measured

pointwise reconstruction residual is below 5×10^{-14}. Acquisition costs satisfy the expected-budget constraint up to the stated numerical tolerances. This is not an almost-sure spending cap or a rigorous enclosure of floating-point error.

T FIXED-SCOPE MECHANISM EXPERIMENTS AND STRUCTURAL CONTROLS

These controls use prespecified measurement scopes, without paid reference fitting or support discovery. Risk-directed acquisition here denotes maximin scope coverage, not the learned-reference optimizer.

T.1 WHAT EACH EXPERIMENT MEASURES

The parity and multi-policy experiments evaluate the MSE and interval coverage of unbiased *fixed-risk estimators*. The hidden-state and image experiments evaluate the *target Brier risk* of fitted predictors. Their standard errors summarize complete repetitions, not uncertainty for individual predictions. Budgets count additional coordinates, excluding native observations and the eventual label.

Table 21: Fixed-scope prediction control at cost one per native-missing record. The relevant six-coordinate union is given in advance, and learning that support is not charged. Entries are mean $\pm$ standard error over 40 repetitions and are reported separately from the paid-pilot study.

Acquisition / estimation	Hidden-state	Digits (semi-synthetic)
Maximin scope coverage	0.173436 ± 0.000075	0.230529 ± 0.000447
Uniform target bundles	0.174787 ± 0.000119	0.234980 ± 0.000534
Full-record audits	0.177617 ± 0.000311	0.264626 ± 0.002227
MI-guided bundles	0.182203 ± 0.000530	0.259159 ± 0.001472
Independent coordinates	0.180643 ± 0.000235	0.247350 ± 0.000847
Native MAR estimator	0.180773 ± 0.000238	0.249962 ± 0.000924
Joint two-state model (no audit)	0.176389 ± 0.000145	0.249647 ± 0.000255

T.2 PARITY SCALING AND A DENSE NEGATIVE CONTROL

We evaluate the family P_θ in (28) with $k = 3$, $b = 0.75$, $n = 4{,}000$, $\gamma = 0.8$, and $d \in \{8, 16, 32, 64, 128\}$. It uses $\theta \in \{0, 0.6\}$ and 5,000 independent repetitions of each configuration. The second parameter lies outside the interval used for the numerical constants in Theorem 10, but remains inside the valid full-support parity family. The exact unbiasedness and variance formula hold for all $|\theta| < 1$.

An audit has three possible sufficient-statistic outcomes, namely no audit, audited positive parity, and audited negative parity. Their probabilities are $1 - p$, $p(1 + \theta)/2$, and $p(1 - \theta)/2$. Sampling their multinomial counts gives exactly the same estimator distribution as generating the full feature vectors and subsequently auditing them. It is not a normal approximation. If N_+ and N_- are the last two counts, then

$$\widehat{R} = \frac{1}{2} + \frac{\gamma}{2np}(N_+ - N_-), \qquad s_W^2 = \frac{\gamma^2}{4p^2(n-1)}\left(N_+ + N_- - \frac{(N_+ - N_-)^2}{n}\right).$$

The reported interval is $\widehat{R} \pm 1.96\sqrt{s_W^2/n}$. Its coverage is measured against the exact $R_\theta = 1/2 + \gamma\theta/2$. The targeted probability is $p = b/3$, and the full-record probability is $p = b/d$. Average acquisition cost is the audit count times its bundle size divided by n. Complete results are in the supplement.

Across these configurations, observed Wald coverage ranges from 0.9428 to 0.9538. These values illustrate finite-sample calibration at the tested budgets; they are not a proof of uniform normal approximation. We also consider a dense negative control with $K = [d]$ for $d \in \{4, 8, 16, 32\}$. There, targeted and full-record audit probabilities coincide at b/d and their variance ratio is exactly one. This dense control is evaluated analytically rather than by simulation.

Table 22: Parity results. Mean squared errors are multiplied by 10^4. Coverage is the fraction of nominal 95% Wald intervals containing the exact target risk. Each row uses 5,000 repetitions.

θ	d	Targeted MSE $\times 10^4$	Full MSE $\times 10^4$	Targeted coverage	Full coverage
0.0	8	1.589	4.279	0.9538	0.9494
0.0	16	1.655	8.592	0.9476	0.9460
0.0	32	1.611	17.058	0.9454	0.9516
0.0	64	1.617	34.811	0.9526	0.9476
0.0	128	1.606	68.932	0.9474	0.9428
0.6	8	1.448	4.052	0.9528	0.9538
0.6	16	1.437	8.527	0.9528	0.9452
0.6	32	1.467	16.869	0.9518	0.9494
0.6	64	1.444	34.095	0.9492	0.9482
0.6	128	1.406	67.071	0.9530	0.9454

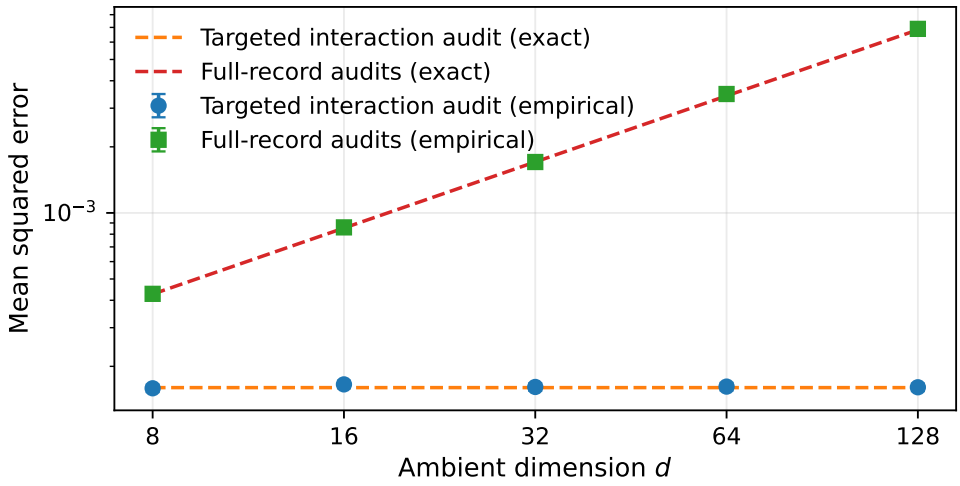

Figure 5: Parity mechanism study at $\theta = 0$, with 5,000 repetitions of 4,000 records at each dimension. Targeted acquisition purchases the specified three-coordinate interaction, whereas full-record acquisition pays for all coordinates. The comparison illustrates Theorem 10, not support discovery.

T.3 Overlapping interactions and multiple target losses

The second experiment uses $d = 8$ with the five one-indexed interactions

$$\{1,2\}, \quad \{2,3,4\}, \quad \{4,5\}, \quad \{5,6,7\}, \quad \{1,7,8\}.$$

The three losses have constant coefficient $1/2$ and nonconstant coefficient matrix

$$\begin{pmatrix} 0.24 & 0.08 & 0.04 & 0 & 0 \\ 0 & 0.06 & 0.19 & 0.09 & 0 \\ 0.06 & 0 & 0 & 0.08 & 0.21 \end{pmatrix}. \tag{55}$$

Each row has absolute coefficient sum at most 0.36, so all losses lie in $[0, 1]$. Each can also be represented as a target-policy loss using a binary mask-only predictor and the loss value as the probability of its loss-one mask. Native records are empty. Acquisition costs equal bundle cardinality, and the budget is 1.5 per record.

The optimized distribution minimizes the largest uniform-reference variance over the union menu. The uniform target-bundle comparator assigns equal probability to the five required interaction bundles, with the total acquisition probability scaled to meet the same expected cost. Full-record acquisition chooses all eight coordinates with probability $1.5/8$. These comparators use the same score (7); the difference is the audit distribution. The optimization uses SLSQP with analytic derivatives, an epigraph variable, tolerance 10^{-12}, and at most 3,000 iterations. Required coverages are constrained numerically to exceed 10^{-9}. The design is then compressed by linear-dependence elimination.

Evaluation first uses the uniform distribution on the 256 completions. The shifted evaluation distribution is

$$P_1(x) \propto \exp\left(0.6x_1x_2 + 0.4x_3x_4 + 0.3x_5x_6x_7 + 0.2x_8\right). \tag{56}$$

Its normalizing constant, target means, and estimator variances are computed by exhaustive enumeration. The same uniform-reference design is used unchanged on this dependent law. For each design and distribution, 2,000 repetitions of 4,000 records are simulated through the exact joint state/action distribution. Sample-variance Wald intervals use the same convention as the parity experiment.

Table 23: Multi-policy risk estimation. Empirical and exact mean squared errors are multiplied by 10^4. Coverage is marginal for each target, not simultaneous across targets. Each entry uses 2,000 repetitions.

Law	Design	Target	MSE $\times 10^4$	Exact $\times 10^4$	Coverage
Reference	Reference-variance design	1	0.723	0.726	0.9475
Reference	Reference-variance design	2	0.728	0.726	0.9550
Reference	Reference-variance design	3	0.747	0.726	0.9470
Reference	Uniform target bundles	1	1.454	1.421	0.9485
Reference	Uniform target bundles	2	1.050	1.036	0.9485
Reference	Uniform target bundles	3	1.124	1.172	0.9560
Reference	Full-record audits	1	0.820	0.875	0.9535
Reference	Full-record audits	2	0.671	0.637	0.9460
Reference	Full-record audits	3	0.703	0.721	0.9500
Dependent	Reference-variance design	1	0.677	0.685	0.9505
Dependent	Reference-variance design	2	0.694	0.725	0.9475
Dependent	Reference-variance design	3	0.721	0.734	0.9490
Dependent	Uniform target bundles	1	1.385	1.380	0.9465
Dependent	Uniform target bundles	2	1.023	1.034	0.9490
Dependent	Uniform target bundles	3	1.185	1.164	0.9535
Dependent	Full-record audits	1	0.867	0.833	0.9500
Dependent	Full-record audits	2	0.683	0.636	0.9375
Dependent	Full-record audits	3	0.750	0.734	0.9530

The optimized reference variances agree to numerical precision across the three targets. This is consistent with the minimax-over-targets objective, but need not occur in every problem. Some targets have lower variance under full-record audits. The dependent-law run is consistent with the all-laws expectation identity and shows the effect of a misspecified variance reference. Marginal Wald coverage ranges from 0.9375 to 0.9560 across the multi-policy configurations.

T.4 Independent hidden-state episodes

Each synthetic record is an independent two-time-point episode. Its hidden initial state S_0 is uniform on $\{-1, 1\}$, and S_1 equals S_0 with probability 0.9 and $-S_0$ otherwise. The label is $Y = (S_1 + 1)/2$. Conditional on S_0, the first six coordinates of $X \in \{-1, 1\}^{32}$ are independent noisy state measurements with error probabilities

$$(0.10, 0.20, 0.25, 0.10, 0.20, 0.30).$$

The other 26 coordinates are independent fair signs, independent of the state and label. A native completeness indicator D gives $M = [32]$ when $D = 1$ and $M = \emptyset$ otherwise. In the self-masking mechanism experiment,

$$\mathbb{P}(D = 1 \mid S_0, X) = \mathrm{sigmoid}(-1.5 + S_0 + 1.2X_1X_4 - 0.7X_5). \tag{57}$$

The additional dependence on noisy emissions makes the state-only joint-mask model misspecified. In the correctly specified control, the native rule is instead

$$\mathbb{P}(D = 1 \mid S_0, X) = \mathrm{sigmoid}(-1.5 + S_0). \tag{58}$$

The complete feature/label law and all target policies are identical between these two experiments. There are 4,000 training episodes and 500 independent pilot episodes per repetition. Training records and native masks are regenerated for every repetition. The pilot is used only to choose the mutual-information-guided comparator's mixture and has its own native missingness. It does not furnish fully observed hidden training data to any learner.

The target policies have common core $C = \{1, 2, 3\}$ and gates $G = \{4, 5, 6\}$. Policy j observes exactly the core with probability

$$\pi_j(C \mid x) = 0.5 + 0.35x_{3+j}, \qquad \pi_j(\emptyset \mid x) = 1 - \pi_j(C \mid x). \tag{59}$$

The gate controls the mask but is not itself supplied to the predictor. Both probabilities lie in $\{0.15, 0.85\}$. Thus the predictor's input has nine states, one for the empty mask and eight for the binary core pattern. All policies and scopes are fixed before data generation.

For synthetic evaluation, nuisance coordinates can be marginalized out. The exact law on $(X_1, \ldots, X_6, Y)$ has 128 states and is obtained by summing over S_0,

$$p(x, y) = \frac{1}{2} \sum_{s \in \{-1,1\}} \left[\prod_{r=1}^{6} p_r(x_r \mid s) \right] p_Y(y \mid s), \quad p_r(s \mid s) = 1 - e_r, \quad p_Y((s+1)/2 \mid s) = 0.9.$$

Every reported synthetic prediction risk averages exactly over this law and the known target kernel, without test-set Monte Carlo noise. The average of the three target Bayes risks is 0.1722750868.

T.5 Semi-synthetic handwritten digits

We use the 1,797 8×8 images bundled with scikit-learn, a subset of the Optical Recognition of Handwritten Digits dataset (Alpaydin & Kaynak, 1998; scikit-learn developers, 2026). Each intensity is mapped to $+1$ when it exceeds 8 and to -1 otherwise. The binary label is one for original digit labels 5 through 9 and zero for labels 0 through 4. Six fixed zero-indexed pixel locations,

$$19,\ 27,\ 35,\ 20,\ 28,\ 36,$$

are moved to the first six feature positions in that order, with the remaining positions appended in ascending order. The first three form the common core and the next three are the gates. This ordering is prespecified rather than selected using the test set.

A stratified split with random state 1701 assigns 1,200 images to a development pool and 597 to testing. A second stratified split with random state 1702 assigns 200 development images to the pilot and 1,000 to training. These image partitions remain fixed across all 40 repetitions. The split indices are in the supplement. The native completeness probability is artificially imposed as

$$\mathbb{P}(D = 1 \mid X, Y) = \mathrm{sigmoid}(-1.2 + 1.2X_1 - 0.8X_5 + 1.3Y). \tag{60}$$

The outcome-dependent rule is a retrospective censoring device. It is not a prospective claim that a sensing system knows the later outcome. All outcomes remain recorded. Native masks and audit randomization are regenerated in each repetition for the fixed training and pilot images. The target kernels remain (59).

The image distribution is neither uniform nor full-support on the binary cube, neither of which is needed for sufficient scope weighting. The unrestricted model's necessity result is not asserted on this empirical image support. Nor does the experiment establish an acquisition advantage for every representation. Many ambient pixels are irrelevant to the prescribed core/gate policies, and full-record audits still pay for them.

Held-out Brier risk is the average over the 597 test images and the *exact* target observation probabilities. The reference predictor is the same smoothed table learner fitted on all 1,000 complete training images. Its average held-out risk is 0.2254326450. These complete training values are withheld from every audit-based or native-data learner, so the reference is an inaccessible full-measurement comparator. It is not a population Bayes predictor. Standard errors across repetitions are conditional on this fixed image sample and do not quantify dataset-sampling or writer-level generalization uncertainty.

T.6 Acquisition policies and shared learning procedure

All acquisition policies receive the same native training records and known target scopes $S_j = C \cup \{3 + j\}$. Each is applied only to native-missing records. For a missing record and budget $b \in \{0.5, 1, 2\}$, the policies are as follows. Risk-directed acquisition measures $C \cup G$ with probability $b/6$. Full-record acquisition measures all d features with probability b/d. Uniform target-bundle

acquisition chooses each of the three scopes with probability $b/12$. Independent acquisition purchases each of the d coordinates independently with probability b/d. Unused probability mass in the bundle policies is assigned to the empty audit. Their expected cost is exactly b, although realized cost fluctuates.

The MI-guided comparator estimates the mutual information between (X_C, X_{3+j}) and Y for each j, using only natively complete pilot records. Its 16×2 contingency table receives pseudocount $1/2$ per cell. If j^* maximizes this empirical mutual information, the bundle mixture is

$$w_j = 0.9\mathbf{1}\{j = j^*\} + 0.1/3.$$

It selects scope S_j with probability $(b/4)w_j$ and otherwise performs no audit. Every scope has positive probability, with more weight on the most predictive scope under the observed pilot law. Native selection can bias the pilot estimate, and the native-missing stratum has no feature-valued context. This is an information-guided objective baseline, not a replication of the estimator, meta-learning system, or evaluation protocol of Kobayashi et al. (2025).

For target j, a record is eligible when it is natively complete or its entire scope is acquired. Native-complete records receive weight one. A missing record receives weight $\mathbf{1}\{S_j \subseteq A\}/p_{S_j}$, with exact coverage probabilities from the acquisition design. For eligible records, the target-observed probability is $r_j(x) = 0.5 + 0.35x_{3+j}$. The estimator adds weighted mass $1 - r_j(x)$ to the empty observation/label cell and weighted mass $r_j(x)$ to its core-pattern/label cell. These are the nonnegative scope scores in (51). The learner uses (52) with $a = 1/2$.

The table learner receives only native or acquired feature values. A record contributes to a target table only when its entire core and gate are observed. The common smoothing causes the ratio estimator to be biased at finite sample size, but its unnormalized target table is design-unbiased.

The native MAR comparator uses only complete native records with constant weight $1/\hat{r}$, where $\hat{r}$ is the empirical complete-record fraction. It uses no audit budget. With no other pre-audit covariate available in this experiment, this is an ignorable-selection benchmark rather than a fitted correct MNAR propensity model. Its finite-sample smoothing differs slightly from the no-eligible-audit outcome of an audit method because the overall weight scale differs. This explains why the independent-coordinate and native-MAR predictions can be close without being identical.

Table 24: Realized audit cost per missing record and the mean worst-policy excess at $b = 1$. Worst-policy excess is maximized within each repetition before averaging. The synthetic reference is Bayes; the digit reference is the complete-training-data table predictor.

Method	HMM cost	HMM worst excess	Digits cost	Digits worst excess
Maximin scope coverage	0.998	0.001493	1.015	0.006800
Uniform target bundles	0.995	0.003529	0.997	0.014663
Full-record audits	1.005	0.006647	0.982	0.048605
MI-guided bundles	0.998	0.018484	1.011	0.060213
Independent coordinates	1.003	0.011824	0.992	0.024549
Native MAR estimator	0.000	0.011984	0.000	0.027617

T.7 A JOINT-MODEL COMPARATOR AND ITS CORRECTLY SPECIFIED CONTROL

The joint comparator is a two-component Bernoulli latent-class model for the first six features, label, and native completeness indicator. Given latent state $s \in \{1, 2\}$, it assumes independent feature emissions, a Bernoulli outcome, and a Bernoulli native mask. Its observed-data log likelihood includes feature terms only on complete native records. Empty records contribute their observed label and mask but no hidden feature values. The learner is told the latent dimension and the six relevant coordinates, providing it a favorable structural specification.

The implementation uses four initializations, 250 EM iterations per initialization, and selects the fit with highest recorded observed-data log likelihood. The first outcome-probability initialization is $(0.2, 0.8)$; the remaining initialization distributions are specified in the supplement. Updates use 0.01 numerator and 0.02 denominator pseudocounts for each Bernoulli parameter, with clipping to $[10^{-5}, 1 - 10^{-5}]$. The fitted complete-data distribution is enumerated over the six features and label, then combined with each known target policy to form its prediction table.

The model is correctly specified for (58), but not for (57) or the artificial image law. Table 25 reports the state-only control. Without audits, the joint model has mean risk 0.172714 ± 0.000045 and outperforms the audit learners. Its valid structural assumptions supply information excluded from the unrestricted analysis. Under self-masking it instead gives 0.176389 ± 0.000145; on digits it gives 0.249647 ± 0.000255. These are fixed-scope controls and are not pooled with the paid-pilot comparisons in Table 3.

Table 25: Correctly specified state-only native-mask control. The complete-data law and target policies are unchanged. The joint model uses no additional measurements, while audit methods have budget one per missing record. Entries are mean $\pm$ standard error over 40 repetitions.

Acquisition / estimation	Target Brier risk	Audit cost
Maximin scope coverage	0.173318 ± 0.000077	1.001
Uniform target bundles	0.174527 ± 0.000110	1.010
Full-record audits	0.178091 ± 0.000390	0.976
MI-guided bundles	0.184567 ± 0.000865	0.999
Independent coordinates	0.222251 ± 0.001130	1.005
Native MAR estimator	0.222556 ± 0.001136	0.000
Joint two-state model	0.172714 ± 0.000045	0.000

T.8 Neural predictors on the same acquired information

At budget one, the six acquisition/native-data methods supply their weighted tables to separate small neural networks, one per repetition and target. The input is four-dimensional: an empty mask is encoded as $(0, 0, 0, 0)$, and an observed core is encoded as $(1, x_1, x_2, x_3)$. Each network has one 16-unit ReLU hidden layer and a scalar sigmoid output. It minimizes the weighted Brier loss on the nine observation states after adding the same $1/2$ pseudocount per outcome used by the table learner. Table masses are normalized within each network's objective.

Networks use 1,200 full-batch Adam steps at learning rate 0.025 in float64. Initialization uses seed 781, Gaussian weight scales 0.3 and 0.2 in the two layers, and zero biases. Each network has separate parameters. Batched evaluation uses the final forward pass and exact-population or held-out-image risk. Initialization, stopping, hyperparameters, and architecture are fixed without test data. The numerical environment is in the supplement.

Table 26: Neural prediction from the same weighted target tables at budget one. Entries are mean target Brier risk $\pm$ standard error across 40 repetitions. This experiment changes the predictor, not the information supplied by acquisition.

Acquisition / estimation	Hidden-state	Digits (semi-synthetic)
Maximin scope coverage	0.173436 ± 0.000075	0.230525 ± 0.000447
Uniform target bundles	0.174787 ± 0.000119	0.234973 ± 0.000533
Full-record audits	0.177670 ± 0.000307	0.264720 ± 0.002230
MI-guided bundles	0.182208 ± 0.000531	0.259202 ± 0.001478
Independent coordinates	0.180643 ± 0.000235	0.247350 ± 0.000846
Native MAR estimator	0.180773 ± 0.000238	0.249953 ± 0.000925

This nine-state, small-network experiment compares optimization on fixed acquired information. It does not test high-dimensional continuous representations or discovery of unobserved target scopes.

T.9 Randomization

Synthetic and image repetitions use NumPy generators with seeds $62000 + r + 1000d_0$, where $r \in \{0, \dots, 39\}$ and $d_0 \in \{0, 1, 2\}$ indexes the two datasets and the structural control. Audit generators use $950000 + r + 100m_0 + 10000b_0 + 1000000d_0$, where m_0 is the method index and b_0 the budget index. The joint-model seed is $90000 + r$. Further seed specifications are in the supplement.

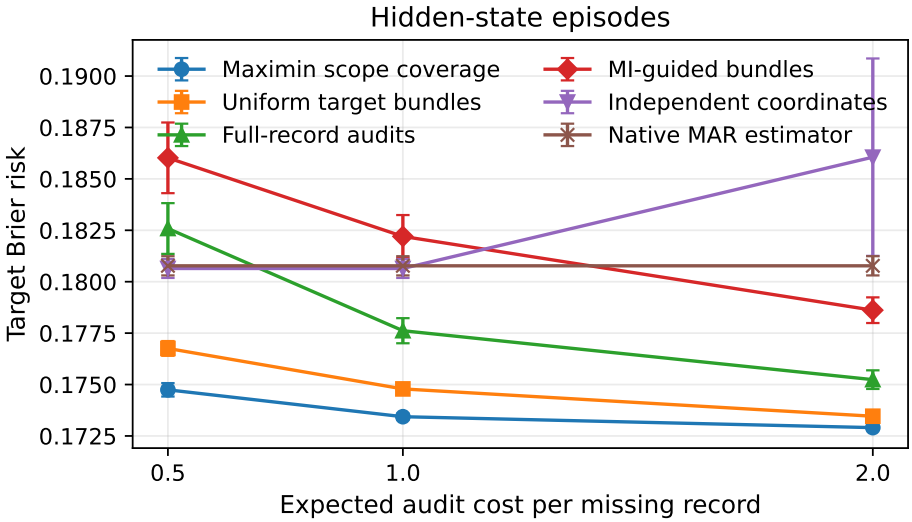

Figure 6: Target Brier risk versus expected audit cost for the hidden-state experiment. Error bars are 1.96 times the standard error of the 40 repetition-level mean risks. The native estimator uses zero additional measurements and is repeated horizontally only as a reference.

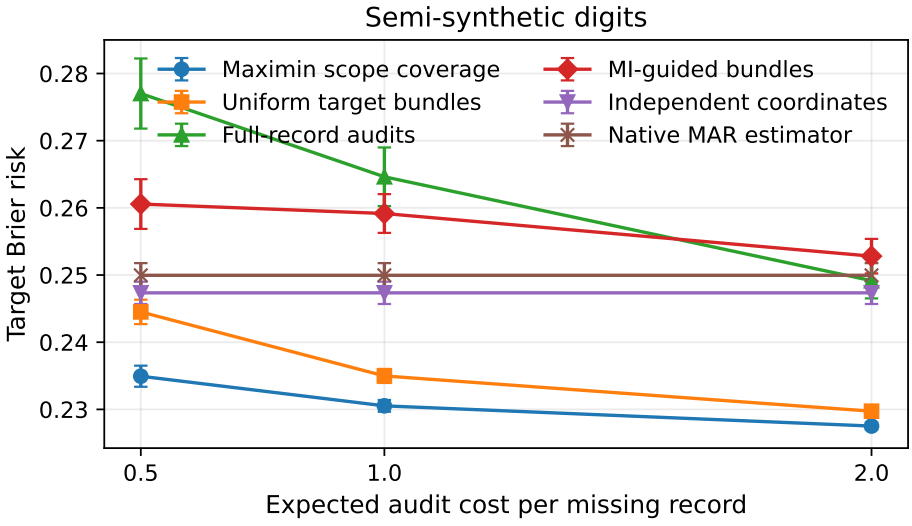

Figure 7: The analogous budget comparison on semi-synthetic digits. Uncertainty is conditional on the fixed images and split. Independent coordinate acquisition almost never covers a four-feature target scope at these budgets, so buying more isolated coordinates need not improve the observed learning curve.