THE PRICE OF HIDDEN ASSIGNMENTS: CALIBRATION AND MINIMAX RATES IN POOLED CATE LEARNING

Anonymous authors
Paper under double-blind review

ABSTRACT

How accurately can the conditional average treatment effect (CATE) be estimated from noisy pooled treatment reports? We study randomized experiments with G fixed-size pools, linked individual covariates and outcomes, hidden treatment assignments, known assignment probabilities, and a common unknown nondifferential reporting channel. Validation of m uniformly selected pools reveals their true OR states: whether any member received treatment. Under fixed overlap and channel bounds, pooling with a known channel preserves the usual CATE minimax rate. For an unknown channel, assuming smoothness only of the CATE, we establish matching classwise bounds for expected integrated squared risk, of order $G^{-\alpha} + \{G(1 + m + G\sigma_0)\}^{-1/2}$, where $\alpha = 2\beta/(2\beta + d)$, $\beta \in (0, 1]$ is CATE Hölder smoothness, d is covariate dimension, and σ_0 is the class-specified propensity-variance lower bound. The two terms quantify the regression error and the additional calibration cost. Validation and propensity heterogeneity jointly determine the calibration information scale $1 + m + G\sigma_0$, yielding distinct validation scales for improving on the no-validation minimax rate and recovering the known-channel rate. The lower bound retains treatment-effect information in complete validation records. An upper-bound construction nevertheless attains the same minimax rate using only covariates, true OR states, and noisy reports from a reserved validation subset, without its outcomes, and estimates CATE from the remaining pools. Finite-budget diagnostics show that qualification checks and clipping can suppress calibration gains in this construction. Across all 90 prespecified settings, its mean risk exceeded that of the zero-effect estimator.

1. INTRODUCTION

Estimating a conditional average treatment effect (CATE) typically requires knowing whether each individual received treatment. We consider a randomized experiment with incomplete treatment information: individual covariates remain linked to outcomes, but treatment assignments are withheld. Routine treatment information consists only of noisy logical OR reports for fixed-size pools. Each report indicates, possibly incorrectly, whether at least one pool member received treatment. Random validation reveals the true OR status of selected pools without directly releasing individual treatment labels. How accurately can treatment effects be estimated from these data?

This question is motivated by data-release settings in which individual treatment labels are withheld while aggregate reports and limited validation information remain available. We distinguish the cost of withholding individual treatment labels from the additional cost of reporting errors. Associations between outcomes and pool-level reports depend on both treatment effects and reporting errors.

The known-channel benchmark separates the cost of pooling from the additional cost of channel uncertainty. With fixed pool size, known propensity scores, and fixed overlap and channel bounds, an inverse response preserves the usual CATE minimax rate. Pooling therefore need not change the rate at which risk decreases, even though it conceals individual labels. When the channel is unknown, however, small changes in treatment effects can be masked by changes in reporting errors. The risk then includes an additional cost of channel calibration.

Suppose there are G pools, of which m are selected uniformly for validation, independently of the study variables. Let σ_0 denote the lower bound on propensity score variance defining the model class. Proposition 1 establishes matching classwise bounds for expected integrated squared risk:

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.

$$\inf_{\widehat{\tau}} \sup_{P \in \mathcal{P}(\sigma_0)} R_P(\widehat{\tau}) \asymp G^{-\alpha} + \{G(1 + m + G\sigma_0)\}^{-1/2}, \qquad \alpha = \frac{2\beta}{2\beta + d}.$$

Here, d is the covariate dimension and $\beta \in (0, 1]$ is the CATE's Hölder smoothness. The two terms represent regression difficulty and channel calibration cost. Validation and variation in propensity scores both provide calibration information, with an effective scale at the rate level of

$$I_{\mathrm{cal}} \asymp 1 + m + G\sigma_0.$$

Validation is therefore not the only source of channel calibration. When propensity score variation already supplies substantial information, validation has less scope to change the overall rate.

Whether the validation budget is sufficient depends on the objective: improving the original convergence rate or recovering the regression rate attainable with a known channel. Write $D_0 = 1 + G\sigma_0$ for the existing calibration scale and $H_{\mathrm{rate}} = G^{2\alpha-1}$ for the recovery scale. Validation changes the overall risk rate only when existing calibration information is far below the recovery threshold and the validation size is much larger than D_0. Once $D_0 + m \gtrsim H_{\mathrm{rate}}$, the total risk attains the standard regression rate.

For example, when $d = \beta = 1$ and $\sigma_0 = 0$, we have $\alpha = 2/3$ and $H_{\mathrm{rate}} = G^{1/3}$. Taking $m \asymp G^{1/6}$ improves the calibration term from $G^{-1/2}$ to $G^{-7/12}$, but does not yet attain $G^{-2/3}$. Recovering that regression rate requires $m \gtrsim G^{1/3}$. Improving the convergence rate and recovering the regression rate thus impose different validation requirements.

We establish these matching risk bounds and use them to characterize the validation scales required for improvement and recovery. Under the stated model conditions, the results require smoothness only of the CATE, not of the two outcome regression functions separately.

Outcomes in validated pools contain treatment-effect information, so the lower bound must account for the complete validated observation law. The upper-bound construction can nevertheless discard outcomes from a reserved validation subset and use only its (X, T, S) observations for calibration. Estimating the effect from the remaining pools still attains the same optimal rate. Attaining the optimal rate does not require using all available information, nor does it guarantee the best finite-sample performance. We therefore conclude with finite-budget calculations examining how the rate-optimal construction performs at specific sample sizes.

The analysis connects CATE theory with observed individual treatment labels, research on pooling and noisy reports, and the theory of calibration with unknown operators. Our focus is the full observation experiment in which these features coexist, together with its risk, validation requirements, and finite-budget performance.

2. Observation model and risk

We observe G disjoint pools, each containing a fixed number $k \geq 2$ of individuals. All kG individuals are independent and identically distributed. Each individual has covariates $X \in [0, 1]^d$ and bounded potential outcomes, $|Y(a)| \leq B$. Conditional on the covariates, treatment assignments are independent across individuals, with

$$A \mid X = x \sim \mathrm{Bernoulli}(e(x)),$$

where $e(x) \in [\eta, 1 - \eta]$ is known and $0 < \eta \leq 1/2$. Treatment assignment is independent of potential outcomes conditional on X. Consistency and no interference link observed outcomes to potential outcomes.

Individual treatment assignments are not released directly. Each pool provides only a report S on its true OR status $T = \max_i A_i$. All pools share an unknown sensitivity s and specificity z, satisfying

$$\kappa \leq s, z \leq 1 - \kappa, \qquad s + z - 1 \geq \kappa, \qquad 0 < \kappa < 1/3.$$

Conditional on T, the report has no further dependence on any member variables. Reporting errors are independent across pools. These assumptions specify a common, nondifferential reporting mechanism. The channel contrast is also bounded away from zero as the sample size increases.

The researcher observes linked individual covariates and outcomes, pool-level reports, and random validation results. Exactly m pools, with $0 \le m \le G$, are selected uniformly from all pools, independently of the study variables. Validation reveals their true OR statuses without error. All pools used for reference estimation, fitting, and calibration count toward the total budget G, and each revealed status counts toward m.

An estimator may use the supplied propensity score function and fixed inputs $(k, B, \eta, \kappa, d, \beta, L_\tau)$. It does not know the channel, outcome means, covariate distribution, or propensity score variance. It must be a measurable function of the observed records, validation indices, supplied inputs, and parameter-independent external randomization. It cannot use hidden treatment assignments or unknown distribution parameters. The minimax infimum permits a separate rule for each declared class index σ_0. The upper construction is stronger: one rule works without receiving either the actual propensity variance or σ_0. Implicit constants depend only on the fixed inputs and are uniform over G, m, σ_0. We use $a \asymp b$ for matching bounds up to such constants and $a = o(b)$ for $a/b \to 0$ along admissible sequences.

The target is

$$\tau(x) = \mathbb{E}[Y(1) - Y(0) \mid X = x],$$

which satisfies

$$|\tau(x) - \tau(x')| \le L_\tau \|x - x'\|^\beta, \qquad 0 < \beta \le 1.$$

Only the CATE must be smooth; the covariate law, propensity score function, and both outcome regressions may be nonsmooth. Risk is measured by expected integrated squared error:

$$R_P(\widehat{\tau}) = \mathbb{E}_P \|\widehat{\tau} - \tau\|^2_{L^2(P_X)}.$$

The model class $\mathcal{P}(\sigma_0)$ consists of distributions satisfying the preceding conditions and $\mathrm{Var}(e(X)) \ge \sigma_0$. The propensity score function associated with each distribution is supplied to the estimator as known information. We assume

$$0 \le \sigma_0 \le \min\{1/16, (1/2 - \eta)^2\}.$$

Our risk bounds characterize the worst case over this model class; they do not imply that every fixed distribution is equally difficult. Appendix K explains the boundaries of the theorem, including growing pool size, vanishing overlap, and weakening channel contrast.

3. The cost of hidden treatment assignments

3.1 The main minimax risk law

Proposition 1 separates the regression error present under a known channel from the additional calibration error caused by channel uncertainty.

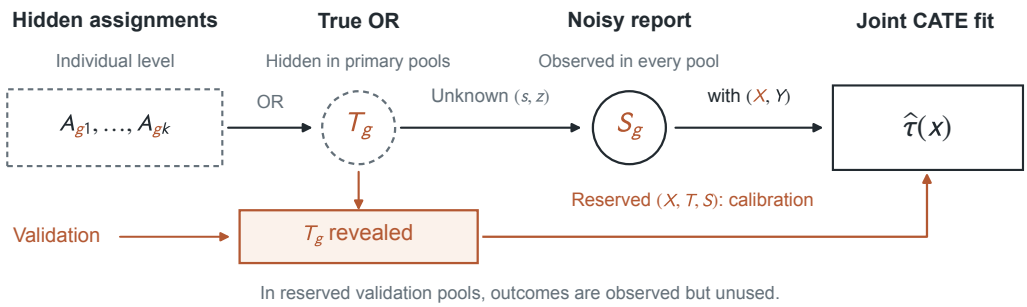


Figure 1: Observation mechanism for fixed-size pools. Individual covariates and outcomes are observed, whereas treatment assignments are hidden. All pools release noisy OR reports; validated pools additionally reveal their true OR statuses. The upper-bound construction calibrates the channel using (X, T, S) from a reserved validation subset and then estimates the CATE using linked records from the remaining pools.

Proposition 1 (Minimax risk). Let

$$r(G, m, \sigma_0) = G^{-\alpha} + \{G(1 + m + G\sigma_0)\}^{-1/2}, \qquad \alpha = \frac{2\beta}{2\beta + d}.$$

Under the conditions in Section 2, there exist positive constants C_{low} and C_{up}, independent of G, m, σ_0, such that

$$C_{\text{low}}\, r(G, m, \sigma_0) \le \inf_{\widehat{\tau}} \sup_{P \in \mathcal{P}(\sigma_0)} R_P(\widehat{\tau}) \le C_{\text{up}}\, r(G, m, \sigma_0). \tag{1}$$

The proof follows these two sources of difficulty. A Hölder hypercube construction with a fixed channel yields the regression lower bound. For the calibration term, we choose two models with the same supplied propensity score function and actual propensity variance σ_0. Their CATEs differ, but their observed-data distributions are close. We then separately quantify how well validated and unvalidated pools distinguish the models. Product likelihoods and Le Cam's method turn these calculations into a risk lower bound. The upper bound attains the same rate through sample splitting that respects the total budget, observable identities, moment bounds, and uniform inversion.

3.2 Information benchmarks

To interpret the risk bound, consider how risk changes when the researcher has more or less information. Under the same total pool budget, smoothness conditions, and loss, observing all individual treatment labels yields the rate $G^{-\alpha}$. When labels are hidden but the reporting channel is known, an inverse-transformed response still attains the same rate. The rates coincide, although the risk constants need not.

Retaining only the $m \ge 1$ validated pools allows a bounded response with conditional mean equal to the CATE. The minimax risk for this restricted experiment has order $m^{-\alpha}$. This approach discards unvalidated pools, however, and therefore does not characterize the best performance available from jointly using all the data.

For the known-channel benchmark, write $r_e = \mathbb{E}[1 - e(X)]$ and $b = (s + z - 1)r_e^{k-1}$. The inverse response

$$Z_{\text{known}} = \frac{Y\{(S - s)/b + 1 - e(X)\}}{e(X)(1 - e(X))}$$

has conditional mean $\tau(X)$. An independent fold estimates r_e with $O(G^{-1})$ mean squared error, and its pools count toward G. Each validated pool also gives

$$Z_{\text{val}} = \frac{Y_1}{e(X_1)} \left[1 - \frac{\mathbf{1}\{T = 0\}}{\prod_{i=1}^{k}(1 - e(X_i))}\right],$$

whose conditional mean is $\tau(X_1)$. Fixed overlap and pool size keep both responses bounded. Cell regression yields the upper rates; the corresponding full-label hypercube gives the matching lower orders.

These comparisons describe the roles of channel knowledge, validated records, and unvalidated records under different observation conditions. Inverse responses have precedents in group-testing research, including Delaigle and Tan (2024). We use this construction to distinguish the consequences of pooling from those of an unknown reporting channel. Appendices A and J give the benchmark constructions.

3.3 Why channel calibration is difficult

Channel calibration is particularly difficult when treatment effects are weak. Conditional on the independent pilot and normalization results, the observable statistics U, V in Appendix B satisfy

$$\mathbb{E}U = bt, \qquad \mathbb{E}V = bt^2,$$

where b is a channel factor and t is an effect coefficient. As t approaches zero, $\mathbb{E}V = bt^2$ becomes less sensitive to parameter changes. These moments then provide less information for distinguishing effect changes from channel changes.

The lower-bound construction exploits this feature. We perturb the effect and the channel jointly so that first-order changes in the joint observed-data distribution of unvalidated pools cancel. The two models then have different CATEs but generate very similar observations, making them difficult for any estimator to distinguish reliably.

The complete validated law makes this distinction explicit. In the lower-bound family, let $Y = BR$ with $R \in \{-1, 1\}$ and

$$\Pr(R = r' \mid A = a, X) = \tfrac{1}{2} + \theta(a - e(X))r'.$$

The marginal distribution of R is uniform for every θ. Given a pool's (X, R), set

$$v_\theta = \prod_{i=1}^{k}(1 - e_i)(1 - 2e_i\theta R_i).$$

The four conditional (T, S) probabilities are zv_θ, $(1-z)v_\theta$, $(1-s)(1-v_\theta)$, and $s(1-v_\theta)$. Holding $(s + z - 1)\theta$ fixed cancels the first-order effect–channel term in the unvalidated report law. The lower-bound calculation still retains all four probabilities for validated pools; replacing them by calibration-only marginals would discard effect information.

Validation helps distinguish the models, but its outcomes also contain treatment-effect information. Its contribution must therefore be calculated from the complete validated observation law. Accounting separately for validated and unvalidated pools yields the remaining testing difficulty represented by the calibration term in the risk bound.

3.4 Nonsmooth nuisance regressions and a common projection

The upper bound must address another issue: a smooth CATE does not imply that the two outcome regression functions are easy to estimate. Approximating them separately and then taking their difference may allow approximation errors to obscure the effect of interest.

A common projection makes the cancellation explicit. Write $\mu = \mathbb{E}[Y \mid X]$, $\bar{e} = \mathbb{E}e(X)$, and $M_0 = \mu(e - \bar{e})/\{e(1 - e)\}$. For the same cell projection Π,

$$f = \Pi(\tau + M_0), \qquad M = \Pi M_0, \qquad f - M = \Pi\tau.$$

Only the approximation error of τ requires smoothness control. The sampling errors of both fitted components remain and are bounded separately.

For $G \geq 44$, reserve $\ell = \min\{m, \lfloor G/4 \rfloor\}$ validated pools and split the remainder into eleven folds of size $q = \lfloor (G - \ell)/11 \rfloor$. The reserve estimates the channel factor b using only (X, T, S), giving

a bounded estimate b_{cal} with mean squared error $O(1/\ell)$ when $\ell \geq 1$. Empty status categories use the prescribed default. When $\ell = 0$, the calibration penalty has zero weight. The moment objective adds

$$\frac{\ell}{q}(b' - b_{\mathrm{cal}})^2.$$

Uniform inversion controls coefficient error, while an energy-weighted bound controls the contribution of channel error to CATE risk. With zero propensity variance and no validation, a quartic bound avoids division by zero calibration weight. The resulting bound is

$$\mathcal{A}_\tau + C\left\{\frac{J}{q+1} + [q(1+\ell+q\sigma)]^{-1/2}\right\},$$

where J is the cell count, $\sigma = \mathrm{Var}(e(X))$, and $\mathcal{A}_\tau$ is the Hölder approximation error. Choosing $J \asymp G^{d/(2\beta+d)}$ gives Proposition 1; for $G < 44$, the procedure returns zero. Every fold counts toward G. Section 6 examines the finite-budget behavior of this construction.

4. How much validation is enough?

Proposition 1 distinguishes the validation scales needed to improve the no-validation rate and to recover the known-channel regression rate. Write $D_0 = 1 + G\sigma_0$ and $H_{\mathrm{rate}} = G^{2\alpha-1}$.

4.1 When does validation improve the convergence rate?

Here, improvement means that the ratio of the minimax risk with validation to that without validation tends to zero. A reduction in the risk constant alone does not qualify. This occurs if and only if

$$D_0 = o(H_{\mathrm{rate}}), \qquad D_0 = o(m).$$

The first condition requires the calibration term to dominate total risk in the absence of validation. The second requires the validation size to be much larger than the amount of existing calibration information. If the standard regression rate is already attained without validation, additional validation may reduce the calibration term but cannot further improve the overall risk rate.

4.2 When is the regression rate recovered?

Even after validation improves the risk rate, the calibration term may remain larger than the regression term. Recovering the regression rate attainable with a known channel requires

$$D_0 + m \gtrsim H_{\mathrm{rate}}.$$

The calibration term is then at most of the same order as the regression term, but the two may still have the same order. The calibration term becomes negligible relative to the regression term only under the stronger condition $H_{\mathrm{rate}} = o(D_0 + m)$.

In the example with $d = \beta = 1$ and $\sigma_0 = 0$, $m \asymp G^{1/6}$ improves the rate, whereas $m \gtrsim G^{1/3}$ recovers $G^{-2/3}$. Validating a fixed proportion of pools recovers the regression rate, although a smaller validation subset may suffice.

Rate improvement also requires $2\beta > d$. Within our range $\beta \leq 1$, this requires $d = 1$ and $\beta > 1/2$. The lower bound on propensity score variance must also be small enough for $D_0 = o(H_{\mathrm{rate}})$ to hold.

For $d = \beta = 1$, using only $m \asymp G^{1/3}$ validated pools gives risk of order $m^{-2/3} = G^{-2/9}$. Jointly using all pools instead reaches $G^{-2/3}$. The validation-only benchmark therefore cannot substitute for the complete experiment when measuring the value of validation.

Table 1: Validation regimes for regular asymptotic sequences. Improvement means a vanishing minimax-risk ratio relative to no validation.

Regime	Risk order	Role of validation
$D_0 \gtrsim H_{\mathrm{rate}}$	$G^{-\alpha}$	Regression rate already attained.
$D_0 = o(H_{\mathrm{rate}}),\ m = O(D_0)$	$(GD_0)^{-1/2}$	No change in order.
$D_0 = o(m),\ m = o(H_{\mathrm{rate}})$	$(Gm)^{-1/2}$	Improvement before recovery.
$D_0 = o(H_{\mathrm{rate}}),\ m \gtrsim H_{\mathrm{rate}}$	$G^{-\alpha}$	Known-channel rate recovered.

5. Does attaining the optimal rate require all validation information?

Attaining the optimal rate does not require using outcomes from the reserved validation subset. These outcomes contain treatment-effect information, but channel calibration alone suffices for the rate bound.

The upper-bound construction retains only (X, T, S) from the reserved validation subset and discards its outcomes. The CATE is then estimated using (X, Y, S) from the remaining pools. It nevertheless attains the minimax risk rate of the full experiment. Conditional on covariates, the distribution of (T, S) in the reserved subset does not depend on the potential-outcome distribution and can be used separately for channel calibration.

The omitted outcomes remain informative. In a scalar example, two models with different treatment effects can have identical marginal distributions of pool-level statuses but different joint distributions with outcomes. The lower bound must therefore retain the complete validated observation law. Equal risk rates do not imply equivalent experiments. Under the specified class and loss, outcomes from the reserved validation subset are unnecessary for attaining the optimal rate. This rate result does not determine whether including those outcomes would improve risk constants.

6. Why asymptotic information gains may not yield finite-budget benefits

An optimal rate does not fully describe performance at a given budget. The construction also involves qualification checks, inversion, and clipping. Validation information may enter the computation without changing the final prediction; even when predictions change, risk need not decrease. This section examines the specific construction used to prove the upper bound.

6.1 Qualification checks and clipping can erase calibration effects

Across 84 exact comparisons in two prespecified scalar families, the construction matches its comparator with the same sample split and no calibration penalty. In every case, its risk exceeds that of the zero-effect estimator.

These cases satisfy a finite-budget condition under which predictions after qualification checks and clipping are algebraically independent of the calibration term. The penalty enters the computation but cannot affect the final prediction under this condition. This condition need not continue to hold as the budget increases. Numerical Supplement S1 gives the condition and exact calculations.

6.2 Finite-budget Monte Carlo diagnostics

The confirmatory simulations vary the number of pools, validation size, and propensity variance across two CATE scenarios, covering 90 settings with 200 replicates per setting.

To isolate the effect of the calibration penalty, we compare the construction with a version that removes only this penalty, using identical sample splits and moment statistics. We define the paired risk difference as the risk with the penalty minus the risk without it; negative values favor calibra-

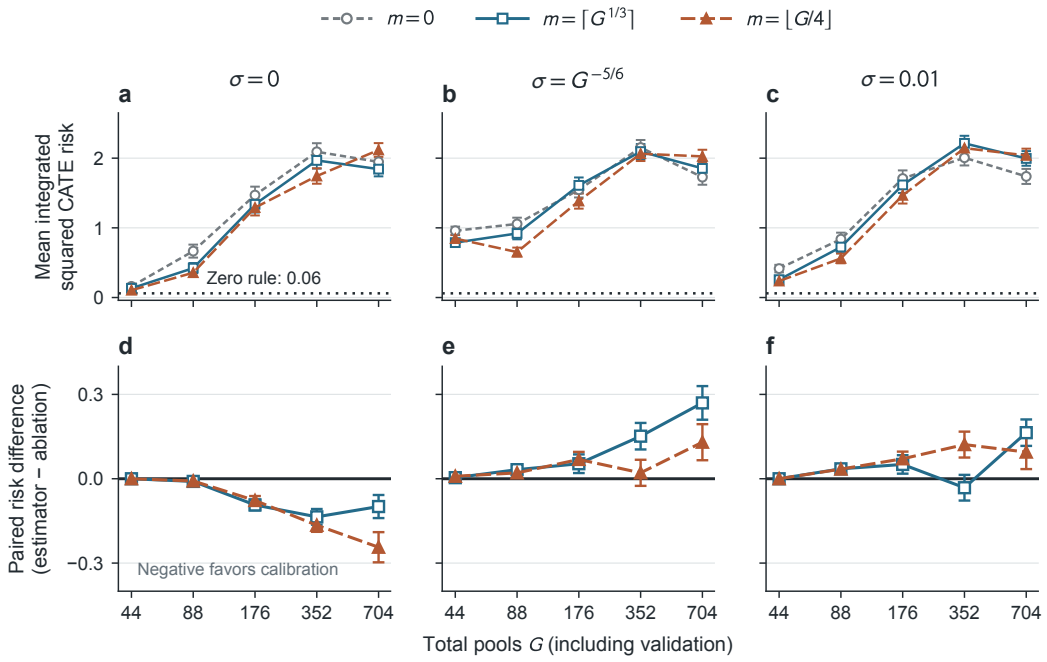


Figure 2: Fixed-budget results for the rough-baseline CATE scenario. The top row shows the estimator's mean integrated squared risk; the reference line indicates the zero-effect estimator's risk of 0.06. The bottom row shows paired risk differences between the estimator and a version using the same sample split but omitting the calibration penalty. Negative values favor calibration. Each point averages 200 replicates. Error bars represent one Monte Carlo standard error, with paired standard errors used in the bottom row. Validated pools count toward G. When $m = 0$, the risk difference is identically zero and is omitted from the bottom row.

tion. Figure 2 shows these differences and their Monte Carlo standard errors for the rough-baseline scenario. Numerical Supplement S6 provides detailed results.

Among the 12,000 runs with positive validation size, the calibration penalty changed the final clipped predictions in 2,681 runs (22.34%). However, across all 90 settings, the construction's mean integrated squared risk exceeded that of the corresponding zero-effect estimator. Further diagnostics revealed substantial instability before clipping and a long right tail in coefficient errors.

The paired comparison measures the penalty's contribution to risk, whereas the zero-effect benchmark assesses the performance of the complete estimator. A change in predictions alone does not establish that the penalty reduces risk.

6.3 A SCALAR FINITE-BUDGET LOWER BOUND

Examining a particular estimator does not establish what risk is attainable from the available information. We therefore compute upper and lower bounds on minimax risk in the scalar full-observation experiment of Appendix H. At $G = 44, m = 0, k = 2, e = 1/2$, the minimax squared-CATE risk R^* satisfies

$$0.946362654615 \leq R^* \leq \frac{2779}{2823} = 0.984413744244\ldots.$$

This interval pertains to the scalar parameter family and effect scale specified in Appendix H. These differ from the simulation scenarios in Section 6.2, so the numerical values should not be compared directly with the risks in Figure 2. The interval bounds finite-budget risk in that scalar experiment,

but it neither identifies the exact minimax risk nor directly establishes an equivalent limitation for the full nonparametric model class.

The simulations assess the performance achieved by a specific method. This calculation instead starts from the full observed-data distribution and bounds the optimal risk in the scalar experiment. Model specifications and computational details appear in Appendix H and Numerical Supplement S2.

7. Related work

Minimax CATE theory with observed individual treatment labels provides the regression benchmark for our analysis. Kennedy et al. (2024) study minimax rates for heterogeneous causal effects with a different loss and different nuisance smoothness assumptions. Nie and Wager (2021) and Foster and Syrgkanis (2023) establish learning guarantees using individual treatment information. Gao and Han (2020) allow effects smoother than the outcome regressions. These studies support learning that focuses on the effect itself, but their available treatment information differs from ours. Their risk results do not directly apply to an unknown pooled reporting channel.

Pooling and group-testing research provides misclassification corrections for prevalence estimation, inverse-transformed responses, and regression models with imperfect tests (Zhang et al., 2014; Delaigle and Hall, 2015; Delaigle and Tan, 2024). The misclassification literature also examines identification under other observation schemes and assumptions (Lewbel, 2007; Chen et al., 2008; Di-Traglia and Garcia-Jimeno, 2019; Kasahara and Shimotsu, 2022). For misclassified treatment labels, Braun et al. (2017) use gold-standard labels from a validation sample to adjust propensity score analyses. Shang et al. (2026) model true individual treatment as a latent variable and allow available validation information to be incorporated. Both studies concern misclassified individual treatment labels, whereas routine reports and validation results in our setting are at the pool level.

Unknown-operator theory motivates auxiliary-sample calibration and joint target–nuisance perturbations (Johannes, 2009; Hoffmann and Reiss, 2008; Johannes and Schwarz, 2013). These ideas help analyze ambiguity between treatment effects and the reporting channel. Our bounds must retain the complete joint law of individual outcomes, noisy OR reports, and validated OR states. Appendix C compares the observation laws, parameter classes, and losses; the cited results do not directly supply the risk law for this experiment.

8. Discussion

Under our assumptions, fixed-size pooling with a known reporting channel preserves the usual CATE regression rate. An unknown channel adds a calibration cost. Validation and propensity score heterogeneity jointly supply calibration information, and the amount of existing information determines whether additional validation can change the overall risk rate.

These conclusions rely on fixed pool size, known propensity scores, fixed overlap and channel bounds, and a nondifferential reporting mechanism. Smoothness is required only of the CATE, allowing the two outcome regression functions to be individually nonsmooth, but the remaining assumptions must still hold. The current risk formula does not apply directly when pool size grows, overlap deteriorates, or channel contrast weakens.

These are classwise rate comparisons, not exact finite-sample risk guarantees. Equal rates do not imply equivalent experiments or equal risk constants, and improvements for a particular distribution need not follow the minimax ordering.

The finite-budget diagnostics show that a rate-optimal construction can still have substantial finite-sample risk. Improving estimation, stabilizing inversion, and reducing risk under the same observation and validation budgets remain open tasks. Whether true OR statuses can be validated without error, and at what cost, must also be assessed in specific applications.

AI USE STATEMENT

Generative AI tools assisted with developing the theoretical model, formulating mathematical claims, drafting proof arguments, designing and reviewing methodology and deterministic checks, implementing methods and editing analysis software, revising scientific figures, searching and summarizing literature, translating and revising the manuscript, and interpreting results. The authors reviewed the AI-assisted text, code, and outputs against the current source records and deterministic checks. The proof checks reported in the accompanying materials are author-side checks assisted by automated algebra and replay scripts; they are not independent mathematical certification. The authors take responsibility for the final text, mathematics, software, and reported results.

REPRODUCIBILITY STATEMENT

Appendix B gives the proof and the accompanying records document author-side checks of Proposition 1. The anonymous package includes code and dependencies for the 84 exact-risk cases and finite-budget certificate, the frozen 18,000-run data, and 12,000 retained positive-validation predictions. Its documented commands were executed from a clean extraction: exact results were recomputed, saved summaries checked, 26 audit rows regenerated, and 1,200 existing confirmation seeds replayed for the targeted error-path diagnostic. Numerical Supplement S4 specifies the commands and scope. These checks are distinct from independent replication or a full rerun of the Monte Carlo study.

REFERENCES

Chen, Xiaohong, Yingyao Hu, and Arthur Lewbel. Nonparametric identification of regression models containing a misclassified dichotomous regressor without instruments. Economics Letters 100(3), 381–384 (2008). https://doi.org/10.1016/j.econlet.2008.03.015.

Braun, Danielle, Malka Gorfine, Giovanni Parmigiani, Nils D. Arvold, Francesca Dominici, and Corwin Zigler. Propensity Scores with Misclassified Treatment Assignment: A Likelihood-Based Adjustment. Biostatistics 18(4), 695–710 (2017). https://doi.org/10.1093/biostatistics/kxx014.

Delaigle, Aurore, and Peter Hall. Nonparametric methods for group testing data, taking dilution into account. Biometrika 102(4), 871–887 (2015). https://doi.org/10.1093/biomet/asv049.

Delaigle, Aurore, and Ruoxu Tan. Group Testing Regression Analysis with Missing Data and Imperfect Tests. Statistica Sinica 34, 201–228 (2024). https://doi.org/10.5705/ss.202021.0382.

DiTraglia, Francis J., and Camilo García-Jimeno. Identifying the effect of a mis-classified, binary, endogenous regressor. Journal of Econometrics 209(2), 376–390 (2019). https://doi.org/10.1016/j.jeconom.2019.01.007.

Foster, Dylan J., and Vasilis Syrgkanis. Orthogonal statistical learning. The Annals of Statistics 51(3), 879–908 (2023). https://doi.org/10.1214/23-AOS2258.

Gao, Zijun, and Yanjun Han. Minimax Optimal Nonparametric Estimation of Heterogeneous Treatment Effects. Advances in Neural Information Processing Systems 33, 21751–21762 (2020). https://proceedings.neurips.cc/paper/2020/hash/f75b757d3459c3e93e98ddab7b903938-Abstract.html.

Hoffmann, Marc, and Markus Reiss. Nonlinear estimation for linear inverse problems with error in the operator. The Annals of Statistics 36(1), 310–336 (2008). https://doi.org/10.1214/009053607000000721.

Johannes, Jan, and Maik Schwarz. Adaptive Gaussian Inverse Regression with Partially Unknown Operator. Communications in Statistics—Theory and Methods 42(7), 1343–1362 (2013). https://doi.org/10.1080/03610926.2012.731548.

Johannes, Jan. Deconvolution with unknown error distribution. The Annals of Statistics 37(5A), 2301–2323 (2009). https://doi.org/10.1214/08-AOS652.

Kasahara, Hiroyuki, and Katsumi Shimotsu. Identification of regression models with a misclassified and endogenous binary regressor. Econometric Theory 38(6), 1117–1139 (2022; online 2021). https://doi.org/10.1017/S0266466621000451.

Kennedy, Edward H., Sivaraman Balakrishnan, James M. Robins, and Larry Wasserman. Minimax rates for heterogeneous causal effect estimation. The Annals of Statistics 52(2), 793–816 (2024). https://doi.org/10.1214/24-AOS2369.

Lewbel, Arthur. Estimation of Average Treatment Effects with Misclassification. Econometrica 75(2), 537–551 (2007). https://doi.org/10.1111/j.1468-0262.2006.00756.x.

Nie, Xinkun, and Stefan Wager. Quasi-oracle estimation of heterogeneous treatment effects. Biometrika 108(2), 299–319 (2021; online 24 September 2020). https://doi.org/10.1093/biomet/asaa076.

Shang, Yimeng, Yu-Han Chiu, and Lan Kong. A Latent Variable Approach for Causal Effect Estimation Under Misclassified Treatment Assignment. Statistics in Medicine 45(8–9), e70528 (2026). https://doi.org/10.1002/sim.70528.

Zhang, Zhiwei, Chunling Liu, Sungduk Kim, and Aiyi Liu. Prevalence estimation subject to misclassification: the mis-substitution bias and some remedies. Statistics in Medicine 33(25), 4482–4500 (2014). https://doi.org/10.1002/sim.6268.

Appendix organization. Appendices A–B contain the information benchmarks and the complete proof of Proposition 1. C–D give literature comparisons, reductions and budget corollaries; E–K give scope and auxiliary bounds. Within L, L1–L2 are *theoretical diagnostics*, not prerequisites of Proposition 1; L3 is *proof support* for the function-risk algebra in P9. Numerical Supplement S1–S7 is provided in a separate PDF. S1–S4 document exact calculations and reproduction; S5–S7 contain the plug-in, fixed-budget and covariance diagnostics.

Appendix A: Information benchmarks

Use the model and loss of Section 2 with fixed $k \geq 2$. The same G pools contain $n = kG$ members; $e \in [\eta, 1-\eta]$ is supplied, outcomes are bounded, only τ is β-Hölder, and P_X is arbitrary. Channel parameters (s, z) obey the compact bounds of Section 2. Define $c = s + z - 1$, $r_e = \mathbb{E}[1 - e(X)]$, and $b = cr_e^{k-1}$.

Full labels

Observe (X, A, Y) for all members. The response $Z_{\text{full}} = Y(A - e(X))/\{e(X)(1 - e(X))\}$ has conditional expectation $\tau(X)$ and absolute value at most B/η. The ordinary cell-mean histogram with J cells therefore has expected error from $\Pi\tau$ at most $3(B/\eta)^2 J/(kG + 1)$. Add Hölder approximation $O(J^{-2\beta/d})$. Taking J of order $G^{d/(2\beta+d)}$ yields $G^{-\alpha}$, where $\alpha = 2\beta/(2\beta + d)$. The existing full-label Hölder hypercube in Appendix B gives the matching lower bound at the same supplied step e and variance floor σ_0. This is reuse of an existing proof ingredient, not a new minimax principle. The reduction from full labels to the pooled experiment is a parameter-independent simulation kernel: generate the OR state and then the fixed-channel report using only the supplied e, the fixed channel, and fresh parameter-free randomness.

Known-channel pools

Observe original linked (X, Y, S) and additionally receive (s, z) as oracle inputs; P_X and therefore r_e remain unknown. Write $m(x) = \mathbb{E}[Y \mid X = x]$ and $\mu_0(x) = \mathbb{E}[Y(0) \mid X = x]$. Marginalizing the other pool members gives $\mathbb{E}[YS \mid X = x] = sm(x) - b(1 - e(x))\mu_0(x)$. Consequently the response $Z_b = Y\{(S - s)/b + 1 - e(X)\}/\{e(X)(1 - e(X))\}$ has conditional mean $\{m(x) - \mu_0(x)\}/e(x) = \tau(x)$. Its absolute value is bounded by a fixed constant, since $b \geq \kappa\eta^{k-1}$. No regression smoothness of m or μ_0 is used. This identity reformulates the existing pooled identities; it is not claimed as new identification.

To avoid supplying the unknown r_e, split pools into $N_0 = \lfloor G/2 \rfloor$ reference pools and $N_1 = G - N_0$ learning pools for $G \geq 2$. Estimate r_e using the first member of each reference pool and set $\widehat{b} = c\widehat{r}_e^{k-1}$. Each summand lies in $[\eta, 1 - \eta]$, so $\widehat{b} \geq \kappa\eta^{k-1}$; no estimated denominator approaches zero. Independence is by complete pools.

Given $\widehat{r}_e$, the learning response with $\widehat{b}$ has conditional mean $\widehat{g}(x) = \tau(x) + (1 - b/\widehat{b})\mu_0(x)/e(x)$. Its bound is uniform in $\widehat{r}_e$, and $\mathbb{E}\|\widehat{g} - \tau\|^2 \leq (B/\eta)^2\mathbb{E}(1 - b/\widehat{b})^2 \leq C/N_0$, because the map

$r_e \mapsto c r_e^{k-1}$ is Lipschitz with fixed constants and $\mathbb{E}(\widehat{r}_e - r_e)^2 \le 1/(4N_0)$. The unknown-mass histogram bound, projection contraction, and a square inequality give

$$\mathbb{E}\|\widehat{\tau} - \tau\|^2 \le C\{J/(N_1+1) + 1/N_0 + J^{-2\beta/d}\}.$$

Choose the same order of J and clip to $[-2B, 2B]$; $G = 1$ uses zero. Since $\alpha < 1$, the rate is $G^{-\alpha}$ uniformly over the permitted σ_0. For the lower bound, fix a legal channel in the existing Hölder hypercube. Full labels plus this fixed channel simulate S. Any faster known-channel pooled rule would contradict the full-label hypercube lower bound. Thus the oracle pooled risk has matching order $G^{-\alpha}$ at fixed k.

Unknown-channel pools

Proposition 1 supplies $R_{\text{unknown}} \asymp G^{-\alpha} + (G + G^2\sigma_0)^{-1/2}$. Known channel is more informative than unknown channel by dropping its supplied parameters. The full-label versus known-channel lower-bound comparison is performed at a fixed channel; it does not imply a universal simulation map that remains valid when the channel parameters are unknown. Equality of minimax orders does not imply equivalence of experiments, finite-risk equality or equality of optimal constants.

At fixed k, known-channel pooling entails no extra order relative to full labels under this model. The extra term is associated with unknown-channel estimation, not an unavoidable fixed-k pooling penalty. For $d \ge 2\beta$ it is masked by the regression term; otherwise it can dominate.

Why validation requires the complete law

For m validated study pools observing true OR T, validated (X, Y, T) rows already provide noiseless pooled information. They are not channel-only observations, so a valid lower bound must use the full augmented likelihood of (X, Y, S, T). Appendix B does so. Its upper construction charges validation and learning to the same G pools and uses only a reserved outcome-free validation block for channel calibration.

Appendix B: Random validation has a uniform CATE calibration threshold

The proof first constructs a complete-law two-point lower bound and a fixed-channel Hölder hypercube. It then gives a sample-split estimator, controls channel calibration uniformly over vanishing propensity variance, and transfers the resulting scalar inversion bound to integrated CATE risk. The argument keeps all sample charges explicit and does not assume smooth arm-specific outcome regressions.

P1. Experiment, quantifiers and conclusion

Fix an integer $k \ge 2$, $B > 0$, $0 < \eta \le 1/2$, $0 < \kappa < 1/3$, an integer $d \ge 1$, $0 < \beta \le 1$, and $L_\tau > 0$. There are $G \ge 1$ independent pools of k iid members. Covariates lie in $[0,1]^d$ with arbitrary law P_X. A supplied measurable propensity e lies in $[\eta, 1-\eta]$. Conditional randomization is independent of the bounded potential outcomes, $|Y(a)| \le B$; assignments are independent across members and there is no interference. Only the CATE τ is assumed β-Hölder:

$$|\tau(x) - \tau(x')| \le L_\tau \|x - x'\|^\beta.$$

The observation contains every linked (X, Y) and a pool bit S. Conditional on the true OR $T = \max_i A_i$, S has sensitivity s and specificity z and no further dependence on member variables. Independent pool errors use a common unknown channel with $\kappa \le s, z \le 1-\kappa$ and $c = s + z - 1 \ge \kappa$.

An observed uniform subset I of exactly m pools is drawn independently of all study variables, where $0 \le m \le G$. The true T is additionally observed on I; individual A remains hidden. The same study budget G pays for every fold and reference; m counts additional validation measurements. No population covariate law or channel parameter is supplied to the procedure. The fixed class inputs $(k, B, \eta, \kappa, d, \beta, L_\tau)$ and the function e are available to the procedure, but σ_0 is not.

For $0 \le \sigma_0 \le \min\{1/16, (1/2-\eta)^2\}$, let $\mathcal{P}(\sigma_0)$ include every such law and supplied function e with $\sigma = \mathrm{Var}(e(X)) \ge \sigma_0$. Risk is the training expectation of integrated squared error at a fresh covariate. Estimators follow the information-set and parameter-independent external-randomization convention of Section 2; the minimax infimum may be indexed by the declared class σ_0. Write $\alpha = 2\beta/(2\beta+d)$.

Proposition. There are positive finite constants $C_{\mathrm{low}}, C_{\mathrm{up}}$, depending only on the fixed class inputs, such that simultaneously for every G, m, σ_0,

$$\begin{aligned} C_{\mathrm{low}} r(G,m,\sigma_0) \le \inf_{\widehat{\tau}} \sup_{P \in \mathcal{P}(\sigma_0)} R_P(\widehat{\tau}) \le C_{\mathrm{up}} r(G,m,\sigma_0), \\ r(G,m,\sigma_0) = G^{-\alpha} + \{G(1+m+G\sigma_0)\}^{-1/2}. \end{aligned} \tag{P1}$$

The upper procedure below does not receive σ or σ_0. The lower bound is classwise: its supplied propensity can depend on σ_0, and its two alternative laws can depend on G, m, σ_0. No pointwise lower bound for each fixed e or law follows. Constants are not uniform as k grows or κ vanishes.

P2 proves the complete-law calibration lower bound, and P3 supplies the regression lower bound. P4 allocates the upper construction's folds; P5–P6 establish its observable identities and moment error bounds. P7–P8 control augmented inversion, including the zero-variance endpoint, and P9 transfers the resulting scalar bound through the common projection to integrated CATE risk.

P2. The complete validated law supplies the channel lower bound

Take uniform X and the known two-valued propensity $e = 1/2 \pm \sqrt{\sigma_0}$ on the two halves of the cube. Because $\sigma_0 \le \min\{1/16, (1/2-\eta)^2\}$, it lies in $[1/2-\sqrt{\sigma_0}, 1/2+\sqrt{\sigma_0}] \subseteq [\eta, 1-\eta]$ and has variance exactly σ_0. This chooses the boundary of the class $\sigma \ge \sigma_0$; both alternatives retain this same e. Let $Y = BR$ with $R \in \{-1, 1\}$ and

$$\Pr(R = r' \mid A = a, X) = \tfrac{1}{2} + \theta(a - e(X))r', \qquad |\theta| \le \tfrac{1}{2}.$$

Independent potential outcomes with these marginals realize the causal model. The CATE is $2B\theta$, and the observed member marginal has independent uniform R, independent of X, at every θ. Conditional on the complete (X, R) in a pool,

$$v_\theta = \Pr(T = 0 \mid X, R) = \prod_{i=1}^{k} (1 - e_i)(1 - 2e_i\theta R_i).$$

In order $(T, S) = (0,0), (0,1), (1,0), (1,1)$, the four validated probabilities are

$$z v_\theta, \quad (1-z)v_\theta, \quad (1-s)(1-v_\theta), \quad s(1-v_\theta). \tag{P2}$$

Marginalizing T gives $\Pr(S = 1 \mid X, R) = s - c v_\theta$. This law includes every observed outcome and covariate. Validated T carries effect information even when the outcome marginal is common; the proof does not discard that information.

Set $c_0 = (1-\kappa)/2$, $s_0 = z_0 = (3-\kappa)/4$, and $h_{\max} = (1-3\kappa)/8$. Compare (s_0, z_0, θ) with

$$\begin{aligned} c_1 = c_0 + h, \quad s_1 = s_0 + 2^{-k}h, \quad z_1 = z_0 + (1-2^{-k})h, \\ \theta_1 = a\theta, \qquad a = \frac{c_0}{c_0 + h}, \qquad 0 \le h \le h_{\max}. \end{aligned}$$

At $h = 0$ the alternatives coincide. At the other endpoint, $z_1 \le z_0 + h_{\max} = (7-5\kappa)/8 \le 1-\kappa$; $s_1 \le z_1$ and both coordinates increase from $s_0 = z_0 \ge \kappa$. Also $c_0 \ge \kappa$ and $c_1 \ge c_0$. Thus the entire closed interval is legal. Sign orthogonality conditional on X gives the exact squared distance between their unvalidated report probabilities:

$$D_U = h^2\big[(1/4+\sigma_0)^k - (1/4)^k\big] + c_0^2 \sum_{j=2}^{k} \binom{k}{j} (4\theta^2)^j (1-a^{j-1})^2 (1/4+\sigma_0)^{k-j} (1/4-\sigma_0)^{2j}. \tag{P3}$$

To derive P3, expand $\prod_i (1-2e_i\theta R_i)$ by subsets of signs. The singleton coefficients cancel because $c_1\theta_1 = c_0\theta$. The constant coefficient differs by $h\{\prod_i(1-e_i) - 2^{-k}\}$. Different sign subsets are orthogonal. Finally use $\mathbb{E}(1-e)^2 = 1/4+\sigma_0$ and $e(1-e) = 1/4-\sigma_0$ for the two-valued propensity.

Let $A_{\max} = 5/16$, $\theta_{\max} = 1/2$, and

$$C_1 = kA_{\max}^{k-1},$$
$$C_2 = \sum_{j=2}^{k} \binom{k}{j} 4^j (j-1)^2 \theta_{\max}^{2j-4} (1/16)^j A_{\max}^{k-j}.$$

Since $1-a^{j-1} \le (j-1)h/c_0$, P3 is at most $h^2(C_1\sigma_0 + C_2\theta^4)$. Both Bernoulli report probabilities lie in $[\kappa, 1-\kappa]$. Their one-pool chi-square divergence $\chi^2(U_0\|U_1)$, after averaging over the common (X, R) marginal, is therefore at most that envelope divided by $\kappa(1-\kappa)$.

For the validated law, each factor is $(1-e_i)(1-2e_i\theta R_i)$. Since $e_i \in [1/4, 3/4]$ and $|\theta| \le 1/2$, its minimum is at least $(1-e_i)^2 \ge 1/16$ and its maximum is at most $1-e_i^2 \le 15/16$. Hence v_θ and $1-v_\theta$ are at least 16^{-k}. Every atom of P2 is at least $\kappa 16^{-k}$. Differentiating v_θ gives $|\partial_\theta v_\theta| \le k/2$. Moreover $|\theta - \theta_1| \le h/(2c_0)$, and both channel coordinates change by at most h. Each atom therefore changes by at most $h\{1 + k/(4c_0)\}$, yielding

$$\chi^2(V_0\|V_1) \le C_V h^2, \qquad C_V = \frac{4\cdot 16^k}{\kappa}\left(1 + \frac{k}{4c_0}\right)^2. \tag{P4}$$

Given the observed subset I, independence gives the exact product identity

$$1 + \chi^2(P_0^{(G,m)}\|P_1^{(G,m)}) = \big(1+\chi^2(U_0\|U_1)\big)^{G-m}\big(1+\chi^2(V_0\|V_1)\big)^m. \tag{P5}$$

The subset probabilities are parameter-free. Every subset of size m has the same conditional divergence, so averaging over I preserves P5.

Put $z_G = \sigma_0 + (m+1)/G$, $\epsilon = 1/10$, and

$$C_{\mathrm{div}} = \frac{C_1 + C_2\epsilon^2}{\kappa(1-\kappa)} + C_V, \qquad d_0 = \min\left\{h_{\max}, \sqrt{\frac{\log(5/4)}{C_{\mathrm{div}}}}\right\},$$
$$\theta^2 = \epsilon\sqrt{z_G}, \qquad h = \frac{d_0}{\sqrt{Gz_G}}.$$

Here $1/G \le z_G \le 33/16$, $\theta^2 < 1/4$, and $h \le h_{\max}$. The product exponent in P5 is at most

$$h^2\left[\frac{Gz_G(C_1 + C_2\epsilon^2)}{\kappa(1-\kappa)} + mC_V\right] \le d_0^2 C_{\mathrm{div}} \le \log(5/4).$$

because $\sigma_0 \le z_G$ and $m \le Gz_G$. Chi-square is at most $1/4$ and total variation at most $1/4$. The two constant CATEs have squared separation

$$\Delta^2 = \frac{4B^2\theta^2h^2}{(c_0+h)^2} \ge \frac{4B^2\epsilon d_0^2}{G\sqrt{z_G}}.$$

Decode the closer constant from any estimated function. An error entails loss at least $\Delta^2/4$, and average testing error is at least $3/8$. This proves a fixed positive multiple of $\{G(1+m+G\sigma_0)\}^{-1/2}$ uniformly in all three varying inputs. Constant CATEs are Hölder for any $L_\tau > 0$.

P3. A FIXED-CHANNEL HYPERCUBE RETAINS THE REGRESSION LOWER BOUND

Use the same uniform X and supplied two-valued e. Fix one legal channel. For $Y \in \{-B, B\}$ and $|\tau| \le B/2$, set $\Pr(Y = B \mid A, X) = 1/2 + \tau(X)(A - e(X))/(2B)$. These arm laws give CATE τ and probabilities in $[1/4, 3/4]$. Full individual labels can generate both T and S and the independent validation subset at this fixed channel. Thus a lower bound for kG full-label members also bounds our experiment.

Divide the cube into j^d subcubes Q_v of width $h = 1/j$. For independent signs ω_v, use $\tau_\omega(x) = a_0\omega_v \operatorname{dist}(x, Q_v^c)^\beta$ within Q_v, where $0 < a_0 \le \min(B/2, L_\tau/2)$. These functions meet the difference Hölder bound, including across cell boundaries. Let $C_{\mathrm{bd}} = \int_{[0,1]^d} \operatorname{dist}(u, ([0,1]^d)^c)^{2\beta}\, du > 0$. Flipping one sign gives squared separation $4a_0^2 C_{\mathrm{bd}} h^{2\beta+d}$.

For completeness, the cross-cell Hölder bound follows from the line segment between two points. If x and y lie in different cells, that segment crosses the boundary of each relevant cell, so their two boundary distances sum to at most $\|x - y\|$. For $0 < \beta \le 1$,

$$a^\beta + b^\beta \le 2^{1-\beta}(a+b)^\beta \le 2(a+b)^\beta.$$

Thus opposite signs across cells change τ by at most $2a_0\|x-y\|^\beta \le L_\tau\|x-y\|^\beta$. Within one cell, the distance function is one-Lipschitz and $|a^\beta - b^\beta| \le |a-b|^\beta$. This verifies the global difference bound.

The one-member KL is at most $\|\tau - \tau'\|^2_{L^2(P_X)}/(3B^2)$ by the Bernoulli KL inequality and $\mathbb{E}[(A - e)^2 \mid X] \le 1/4$. Choose $j = \lceil C(kG)^{1/(2\beta+d)} \rceil$ with fixed C large enough that neighboring full-data KL is at most $1/8$. Their total variation is at most $1/4$. Nearest-sign decoding on each subcube has average error at least $3/8$; summing its squared losses gives a positive fixed multiple of $h^{2\beta}$, hence $G^{-\alpha}$. Combining this lower bound with P2 by $\max(u, v) \ge (u+v)/2$ proves the lower half of P1.

P4. THE UPPER CONSTRUCTION CHARGES CALIBRATION AND LEARNING TO THE POOL BUDGET

For $G < 44$ use zero. For $G \ge 44$ reserve $\ell = \min\{m, \lfloor G/4 \rfloor\}$ validated pools by a deterministic index rule within I. Conditional on I, their observations are iid and independent of all remaining pools. Set $q = \lfloor (G - \ell)/11 \rfloor$. Use $11q$ remaining pools and leave at most ten unused. The bounds

$$\frac{G}{22} \le q \le \frac{G}{11}, \quad \ell \le m, \quad \ell \ge \frac{m}{8} \text{ if } m \ge 1, \quad \frac{\ell}{q} \le \frac{11}{2}. \tag{P6}$$

follow from $G - \ell \ge 3G/4$, $\lfloor x \rfloor \ge x - 1$, and $m \le G$. All constants below use these inequalities only; no limit fixing m/G is taken.

Write $r_e = \mathbb{E}[1 - e(X)]$, $b = cr_e^{k-1}$, $b_\star = \kappa\eta^{k-1}$, $H_Y = B/\{\eta(1-\eta)\}$, and $K = H_Y/b_\star$. On the reserve estimate s from reports with $T = 1$, z from one minus reports with $T = 0$, and r_e from first-member $1 - e(X)$. Use $1/2$ for an empty status category. Project $(\widehat{s} + \widehat{z} - 1)\widehat{r}_e^{k-1}$ onto $[b_\star, 1]$ to obtain b_{cal}. With $p_0 = r_e^k$ and $p_1 = 1 - p_0$, both category probabilities are bounded below by fixed positive constants.

When $\ell \ge 1$, define $\mathbf{1}\{N > 0\}/N = 0$ at $N = 0$. For $N \sim \mathrm{Bin}(\ell, p)$, the cell-mean MSE is at most $\frac{1}{4}\mathbb{E}[\mathbf{1}\{N > 0\}/N] + \frac{1}{4}\Pr(N = 0) \le 3/\{4(\ell+1)p\}$. This follows from $\mathbf{1}\{N > 0\}/N \le 2/(N+1)$ and $p(1-p)^\ell \le 1/(\ell+1)$. Also $\mathbb{E}(\widehat{r}_e - r_e)^2 \le 1/(4\ell)$. Lipschitz continuity and projection give

$$\begin{aligned} \mathbb{E}(b_{\mathrm{cal}} - b)^2 &\le 3\big\{\mathbb{E}(\widehat{s} - s)^2 + \mathbb{E}(\widehat{z} - z)^2 + (k-1)^2\mathbb{E}(\widehat{r}_e - r_e)^2\big\} \\ &\le C/\ell. \end{aligned} \tag{P7}$$

This step allows dependence among the three estimates. At $\ell = 0$ choose any b_{cal} in the interval and give its penalty zero weight.

Allocate the remaining folds as follows: propensity variance $2q$; report mean q; untreated-propensity mean q; shared pilot/correction q; normalization q; report-propensity calibration $2q$; primary/reference/third outcome moments $3q$. The index-only allocation, conditional on I, is:

Block	Pools	Output and upstream inputs
Reserve	ℓ	b_{cal} from (X, T, S); independent of the eleven folds
1–2	$2q$	X_v from propensity differences
3	q	$\widehat{p}$ from reports
4	q	$\widehat{r}_e$ from supplied propensities
5	q	$F, \widehat{M}$ using $\widehat{p}, \widehat{r}_e$; these fits may be dependent
6	q	$\widehat{\nu}$ using the fixed fit F
7–8	$2q$	$\bar{A}$ from report–propensity products, independent of X_v
9–11	$3q$	$\bar{U}, \bar{V}$ using $F, \widehat{\nu}$; the two moments may be dependent

The seven fold charges sum to

$$2q + q + q + q + q + 2q + 3q = 11q.$$

Together with (P6), this gives the explicit bounds

$$\frac{1}{484} G(1 + m + G\sigma) \le q(1 + \ell + q\sigma) \le \frac{1}{11} G(1 + m + G\sigma).$$

These constants do not depend on G, m, or σ. Only the reserve uses additional T. The construction deliberately uses at most $\ell = \min\{m, \lfloor G/4 \rfloor\}$ validation labels; the remaining labels may be unused. The rate is indexed by the acquired budget m, since $\ell \ge m/8$ for $m \ge 1$. Unused outcomes or validation labels are an allowed choice of an upper-bound estimator, not a reduction of the lower experiment.

P5. Rough arm means cancel after projection

Let $\mu_a(x) = \mathbb{E}[Y(a) \mid X = x]$ for $a \in \{0, 1\}$, $\mu(x) = \mathbb{E}[Y \mid X = x]$, $\bar{e} = \mathbb{E}[e(X)]$, and

$$M_0(x) = \frac{\mu(x)(e(x) - \bar{e})}{e(x)(1 - e(x))}, \qquad f_0 = \tau + M_0.$$

Lemma B.1 (observable moment identities). With $p = \mathbb{E}S$, the three observable identities are

$$\mathbb{E}\left[\frac{Y(S - p)}{e(1 - e)} \middle| X \right] = b f_0(X), \tag{P8a}$$

$$\mathbb{E}\left[\frac{Y(e + r_e - 1)}{e(1 - e)} \middle| X \right] = M_0(X), \tag{P8b}$$

$$\mathrm{Cov}(S, e) = b\sigma. \tag{P8c}$$

For (P8a), marginalizing the other iid members gives $\mathbb{E}(S - p \mid A, X) = c r_e^{k-1}(A - \bar{e})$, and $\mathbb{E}[Y\mathbf{1}\{A = 0\} \mid X] = (1 - e)\mathbb{E}[Y(0) \mid X]$. The treatment and control terms reduce to $b\{\tau + M_0\}$. For (P8b), expand $e + r_e - 1 = e - \bar{e}$ and use $\mathbb{E}[Y \mid X] = \mu(X)$; for (P8c), use $\mathbb{E}[S \mid X] = s - c r_e^{k-1}(1 - e(X))$. For clarity, the first identity follows by expanding

$$\begin{aligned} \mathbb{E}[Y(S - p) \mid X = x] &= b\,\mathbb{E}\{Y(A - \bar{e}) \mid X = x\} \\ &= b\{e(x)(1 - \bar{e})\mu_1(x) - (1 - e(x))\bar{e}\,\mu_0(x)\} \\ &= b\,e(x)(1 - e(x))\{\tau(x) + M_0(x)\}. \end{aligned}$$

The identities themselves use only bounded potential outcomes, conditional randomization, and the common reporting channel. Smoothness enters only when bounding the projection error; no smoothness is imposed on μ_0 or μ_1. For a fixed partition projection Π, put $f = \Pi f_0$ and $M = \Pi M_0$, so $f - M = \Pi\tau$ exactly.

Use the first member of each pool for every scalar covariate, histogram, and normalization row. The paired outcome moments below use the first two members of their primary and reference pools. Estimate p, r_e on their separate folds. On the shared learning fold fit cell means F and $\widehat{M}$ to the responses defined in (P8a)–(P8b) with these estimates. Return zero on empty cells. Both responses have absolute value at most H_Y. For n iid bounded rows Z and any fixed J-cell partition, the cell histogram satisfies $\mathbb{E}\|\operatorname{hist} - \Pi\mathbb{E}(Z \mid X)\|^2 \le 3H_Y^2 J/(n+1)$. For a cell of mass a, condition on its binomial count and use $a\mathbb{E}[1/(N+1)] \le (n+1)^{-1}$ and $a\Pr(N=0) \le (n+1)^{-1}$.

Conditional on $\widehat{p}, \widehat{r}_e$, the two regression perturbations have squared norms at most $H_Y^2(\widehat{p}-p)^2$ and $H_Y^2(\widehat{r}_e - r_e)^2$, respectively. Averaging over their separate folds bounds each expected squared norm by $H_Y^2/(4q)$. Projection contraction yields

$$\max\Big\{\mathbb{E}E_F, \mathbb{E}\|\widehat{M} - M\|^2\Big\} \le \frac{6H_Y^2 J}{q+1} + \frac{H_Y^2}{2q}, \qquad E_F = \|F - bf\|^2. \tag{P9}$$

No independence between these two errors is needed. Also $|F| \le H_Y$ and $|f| \le K$. Take regular cells of width $h = 1/\lceil G^{1/(2\beta+d)}\rceil$ and $J = h^{-d}$. The sole approximation error is $\mathcal{A}_\tau = \|\Pi\tau - \tau\|^2 \le L_\tau^2 d^\beta h^{2\beta}$.

P6. Normalization and complete outcome moments

Let $\nu = \|F\|^2$ and estimate $\widehat{\nu}$ by averaging $F(X)^2$ on its independent normalization fold. The implemented qualification rule is the observable condition $\widehat{\nu} \ge H_Y^2/q$. On acceptance the pre-correction fit is $F/\widehat{b}$; otherwise it is defined to be zero and the final output is $\operatorname{clip}(-\widehat{M}, [-2B, 2B])$. The events $\mathcal{E}_{\text{energy}}$, $\mathcal{E}_{\text{norm}}$, and $\mathcal{N}$ introduced below are proof events only and are never evaluated by the algorithm. Conditional on F with $\nu \ge 4H_Y^2/q$, Bernstein gives

$$\Pr(\widehat{\nu} \notin [\nu/2, 2\nu] \mid F) \le 2\exp(-c_{\mathrm{B}} q\nu/H_Y^2). \tag{P10}$$

$c_{\mathrm{B}} > 0$ is a fixed Bernstein constant, distinct from the channel contrast $c = s + z - 1$.

Since $|F| \le H_Y$, the proof event $\nu \ge 4H_Y^2/q$ is empty for $q < 4$; the small-energy bound below covers that case. More generally, if $\rho = \sup_x \eta(1-\eta)/[e(x)(1-e(x))]$, qualification requires $q\rho^2 \max(\widehat{p}, 1-\widehat{p})^2 \ge 1$ by the pilot-response bound. This is a necessary condition, not a replacement threshold.

On the good event, qualification holds and $\phi = F/\sqrt{\widehat{\nu}}$ has squared norm in $[1/2, 2]$. In every branch, the pre-correction function is $F/\widehat{b}$ with $\widehat{b} \ge b_\star$ or zero; its squared error from f is at most $C(E_F + \nu)$. Thus small ν and bad normalization contribute at most $C\mathbb{E}E_F + C/q$. This follows from $\nu\exp(-c_{\mathrm{B}} q\nu/H_Y^2) \le C/q$, even when the fit depends on normalization. Below all expectations condition on F and a good normalization.

Let $t = \langle f_0, \phi\rangle = \langle f, \phi\rangle$, so $|t| \le \sqrt{2}K$. For $g(W) = Y\phi(X)/\{e(X)(1-e(X))\}$ and independent reference members, form $D_i = g(W_i) - g(W_i')$, $i = 1, 2$. Use a primary pool for S and W_i, a reference pool for W_i', and a third pool for X''. Their moment responses are indexed by $j = 1, \ldots, q$: the first two members of primary pool j and reference pool j supply D_{1j}, D_{2j}, and the first member of third pool j supplies X_j''. More explicitly, writing these pools as P_j, R_j, C_j,

$$D_{1j} = g(P_j[1]) - g(R_j[1]), \qquad D_{2j} = g(P_j[2]) - g(R_j[2]),$$
$$U_j = \tfrac{1}{2}S(P_j)(D_{1j} + D_{2j}), \qquad V_j = -(1 - e(C_j[1]))S(P_j)D_{1j}D_{2j}.$$

The three pool roles are disjoint, and different j use disjoint pools. Thus the D_{ij} are independent and centered before multiplication by S_j; after observing S_j, no independence between U_j and V_j is asserted or needed. Their moment responses are

$$U = \frac{S}{2}(D_1 + D_2), \qquad V = -(1 - e(X''))SD_1D_2,$$
$$\mathbb{E}U = bt, \quad \mathbb{E}V = bt^2, \quad \operatorname{Var}(U) \le 2H_Y^2, \quad \operatorname{Var}(V) \le 16H_Y^4. \tag{P11}$$

For completeness, fix the independent pilot and normalization results, so g and ϕ are fixed, and let $\pi_0(W) = \Pr(A = 0 \mid W)$. This posterior is only a proof device: the estimator uses observable U, V and never estimates π_0. Then $\mathbb{E}\pi_0 = r_e$, and direct conditioning on A gives $\mathbb{E}[(\pi_0 - r_e)g] = -t$. Since $\mathbb{E}[S \mid W_1, \ldots, W_k] = s - c\prod_{i=1}^{k} \pi_0(W_i)$ and each reference member is independent of its primary member,

$$\mathbb{E}[SD_1] = -cr_e^{k-1}\mathbb{E}[(\pi_0 - r_e)g] = cr_e^{k-1}t = bt,$$
$$\mathbb{E}[SD_1D_2] = -cr_e^{k-2}\{\mathbb{E}[(\pi_0 - r_e)g]\}^2 = -cr_e^{k-2}t^2.$$

The third-pool covariate X'' is independent and satisfies $\mathbb{E}\{1 - e(X'')\} = r_e$; hence $\mathbb{E}V = -r_e\mathbb{E}[SD_1D_2] = cr_e^{k-1}t^2 = bt^2$. These equalities are conditional on the independent pilot and normalization results. They do not assume that the two response differences remain independent after observing S. Before multiplication by S, D_1, D_2 are independent, with second moments at most $2\mathbb{E}g^2 \le 4H_Y^2$. This proves the variance bounds. Dependence of U and V within a triple causes no problem.

Here are the constants explicitly. Since $S \le 1$,

$$\mathbb{E}U^2 \le \tfrac{1}{4}\mathbb{E}(D_1 + D_2)^2.$$

The unweighted D_1, D_2 are independent and centered, so their cross term vanishes. Hence $\mathbb{E}U^2 \le \{\mathbb{E}D_1^2 + \mathbb{E}D_2^2\}/4 \le 2H_Y^2$. Similarly, remove S and the third factor before bounding V^2. Independence gives $\mathbb{E}(D_1^2D_2^2) \le 16H_Y^4$. These second-moment bounds also bound the corresponding variances. The argument does not assert conditional independence after observing S.

On separate folds average $A_e = S\{e(X) - e(X')\}$ and $Z_e = \{e(X) - e(X')\}^2/2$. Write their averages $\bar{A}, X_v$. Their means are $b\sigma, \sigma$ and variances at most $2\sigma/q, \sigma/(2q)$, respectively. Also $Z_e \in [0, 1/2]$. All these folds are independent of the reserve, pilot and normalization; $\bar{A}$ is independent of X_v.

P7. Augmented discrepancy and a feasible grid approximation

Set $\omega_{\mathrm{cal}} = \ell/q$ and minimize

$$Q_+(b', t') = (\bar{U} - b't')^2 + (\bar{V} - b't'^2)^2 + \frac{(\bar{A} - b'X_v)^2}{X_v + 1/q} + \omega_{\mathrm{cal}}(b' - b_{\mathrm{cal}})^2. \tag{P12}$$

The additive $1/q$ in $X_v + 1/q$ is deterministic denominator regularization. It keeps the objective defined when the empirical propensity variance is zero and charges the resulting bias in (P13). The finite grid uses $b' = b_\star + (1 - b_\star)i/q$, $i = 0, \ldots, q$, and $t_0 = KH_Yj/q^2$, $j \in \mathbb{Z}$, $|j| \le q^2$. Retain points with $t_0^2 \le 4K^2\widehat{\nu}$ and set $t' = t_0/\sqrt{\widehat{\nu}}$. There are at most $(q + 1)(2q^2 + 1)$ points. Ties are resolved by a fixed ordering. The grid depends only on observed data.

Let $\mathscr{F}_{\mathrm{pilot}} = \sigma(F, \widehat{\nu})$ and define the proof event

$$\mathcal{N} = \{\nu \ge 4H_Y^2/q,\ \nu/2 \le \widehat{\nu} \le 2\nu\}.$$

Conditional on the observed validation index set, the index-only reserve and the eleven disjoint folds are independent. The pilot bundle (including $\widehat{M}$) uses folds 3–5; normalization uses fold 6; X_v uses 1–2; $\bar{A}$ uses 7–8; and $\bar{U}, \bar{V}$ share folds 9–11. The last two moments need not be independent of each other. Conditioning on F does not in general fix $\widehat{M}$, and no such claim is needed below. On $\mathcal{N}$, qualification holds and the true (b, t) belongs to the continuous feasible region. It need not be a grid point. The true point is used only as a comparison point in the proof; it is never evaluated or supplied to the grid-search algorithm. Given $\mathscr{F}_{\mathrm{pilot}}$ and X_v, the unused fold data and reserve remain independent of the pilot and normalization data. On $\mathcal{N}$,

$$\mathbb{E}[\bar{U} \mid \mathscr{F}_{\mathrm{pilot}}, X_v] = bt, \quad \mathrm{Var}(\bar{U} \mid \mathscr{F}_{\mathrm{pilot}}, X_v) \le 2H_Y^2/q,$$
$$\mathbb{E}[\bar{V} \mid \mathscr{F}_{\mathrm{pilot}}, X_v] = bt^2, \quad \mathrm{Var}(\bar{V} \mid \mathscr{F}_{\mathrm{pilot}}, X_v) \le 16H_Y^4/q,$$
$$\mathbb{E}[\bar{A} \mid \mathscr{F}_{\mathrm{pilot}}, X_v] = b\sigma, \quad \mathrm{Var}(\bar{A} \mid \mathscr{F}_{\mathrm{pilot}}, X_v) \le 2\sigma/q.$$

We do not condition on b_{cal}; reserve independence and (P7) give $\omega_{\mathrm{cal}}\mathbb{E}[(b_{\mathrm{cal}}-b)^2 \mid \mathscr{F}_{\mathrm{pilot}}, X_v] \le C/q$ for $\ell \ge 1$, and zero penalty at $\ell = 0$. Thus

$$\mathbb{E}[Q_+(b,t) \mid \mathscr{F}_{\mathrm{pilot}}, X_v] \le \frac{C}{q} + \frac{2\sigma/q + (X_v-\sigma)^2}{X_v + 1/q}. \tag{P13}$$

Splitting at $X_v \ge \sigma/2$ and applying Bernstein's inequality gives

$$\mathbb{E}\frac{\sigma}{X_v+1/q} \le C, \qquad \mathbb{E}\frac{(X_v-\sigma)^2}{X_v+1/q} \le C/q.$$

For $\sigma > 0$, let $B_v = \{X_v < \sigma/2\}$ and $x = q\sigma$. Bernstein gives $\Pr(B_v) \le e^{-3x/14}$. On B_v^c the denominator is at least $\sigma/2$; on B_v it is at least $1/q$. Since $\mathrm{Var}(X_v) \le \sigma/(2q)$ and $|X_v - \sigma| \le \sigma$ on B_v, the respective expectations are at most $2 + xe^{-3x/14}$ and $q^{-1}(1 + x^2 e^{-3x/14})$. Both numerators vanish almost surely at $\sigma = 0$. This makes the uniformity of the integrated C/q bound explicit.

To construct the grid approximation, round $t_0 = \langle f_0, F\rangle = \langle f, F\rangle$ toward zero and round b to its grid. Since $|t_0| = |\langle f, F\rangle| \le KH_Y$ and, on $\mathcal{N}$, $|t_0| \le \sqrt{2}K\sqrt{\widehat{\nu}}$, rounding preserves feasibility. The normalized coefficient error is at most $Kq^{-3/2}$, and the channel error is at most q^{-1}. Write $Q_+(b',t') = \|y - M(b',t')\|^2$, where

$$y = \left(\bar{U}, \bar{V}, \frac{\bar{A}}{\sqrt{X_v + 1/q}}, \sqrt{\omega_{\mathrm{cal}}}\, b_{\mathrm{cal}}\right),$$
$$M(b',t') = \left(b't', b't'^2, \frac{b'X_v}{\sqrt{X_v+1/q}}, \sqrt{\omega_{\mathrm{cal}}}\, b'\right).$$

The rounded point has squared map discrepancy at most C/q^2, because $X_v^2/(X_v + 1/q) \le 1/2$ and $\omega_{\mathrm{cal}} \le 11/2$. Grid optimality and the triangle inequality therefore give

$$\begin{aligned} D = \|M(\widehat{b},\widehat{t}) - M(b,t)\|^2 &\le 6Q_+(b,t) + C/q^2, \\ D = d_a^2 + d_2^2 + (w_v + \omega_{\mathrm{cal}})d_c^2, \qquad w_v &= \frac{X_v^2}{X_v + 1/q}, \\ d_a = \widehat{b}\widehat{t} - bt, \quad d_2 = \widehat{b}\widehat{t}^{\,2} - bt^2, \quad d_c &= \widehat{b} - b. \end{aligned} \tag{P14}$$

In particular, $\mathbb{E}[D\mathbf{1}_{\mathcal{N}}] \le C/q$. The complement of $\mathcal{N}$ is controlled directly in function risk by P6; no objective bound is required there.

P8. Uniform inversion, including zero variance

Quartic inversion. Put $E_t = (\widehat{t} - t)^2$ and note that $E_t^2 \lesssim D$ is a quartic bound on the coefficient error; this is precisely the device used at the zero-calibration endpoint. Then $w = w_v + \omega_{\mathrm{cal}}$, so $0 \le w < 6$. The exact identities

$$d_2 - \widehat{t}\,d_a = bt(\widehat{t} - t), \qquad d_2 - t d_a = \widehat{b}\widehat{t}(\widehat{t} - t) \tag{P15}$$

imply $E_t^2 \le CD$: square both identities, use $b, \widehat{b} \ge b_\star$ and $|t|, |\widehat{t}| \le 2K$, then use $(\widehat{t} - t)^2 \le 2(t^2 + \widehat{t}^{\,2})$. For $w > 0$ they also imply $E_t \le CD/\sqrt{w}$. If $t^2 \ge \sqrt{w}$, the first identity suffices. Otherwise use $\widehat{b}(\widehat{t} - t) = d_a - t d_c$ and $t^2 d_c^2 \le D/\sqrt{w}$, absorbing D since $w < 6$.

Calibrated inversion. Set $\omega = \sigma + \omega_{\mathrm{cal}}$. If $q\omega \le 2$, the quartic bound gives $\mathbb{E}[E_t\mathbf{1}_{\mathcal{N}}] \le C/\sqrt{q}$. If $q\omega > 2$ and $\omega_{\mathrm{cal}} \ge \sigma$, then $w \ge \omega/2$, giving $\mathbb{E}[E_t\mathbf{1}_{\mathcal{N}}] \le C/(q\sqrt{\omega})$. These two cases include $\sigma = 0$.

It remains to consider $q\omega > 2$ and $\sigma > \omega_{\mathrm{cal}}$, hence $q\sigma > 1$. On $\mathcal{E}_v = \{X_v \ge \sigma/2\}$, $w_v \ge \sigma/6$, so $w \ge \omega/6$ and the same expectation bound follows. On $\mathcal{E}_v^c$, Bernstein's inequality gives $\Pr(\mathcal{E}_v^c) \le \exp(-3q\sigma/14)$. The right side of (P13) is at most $C/q + 2\sigma + q\sigma^2$ on this event. With $x = q\sigma$, conditioning as in (P13) and then integrating yields

$$\begin{aligned} \mathbb{E}[D\mathbf{1}_{\mathcal{N}\cap\mathcal{E}_v^c}] &\le (C/q)(1 + x + x^2)e^{-c_{\mathrm{tail}}x}, \\ \mathbb{E}[E_t\mathbf{1}_{\mathcal{N}\cap\mathcal{E}_v^c}] &\le \sqrt{C\mathbb{E}[D\mathbf{1}_{\mathcal{N}\cap\mathcal{E}_v^c}]\Pr(\mathcal{E}_v^c)} \\ &\le Cq^{-1/2}(1+x)e^{-c_{\mathrm{tail}}x} \le \frac{C}{q\sqrt{\sigma}} \le \frac{C'}{q\sqrt{\omega}}. \end{aligned} \tag{P16}$$

Here $c_{\mathrm{tail}} > 0$ is a fixed tail constant. The last inequality uses $\sigma \le \omega < 2\sigma$. Both probability factors in Cauchy–Schwarz are retained; no independence between the selected fit and X_v is assumed. By (P6) and the propensity bound, $\omega \le 1/4 + 11/2 < 6$. This bound depends only on the fixed overlap and fold constants, not on σ, σ_0, or the CATE energy. Thus the C/q contribution from d_a^2 is absorbed into $C/(q\sqrt{\omega})$ when $q\omega > 2$; when $q\omega \le 2$, it is at most $C/\sqrt{q}$. All cases give

$$\mathbb{E}[(E_t + d_a^2)\mathbf{1}_{\mathcal{N}}] \le \frac{C}{\sqrt{q}\sqrt{1+\ell+q\sigma}}. \tag{P17}$$

This covers arbitrary sequences of m and σ. At $\sigma = \omega_{\mathrm{cal}} = 0$, $X_v = 0$ almost surely, and only the quartic bound is used; no division by w or σ occurs. For $q\omega \le 2$, $1 + q\omega \le 3$; for $q\omega > 2$, $1 + q\omega \le 3q\omega/2$. These comparisons give (P17) with the same fixed constants.

P9. From scalar error to integrated CATE risk

Define the pre-correction fit by

$$\widetilde{f} = \begin{cases} F/\widehat{b}, & \widehat{\nu} \ge H_Y^2/q, \\ 0, & \widehat{\nu} < H_Y^2/q. \end{cases}$$

The implemented estimator in both branches is $\widehat{\tau} = \mathrm{clip}(\widetilde{f} - \widehat{M}, [-2B, 2B])$. On $\mathcal{N}$, define $u = F/\sqrt{\nu}$ and $r_\phi = \|\phi\| = \sqrt{\nu/\widehat{\nu}} \in [1/\sqrt{2}, \sqrt{2}]$. Since ϕ lies in the fitted cell space, $t = \langle f, \phi\rangle = r_\phi\langle f, u\rangle$, and

$$|\langle F - bf, u\rangle| = |\sqrt{\nu} - bt/r_\phi| \le \sqrt{E_F}.$$

Hence $\sqrt{\nu} \le \sqrt{2}|t| + \sqrt{E_F}$ and $\nu \le 4t^2 + 2E_F$. Using $|d_c| \le 1$, $td_c = d_a - \widehat{b}(\widehat{t} - t)$, and $F/\widehat{b} - f = (F - bf)/b - d_c F/(\widehat{b}b)$ gives

$$\begin{aligned} \nu d_c^2 &\le 2E_F + 8d_a^2 + 8E_t, \\ \|F/\widehat{b} - f\|^2 &\le \frac{2E_F}{b_\star^2} + \frac{2\nu d_c^2}{b_\star^4}. \end{aligned} \tag{P18}$$

Random fitted-function risk. Write $\mathcal{E}_{\mathrm{norm}} = \{\nu/2 \le \widehat{\nu} \le 2\nu\}$, $\mathcal{E}_{\mathrm{qual}} = \{\widehat{\nu} \ge H_Y^2/q\}$, and $\mathcal{E}_{\mathrm{energy}} = \{\nu \ge 4H_Y^2/q\}$, and set $\widetilde{f} = \mathbf{1}_{\mathcal{E}_{\mathrm{qual}}} F/\widehat{b}$, with value zero on $\mathcal{E}_{\mathrm{qual}}^c$. The three proof events $\mathcal{E}_{\mathrm{energy}} \cap \mathcal{E}_{\mathrm{norm}}$, $\mathcal{E}_{\mathrm{energy}}^c$, and $\mathcal{E}_{\mathrm{energy}} \cap \mathcal{E}_{\mathrm{norm}}^c$ are disjoint and exhaustive; they do not change the qualification rule. For every output branch,

$$\Omega = (\mathcal{E}_{\mathrm{energy}} \cap \mathcal{E}_{\mathrm{norm}}) \,\dot\cup\, \mathcal{E}_{\mathrm{energy}}^c \,\dot\cup\, (\mathcal{E}_{\mathrm{energy}} \cap \mathcal{E}_{\mathrm{norm}}^c).$$

The first event is the qualified, normalized branch; the other two are handled by the direct all-branch bound below. Since $b^2\|f\|^2 \le 2\|F - bf\|^2 + 2\|F\|^2$, both the qualified and unqualified branches obey

$$\|\widetilde{f} - f\|^2 \le (4E_F + 6\nu)/b_\star^2.$$

On $\mathcal{E}_{\mathrm{energy}} \cap \mathcal{E}_{\mathrm{norm}} = \mathcal{N}$, $\widehat{\nu} \ge \nu/2 \ge 2H_Y^2/q$, so qualification holds. Equations (P17)–(P18) give

$$\mathbb{E}[\|\widetilde{f} - f\|^2\mathbf{1}_{\mathcal{N}}] \le C\mathbb{E}[E_F] + C\mathbb{E}[(d_a^2 + E_t)\mathbf{1}_{\mathcal{N}}] \le C\mathbb{E}[E_F] + \frac{C}{\sqrt{q(1+\ell+q\sigma)}}.$$

On $\mathcal{E}_{\mathrm{energy}}^c$, the all-branch inequality, regardless of qualification, gives

$$\mathbb{E}[\|\widetilde{f} - f\|^2\mathbf{1}_{\mathcal{E}_{\mathrm{energy}}^c}] \le C\mathbb{E}[E_F] + C/q.$$

Finally, condition on F on $\mathcal{E}_{\mathrm{energy}} \cap \mathcal{E}_{\mathrm{norm}}^c$. By (P10), $\Pr(\mathcal{E}_{\mathrm{norm}}^c \mid F) \le 2e^{-c_{\mathrm{B}} q\nu/H_Y^2}$, and hence

$$\begin{aligned} \mathbb{E}[\nu\mathbf{1}_{\mathcal{E}_{\mathrm{energy}} \cap \mathcal{E}_{\mathrm{norm}}^c}] &\le 2\mathbb{E}[\nu\mathbf{1}_{\mathcal{E}_{\mathrm{energy}}} e^{-c_{\mathrm{B}} q\nu/H_Y^2}] \le C/q, \\ \mathbb{E}[E_F\mathbf{1}_{\mathcal{E}_{\mathrm{energy}} \cap \mathcal{E}_{\mathrm{norm}}^c}] &\le \mathbb{E}[E_F]. \end{aligned}$$

The all-branch inequality therefore charges this third event by $C\mathbb{E}[E_F] + C/q$, including either qualification output and $F = 0$. Adding the three disjoint contributions yields

$$\mathbb{E}\|\widetilde{f} - f\|^2 \le C\mathbb{E}[E_F] + \frac{C}{\sqrt{q(1+\ell+q\sigma)}} + \frac{C}{q}.$$

Using (P9), $\mathbb{E}[E_F] \le CJ/(q+1) + C/q$, and absorbing q^{-1} into the histogram term, gives

$$\mathbb{E}\|\widetilde{f} - f\|^2 \le C\left\{\frac{J}{q+1} + \frac{1}{\sqrt{q(1+\ell+q\sigma)}}\right\}.$$

The nuisance estimate $\widehat{M}$ is not treated as fixed by $\mathscr{F}_{\mathrm{pilot}}$; its (P9) bound is applied after integrating over the pilot bundle.

Projection and clipping. Let $\widetilde{f}$ denote the pre-correction fit, including zero on qualification failure. The following is the deterministic transfer in Lemma L.3. Since both $\widetilde{f} - \widehat{M}$ and $f - M = \Pi\tau$ lie in the common cell space, orthogonality gives

$$\begin{aligned}\|\widetilde{f} - \widehat{M} - \tau\|^2 &= \|\Pi\tau - \tau\|^2 + \|(\widetilde{f} - f) - (\widehat{M} - M)\|^2 \\ &\le \mathcal{A}_\tau + 2\|\widetilde{f} - f\|^2 + 2\|\widehat{M} - M\|^2.\end{aligned}$$

The preceding event bounds and (P9) control the two random terms. By the bounded-potential-outcome assumption in Section 2, $|\tau(x)| \le 2B$; projection onto $[-2B, 2B]$ therefore contracts pointwise squared loss. Absorbing $1/q$ into the histogram term yields

$$R_P(\widehat{\tau}) \le \mathcal{A}_\tau + C\left\{\frac{J}{q+1} + \frac{1}{\sqrt{q(1+\ell+q\sigma)}}\right\}. \tag{P19}$$

Budget conversion. The explicit P4 bounds give

$$\frac{1}{\sqrt{q(1+\ell+q\sigma)}} \le \frac{22}{\sqrt{G(1+m+G\sigma)}} \le \frac{22}{\sqrt{G(1+m+G\sigma_0)}}.$$

This holds uniformly, including $m = 0$. The partition choice gives $\mathcal{A}_\tau + J/(q+1) \le CG^{-\alpha}$. Together with monotonicity in σ, it gives the class bound at σ_0. For $1 \le G < 44$, the zero estimator has risk at most $4B^2$, while $r(G, m, \sigma_0) \ge G^{-\alpha} \ge 43^{-\alpha}$. Thus the fixed choice $C_{\mathrm{up}} \ge 4B^2\, 43^\alpha$ covers all small budgets. Only the CATE smoothness was used.

P10. What the threshold means

For sequences in the stated class, the full-label regression order is attainable if and only if

$$1 + m + G\sigma_0 \gtrsim G^{2\alpha-1}. \tag{P20}$$

Here $\gtrsim$ means a positive eventual lower multiplicative bound, not a unit-constant finite-sample cutoff. At $m = 0$, P1 recovers the original result; at $m = G$ its extra term is $O(G^{-1})$. At $\sigma_0 = 0$ it becomes $\{G(1+m)\}^{-1/2}$. For $\alpha \le 1/2$ no validation is needed for this rate order. Under $\beta \le 1$, the leading calibration regime is restricted to $d = 1$, $\beta > 1/2$.

P1 is an asymptotic statement for the explicit observation model above. It does not provide sharp finite-risk constants or imply that the rate-attaining construction improves on simpler estimators at a given budget. Section 6 and Appendix H give separate finite-sample limitations.

Appendix C: Prior results, remaining proof obligations, and calibration reduction

We use a calibration reduction to construct the upper bound and then compare it with related inverse-problem results. The complete validation law depends on the treatment effect, whereas discarding validation outcomes isolates an independent calibration block. This reduction supplies the upper construction; the lower-bound argument remains in Appendix B.

1. Fixed experiment and the direction of a reduction

Use the experiment and parameter class of Proposition 1. The class inputs k, η, κ stay fixed while $\mathrm{Var}(e(X))$ may approach zero. The supplied propensity is the same under both lower-bound alternatives; weak overlap, unknown propensity, and deteriorating channel contrast are different problems. Arm-specific regressions remain unrestricted.

Let $\mathcal{E}$ denote the complete G-pool experiment, including the observed independent uniform m-subset I with true OR states T. A parameter-free Markov kernel that preserves the target and sample budget can reduce $\mathcal{E}$ to a less informative experiment $\mathcal{E}_{\mathrm{red}}$. Writing $M(\cdot)$ for classwise minimax risk gives

$$M(\mathcal{E}) \le M(\mathcal{E}_{\mathrm{red}}).$$

An upper-bound estimator for the reduction is therefore feasible in the complete experiment. The reverse transfer fails for lower bounds: discarded validated outcomes may contain information about the effect.

2. An exact reduction provides a nuisance-only calibration block

For $G \ge 44$, reserve $\ell = \min\{m, \lfloor G/4 \rfloor\}$ pools using the index-only rule of Appendix B, P4. The reduced experiment retains

- (X, T, S), but not Y, on the reserve;
- (X, Y, S), but no further T, on the remaining pools;
- I and the supplied e, which specify the allocation.

This discarding map is deterministic and parameter-free. Conditional on I, the reserve and remaining pools are independent. For reserve covariates, write $q_X = \prod_{i=1}^{k}\{1 - e(X_i)\}$. In the order $(T, S) = (0,0), (0,1), (1,0), (1,1)$, the four probabilities are

$$zq_X, \quad (1-z)q_X, \quad (1-s)(1-q_X), \quad s(1-q_X). \tag{R1}$$

They contain no potential-outcome law or CATE. Marginally, $\Pr(T = 0) = r_e^k$ with $r_e = \mathbb{E}[1 - e(X)]$; the retained covariates let the estimator learn r_e without knowing P_X. Thus the reserve forms an independent calibration block, even though the complete validated law depends on the effect.

The Appendix B estimator uses only the retained data. Combining its upper bound with the complete-law lower bound gives

$$C_{\mathrm{low}} r(G, m, \sigma_0) \le M(\mathcal{E}) \le M(\mathcal{E}_{\mathrm{red}}) \le C_{\mathrm{up}} r(G, m, \sigma_0). \tag{R2}$$

The reduction preserves the minimax order for this class and loss, without equating the experiments or their constants. At $m = G$, the reserve discards outcomes from at most $G/4$ pools and leaves a learning block. For $G < 44$, the zero upper rule covers the finite set of budgets.

3. Discarding outcomes really can remove information

In the scalar lower family, $Y = BR$ with uniform signs R. The complete conditional law uses $v_\theta = \prod_i (1 - e_i)(1 - 2e_i\theta R_i)$. Averaging over R gives $\mathbb{E}_R v_\theta = q_X$, recovering (R1). The reserve marginal is effect-free, while its joint law with outcomes is not.

For one validated pool, take $k = 2$, $e_i = 1/2$, and $s = z = 3/4$. The (T, S) marginal agrees at $\theta = 0$ and $\theta = 1/4$, whereas

$$\Pr(R_1 = R_2 = 1, T = 0, S = 0) = \frac{3}{64}(1 - \theta)^2$$

differs. A parameter-free kernel cannot reconstruct both full laws from an identical (T, S) input marginal. This witness concerns the discarded block; it makes no claim about kernels using the remaining study pools or about asymptotic equivalence.

4. Comparison with related observation models

The comparison separates established components from the calculations required for this observation contract:

Literature	Established component	Remaining task for Proposition 1
Smooth CATE	Different smoothness of effect and baseline; risk analysis	Integrated risk with hidden assignments and an unknown pooled channel.
Misclassification	Identification under exclusion, relevance and rank	Uniform risk dependence on G, m, σ_0.
Group testing	Inverse pooled responses, regression and assay calibration	Matching bounds under the complete linked-observation contract.
Unknown operators	Auxiliary information and operator error in risk	Nonlinear OR inversion and full validated-law information bounds.
Present work	Fixed-size pooled OR reports, linked (X, Y) records, an unknown common channel, arbitrary arm regressions, and revealed OR states	Classwise minimax risk under the combined observation contract.

None of these lines of work simultaneously assumes fixed-size pooled OR reports, linked (X, Y) records, a common unknown nondifferential channel, arbitrary arm-specific regressions, and a random subset of revealed OR states. The contribution here is the joint risk law under this combined observation contract, rather than a new identification result for any one of these component models.

With observed individual treatments, Gao and Han (2020, Assumptions 1–3 and Theorems 1–2) already separates a smooth effect from a less-smooth Hölder baseline. Its displayed minimax losses are expected L^1 and $L^1(g_0)$ error, with rates depending on covariate matching or density ratios and outcome noise. Kennedy et al. (2024, Theorem 1) likewise imposes nuisance smoothness and studies pointwise CATE loss. These statements do not supply the integrated squared-risk result for an unknown pooled channel.

There is positive overlap with Kasahara and Shimotsu (2022): for one preselected member per independent pool, S is a treatment proxy with $\Pr(S = 1 \mid A = 1, X) = s$ and $\Pr(S = 1 \mid A = 0, X) = s - cr_e^{k-1}$. Some subclasses satisfy that paper's exclusion, relevance and rank conditions, including exogenous-treatment subclasses. Constant propensity removes the required relevance, and unrestricted arm means need not satisfy the exclusion. This overlap precludes a blanket identification-priority claim; fixed-law identification does not give the uniform transition in Proposition 1.

Group-testing work also treats assay calibration and inverse responses (Delaigle and Hall, 2015); see also Delaigle and Tan (2024, Eq. (3.3)). The table records the specific results and assumptions used in this comparison; it is not a survey of the full literature.

Reference and result	Input and target	Established component	Remaining transfer requirement
Johannes (2009), ğ4, Corollaries 4.4–4.5 and Remark 4.3, PDF pp11–12, DOI 10.1214/08-AOS652	Y=X+epsilon, independent error sample, density/derivative Sobolev risk and Fourier decay class	Auxiliary-sample necessary/sufficient orders, two-sample lower-bound logic, regularization tools	A target-preserving convolution representation with the required error-density and smoothness classes. Binary OR calibration is not that representation. Replacing n,m or exponents alone proves nothing.

Reference and result	Input and target	Established component	Remaining transfer requirement
Johannes–Schwarz (2013), model (1.1), Theorems 2.3/2.5, PDF pp2–4; lower proof Step B pp9–10, DOI 10.1080/03610926.2012.731548	Independent Gaussian target and operator sequences; known bases and ellipsoid/operator classes	Separate information terms, joint perturbation of target and operator, identical-main-image lower pairs	A verified experiment mapping or uniform approximation, known bases/classes and suitable loss comparison. The current moments include (bt,btš) and dependent estimated noise, not the stipulated Gaussian sequence experiment.
Hoffmann–Reiss (2008), model (1.1)–(1.2), Assumption 3.1, Theorem 5.1, PDF pp1–2,4–5,8, DOI 10.1214/009053607000000721	g_epsilon=Kf+epsilon W and noisy operator K_delta; operator regularity and Besov risk	Operator estimation as a separate cost, thresholding and regularization, two-noise rate comparison	A linear-operator model satisfying the stated regularity, Gaussian observation structure and target class. No such map has been established for the full pooled causal law.
Delaigle–Hall (2015), ğ3.1 and Theorem 4, PDF pp4–5,9, DOI 10.1093/biomet/asv049	Biomarker group testing with positive/negative training samples comparable to study size; pointwise prevalence behavior	Parametric assay estimation and oracle equivalence under the paper's assumptions	Different sample regime, CATE target, rough-arm class and a matching necessity result. The inspected theorem is not that result.
Zhang et al. (2014), ğ3.1–3.2, DOI 10.1002/sim.6268	Validated group specimens and joint prevalence/channel likelihood with study/validation budgets	The validation interface, likelihood factorization and budget separation	A uniform heterogeneous-function squared-risk result at a vanishing calibration witness, rather than local Fisher-information analysis.
Brenner Miguel–Johannes–Siebel (2024), DOI 10.1214/24-EJS2314	Y=X*U, independent multiplicative-error sample, Mellin methods and weighted density/survival loss	Recent Two-sample unknown-error risk analysis for a multiplicative convolution model	A multiplicative convolution representation and the paper's transform assumptions; these do not yield the pooled-CATE lower bound.

5. Relation to joint perturbation arguments

The lower proof Step B of Johannes–Schwarz explicitly varies f and A together while preserving A_1 f_1=A_-1 f_-1. Only the auxiliary operator observation then distinguishes that pair. Thus joint perturbation of the target and nuisance is an established lower-bound device.

For the current scalar pair, c1=c0+h and theta1=c0*theta/(c0+h) preserve c*theta, but do not preserve the entire unvalidated law. The degree-two sign coefficient has squared gap

```
(c0*theta^2-c1*theta1^2)^2
   = c0^2*theta^4*h^2/(c0+h)^2.                              (R3)
```

At nonconstant e, a constant coefficient also contributes order $h^2\sigma_0$. The complete-law exponent is bounded by a fixed multiple of

```
h^2[(G-m)(sigma0+theta^4)+m].                                    (R4)
```

The nonzero residual is exactly the extra information an argument based on identical main images would ignore. Establishing R4 for all observed X,Y,S,T and transferring the moment fit to the full CATE function remain local model-specific work. Similarity of the perturbation method does not by itself establish either step or determine the contribution's novelty.

6. Scope of the comparison

The model-specific result is the uniform full-function risk law $G^{-\alpha}+\{G(1+m+G\sigma_0)\}^{-1/2}$ in the declared causal experiment. Its threshold is an algebraic consequence, and the calibrated extension reuses standard minimum-distance, sample-splitting, inversion and testing arguments. The complete-law calculation and rough-nuisance cancellation are the model-specific steps. The comparisons above do not supply the full result by direct substitution and do not rule out a future reduction under a different representation.

Appendix D: The calibration threshold has a precise total-budget interpretation

1. The result inherits the existing local risk theorem

Corollary D.1 (total-budget consequence). The following statements are conditional on Proposition 1. The varying inputs are N, λ, and σ_0; the constants remain uniform in those inputs and depend only on the fixed class inputs.

Write $\alpha = 2\beta/(2\beta+d) \in (0,1)$. Proposition 1 gives

$$M(G,m,\sigma_0) \asymp G^{-\alpha} + \{G(1+m+G\sigma_0)\}^{-1/2},$$

with constants uniform in G, m, σ_0 at fixed class inputs. All assumptions of that proposition, including fixed pool size and random validation sampling, remain in force.

Give each study pool unit cost and each true-OR validation additional price $\lambda \geq 1$. For integer budget $N \geq 2$, define

$$R_{\text{budget}} = \inf_{\substack{G\geq 1,\ 0\leq m\leq G \\ G,m\in\mathbb{Z},\ G+\lambda m\leq N}} M(G,m,\sigma_0).$$

The allocation is deterministic; data-dependent allocation lies outside this comparison. Uniformly in N, λ, σ_0,

$$R_{\text{budget}} \asymp N^{-\alpha} + \{N(1+N/\lambda+N\sigma_0)\}^{-1/2}. \tag{D.1}$$

The constants are independent of those three varying inputs. Equation (D.1) identifies an order, not exact optimal allocations or finite-risk constants.

2. A feasible allocation proves the upper comparison

Every feasible allocation has $G \leq N$ and $m \leq N/\lambda$. Set $H_{\text{cap}} = N(1+N/\lambda+N\sigma_0)$. Then

$$G^{-\alpha} \geq N^{-\alpha}, \qquad G(1+m+G\sigma_0) \leq H_{\text{cap}},$$

so the uniform lower constant in Proposition 1 gives the lower half of (D.1) after taking the infimum.

For the upper half, choose the feasible witness

$$G_{\text{half}} = \lfloor N/2 \rfloor, \qquad m_{\text{half}} = \lfloor N/(2\lambda) \rfloor.$$

Indeed, $m_{\text{half}} \leq G_{\text{half}}$ and $G_{\text{half}} + \lambda m_{\text{half}} \leq N$. For $N \geq 2$, $G_{\text{half}} \geq N/3$; for every $x \geq 0$, $1 + \lfloor x/2 \rfloor \geq (1+x)/3$. Hence

$$1 + m_{\text{half}} + G_{\text{half}}\sigma_0 \geq \tfrac{1}{3}(1+N/\lambda+N\sigma_0),$$
$$G_{\text{half}}(1+m_{\text{half}}+G_{\text{half}}\sigma_0) \geq H_{\text{cap}}/9.$$

The regression proxy changes by at most $3^\alpha \leq 3$. Applying the uniform upper constant proves (D.1). This half allocation is an existence witness, not a claim of numerical optimality.

3. Fixed validation prices do not create a new rate exponent

The budget-normalized full-label reference has order $N^{-\alpha}$ when one study pool and its k members carry the stipulated unit cost. Equation (D.1) reaches that order if and only if

$$1 + N/\lambda + N\sigma_0 \gtrsim N^{2\alpha-1}. \tag{D.2}$$

At fixed finite λ and $\sigma_0 = 0$, the calibration term in (D.1) is $O(N^{-1})$, so the total-budget exponent remains α. For $\alpha > 1/2$, the local threshold uses m of order $G^{2\alpha-1}$. A plug-in bound with remainder $1/m$ instead asks for m of order G^{α}. Both budgets are sublinear and recover the same total-budget exponent at fixed validation price. The distinction is how much validation suffices within the pooled experiment, not a new exponent over all calibrated competitors.

4. Growing validation prices expose the threshold's consequence

Let $\lambda_N \asymp N^{\rho}$ with $\rho \geq 0$ and first take $\sigma_0 = 0$. Equation (D.1) becomes

$$R_{\text{budget}} \asymp N^{-r}, \qquad r = \min\{\alpha, \max(1/2, 1 - \rho/2)\}. \tag{D.3}$$

For $\alpha > 1/2$, the full-label exponent is attainable precisely when $\rho \leq 2 - 2\alpha$. At $\alpha = 2/3$, the boundary is $\rho = 2/3$. Below that boundary the minimum sufficient validation cost is sublinear in N; at the boundary a suitable constant fraction suffices. These are rate comparisons, not statements about leading finite-risk constants.

For $\rho \geq 1$, zero validation already attains the channel exponent $1/2$. If $\lambda_N > N$, the integer budget cannot buy one validation; the 1 in (D.1) covers this endpoint. If instead $\sigma_0 = cN^{-\gamma}$ with fixed legal $c > 0$ and $\gamma \geq 0$, then

$$r = \min\{\alpha, \tfrac{1}{2}\max(1, 2 - \rho, 2 - \gamma)\}. \tag{D.4}$$

Propensity variation can supply calibration information. When $\sigma_0 \equiv 0$, omit the $2 - \gamma$ term. If $\alpha \leq 1/2$, regression determines the order even at $m = 0$. These comparisons keep pool size and channel contrast fixed.

5. Interpretation and limitations

(D.1) is an algebraic consequence of Proposition 1. The cost-normalized analysis translates the validation threshold into an allocation comparison under a stipulated price model; it does not establish an empirical collection-cost advantage. Growing validation prices provide a sensitivity analysis rather than an empirically established cost model. Because the relative costs of collecting outcomes, forming reports, retaining individual labels and acquiring true-OR validation are unspecified, this calculation does not show that validated pooling is preferable to individual-label observation at equal actual cost.

Appendix E: A boundary for outcome-dependent reporting errors

We now allow outcome-dependent reporting errors to mark the boundary of the nondifferential model in Proposition 1. The calculations below are scope results, not corrections to the main theorem.

1. Independent channel variation averages out exactly

Let each pool draw an unobserved iid channel $H_g = (s_g, z_g)$ independently of its assignments, covariates, and potential outcomes. Conditional on H_g and the true OR T_g, generate the report independently across pools. Marginalizing H_g gives sensitivity $\mathbb{E}s_g$ and specificity $\mathbb{E}z_g$. These probabilities have no residual dependence on member variables. Independence across pools yields exactly the original product experiment with the averaged channel. The compact channel class is convex, so its bounds and positive contrast persist under averaging.

Thus independent iid random channel variation alone need not violate the model. This statement does not cover fixed nonidentical pool channels, shared batch effects, covariate-dependent channels or outcome-dependent errors. It does not imply that an observed H_g can be discarded without loss.

2. A PRECISE ENLARGED SCALAR CLASS

Take $k = 2$, $e = 1/2$, $B = 1$, and $\kappa = 1/10$. Covariates are uninformative and constant. Members and pools are independent. Treatments are independent Bernoulli$(1/2)$ draws, independent of potential outcomes. For $\theta \in [-1/2, 1/2]$, choose independent potential signs with marginals

$$\Pr\{R(a) = r\} = \frac{1}{2} + \theta(a - 1/2)r, \qquad r \in \{-1, 1\}.$$

Set $Y = R = R(A)$. Then $\tau = \mathbb{E}\{R(1) - R(0)\} = 2\theta$. Randomization, consistency, bounded outcomes and no interference all hold. The observed R_i are iid uniform signs under every θ. Write $R = (R_1, R_2)$, $T = \max(A_1, A_2)$, and

$$v_\theta(R) = \Pr(T = 0 \mid R) = \frac{(1 - \theta R_1)(1 - \theta R_2)}{4}.$$

The unknown baseline channel obeys $s, z \in [1/10, 9/10]$ and $c = s + z - 1 \geq 1/10$. Enlarge its conditional report law to

$$\begin{aligned} &\Pr(S = 1 \mid T = 1, R) = s + \rho_1(R), \\ &\Pr(S = 1 \mid T = 0, R) = 1 - z + \rho_0(R), \\ &\qquad |\rho_0(R)|, |\rho_1(R)| \leq \epsilon, \qquad 0 \leq \epsilon \leq 1/160. \end{aligned} \tag{E.1}$$

Require actual sensitivities/specificities also to satisfy the same interior and contrast bounds at each R. Report randomization is independent across pools. The estimator receives epsilon and the fixed class inputs, but not the baseline channel or the perturbations. Each alternative has its own unknown baseline; this is not a supplied nominal-channel problem.

Let $R_{G,m}(\epsilon)$ denote scalar CATE minimax squared risk from all (R, S) records and exact T on a uniform independent subset of m pools. The observation includes the subset indices, never individual assignments.

3. AN EXACTLY INDISTINGUISHABLE UNVALIDATED PAIR

Start with $s_0 = z_0 = 3/4$, $c_0 = 1/2$, and $\theta_0 = \theta$, $|\theta| \leq 1/2$. For $0 < h \leq 1/8$, define

$$\begin{gathered} c_1 = \frac{1}{2} + h, \qquad s_1 = \frac{3}{4} + \frac{h}{4}, \qquad z_1 = \frac{3}{4} + \frac{3h}{4}, \qquad \theta_1 = \frac{\theta}{1 + 2h}, \\ \rho_0(R) = \rho_1(R) = \rho(R) = -\frac{h\theta^2 R_1 R_2}{4(1 + 2h)}. \end{gathered} \tag{E.2}$$

Model 0 has zero perturbation. Model 1 uses (E.2). Both models have the same unconditional outcome law. Without rho, the conditional report probability is $p_j(R) = s_j - c_j v_{\theta_j}(R)$.

Expansion shows the constant and linear sign coefficients agree, and

$$p_0(R) - p_1(R) = -\frac{h\theta^2 R_1 R_2}{4(1 + 2h)} = \rho(R). \tag{E.3}$$

Consequently the full observable law of (R1,R2,S) agrees exactly in the two models after applying rho. This is equality of every atom, not only of a selected set of moments. It persists for arbitrarily many pools.

The perturbation amplitude is $h\theta^2/\{4(1 + 2h)\} \leq h/16 \leq 1/128$. The actual sensitivity is $s_1 + \rho$ and specificity is $z_1 - \rho$; their contrast remains c_1. Both stay in $[1/10, 9/10]$ under the stated bounds. To meet any smaller ϵ, choose θ accordingly.

In particular, take $h = 1/8$ and $\theta^2 = 40\epsilon$. This is feasible for $\epsilon \leq 1/160$, has perturbation amplitude exactly ϵ, and yields

$$(\tau_0 - \tau_1)^2 = \frac{4\theta^2}{25} = \frac{32}{5}\epsilon. \tag{E.4}$$

Under identical observable laws, the average squared error for the two targets is at least one quarter of their squared separation (complete the square around their midpoint). Thus, for every integer $G \geq 1$,

$$R_{G,0}(\epsilon) \geq \frac{8}{5}\epsilon. \tag{E.5}$$

The CATE separation is of order $\sqrt{\epsilon}$, and the squared-risk floor is of order ϵ. Both the unknown baseline and the weak target matter. A fixed supplied nominal channel is a different information class.

For example, $\epsilon = 1/1000$ gives $\theta_0 = 1/5$ and $\theta_1 = 4/25$, hence $\tau_0 = 2/5$ and $\tau_1 = 8/25$. A reporting perturbation of 0.1 percentage point around model 1's unknown baseline makes its unvalidated law identical to model 0's. Their CATEs differ by 0.08; the maximum squared risk is at least $1/625 = 0.0016$ for every G. The two baseline channels differ; the example does not assert this loss under a common baseline already known to the analyst.

4. Matching scalar risk at zero validation

In the enlarged scalar class,

$$R_{G,0}(\epsilon) \asymp G^{-1/2} + \epsilon \quad (G \geq 1,\ 0 \leq \epsilon \leq 1/160). \tag{E.6}$$

Constants depend only on the fixed class bounds. This concerns the binary scalar class defined above, not general heterogeneous CATE.

Upper bound

Set $b = c/4 \in [1/40, 1/5]$. For each observed pool form $U = S(R_1 + R_2)/2$ and $V = -SR_1R_2$. Under the baseline law, $\mathbb{E}U = b\theta$ and $\mathbb{E}V = b\theta^2$. Condition (E.1) perturbs the conditional mean of S given R by at most ϵ. Uniform independent signs then give

$$|\mathbb{E}U - b\theta| \leq \epsilon/2, \qquad |\mathbb{E}V - b\theta^2| \leq \epsilon, \qquad \mathrm{Var}(U), \mathrm{Var}(V) \leq 1. \tag{E.7}$$

Fit $(\widehat{b}, \widehat{\theta})$ by minimum distance between the sample means of (U, V) and $(b't', b't'^2)$ over endpoint-including grids of mesh at most $1/G$ on $[1/40, 1/5] \times [-1/2, 1/2]$. Break ties by fixed grid index. Set

$$D = (\widehat{b\theta} - b\theta)^2 + (\widehat{b\theta^2} - b\theta^2)^2.$$

If ξ is the sample-mean vector minus $(b\theta, b\theta^2)$, then $\mathbb{E}\|\xi\|^2 \leq 2/G + 5\epsilon^2/4$. A nearest grid point is within C/G in moment space. Minimum-distance optimality and the triangle inequality yield $\mathbb{E}D \leq C(G^{-1} + \epsilon^2)$.

Writing $d_1 = \widehat{b\theta} - b\theta$ and $d_2 = \widehat{b\theta^2} - b\theta^2$, exact algebra gives

$$d_2 - \widehat{\theta} d_1 = b\theta(\widehat{\theta} - \theta), \qquad d_2 - \theta d_1 = \widehat{b\theta}(\widehat{\theta} - \theta).$$

On the compact parameter rectangle these identities imply $(\widehat{\theta} - \theta)^4 \leq CD$. Cauchy–Schwarz therefore bounds the risk of $2\widehat{\theta}$ by $C(G^{-1/2} + \epsilon)$. This is an order bound, not a sharp finite-risk comparison.

Clean-model lower bound

The enlarged class includes zero-perturbation channels. Use the two baselines in (E.2) with $h = 1/8$, zero perturbation in both, and $\theta^2 = 1/(4\sqrt{G})$. For one unvalidated pool the conditional Bernoulli probability gap has magnitude $\theta^2/40$. Every baseline report probability lies in $[1/10, 9/10]$, so

$$\chi^2(P_0 \| P_1) \leq \frac{\theta^4}{144} = \frac{1}{2304G}.$$

The G-pool divergence is at most $e^{1/2304} - 1 < 1/16$, hence total variation is below $1/4$. The squared CATE gap is $1/(25\sqrt{G})$, and nearest-target testing gives risk at least $3/(800\sqrt{G})$. Combining this with (E.5) proves the lower half of (E.6).

5. Exact true-OR validation breaks the indistinguishability

(E.3) does not equate the validated joint law of (R, T, S). In fact, $\mathbb{E}[(1-T)R_1] = -\theta/4$. For $m \geq 1$, average the bounded response $H_T = -8(1-T)R_1$ over validated pools and project the result onto $[-1, 1]$. It has mean τ and $\mathbb{E}H_T^2 = 16$. Uniform independent validation supplies m iid copies, so squared risk is at most $16/m$ for any reporting kernel. The response does not use S. Zero estimation has risk at most one; hence $R_{G,m}(\epsilon) \leq \min\{1, 16/m\}$.

For a lower bound with m validated pools, use (E.2) with

$$h = \frac{1}{256\sqrt{m+1}}, \qquad \theta^2 = \min\left\{\frac{1}{4}, \frac{4\epsilon}{h}\right\}. \tag{E.8}$$

The unvalidated laws still coincide and the perturbation is at most ϵ. Conditional on R, every validated (T, S) atom in model 1 has probability at least $1/128$: $v_{\theta_1} \geq 1/16$, $1 - v_{\theta_1} \geq 7/16$, and each report probability is at least $1/8$. Moreover,

$$|\theta - \theta_1| \leq h, \qquad |v_\theta - v_{\theta_1}| \leq 3h/4,$$
$$|s_0 - (s_1 + \rho)| \leq 5h/16, \qquad |z_0 - (z_1 - \rho)| \leq 13h/16.$$

Thus every conditional validated atom changes by less than $2h$, and $\chi^2(V_0 \| V_1) \leq 2048h^2$. Conditioning on the observed validation indices, complete-data divergence is at most $(1+2048h^2)^m - 1 \leq e^{1/32} - 1 < 1/16$, independently of G. The same bound holds after averaging over the indices. Meanwhile,

$$(\tau_0 - \tau_1)^2 = \frac{16h^2\theta^2}{(1+2h)^2} \geq 10 \min\left\{\frac{h^2}{4}, 4h\epsilon\right\}.$$

The two-point argument proves, for a fixed $c_0 > 0$,

$$0 \leq m \leq G, \qquad 0 \leq \epsilon \leq 1/160,$$

with c_0 independent of G, m, and ϵ,

$$R_{G,m}(\epsilon) \geq c_0 \min\left\{\frac{1}{m+1}, \frac{\epsilon}{\sqrt{m+1}}\right\}. \tag{E.9}$$

No matching rate is claimed for every joint sequence (G, m, ϵ). For fixed $\epsilon > 0$ and bounded m, increasing G alone cannot give uniform consistency in this class. For fixed $\epsilon > 0$ and $m \to \infty$, risk has order m^{-1} uniformly over $G \geq m$, eventually. The latter follows because the first term in (E.9) is eventually the smaller one and the validated-only estimator supplies 16/m.

6. Implications and limits

The unknown-channel theorem is sensitive to a specific violation of nondifferential reporting. A small adversarial dependence on outcomes can erase the second-order information used for self-calibration. Independent random channel heterogeneity is a different case and can average out exactly.

(E.5) also gives a lower bound for an enlarged heterogeneous-CATE class that contains this constant-effect, constant-propensity, k=2 submodel. It applies at sigma0=0 and kappa=1/10 (or a weaker bound), not at every supplied design. It supplies no full-function robust upper bound, optimal validation schedule under shrinking epsilon, confidence interval, practical estimator improvement or claim that such reporting dependence occurs in an actual institution.

The simulations in Section 6 use a common nondifferential channel. They do not test this observationally equivalent differential-error construction, which limits the model without supplying a robust CATE estimator.

APPENDIX F: THE SCALAR BOUNDARY AND CONTAMINATION MODULI

Appendix E's scalar boundary can be compared with a standard contamination modulus. The structured reporting model remains distinct from unrestricted contamination.

1. CONTAMINATION DECISION THEORY

Chen, Gao, and Ren, *A general decision theory for Huber's epsilon-contamination model*, Electronic Journal of Statistics (2016), DOI 10.1214/16-EJS1216.

Equation (1), PDF p. 1, defines observations from $(1-\alpha)P_\xi + \alpha Q$ with arbitrary Q. Theorem 3.2, PDF p. 6, recalls a lower bound from the authors' earlier work: the clean-data difficulty combines with the loss modulus

$$\omega(\alpha) = \sup\left\{ L(\xi, \xi') : \mathrm{TV}(P_\xi, P_{\xi'}) \le \frac{\alpha}{1-\alpha} \right\}.$$

It gives a lower-bound template for general losses, not a universal matching upper bound for every target. Theorem 3.1 primarily controls total variation. Section 5 and Theorem 5.1 explicitly distinguish losses not equivalent to TV; the paper's worked example has a contamination exponent different from two. Thus neither an irreducible contamination term nor a nonquadratic dependence on its size is, by itself, an original principle of pooled-CATE.

2. THE CLEAN POOLED SCALAR FAMILY HAS A LINEAR SQUARED-TARGET MODULUS

Use the clean scalar class of Appendix E, with $k = 2$, $e = 1/2$, $|\theta| \le 1/2$, $\tau = 2\theta$, and unknown channel bounded by $\kappa = 1/10$. Write $\xi = (\theta, s, z)$, let P_ξ be the full (R_1, R_2, S) law, and define

$$\Omega(\delta) = \sup\{(\tau - \tau')^2 : \mathrm{TV}(P_\xi, P_{\xi'}) \le \delta\}.$$

The clean moments of $U = S(R_1 + R_2)/2$ and $V = -SR_1R_2$ are $(b\theta, b\theta^2)$, where $b = (s + z - 1)/4 \ge 1/40$. Both statistics are bounded by one. Total variation at most δ changes each moment by at most 2δ. The inversion inequality of Appendix E gives $(\theta - \theta')^4 \le 8000(d_1^2 + d_2^2)$, hence

$$(\tau - \tau')^4 \le 1{,}024{,}000\delta^2, \qquad (\tau - \tau')^2 \le 1024\delta. \tag{F.1}$$

For a lower modulus, use the two clean channels in (E.2) with $h = 1/8$ and $\theta' = 4\theta/5$, before adding differential errors. Their conditional Bernoulli report probabilities differ by $-\theta^2 R_1 R_2/40$. The sign marginal is common, so

$$\mathrm{TV}(P_0, P_1) = \frac{\theta^2}{40}, \qquad (\tau_0 - \tau_1)^2 = \frac{32}{5}\,\mathrm{TV}(P_0, P_1). \tag{F.2}$$

Set $\theta^2 = 40\delta$ for $0 \le \delta \le 1/160$. Together with (F.1), this yields

$$\frac{32}{5}\delta \le \Omega(\delta) \le 1024\delta. \tag{F.3}$$

This explicit computation explains why the squared-risk perturbation term is linear in epsilon. It is a model-specific evaluation of an established modulus, rather than a new robust-decision principle.

3. STRUCTURED DIFFERENTIAL ERRORS EMBED IN A HUBER NEIGHBORHOOD

Let P_ξ^ρ be an admissible differential-report law within ϵ of its unknown baseline P_ξ, as in Appendix E. Averaging over T conditional on R changes $\Pr(S = 1 \mid R)$ by at most ϵ. Thus

$$\mathrm{TV}(P_\xi^\rho, P_\xi) \le \epsilon. \tag{F.4}$$

Every baseline conditional report probability is at least $\kappa = 1/10$, and, atomwise, $P_\xi^\rho(o) \ge (1 - \epsilon/\kappa)P_\xi(o)$. For $\epsilon > 0$, put $\alpha = \epsilon/\kappa = 10\epsilon \le 1/16$ and define

$$Q(o) = \frac{P_\xi^\rho(o) - (1-\alpha)P_\xi(o)}{\alpha}.$$

The entries of Q are nonnegative and sum to one. Hence

$$P_\xi^\rho = (1-\alpha)P_\xi + \alpha Q. \tag{F.5}$$

This is a same-observation distribution-class inclusion. It is not equality of the classes: the structured model fixes the uniform outcome-sign marginal, whereas arbitrary Huber contaminations can change it. The differing baseline channels in the lower pair remain unknown in both classes. For the same loss and observation space, the inclusion has the minimax direction

$$\mathcal{P}_{\text{structured}} \subseteq \mathcal{P}_{\text{Huber}} \quad \Longrightarrow \quad R^*(\mathcal{P}_{\text{structured}}) \leq R^*(\mathcal{P}_{\text{Huber}}).$$

Thus a Huber upper bound transfers to the structured class; a Huber lower bound does not transfer without an admissible structured pair.

4. The direction of the argument matters

A lower bound for the larger Huber class does not automatically lower-bound the structured subclass. The construction in Appendix E supplies exactly the missing step: it realizes the modulus-sized target separation inside the allowed conditional reporting perturbations, while equating the complete observed laws. That realization must remain in the manuscript.

The upper bound transfers at the level of bounded moments. Under arbitrary Huber contamination, U and V have mean shifts at most 2α. Finite-grid fitting and (F.1) give target risk $O(G^{-1/2} + \alpha)$. Together with the clean lower bound and (F.3), the Huber class has this scalar order for $0 \leq \alpha \leq 1/16$. For larger α in that interval, use $\delta = \min\{\alpha/(1-\alpha), 1/160\} \geq \alpha/10$ in (F.3); only the constant changes. Setting $\alpha = 10\epsilon$ yields an upper bound for the structured subclass. Its sharper moment biases and its own exact lower pair recover $G^{-1/2} + \epsilon$ without treating the classes as equivalent.

This proves a concrete relationship beyond a verbal analogy. It does not show that the earlier paper states the pooled model, the conditional-error class, the exact constants or the true-OR validation result.

5. Validation remains a different observation problem

Huber's single contaminated iid stream does not directly describe observations with exact T on an independent selected subset of m pools. The validation proof in Appendix E uses an error-free target response on those pools and a complete joint-law lower bound with different counts for the two observation types. Theorems 3.1–3.2 do not supply those steps by replacing n with m. However, unbiased estimation from a validation sample and two-point testing are standard techniques; the existence of these local steps does not establish their novelty or importance as an independent paper.

6. Ura is a direct differential-error comparator, with a different target

Ura (2018), DOI 10.3982/QE886, explicitly allows the measured treatment to depend on the outcome conditional on the true treatment (journal p1337, PDF p4). Assumption 1 (journal p1343, PDF p10) imposes instrument exclusion, true-treatment monotonicity and a nondegenerate binary instrument. Theorem 4 (journal p1345, PDF p12) gives a sharp LATE identified set using differences across instrument groups and a total-variation quantity. Here that TV compares instrument-conditional distributions; (F.1)–(F.3) compare different structural laws. These are not interchangeable uses of TV.

Our model supplies randomization propensity and noisy pool OR reports, uses a scalar randomized-effect target, and has no nondegenerate observed instrument in its constant-propensity submodel. Ura's theorem is not a rate theorem for this experiment. General statements that outcome-dependent misclassification creates identification problems, or that additional measurements can tighten identification, have direct precedent and are not innovations of this paper.

7. Scope of the comparison

The differential-error result is retained as a limitation because it constructs observational equivalence inside the stated enlarged class. The modulus calculation places its scalar rate within established

robust decision theory; it does not by itself extend to heterogeneous CATE or establish a separate general methodology.

APPENDIX G: TWO EXPLICIT MAPPINGS DELIMIT WHAT THIS COMPARISON ESTABLISHES

We compare two latent-variable formulations with the scalar submodel of Appendix E. These mappings clarify which parts of the related models can be carried over; the proof of Proposition 1 remains in Appendix B. Assignments are independent Bernoulli$(1/2)$ draws, T is their true OR, and $R_1, R_2 \in \{-1, 1\}$ satisfy $\mathbb{E}[R_i \mid A_i] = \theta(2A_i - 1)$, so $\tau = 2\theta$. In the clean model, write $s = \Pr(S = 1 \mid T = 1)$, $z = \Pr(S = 0 \mid T = 0)$, and $c = s + z - 1$. Outcome draws are independent given assignments.

HU–SHIU'S DENSITY MODEL ADMITS A MEANINGFUL ATTEMPTED MAPPING

Set $X^* = T$, $X = S$, $Y = R_1$, and $Z = R_2$, recoding signs if required. The observable triples coincide. Moreover, $\tau = -2\mathbb{E}[R_1 \mid T = 0]$, so a bound on the latent conditional distribution would bear directly on our target. The mere absence of a separately named instrument does not exclude comparison.

Conditioning on $T = 1$ leaves assignments $(1, 0)$, $(0, 1)$, and $(1, 1)$ equally likely. Consequently,

$$\mathbb{E}[R_i \mid T = 0] = -\theta, \qquad \mathbb{E}[R_i \mid T = 1] = \theta/3,$$
$$\mathbb{E}[R_1 R_2 \mid T = 1] = -\theta^2/3, \quad \mathrm{Cov}(R_1, R_2 \mid T = 1) = -4\theta^2/9.$$

For binary sign variables, their joint mass minus the product of marginal masses equals one quarter of their covariance times r1*r2. Clean reporting makes S independent of (R_1, R_2) given T. The Hu–Shiu Assumption 2.1 residual under this mapping is

$$d_0(r_1, u, r_2, 1) = -\Pr(S = u \mid T = 1)\frac{\theta^2 r_1 r_2}{9}, \qquad d_0(r_1, u, r_2, 0) = 0.$$

Their h must bound both the pointwise residual and its absolute sum over Y. In this clean model the smallest such value is $h_{\min} = 2\max\{s, 1 - s\}\theta^2/9$. Thus our $\epsilon = 0$ does not imply their $h = 0$ when $\theta \neq 0$. Our ϵ bounds a conditional *reporting-probability* departure from an unknown common channel; their h bounds a joint *factorization* departure. They cannot simply be substituted for each other.

With rows $S = 0, 1$ and columns $R_2 = -1, 1$, the observed matrix satisfies $\det M_{S,R_2} = c\theta/8$.

The difference between P(R1=r|T=1) and P(R1=r|T=0) is 2theta*r/3. The absolute conditional-density eigenvalue gap is 2|theta|/3. At theta=0 both observed invertibility and the distinct-density condition fail. Both deteriorate along weak-effect sequences even with fixed c>0. Hu–Shiu's displayed Proposition 2.1 also imposes a perturbation/eigen-gap condition; checking that inequality uniformly is separate from checking these two necessary conditions.

Their regression subsection has another obstacle for this mapping:

$$\mathbb{E}[R_1 \mid R_2 = r, T = 1, S = u] - \mathbb{E}[R_1 \mid T = 1] = -\frac{4\theta^2 r}{3(3 + \theta r)}.$$

For nonzero theta this violates Assumption 2.7 (exact conditional mean independence) even in our clean model. At theta=0 the other degeneracies remain. This observation applies to the mapped regression result, not to their broader density sensitivity analysis.

Conclusion: the displayed results cannot be invoked as an immediate uniform risk theorem over our class using this mapping. For nonzero theta the density model's partial-independence approach remains relevant; a separate calculation could establish useful pointwise bounds. We have not excluded alternative mappings, proved a universal nontransfer theorem, or established priority.

MOLINARI'S CONSTRAINED-CHANNEL FORMULATION REPRESENTS THE SCALAR LAW

Use eight latent categories $x = (r_1, r_2, t)$ and eight observed categories $w = (r_1, r_2, u)$. Define a column-stochastic matrix

$$\Pi_{w,x} = \mathbf{1}\{(w_1, w_2) = (x_1, x_2)\} \Pr(S = u \mid R_1 = r_1, R_2 = r_2, T = t).$$

Then $P_w = \Pi P_x$ is exact, including outcome-dependent reporting. Each outcome-pair block is a two-by-two reporting matrix. The target is a linear functional of the latent probabilities:

$$\tau(P_x) = -8 \sum_{r_1, r_2, t} (1 - t) r_1 P_x(r_1, r_2, t) = 2\theta.$$

A true-OR validation observation supplies both x and w because the outcomes and S are already observed. Thus neither representing the complete discrete law as a constrained channel nor using observed true/noisy pairs is a new principle here. Molinari's 2005 manuscript explicitly also treats estimated validation restrictions, with n=N/kappa and fixed 0<kappa<infinity in C2.

The specific pooled problem additionally restricts P_x through randomized assignments and the outcome model, ties reporting blocks to a common unknown baseline with a prescribed epsilon budget, and varies validation size m relative to G. Recovering its risk law requires preserving those restrictions and proving risk bounds, beyond writing P_w=Pi P_x. The generic representation alone neither proves nor disproves the model-specific minimax contribution.

DETERMINISTIC VERIFICATION SCOPE

An exact-arithmetic diagnostic checks the displayed identities on 27 clean models (nine signed or zero effect values and three channels): 432 factorization atoms, 27 determinants, 27 density gaps, and 108 conditional means. For each model it additionally checks reporting perturbations of the form ar_1r_2 with $a \in \{-1/160, 0, 1/160\}$: 81 augmented channels, 648 observed atoms, and 81 target functionals. The identities above supply the algebra for all stated parameters; the finite grid is a diagnostic against implementation/algebra mistakes, not a replacement for proof.

APPENDIX H: A FINITE-BUDGET INFORMATION OBSTRUCTION

For a specified scalar experiment, a two-point Bayes–Hellinger argument gives the finite-budget lower bound below. It is a bound for this fixed experiment, rather than an exact minimax calculation.

1. EXPERIMENT, TARGET, AND COMPARATOR

There are $G = 44$ independent pools of size two. Each pool reveals both outcome signs $R_1, R_2 \in \{-1, 1\}$ and a reported OR bit S. Individual assignments and the true OR remain hidden ($m = 0$), and the supplied propensity is $e = 1/2$. The parameter set and target are

$$\theta \in [-1/2, 1/2], \qquad s, z \in [1/10, 9/10], \qquad c = s + z - 1 \geq 1/10, \qquad \tau = 2\theta.$$

The marginal outcome signs are independent and uniform. Their complete reported law satisfies

$$\Pr(S = 1 \mid R_1, R_2) = s - \frac{c}{4}(1 - \theta R_1)(1 - \theta R_2). \tag{H1}$$

Risk is expected squared error for τ. Ancillary covariates, if included, have the same independent law under both alternatives. Write R^* for the minimax risk over all estimators measurable from the full observations, including parameter-independent randomization.

For $U = S(R_1 + R_2)/2$, set $b = c/4$ and $\pi = s - b$. Direct calculation gives

$$\mathbb{E}U = b\theta, \qquad \mathrm{Var}(U) = \frac{\pi}{2} - \frac{b\theta^2}{2} - b^2\theta^2. \tag{H2}$$

For the G independent pools, write

$$\bar{U} = G^{-1} \sum_{g=1}^{G} U_g.$$

Consider the linear family $\mathrm{clip}(a\bar{U}, -1, 1)$ with $a = 80h$. The factor 80 converts the original coefficient for θ to the present target 2θ: at $c = 1/10$ and $e = 1/2$, $b = 1/40$. For its raw, unclipped version,

$$R_{\mathrm{raw}} = \frac{a^2\pi}{2G} + \theta^2 \left[4 - 4ab + a^2 \left\{ \left(1 - \frac{1}{G} \right) b^2 - \frac{b}{2G} \right\} \right]. \tag{H3}$$

At a fixed channel, this is affine in θ^2, so its maximum occurs at $\theta = 0$ or $|\theta| = 1/2$. At fixed θ, it is convex in (s, z): the only quadratic channel coefficient is $a^2(1 - 1/G)\theta^2 \geq 0$. The legal channel set is the triangle with vertices $(1/5, 9/10)$, $(9/10, 1/5)$, and $(9/10, 9/10)$. Its maximum therefore occurs at a vertex. This reduces the continuous class to six corner risks, without approximating it by a grid.

The maximum of those six risks is minimized within the raw linear family at

$$h = \frac{44}{2823}, \qquad a = \frac{3520}{2823}, \qquad U := \sup R_{\mathrm{raw}} = \frac{2779}{2823} = 0.9844137442437123\ldots. \tag{H4}$$

At this coefficient, the $(\theta = 1/2, s = 9/10, z = 1/5)$ quadratic has zero derivative and equals the envelope; the other five values are no larger. Convexity in h proves optimality within the raw family. Clipping contracts distance to every $\tau \in [-1, 1]$, so the feasible clipped rule has maximum risk $B_{\mathrm{linear}} \leq U$. Neither equality after clipping nor optimality among nonlinear rules is claimed. This fixed-budget coefficient does not replace the paper's rate estimator.

2. Lower bound from the full observed experiment

Use $\theta \in \{-1/2, 1/2\}$ with the same channel $s = z = 11/20$. The targets are -1 and 1, and $c = 1/10$. In units of $1/640$, the complete one-pool laws are:

(R_1, R_2, S)	$\theta = 1/2$	$\theta = -1/2$
$(+1, +1, 1)$	87	79
$(+1, +1, 0)$	73	81
$(-1, -1, 1)$	79	87
$(-1, -1, 0)$	81	73
$(+1, -1, 1)$	85	85
$(+1, -1, 0)$	75	75
$(-1, +1, 1)$	85	85
$(-1, +1, 0)$	75	75

These probabilities agree atom by atom with the preserved latent-assignment enumerator. In particular, both S=0 outcome patterns are retained: this is not a lower bound restricted to the four-category coarsening.

Let p_+, p_- denote the one-pool laws and P_+, P_- their G-fold products. For the equal prior on the two alternatives, the posterior mean action is

$$d_B(x) = \frac{P_+(x) - P_-(x)}{P_+(x) + P_-(x)}.$$

Its Bayes risk is

$$B_G = \sum_x \frac{2P_+(x)P_-(x)}{P_+(x) + P_-(x)}. \tag{H5}$$

This follows by minimizing the conditional quadratic. Every estimator's maximum risk dominates this average; external randomization cannot improve it. Cauchy–Schwarz gives

$$\left\{ \sum_x \sqrt{P_+(x)P_-(x)} \right\}^2 \leq B_G \sum_x \frac{P_+(x) + P_-(x)}{2} = B_G.$$

By product independence,

$$R^* \geq B_G \geq A^{2G}, \qquad A = \sum_w \sqrt{p_+(w)p_-(w)} = \frac{1}{2} + \frac{\sqrt{6873} + \sqrt{5913}}{320}. \tag{H6}$$

It is legitimate to use a common fixed channel in a lower-bound pair even though the estimator faces an unknown channel. No actual estimator is given the true channel in the comparator. The argument does not assert that this pair is least favorable or that the bound is sharp.

3. Exact certificate and finite-budget implication

For each rational atom product $x = p_+(w)p_-(w)$, the calculation first constructs rational lower and upper enclosures for $\sqrt{x}$ using $D = 10^{18}$ and $n = \lfloor \sqrt{\lfloor D^2 x \rfloor} \rfloor$. It verifies $(n/D)^2 \leq x < ((n+1)/D)^2$. Only after these square-root enclosures are fixed does it use exact integer arithmetic to sum the rational lower endpoints, raise the sum to power $2G = 88$, and round downward to twelve decimal places. Thus

$$\begin{gathered} R^* \geq L := 0.946362654615, \qquad \frac{19}{20}U = 0.9351930570315268\ldots, \\ L - \frac{19}{20}U = 0.0111695975834732\ldots > 0. \end{gathered} \tag{H7}$$

These are exact rational comparisons, with decimals displayed for readability.

Any feasible named baseline with maximum risk $B_{\text{base}} \leq U$ would need a candidate with maximum risk at most $(19/20)B_{\text{base}} \leq (19/20)U < R^*$ to achieve a 5% improvement. This is impossible. A comparison set containing the clipped linear rule has a best baseline with $B_{\text{base}} \leq U$; ranking its other members is unnecessary for this conclusion. The bound does not substitute U for the actual baseline risk to claim a positive gain. The largest possible relative gain over any such baseline is at most

$$1 - \frac{L}{U} = 0.038653553804194\ldots < 0.05. \tag{H8}$$

This conservative bound does not show that a 3.865% gain is attainable. It rules out a 5% improvement in worst risk for the specified scalar design, channel class, loss, and budget.

4. Verification and limits

The verification uses exact fractions and integer arithmetic. It checks 64 latent-versus-observed atom equalities, eight square-root enclosures, all six raw-risk envelope corners against independent latent-law moments, and the linear-rule coefficient solver's rational bracket. Separate ordered sample enumeration for G=1,2,4 (4,168 samples in total) confirms normalization, the posterior-mean Bayes identity, reflection symmetry, the affinity lower bound, the raw moment-risk formula, clipping contraction and zero risk. These are scoped deterministic checks, not independent human proof review.

The conclusion applies to $k = 2$, $G = 44$, $m = 0$ and the stated scalar model, target, and channel class. It does not cover other budgets, true-OR validation, or the full functional class. The 84-case audit is independent of this calculation.

Appendix I: Two validation scales: changing the order and recovering the benchmark

The matching risk theorem yields two validation scales. This appendix works out their asymptotic consequences under the same experiment and parameter class.

1. Keep the experiment and quantifiers fixed

Let $M(G, m, \sigma_0)$ denote the classwise minimax integrated squared CATE risk in Proposition 1. The pool size and the channel, overlap, and smoothness bounds remain fixed. Along a sequence $G \to \infty$,

the integer $m = m_G$ satisfies $0 \le m \le G$, while $\sigma_0 = \sigma_{0,G}$ may vary within its stated legal range. Set

$$\alpha = \frac{2\beta}{2\beta + d} \in (0,1), \qquad D_0 = 1 + G\sigma_0, \qquad H_{\mathrm{rate}} = G^{2\alpha-1}, r_m = G^{-\alpha} + \{G(D_0 + m)\}^{-1/2}.$$

Proposition 1 gives $C_{\mathrm{low}} r_m \le M(G, m, \sigma_0) \le C_{\mathrm{up}} r_m$, with constants independent of G, m, σ_0. Every comparison below holds at the same G and for the same class $\mathcal{P}(\sigma_0)$, including its supplied propensity functions. The lower bound is classwise, not a bound at each fixed propensity or law. D_0 summarizes the class difficulty; the upper estimator needs neither σ_0 nor a covariate-distribution oracle.

2. An exact ratio for the proxy, not for actual finite risk

Set $t = \sqrt{D_0/H_{\mathrm{rate}}}$ and $v = \sqrt{D_0/(D_0 + m)}$; then $t > 0$ and $0 < v \le 1$. Dividing r_m and r_0 by $(GD_0)^{-1/2}$ gives the exact proxy ratio

$$\frac{r_m}{r_0} = \frac{t + v}{1 + t}. \tag{V1}$$

For every admissible G, m, σ_0,

$$\max\left\{v, \frac{t}{1+t}\right\} \le \frac{t+v}{1+t} \le \min\{1, t + v\}. \tag{V2}$$

Indeed, the difference between the middle expression and v is $t(1-v)/(1+t) \ge 0$; the other bounds follow from $0 < v \le 1$. The proxy ratio equals one at $m = 0$ and is nonincreasing in m.

The experiments have the same monotonicity. Given the m validation indices, select a uniform m'-subset with $m' \le m$ and discard the other revealed T values. For any fixed $J \subseteq \{1, \dots, G\}$ of size m', the resulting probability is

$$\frac{\binom{G-m'}{m-m'}}{\binom{G}{m}\binom{m}{m'}} = \frac{1}{\binom{G}{m'}}.$$

The thinning is independent of the study variables, so the resulting data have the m'-validation law. At fixed G and σ_0, minimax risk cannot increase with m. This does not assert monotonicity of a particular fitted rule or of risk at equal total cost.

The uniform constants of Proposition 1 also give

$$\frac{C_{\mathrm{low}}}{C_{\mathrm{up}}} \frac{r_m}{r_0} \le \frac{M(G, m, \sigma_0)}{M(G, 0, \sigma_0)} \le \frac{C_{\mathrm{up}}}{C_{\mathrm{low}}} \frac{r_m}{r_0}. \tag{V3}$$

Hence, along any admissible sequence,

$$\frac{M(G, m, \sigma_0)}{M(G, 0, \sigma_0)} \longrightarrow 0 \quad \Longleftrightarrow \quad \frac{D_0}{H_{\mathrm{rate}}} \longrightarrow 0 \text{ and } \frac{m}{D_0} \longrightarrow \infty. \tag{V4}$$

By (V2), the proxy ratio vanishes exactly when $t \to 0$ and $v \to 0$; these are the two conditions on the right. Inequality (V3) transfers the equivalence to minimax risk. The criterion does not require polynomial budgets.

3. Separate four regular asymptotic regimes

Regular regime	Risk order	Role of validation
$D_0 \gtrsim H_{\mathrm{rate}}$	$G^{-\alpha}$ at $m = 0$	No change in the minimax order.
$D_0 = o(H_{\mathrm{rate}}),\ m = O(D_0)$	$(GD_0)^{-1/2}$	Added validation does not dominate existing calibration.
$D_0 = o(m),\ m = o(H_{\mathrm{rate}})$	$(Gm)^{-1/2}$	Partial recovery; the risk ratio is of order $\sqrt{D_0/m}$.
$D_0 = o(H_{\mathrm{rate}}),\ m \gtrsim H_{\mathrm{rate}}$	$G^{-\alpha}$	The known-channel order is recovered.

The table covers regular regimes. Equations (V1)–(V4) also cover oscillating sequences. In all cases, the known-channel benchmark is recovered exactly when $D_0 + m \gtrsim H_{\mathrm{rate}}$, up to the constants in Proposition 1. This is not a sharp integer cutoff.

If $\alpha \leq 1/2$, then $H_{\mathrm{rate}} \leq 1 \leq D_0$, so the first row applies. Because $\beta \leq 1$ and d is a positive integer, an order change is possible only for $d = 1$ and $\beta > 1/2$. Even there, fixed positive σ_0 makes D_0 linear in G and again puts the model in the first row. Validation changes the leading rate only when class-level propensity variation is sufficiently weak.

For example, with $\alpha = 2/3$ and $\sigma_0 = 0$, any $m \to \infty$ gives an unbounded improvement over the unvalidated $G^{-1/2}$ order. At $m \asymp G^{1/6}$, the risk order is $G^{-7/12}$, still slower than the known-channel $G^{-2/3}$ order; recovery needs $m \gtrsim G^{1/3}$. These asymptotic statements do not prescribe a budget at $G = 44$.

If $\sigma_0 = cG^{-5/6}$ for fixed legal $c > 0$, then $D_0 \asymp G^{1/6}$. Budgets $m \asymp G^{1/12}$ and $m \asymp G^{1/4}$ give, respectively, the unchanged $G^{-7/12}$ order and the improved $G^{-5/8}$ order. At $m \asymp G^{1/3}$, the order is $G^{-2/3}$. Propensity variation moves the first threshold even though the final benchmark threshold has the same exponent.

4. What the result does and does not supply

The two scales describe when validation changes classwise risk and when it recovers the known-channel order. They are consequences of Proposition 1, not additional identification assumptions or an independent minimax theorem. Neither scale specifies a finite-sample percentage reduction or an institutional validation price. Appendix D handles the stated total-budget comparison; Appendix H records a separate $G = 44$ obstruction to one finite candidate family.

Appendix J: A direct validation-only CATE benchmark

The direct validation-only comparison in the main paper follows from the histogram upper bound and the full-label hypercube lower bound. No new lower-bound construction is needed.

The restricted experiment

Keep only $(X_1, \ldots, X_k, Y_1, \ldots, Y_k, T)$ from each of the m uniformly selected validated pools. Discard S and all unvalidated records. The supplied propensity remains known; P_X and the arm regressions remain unknown. The restricted experiment uses the same acquired (G, m) budget as the joint experiment, but discards most of its observations. For $m \geq 1$, its classwise minimax risk satisfies

$$R_{\mathrm{val}}(m, \sigma_0) \asymp m^{-\alpha}, \qquad \alpha = \frac{2\beta}{2\beta + d}. \tag{J.1}$$

At the endpoint $m = 1$, this means a constant-order risk: bounded clipping gives the upper bound, while a fixed two-point Hölder submodel inside the same class $\mathcal{P}(\sigma_0)$ gives a positive lower bound. For $m \geq 2$, the displayed histogram and hypercube arguments give the stated $m^{-\alpha}$ order. Constants depend only on the fixed inputs of Proposition 1. This bound does not apply to estimators that also use the unvalidated records.

Upper construction without knowing P_X

Use the first prespecified member of each validated pool and set

$$Z = \frac{Y_1}{e(X_1)} \left\{ 1 - \frac{\mathbf{1}\{T = 0\}}{\prod_{i=1}^{k} (1 - e(X_i))} \right\}. \tag{J.2}$$

Conditional on the pool covariates, the weighted $T = 0$ term has outcome mean $\mu_0(X_1)$; hence $\mathbb{E}[Z \mid X_1] = \tau(X_1)$. Overlap and bounded outcomes imply

$$|Z| \le M := \frac{B\{\eta^{-(k-1)} - \eta\}}{\eta(1-\eta)}.$$

The m complete pools give independent regression observations. A regular histogram with J cells has squared approximation error $O(J^{-2\beta/d})$ and cell-mean estimation error $O(M^2 J/(m+1))$, including empty cells under arbitrary P_X. The empty-cell rule and clipping to $[-2B, 2B]$ are those of Appendix A. Taking $J \asymp m^{d/(2\beta+d)}$ proves the upper order without arm-regression smoothness or a covariate-distribution oracle.

Lower comparison

Full (X, A, Y) labels on km independent members determine the validation-only data by grouping and computing T. Fix a legal channel in the original class, although S is discarded. The full-label Hölder hypercube at the same supplied step propensity and σ_0 gives the lower bound $c(km)^{-\alpha}$. Fixed k absorbs its factor $k^{-\alpha}$. A faster validation-only estimator would therefore contradict the full-label lower bound.

What the comparison explains

At $\sigma_0 = 0$ and $m \asymp G^\gamma$ for $0 < \gamma \le 1$, the validation-only risk is of order $G^{-\alpha\gamma}$. Proposition 1 gives the joint experiment exponent $\min\{\alpha, (1+\gamma)/2\}$. For $0 < \gamma < 1$,

$$\min\left\{\alpha, \frac{1+\gamma}{2}\right\} > \alpha\gamma. \tag{J.3}$$

Both $\alpha > \alpha\gamma$ and $(1+\gamma)/2 > \gamma \ge \alpha\gamma$ establish (J.3). At $\gamma = 1$ the risk orders agree. This is an asymptotic comparison at the stated pool budget, not a measured cost advantage.

For $d = \beta = 1$ and $\alpha = 2/3$:

Validation size	Validation-only order	Joint order	Interpretation
$m \asymp G^{1/6}$	$G^{-1/9}$	$G^{-7/12}$	Unvalidated outcome pools remain useful.
$m \asymp G^{1/3}$	$G^{-2/9}$	$G^{-2/3}$	Sublinear validation restores the G-pool order.
$m \asymp G$	$G^{-2/3}$	$G^{-2/3}$	The orders coincide.

The joint method reuses outcome information from all G pools after calibration. The comparison does not imply that true OR labels carry no outcome information. In the finite candidate-selection problem, a rule that improves on zero can still lose to direct regression using (J.2); practical comparisons should include this control at the same acquired budget.

APPENDIX K: WHY THE FIXED-INPUT THEOREM CANNOT BE EXTENDED UNIFORMLY

The following examples show why each restriction in Proposition 1 is needed for a uniform statement. They do not enlarge or modify Proposition 1; they show why its fixed-input restrictions cannot be removed uniformly.

1. A COMMON TESTING LEMMA

Let P_+ and P_- have constant CATE targets B and $-B$, with equal prior weight. For densities p_+, p_- with respect to a common dominating measure, pointwise quadratic minimization gives the optimal Bayes squared risk

$$B^2 \int \frac{2p_+p_-}{p_+ + p_-} \geq B^2 \int \min\{p_+, p_-\} = B^2\{1 - \mathrm{TV}(P_+, P_-)\}. \tag{K.1}$$

The integrand is set to zero where $p_+ + p_- = 0$. The inequality follows from $2ab/(a+b) \geq \min\{a, b\}$. Maximum risk dominates Bayes risk. For constant CATE alternatives and common parameter-free X, the same bound applies to integrated function loss by conditional Jensen.

All examples use $Y = BR$, where $R \in \{-1, 1\}$ and

$$\Pr(R = r \mid A = a, X) = \frac{1}{2} + \theta(a - e)r, \qquad \theta \in \{-1/2, 1/2\}. \tag{K.2}$$

Potential outcomes with these marginal laws realize the randomized causal model. The CATE is $\tau = 2B\theta$. Marginally, R is uniform and independent across members under either alternative. The examples use a common covariate law, constant supplied propensity e, and $\sigma_0 = 0$. Their constant CATE is Hölder for every allowed β, L_τ.

2. GROWING POOLS CAN ERASE INFORMATION EVEN WITH EVERY TRUE OR

Take $e = 1/2$. Conditional on $R_1, \ldots, R_k$, the probability that no member is treated is

$$v_\theta := \Pr(T = 0 \mid R_1, \ldots, R_k) = 2^{-k} \prod_{i=1}^{k} (1 - \theta R_i).$$

The two alternatives have the same marginal law of the R_i. Their complete (R, T) one-pool laws therefore satisfy

$$\mathrm{TV}(P_+, P_-) = \mathbb{E}_R|v_+ - v_-| \leq \mathbb{E}_R(v_+ + v_-) = 2^{1-k}. \tag{K.3}$$

Give the analyst all G true ORs, even if the original budget has $m < G$. A fixed legal channel generates S from T through a common kernel, which adds no information to this enlarged experiment. Product coupling and (K.1) yield, for every $0 \leq m \leq G$,

$$\inf_{\widehat{\tau}} \max_{\theta \in \{-1/2, 1/2\}} R_\theta(\widehat{\tau}) \geq B^2[1 - G2^{1-k}]_+. \tag{K.4}$$

At $k = \lceil 2 \log_2 G \rceil$, this is at least $B^2(1 - 2/G)$. The fixed-k formula at $m = G$ would tend to zero if its constants were incorrectly held uniform in k. This example concerns $\sigma_0 = 0$ and the stated growing-k sequence; it does not classify all growing-pool regimes.

3. VANISHING OVERLAP CAN DEFEAT EVEN INDIVIDUAL LABELS

For constant $e = \eta$, the one-member full-label (A, R) laws in (K.2) have $\mathrm{TV}(P_+, P_-) = \mathbb{E}_A|A - e| = 2e(1 - e) \leq 2e$. Full labels on kG members are more informative than any pooled reports and true-OR validation. With a common legal channel, product coupling and (K.1) give

$$\inf_{\widehat{\tau}} \max_{\theta \in \{-1/2, 1/2\}} R_\theta(\widehat{\tau}) \geq B^2[1 - 2kGe]_+. \tag{K.5}$$

For fixed k and $e = G^{-2}$, the lower bound tends to B^2 even at $m = G$. This is a sequence outside the theorem's fixed positive-overlap class; it does not assert impossibility at fixed $\eta > 0$.

4. Weak reporting contrast matters when m=0

Fix $k = 2$, $e = 1/2$, and $m = 0$. Choose $s = z = (1 + c)/2$ and $\kappa = c$ for $0 < c \le 1/4$. This channel is legal for each c. Conditional on R, $\Pr(S = 1 \mid R) = s - cv_\theta$. The squared difference of these probabilities, averaged over R, is $c^2/8$. Both probabilities lie in $[3/8, 5/8]$, so

$$\chi^2(P_+ \| P_-) \le \frac{8c^2}{15} =: a, \qquad \chi^2(P_+^{\otimes G} \| P_-^{\otimes G}) \le (1 + a)^G - 1. \tag{K.6}$$

If $Ga < 1$, the binomial expansion bounds the last quantity by $Ga/(1 - Ga)$. Using $\mathrm{TV} \le \sqrt{\chi^2}/2$ and setting $c = 1/G$ for $G > 4$ gives

$$\inf_{\widehat{\tau}} \max_{\theta \in \{-1/2, 1/2\}} R_\theta(\widehat{\tau}) \ge B^2 \left\{ 1 - \frac{1}{2}\sqrt{\frac{8}{15G - 8}} \right\}. \tag{K.7}$$

The lower bound tends to B^2. Here κ decreases with G; the example concerns uniformity across channel-contrast bounds, not the fixed-κ theorem. It uses no validation and does not preclude recovery from true-OR labels.

5. The other channel endpoint removes channel uncertainty

The constraints $s, z \le 1 - \kappa$ and $s + z - 1 \ge \kappa$ imply $3\kappa \le 1$. At $\kappa = 1/3$, equality forces $s = z = 2/3$, so the channel is known. Its minimax order is then $G^{-\alpha}$, even at $m = 0$. With $\sigma_0 = 0$ and $\alpha > 1/2$, a uniform $G^{-1/2}$ unknown-channel lower term cannot hold at this endpoint. The lower-pair width in P2 likewise vanishes: $h_{\max} = (1 - 3\kappa)/8 = 0$. For $\kappa > 1/3$ the class is empty. Proposition 1 therefore takes fixed $\kappa < 1/3$.

6. Smoothness and validation restrictions need precise wording

Under $|\tau(x) - \tau(y)| \le L\|x - y\|^\beta$ on the connected cube, $\beta > 1$ forces τ to be constant. Divide a line segment into n equal pieces; the triangle inequality bounds the difference by $L\|x - y\|^\beta n^{1-\beta} \to 0$. Any higher-order Hölder extension must instead define smoothness through derivatives.

Uniform validation indices are sufficient, not necessary. Conditional on any data-independent index set of size m, the iid pools give the same split and product divergence. Selection depending on Y, S, T changes the observation law and is outside the current theorem.

The P8 comparison uses $\sigma \le \omega < 2\sigma$ when $\omega = \sigma + \omega_{\mathrm{cal}}$ and $\sigma > \omega_{\mathrm{cal}}$; equality on the left occurs at $\omega_{\mathrm{cal}} = 0$.

Appendix L: The scalar mechanism and its restricted operator embedding

Reading guide. L1–L2 study a separate Gaussian experiment and its restricted operator embedding. L3 isolates the deterministic function-risk algebra used in P9; the complete argument is in Appendix B. The additional-validation plug-in, fixed-budget construction and covariance-grid diagnostics are in Numerical Supplement S5, S6 and S7, respectively.

L1. A scalar Gaussian diagnostic for the calibration rate

This separate experiment illustrates the geometry behind the rate. Its algebra parallels Appendix B, P7–P8, while its observation model remains distinct from the pooled experiment.

1. Experiment and uniform claim

Fix 0 < l < M, T > 0, and 0 <= A < infinity. Observe

```
Z = (bt, bt^2, sqrt(a)b) + n^(-1/2) epsilon,
epsilon ~ N(0,I_3), b in [l,M], t in [-T,T], a in [0,A].
```

The estimator receives n and a, but not b or t. Define

```
R_n(a) = inf_estimators sup_(b,t) E[(t_hat-t)^2],
r_n(a) = 1/sqrt(n(1+na)).
```

Proposition L.1. There exist $0 < c < C < \infty$, depending only on l,M,T,A, such that c r_n(a) <= R_n(a) <= C r_n(a) for all integers n>=1 and all a in [0,A]. Nondegenerate b and t intervals are essential.

2. A FINITE-GRID MINIMUM-DISTANCE UPPER BOUND

Let f_a(b,t)=(bt,btš,sqrt(a)b). Use equally spaced grids including the endpoints of each parameter interval, with respective meshes at most 1/n. Choose the grid point minimizing ||Z-f_a(b',t')||š, breaking ties by its fixed grid index. This is a measurable finite minimization. A grid point (b*,t*) within 1/n of the truth obeys

```
||f_a(b*,t*)-f_a(b,t)||^2 <= K0/n^2,
K0=(M+T)^2+(2MT+T^2)^2+A.
```

Write D=||f_a(b_hat,t_hat)-f_a(b,t)||š and xi=Z-f_a(b,t). Optimality and the triangle inequality give

```
sqrt(D) <= 2||xi||+sqrt(K0)/n,
E D <= 24/n+2K0/n^2 <= (24+2K0)/n.                    (L.1)
```

For any two points in the rectangle, set u=t_hat, v=b_hat, delta=u-t, d1=vu-bt, d2=vuš-btš, db=v-b, so D=d1š+d2š+a dbš. Exact identities are

```
d2-u d1 = b t delta,
d2-t d1 = v u delta,
v delta = d1-t db.                                  (L.2)
```

Squaring the first two, bounding each by (1+Tš)(d1š+d2š), and using (u-t)š<=2(uš+tš), gives

```
delta^4 <= C4 D, C4=4(1+T^2)/l^2.                   (L.3)
```

For a>0, split at tš=sqrt(a). When tš>=sqrt(a), the first identity yields deltaš <= (1+Tš)D/(lš sqrt(a)). Otherwise the third identity yields

```
delta^2 <= 2(D+t^2D/a)/l^2
        <= 2(1+sqrt(A))D/(l^2 sqrt(a)).
```

Hence, with C2=max(1+Tš,2(1+sqrt(A)))/lš,

```
delta^2 <= C2 D/sqrt(a).                            (L.4)
```

Taking expectations in L.3–L.4 and using (L.1) gives E deltaš <= C min(n^(-1/2),1/(n sqrt(a))) for a>0; at a=0 the first bound applies. For x=na, the ratio of this minimum to r_n(a) lies between 1 and sqrt(2). This proves the uniform upper bound, without an asymptotic normal approximation.

3. A COMPLETE GAUSSIAN-LAW LOWER BOUND

Let b0=(l+M)/2 and fix positive constants

```
d=min((M-l)/4,1/4),
epsilon0=min(T^2/sqrt(A+1),1/4).
```

For each n,a let z=a+1/n and define

```
h=d/sqrt(nz), b1=b0+h,
t0^2=epsilon0 sqrt(z), t0>0,
t1=b0 t0/b1.
```

Both points lie in the parameter rectangle: 0<h<=d, b1<M, t0š<=Tš and 0<t1<=t0. The first mean coordinate is identical. The second mean difference is -b0 h t0š/b1, so the exact Gaussian Kullback–Leibler divergence is

```
KL(P0,P1) = (n h^2/2)[a+(b0/b1)^2 t0^4]
          <= (d^2/2)(a/z+epsilon0^2)
          <= 17/512 < 1/8.                    (L.5)
```

Pinsker's inequality gives total variation <=1/4. The target gap satisfies

```
Delta^2=(t0-t1)^2=t0^2 h^2/b1^2
       =epsilon0 d^2 r_n(a)/b1^2
       >=epsilon0 d^2 r_n(a)/M^2.              (L.6)
```

Convert any estimator into the nearer-target test. A wrong test implies squared error at least Deltaš/4. The average testing error is at least (1-TV)/2, so maximum estimation risk at the pair is at least Deltaš(1-TV)/8 >= 3 Deltaš/32. This proves Proposition L.1 with c=3 epsilon0 dš/(32 Mš). The same pair construction works with both targets negative. It also covers n=1 and a=0.

4. Removing the quadratic moment changes the uniform order

Now observe only Z'=(bt,sqrt(a)b)+n^(-1/2)N(0,I_2), with the same parameter class and information contract. Its minimax squared risk is

```
R'_n(a) asymp 1/(1+na).                       (L.7)
```

For the upper bound, a=0 admits the bounded zero estimator. For a>0, let b_hat be Z'_2/sqrt(a) projected onto [l,M], and let t_hat be Z'_1/b_hat projected onto [-T,T]. Projection contraction gives

```
E(t_hat-t)^2 <= (2/l^2)[1/n+T^2/(na)],
E(t_hat-t)^2 <= 4T^2.
```

Since a<=A, the first expression is at most C/(na). Taking the better bound yields C'/(1+na), uniformly. The bounds apply to the same projected estimator; no unknown parameter chooses the estimator.

For the lower bound, take b0 as above, t0=T/2, h=d/sqrt(1+na), b1=b0+h, t1=b0 t0/b1. The first coordinates agree, KL=na hš/2<=dš/2<1/8, and

```
(t0-t1)^2 >= T^2d^2/[4M^2(1+na)].
```

The same testing argument proves (L.7). At a=0, b0=3/2,t0=1/2 and b1=7/4,t1=3/7 (within l=1,M=2,T=1) give exactly identical laws in the ablated experiment but distinct targets. Restoring btš gives means 3/8 and 9/28 in the second coordinate, which differ by 3/56.

Thus the quadratic coordinate supplies consistency even without direct calibration. This conclusion concerns the Gaussian experiment. Dropping one constructed pooled moment does not remove that information from the full pooled data; (L.7) is not a lower bound for the pooled experiment.

5. Information calculation explains the difficult scale

With t first and b second, the per-n Fisher matrix is J'J where

```
J_t=(b,2bt,0), J_b=(t,t^2,sqrt(a)).
```

Its determinant is bš[t^4+a(1+4tš)]. When tš+t^4+a>0, the Schur complement after projecting out b is

```
I_eff/n = b^2[t^4+a(1+4t^2)]/(t^2+t^4+a).     (L.8)
```

For small positive a and tš of order sqrt(a), (L.8) is of order sqrt(a). The lower pair instead uses tš of order sqrt(a+1/n), which also covers the nonregular boundary. This is an explanation, not the proof of Proposition L.1. The expression is undefined at a=t=0; no regular information claim is made there. Pointwise Cramer–Rao calculations do not prove a uniform minimax bound at a singular nuisance parameter.

6. What the correspondence to pooled CATE establishes

Formally substituting n=G and a=m/G+sigma gives

```
r_G(a) = [G(1+m+G sigma)]^(-1/2).
```

In the pooled upper construction, the actual learning size is q, reserve size l_res and calibration weight l_res/q; its fitted propensity weight is random. Appendix B, P7–P9 handles the resulting random discrepancy, rare underestimation and transfer to the function. In the minimax law the design-class lower bound sigma0 replaces the law-specific sigma.

This diagnostic proves that the numerical scalar rate is shared by a simple nonlinear Gaussian experiment. It does not give a Markov kernel, Le Cam equivalence, a Gaussian approximation for the pooled moments or a theorem from prior literature. It cannot by itself prove the pooled risk result. The full-law lower bound, rough-arm cancellation and functional transfer remain separate steps in the pooled proof.

L2. A curved subexperiment embeds, but the product-class rate does not transfer sharply

This subsection gives a concrete mapping for the Gaussian diagnostic. It does not establish equivalence between Gaussian and pooled observations.

1. The comparison theorem is over a product of parameter classes

Johannes–Schwarz (2013), DOI 10.1080/03610926.2012.731548, observe independent Gaussian sequences Y_j=a_j f_j+sqrt(epsilon)xi_j and X_j=a_j+sqrt(sigma)eta_j. Their model (1.1), classes F_gammaˆr and A_lambdaˆd, and Theorems 2.3/2.5 range over a product of solution and operator classes. Step B of their lower proof varies target and operator while preserving the main image.

2. An explicit positive-calibration mapping exists

Fix b in [1/2,1], t in [-1/2,1/2], n>=1 and a in [0,1]. The diagnostic is

```
Z1=bt+xi1/sqrt(n),
Z2=bt^2+xi2/sqrt(n),
Z3=sqrt(a)b+xi3/sqrt(n),                                    (1)
```

with independent standard Gaussian noises. For a>0, set C=Z3/sqrt(a), so Var(C)=1/(na). Independently generate D~N(0,1/(na)), and output

```
Y1=Z1, Y2=Z2, X1=C+D, X2=C-D.                              (2)
```

Then X1,X2 are independent N(b,2/(na)) variables, independent of Y1,Y2. The inverse kernel retains Y1,Y2 and sets Z3=sqrt(a)(X1+X2)/2. The discarded difference (X1−X2)/2 is independent of the retained data and has a parameter-free distribution. Thus (1) and this finite-coordinate experiment are exactly equivalent by parameter-free randomized kernels. The kernels use only known n,a. At a=0 this construction is unavailable; Z3 is ancillary, not an operator observation with finite noise.

The coordinates take the unknown-operator form with

```
a1=a2=b, (f1,f2)=(t,t^2),
epsilon=1/n, sigma=2/(na).                                 (3)
```

For n>1 and na>2, both noise levels satisfy the theorem's (0,1) range. To embed in its infinite sequence observation form, fix a_j=2^{−(j−2)}, f_j=0 for j>=3 and generate the parameter-free Gaussian tail. One compatible choice is lambda1=lambda2=1, lambda_j=a_jš thereafter, gamma1=gamma2=1, gamma_j=jš thereafter, omega_j=1, radius r=1 and d=4. These weights satisfy the displayed monotonicity requirements; the tied operators are injective and compact, and the curved targets lie in the ellipsoid. This establishes a subexperiment embedding. It does not check every theorem regularity condition, identify its full class with this subclass, or transfer its minimax conclusions.

For two curve points, ordinary coordinate squared loss satisfies

```
(t-t')^2 <= ||(t,t^2)-(t',t'^2)||^2 <= 2(t-t')^2.            (4)
```

Thus loss mismatch alone is not the obstruction here. The restrictions on the parameter class matter.

3. Removing the curve produces a slower experiment

Retain the same tied b and noises but replace (t,tš) by independent u,v in [-1/2,1/2]. Observe

```
Z=(bu,bv,sqrt(a)b)+N(0,I3/n),
```

and estimate u under squared loss. Let R_rect(n,a) be its minimax risk. For constants independent of n and a in [0,1],

```
R_rect(n,a) asymp 1/n+1/(1+na).                              (5)
```

Here is a direct proof for this finite diagnostic, not a quoted theorem from Johannes–Schwarz.

For the upper bound, if na<=1 use zero. Otherwise clip C=Z3/sqrt(a) to [1/2,1] to obtain b_hat, and clip Z1/b_hat to [-1/2,1/2]. Projection gives E(b_hat−b)š<=1/(na), and bounded inverse denominators give

```
E(u_hat-u)^2 <= 8/n+2/(na).
```

The zero branch and 1/(na)<=2/(1+na) for na>1 prove the upper order. No independence between estimated numerator and denominator is required for this inequality (the noises happen to be independent).

For the lower bound, take b0=3/4, u0=1/4, v0=0 and

```
h=1/[8 sqrt(1+na)], b1=b0+h,
u1=b0*u0/b1, v1=0.
```

Both points belong to the rectangle. Their first two mean coordinates agree, so the complete Gaussian-law divergence is

```
KL(P0,P1)=na*h^2/2 <= 1/128.
```

Their target separation satisfies

```
(u0-u1)^2 = (u0*h/b1)^2 >= 1/[784(1+na)].                    (6)
```

Pinsker gives TV<=1/16. Nearest-target decoding yields minimax risk at least (15/128) times the squared separation. Since a<=1 and n>=1, 1/n<=2/(1+na), this also lower-bounds a constant times the right side of (5). This completes both bounds, including a=0; no finite-grid statement is being substituted for the proof.

In contrast, the curved diagnostic in L1 has risk

```
R_curve(n,a) asymp [n(1+na)]^(-1/2).                         (7)
```

For a=nˆ(−1/2), (5) has order nˆ(−1/2), while (7) has order nˆ(−3/4). Their ratio diverges. At a=0 the rectangle is nonidentified in u, whereas the curved experiment's risk decreases at order nˆ(−1/2). The rectangle is already much smaller than the unrestricted infinite product class; this comparison is not an evaluation of that full class's exact minimax rate.

4. The familiar lower-pair cancellation leaves the curve

Starting on the curve at (t,tš), let b1=b0+h and q=b0/b1. Preserving both main images would require the target to become (qt,qtš). But

```
qt^2-(qt)^2 = q(1-q)t^2,                                    (8)
```

which is positive for $h > 0$ and $t \neq 0$. The curve-consistent alternative is (qt,(qt)š); its second mean changes by

```
b1(qt)^2-b0t^2 = -b0*h*t^2/b1.                              (9)
```

The squared discrepancy is order hšt^4, exactly the residual that the curved lower proof charges. Using t=0 to stay on the curve under image cancellation removes the target separation. More generally, bt and btš jointly determine t whenever $bt \neq 0$; when $bt = 0$, positivity of b forces t=0. Therefore exact main-image equality within the curve never produces distinct targets.

Johannes–Schwarz's Step B is a precedent for the joint-perturbation method. It is not itself this tied two-coordinate pair. Equations (8)–(9) analyze the natural attempt to reproduce its image-preserving mechanism here; they do not show that every possible adaptation of its proof fails.

5. What the attempted transfer establishes

An exact observation embedding into a restricted unknown-operator experiment does exist for this Gaussian surrogate at a>0. This embedding does not imply equivalence with the pooled experiment. Applying a larger product-class upper bound loses the curved rate, and its lower bound needs a new admissibility check. The restricted class requires its own analysis. This is a concrete obstruction to direct theorem substitution; it does not establish a literature-priority claim.

The Gaussian surrogate still does not generate the pooled observations. The pooled result additionally requires a complete pooled-law lower bound, charged-fold moment estimates, control of random propensity-variance weights, and a transfer from scalar control to function risk. The formal substitution $n = G$ and $a = m/G + \sigma$ establishes none of these steps.

L3. Deterministic transfer to function risk

Lemma L.3 (deterministic function-risk transfer). The lemma isolates the Hilbert-space step used in Appendix B, P5 and P9. Let $\mathcal{H}$ be a real Hilbert space, $0 < b_{\min} \le b, \widehat{b} \le 1$, and $F, f \in \mathcal{H}$. Write $E = \|F - bf\|^2$ and $\nu = \|F\|^2$. For $\nu > 0$ and a normalizer v with $\nu/2 \le v \le 2\nu$, define $\phi = F/\sqrt{v}$, $t = \langle f, \phi\rangle$, $d_1 = \widehat{bt} - bt$, and $E_t = (\widehat{t} - t)^2$. Then

$$\nu \le 4t^2 + 2E, \qquad \nu(\widehat{b} - b)^2 \le 2E + 8d_1^2 + 8E_t, \tag{10}$$

$$\|F/\widehat{b} - f\|^2 \le \frac{2E}{b_{\min}^2} + \frac{2\nu(\widehat{b} - b)^2}{b_{\min}^4}. \tag{11}$$

To prove (10), put $\delta = \langle F - bf, \phi\rangle$. The identity $bt = \nu/\sqrt{v} - \delta$ and $|\delta| \le \sqrt{E}\sqrt{\nu/v}$ give $\sqrt{\nu} \le \sqrt{v/\nu}\, b|t| + \sqrt{E} \le \sqrt{2}|t| + \sqrt{E}$. Square this inequality, multiply by $(\widehat{b} - b)^2 \le 1$, and use $t(\widehat{b} - b) = d_1 - \widehat{b}(\widehat{t} - t)$. Equation (11) follows by expanding $F/\widehat{b} - f = (F - bf)/b + F(1/\widehat{b} - 1/b)$. At $F = 0$, the error is $\|f\|^2 = E/b^2$, so no direction is needed.

Let Π be an orthogonal projection onto the common fit space. If $f = \Pi(\tau + M_0)$, $M = \Pi M_0$, and $F, \widehat{M}$ lie in this space, then $f - M = \Pi\tau$, and orthogonality gives

$$\|F/\widehat{b} - \widehat{M} - \tau\|^2 \le \|\Pi\tau - \tau\|^2 + 2\|F/\widehat{b} - f\|^2 + 2\|\widehat{M} - M\|^2. \tag{12}$$

Pointwise clipping to an interval containing τ contracts the loss. The projection must be shared: for a two-atom equally weighted law, $\tau = 0$, $M_0 = (1, -1)$, Π the identity, and Π' the projection onto constants, $\|\Pi(\tau + M_0) - \Pi' M_0\|^2 = 1$, while using Π for both components cancels M_0 exactly.

These are deterministic inequalities. The pooled-data proof must still establish the observable moment identities, histogram error bounds, and normalization tails; none follows from the Hilbert-space algebra alone.