SELECTIVE PREFERENCE FEEDBACK CHANGES THE INFORMATION ORDER

Anonymous authors
Paper under double-blind review

ABSTRACT

Can deleting a tie category make preference learning impossible? Under selective response, yes. We study fixed-confidence identification from complete response/nonresponse laws. Binary one-protocol feedback can be nonidentifying; retaining a scored tie yields quartic information and a $\Delta^{-4}\log(1/\delta)$ lower bound, while valid protocols yield quadratic information and a $\Delta^{-2}\log(1/\delta)$ lower bound. After profiling an analytic nuisance manifold, the first normal likelihood jet not reproduced by nuisance or model uncertainty determines the profiled KL order. In a degree-r scored model, we derive a sharp support threshold, coefficient, and finite-gap alias radius at scale Δ^{r+1}. Under imperfect exclusion, arbitrary compact joint validity uncertainty yields a product-space coefficient; when the joint set is convex, maximin protocol design becomes a weighted Chebyshev-center problem, while rectangular budgets alone admit a closed-form pair design. A compact deployment-reference family gives positive, negative, and ambiguous decisions, with separate radii for losing a unique sign and attaining an opposite sign. An exchange-symmetric two-protocol product class exhibits a robust quadratic–quartic crossover; correlated response errors with equal marginals can differ. All sample-complexity statements are lower bounds; we do not claim an attaining same-sample algorithm.

1 INTRODUCTION

Preference data are often observed through a selective reporting channel. A user may answer only when a comparison feels sufficiently clear; a platform may retain a tie category, collapse it into a binary label, or change the prompting protocol. These choices do more than change the number of observations. They can change which likelihood derivative remains after the response mechanism is profiled out, and hence the polynomial order of the information available for the decision. Figure 1 summarizes the three regimes developed in the paper.

The results have three layers. *First*, we give a coordinate-free normal-contact theorem for analytic finite-alphabet experiments with uncertainty on both the true and alternative models. It identifies the profiled KL coefficient as a squared distance in the normal bundle and records the additional localization needed to pass from a local nuisance chart to the global wrong-decision problem. *Second*, we solve a scored exponential-tilt model exactly. The number of distinct positive-mass scores, relative to the response-nuisance degree, determines whether the decision is unidentifiable or only higher-order identifiable; the same calculation gives the leading constant, conditioning laws, and the exact finite-gap robustness radius. *Third*, we study response protocols with imperfect validity. Arbitrary compact joint validity errors yield a product-space coefficient, and convex sets turn maximin protocol design into a weighted Chebyshev-center problem; correlated one-dimensional mechanisms can require three protocols, while rectangular budgets admit a closed-form pair design. Because exclusion violations make the latent target protocol dependent, we also allow a compact family of deployment references fixed independently of the sampling design. The resulting positive, negative, and ambiguous decision regions have exact robustness radii, and profiling the reference contracts the usable protocol margin. A separate symmetric two-protocol product class then yields a quadratic–quartic crossover. This closed form depends on the coupling: correlated response errors with the same protocolwise marginals need not have the same coefficient or optimal allocation.

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

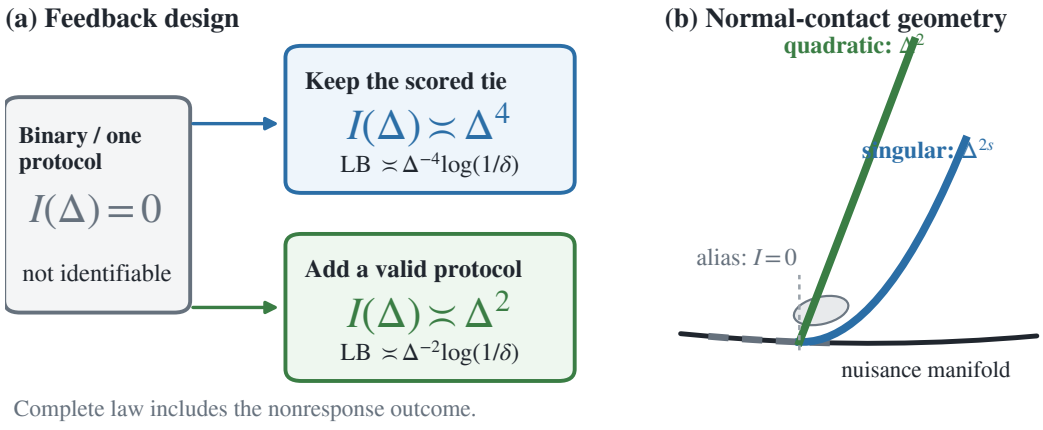


Figure 1: Selective response creates three information orders. Insufficient score resolution permits an exact opposite-decision alias; retaining a scored tie leaves a quartic signal; and a sufficiently valid protocol creates a quadratic transverse signal. All sample scales are fixed-confidence lower bounds.

Our claims are narrower than several nearby literatures. First-order robust testing already projects onto differences of tangent sets; MNAR instruments and sensitivity analysis already address imperfect exclusion; ordinal feedback and ternary preferences already have identification and learning theory; and semiparametric bandit design already provides nuisance-aware lower and upper guarantees (Rieder, 2014; Sun et al., 2018; Chen et al., 2025; Chidambaram et al., 2026; Kim, 2026). The minimax exchange, active-set argument, and sparse-support reduction in our protocol result use classical convex optimal-design tools related to Elfving geometry and conic formulations (Elfving, 1952; Sagnol, 2010). The contribution here is the higher-order, decision-specific geometry of the statistically induced joint validity set in the complete response/nonresponse experiment: an exact support–nuisance boundary, a finite-gap alias topology, correlated-validity design, and a dependence-aware crossover. We do not claim novelty for Blackwell monotonicity, approximate exclusion as a sensitivity device, or convex design machinery by itself. Appendix A gives a source-by-source comparison and makes no claim about unpublished or concurrent work.

2 Complete Experiment and Robust Normal Contact

A model m and a controlled experiment $e \in \mathcal{A}$ determine a positive law P_m^e on a finite alphabet $\mathcal{O}_e$ that includes the observed nonresponse outcome $\perp$. A sequential procedure chooses experiments predictably, stops at τ, and returns $\widehat{d}$. On a compact decision class $\mathcal{M}$, we call it δ-correct if

$$\mathbb{P}_m(\tau < \infty) = 1, \qquad \mathbb{P}_m(\widehat{d} = d(m)) \geq 1 - \delta \quad \forall m \in \mathcal{M}_{\mathrm{uni}}. \tag{1}$$

For strict wrong-decision alternatives $\mathrm{Alt}^\circ_{\mathcal{M}}(m)$, the class-relative information is

$$\Gamma_{\mathcal{M}}(m) = \sup_{w \in \Delta(\mathcal{A})} \inf_{m' \in \mathrm{Alt}^\circ_{\mathcal{M}}(m)} \sum_{e \in \mathcal{A}} w_e D_{\mathrm{KL}}(P_m^e \| P_{m'}^e). \tag{2}$$

The standard stopped-history change of measure gives $\mathbb{E}_m[\tau] \geq \mathrm{kl}(1-\delta, \delta)/\Gamma_{\mathcal{M}}(m)$. When the model class varies with a local parameter, the correctness statement, action set, alternatives, and denominator are indexed by the same row.

We next isolate the local geometry that determines $\Gamma_{\mathcal{M}}(m)$ near a decision boundary. Let $F(t) = S(P_t) = (\sqrt{P_t(o)})_o$ be an analytic signal curve in the positive probability sphere, and let G be a regular analytic chart for the nuisance-observation manifold. Denote by $x_t = G(\xi_H(t))$ the local Hellinger projection and by $R_t = F(t) - x_t$ its residual. If N_t is the normal space to the nuisance manifold inside the probability sphere and $U_t : \mathsf{N}_0 \to \mathsf{N}_t$ is an analytic orthogonal trivialization, assume

$$U_t^{-1} \Pi_{\mathsf{N}_t} R_t = t^s r_s + O(t^{s+1}), \qquad r_s \neq 0. \tag{3}$$

True-side and alternative-side uncertainty are represented by nonempty compact normal cross-sections $\mathcal{B}_{\mathrm{T}}, \mathcal{B}_{\mathrm{A}} \subset \mathsf{N}_0$ and positive laws whose square-root graphs satisfy

$$F(t) + t^s U_t a + O(t^{s+1}), \qquad G(\xi) + t^s U_\xi b + O(t^{s+1}). \tag{4}$$

All remainders are uniform on fixed compact neighborhoods. Complete probabilities are bounded away from zero, the nuisance image is a regular embedded germ, and the base observation fiber is separated from remote nuisance branches. For the one-sided decision profile, the wrong-decision alternatives are taken from the same compact class used in equation 1.

Theorem 1 (A: Robust normal contact)**.** Under the preceding conditions,

$$\mathcal{I}_{\mathrm{loc}}(t) = 2t^{2s}\,\mathrm{dist}^2(r_s, \mathcal{B}_{\mathrm{A}} - \mathcal{B}_{\mathrm{T}}) + O(t^{2s+1}). \tag{5}$$

If the normalized uncertainty sets converge only qualitatively, the same coefficient holds with an $o(t^{2s})$ remainder. An isolated outer-fiber gap gives the same expansion for the unrestricted profile $\mathcal{I}_{\mathrm{glob}}(t)$, and uniform localization at the decision boundary transfers it to the one-sided profile $\mathcal{I}_{\mathrm{dec}}(\Delta)$. If every normal derivative vanishes, analyticity implies exact local observational equivalence. Conversely, a zero leading distance means only that the leading normal jet is cancelled; it does not by itself imply a finite-gap alias.

The proof has three steps. Analytic tubular projection identifies the intrinsic normal jet and is invariant to the nuisance chart and normal frame. Graphical convergence then reduces the robust local problem to the distance from r_s to $\mathcal{B}_{\mathrm{A}} - \mathcal{B}_{\mathrm{T}}$. Finally, every near-minimizer is shown to lie within $O(t^s)$ of the local projection; uniform KL–square-root equivalence supplies the factor two, while the outer-fiber gap rules out a distant branch. Appendices D and C.2 contain the decision-transfer and strict-versus-closed topology arguments.

3 Exact scored support–nuisance phase diagram

Let $y_1 < \cdots < y_J$ be distinct scores with positive weights ω_j, and consider the exponential-tilt family

$$p_u(j) = \omega_j \exp\{uy_j - \psi(u)\}, \qquad \psi(u) = \log \sum_{k=1}^{J} \omega_k e^{uy_k}. \tag{6}$$

A degree-r response-logit nuisance has evaluation vector $q \in V_r := \mathrm{span}\{\mathbf{1}, \mathbf{y}, \ldots, \mathbf{y}^{[r]}\}$. At a category-independent true response intercept a,

$$P_{u,q}(j) = p_u(j)\sigma(a + q_j), \qquad P_{u,q}(\perp) = 1 - \sum_j P_{u,q}(j). \tag{7}$$

The true model is $P_{\Delta,0}$ with $\Delta > 0$, and wrong decisions have $u \le 0$. Let J_{eff} denote the number of distinct positive-mass scores, so duplicate or zero-mass labels do not count as additional resolution. Set $s = r + 1$, $\rho = \sigma(a)$,

$$W = \rho(1-\rho)^2\,\mathrm{diag}(\boldsymbol{\omega}) + \rho^2(1-\rho)\boldsymbol{\omega}\boldsymbol{\omega}^\top, \qquad D_{J,r} = (I - \Pi_r^W)\mathbf{y}^{[s]}, \tag{8}$$

and $C_{J,r} = \frac{1}{2}\beta_s(a)^2\|D_{J,r}\|_W^2$, $A_{J,r} = \sqrt{2C_{J,r}}$, where $\beta_s(a) = G_a^{(s)}(0)/s! > 0$ for $G_a(x) = \mathrm{logit}(\sigma(a)e^x) - a$.

Theorem 2 (B: Exact scored-feedback phase diagram)**.** Fix the compact utility and nuisance class described above.

(i) If $J_{\mathrm{eff}} \le r + 1$, then every sufficiently small $\Delta > 0$ admits a strict wrong-decision model with exactly the same complete law. Hence $I_0(\Delta) = 0$, and no uniformly δ-correct procedure exists for $\delta < 1/2$.

(ii) If $J_{\mathrm{eff}} \ge r + 2$, no positive-gap opposite-decision utility reproduces the complete law. There exist $\Delta_0, K, M > 0$ such that

$$\left|I_0(\Delta) - C_{J,r}\Delta^{2s}\right| \le K\Delta^{2s+1}, \qquad 0 < \Delta \le \Delta_0. \tag{9}$$

Every least-favorable utility satisfies $\Delta - u_\Delta \le M\Delta$ and $-u_\Delta = O(\Delta^2)$; the nuisance minimizer need not be unique. Therefore

$$\mathbb{E}_\Delta[\tau_\delta] \ge \frac{\mathrm{kl}(1-\delta, \delta)}{C_{J,r}\Delta^{2(r+1)} + K\Delta^{2r+3}}, \tag{10}$$

which gives the lower-bound scale $\Omega(\Delta^{-2(r+1)}\log(1/\delta))$.

The strict and closed wrong-decision profiles have the same value. Section 3.1 identifies the exact one-sided finite-gap radius at which that value reaches zero.

The threshold is an interpolation statement. Matching the answered-category probabilities forces the alternative response logit to equal

$$T_{\Delta,u,j} = G_a((\Delta - u)y_j + \psi(u) - \psi(\Delta)) \,. \tag{11}$$

On at most $r+1$ support points, these values can be interpolated by a degree-r polynomial, yielding an exact opposite-decision alias. Above the threshold,

$$T_{\Delta,u} = v_{\Delta,u} + \beta_s(a)(\Delta - u)^s \mathbf{y}^{[s]} + O((\Delta - u)^{s+1}), \qquad v_{\Delta,u} \in V_r, \tag{12}$$

so the first component that cannot be absorbed by the nuisance is $\beta_s(a)(\Delta - u)^s D_{J,r}$. Uniform coercivity of the complete-law Hessian, localization of global minimizers, and profiling over $u \le 0$ then give equation 9.

The coefficient also quantifies how close the experiment is to losing resolution. Up to the fixed response factor,

$$C_{J,r} \;\propto\; \frac{\det H_s}{\det H_{s-1}}, \qquad H_k = [m_{i+j}]_{i,j=0}^{k}, \quad m_\ell = \sum_j \omega_j y_j^\ell. \tag{13}$$

For minimal support, an essential category of mass ε gives $C_{J,r} = \Theta(\varepsilon)$, while an essential score collision of size h gives $C_{J,r} = \Theta(h^2)$. The normalized quantity

$$\chi_{J,r} = \frac{C_{J,r}}{\mathrm{Var}_\omega(Y)^{r+1}} \tag{14}$$

is invariant under positive-affine recoding of the scores with the corresponding reciprocal rescaling of utility. Classical data processing implies that post-processing cannot improve the information order. Merging duplicate scores is lossless, whereas an order-preserving deterministic coarsening to at most $r+1$ non-null contiguous score bins creates an exact local wrong-decision alias. Appendix C.2 gives the interpolation and Rolle arguments, the global identification proof, and the determinant calculation.

3.1 One-sided finite-gap robustness

At the matched law, the square-root derivative with respect to the response logit satisfies $L^\top L = W/4$. Thus a perturbation with $\|g\|_W = \epsilon$ changes the complete law by local KL approximately $\epsilon^2/2$; the radius ϵ is an information-scale misspecification budget. The two-sided local theorem immediately gives, for W-ball radii $\epsilon_{\mathrm{T}}, \epsilon_{\mathrm{A}} = O(\Delta^s)$,

$$I_{\epsilon_{\mathrm{T}},\epsilon_{\mathrm{A}}}(\Delta) = \frac{1}{2}\left(A_{J,r}\Delta^s - \epsilon_{\mathrm{T}} - \epsilon_{\mathrm{A}}\right)_+^2 + O(\Delta^{2s+1}). \tag{15}$$

The exact finite-gap result below is one-sided: the true law remains nominal, and only the wrong-decision nuisance class is enlarged. Let $\mathcal{H} \subset \mathbb{R}^{r+1}$ be compact with the true coefficient in its interior, and define

$$\mathcal{Q}_r^W(\epsilon) = \{V_r\eta + g : \eta \in \mathcal{H},\ \|g\|_W \le \epsilon\}, \qquad \mathcal{K}_\epsilon = [-u_0, u_0] \times \mathcal{Q}_r^W(\epsilon). \tag{16}$$

Writing $\mathcal{Q}_r^0 = V_r\mathcal{H}$, the smallest alternative-side perturbation that can reproduce the true complete law is

$$\epsilon_{\mathrm{alias}}(\Delta) = \min_{-u_0 \le u \le 0} \mathrm{dist}_W(T_{\Delta,u}, \mathcal{Q}_r^0). \tag{17}$$

Local interiority ensures that, for the relevant minimizers, the compact class agrees with the full affine tube. Uniformly over the local parameter rectangle,

$$A_{J,r}\Delta^s - K\Delta^{s+1} \le \epsilon_{\mathrm{alias}}(\Delta) \le A_{J,r}\Delta^s + K\Delta^{s+1}, \tag{18}$$

$$\left| I_\epsilon(\Delta) - \frac{1}{2}\big(\epsilon_{\mathrm{alias}}(\Delta) - \epsilon\big)_+^2 \right| \le K\Delta^{2s+1}. \tag{19}$$

These formulas distinguish three finite-gap cases. Below $\epsilon_{\mathrm{alias}}(\Delta)$ the profile is positive. At equality, the closed profile is zero, but the strict infimum need not be attained if every radius minimizer lies on the decision boundary. Above the radius, a strict opposite-decision alias exists. In particular, a boundary-only zero still precludes uniform error below $1/2$ by continuity of finite adaptive transcript laws.

For fixed Δ, $I_\epsilon(\Delta)$ is nonincreasing and Lipschitz in ϵ. Under any fixed norm equivalent to $\|\cdot\|_W$, the critical radius remains of order Δ^{r+1}, so the exponent is not an artifact of the information-aligned ball. Every almost-surely terminating rule that is correct on the same class $\mathcal{K}_\epsilon$ satisfies

$$\mathbb{E}_{\Delta,\epsilon}[\tau_\delta] \geq \frac{\mathrm{kl}(1-\delta,\delta)}{I_\epsilon(\Delta)}; \tag{20}$$

when $I_\epsilon(\Delta) = 0$, the conclusion is impossibility rather than a formal infinite lower bound. Appendix E proves the compact-tube identity, the uniform KL-to-distance approximation, the nonaligned-norm bounds, and the endpoint transcript-continuity lemma.

4 Protocol restoration under imperfect validity

Consider a platform that randomizes, before the latent judgment, between an optional comparison and an otherwise identical explicitly prompted comparison. The randomization controls assignment, but it does not establish exclusion. We therefore allow protocol-specific effects on both the latent judgment and the response mechanism, and use a separate homogeneous pilot only to calibrate baseline response logits.

For protocol z, set

$$\rho_z = \sigma(a_z), \quad b_z = (1-\rho_z)^{-1}, \quad \lambda_z = \rho_z(1-\rho_z)^2, \quad \alpha_z = \sqrt{\nu\lambda_z}, \tag{21}$$

where $\nu = \sum_j \omega_j (y_j - \bar{y})^2$. Under exact validity, $I_{\mathrm{eff}}(w;a) = \nu \min_c \sum_z w_z \lambda_z (b_z - c)^2$.

Scientific target. Under imperfect exclusion the primitive latent object is the vector V of protocol-potential utilities. A scientific deployment reference $\bar{w}$ is fixed before data collection and independently of the sampling design; the design cannot choose the estimand. Write $V = u_{\bar{w}}\mathbf{1} + \xi_{\bar{w}}$, where $u_{\bar{w}} = \bar{w}^\top V$ and $\bar{w}^\top \xi_{\bar{w}} = 0$. Pairwise contrasts are reference invariant, whereas $u_{\bar{w}}$ changes according to equation 146. For a compact family $\mathcal{W}_0$ of acceptable deployment references, define

$$D_{\mathcal{W}_0}(V) = \begin{cases} +, & \min_{\omega\in\mathcal{W}_0} \omega^\top V > 0, \\ -, & \max_{\omega\in\mathcal{W}_0} \omega^\top V < 0, \\ ?, & \text{otherwise.} \end{cases} \tag{22}$$

The third case is an ambiguous scientific target, not statistical error. Without a fixed reference or anchor protocol, the transformation $(u,\xi) \mapsto (u + c, \xi - c\mathbf{1})$ leaves the complete experiment unchanged: there is no protocol-free scalar decision.

General joint coefficient. After converting direct-effect and response uncertainty to first-order target slopes, let $\mathcal{U} \subset \mathbb{R}^L / \operatorname{span}\{\mathbf{1}\}$ be the actual compact joint uncertainty set. For a design w, define the *fixed-design, fixed-boundary coefficient*

$$\Phi_{\mathcal{U}}(w) = \inf_{\substack{c\in\mathbb{R} \\ v\in\mathcal{U}}} \sum_z w_z \alpha_z^2 (c - b_z - v_z)^2. \tag{23}$$

The same jointly feasible v must be used for every protocol. If $\mathcal{U}$ is convex, the *optimized fixed-boundary coefficient* $\Psi_{\mathcal{U}} := \max_w \Phi_{\mathcal{U}}(w)$ obeys

$$\Psi_{\mathcal{U}} = \min_{c,\, v\in\mathcal{U}} \max_z \alpha_z^2 (c - b_z - v_z)^2, \tag{24}$$

with c in the compact slope range: robust design is a weighted Chebyshev-center problem. The minimax and active-set machinery is classical optimal-design geometry; the new object is the validity set induced by the selective-response experiment. An optimal design uses at most $\dim(\operatorname{aff}\mathcal{U}) + 2$ active protocols; a rank-one correlated class can require three protocols, so pair support is not generic.

Theorem 3 (C: Protocol restoration under joint validity uncertainty and crossover)**.** The following statements hold on the displayed compact row-wise classes.

Joint uncertainty and reference-robust decisions. If the normalized joint target-slope sets converge to $\mathcal{U}$ at Hausdorff rate $O(\Delta)$, then, uniformly in w,

$$K_w^{\mathrm{joint}}(\Delta, 0) = \frac{1}{2}\Delta^2 \Phi_{\mathcal{U}}(w) + O(\Delta^3). \tag{25}$$

Here K_w^{joint} is a fixed-boundary profile. Qualitative convergence gives $o(\Delta^2)$. Suppose additionally that this expansion is uniform for localized wrong utilities, the squared boundary normal distance changes by $O(|u|)$, strict alternatives are dense in the closed boundary slice, and remote strict fibers remain separated. If $\Psi_{\mathcal{U}} > 0$, then every least-favorable wrong utility on coefficient-optimal designs satisfies $-u_{\Delta,w} = O(\Delta^2)$, and the *strict wrong-decision profile* obeys

$$I_w^{\mathrm{joint}}(\Delta) = \frac{1}{2}\Phi_{\mathcal{U}}(w)\Delta^2 + O(\Delta^3), \tag{26}$$

$$I_{\mathrm{joint}}^{\star}(\Delta) := \max_w I_w^{\mathrm{joint}}(\Delta) = \frac{1}{2}\Psi_{\mathcal{U}}\Delta^2 + O(\Delta^3). \tag{27}$$

A zero tangent coefficient alone does not imply a finite-gap alias: the exact joint alias radius separates positive information, a boundary-only zero, and an attained strict alias. The same joint class is used in correctness, profiling, and every denominator. For a robustly positive V, the exact distance to loss of a unique reference-family sign is

$$\rho_{\mathrm{amb}}^{+}(V) = \min_{\omega \in \mathcal{W}_0} \frac{\omega^\top V}{\|\omega\|_*}. \tag{28}$$

The strict opposite-decision distance is a joint cone distance and may be larger. For a deployment-reference family, a binary fixed-confidence lower bound uses the strict-opposite coefficient d_{opp}, not the ambiguity coefficient d_{amb}. Reference profiling contracts the regular protocol margin; in the rectangular two-protocol model,

$$\Theta_{ij}^{\mathrm{ref}} = \frac{\alpha_i \alpha_j}{\alpha_i + \alpha_j}[\underline{m}|b_i - b_j| - r_i - r_j]_+, \tag{29}$$

where $\underline{m} = \min_{\omega \in \mathcal{W}_0} \omega^\top g$.

Rectangular corollary. If validity budgets are protocol-separable, put $r_z = \chi_z + \tau_z/\alpha_z$, $\chi_z = b_z \kappa_z$, and $\mathcal{J}_z = [b_z - r_z, b_z + r_z]$. Then

$$\Phi_{\mathcal{U}_{\mathrm{rect}}}(w) = \min_c \sum_z w_z \alpha_z^2 \operatorname{dist}^2(c, \mathcal{J}_z). \tag{30}$$

Only in this rectangular case is the full-library optimum achieved by a pair. Its square-root allocation is $p_{ij}^{\star} = \alpha_j/(\alpha_i + \alpha_j)$ and its usable margin is

$$\Theta_{ij}^{\mathrm{excl}} = \left[\frac{\alpha_i \alpha_j}{\alpha_i + \alpha_j}|b_i - b_j| - \frac{\alpha_i \alpha_j}{\alpha_i + \alpha_j}(\chi_i + \chi_j) - \frac{\alpha_j \tau_i + \alpha_i \tau_j}{\alpha_i + \alpha_j}\right]_+. \tag{31}$$

In words, the usable margin is calibrated separation minus direct-effect uncertainty minus response-sharing uncertainty. Thus $\max_w \Phi_{\mathcal{U}_{\mathrm{rect}}}(w) = (\max_{i<j} \Theta_{ij}^{\mathrm{excl}})^2$. The corresponding finite-gap radius has the same strict/closed endpoint classification as the general joint radius.

Calibration and critical crossover. The pilot $S_{z,i} \overset{\mathrm{iid}}{\sim} \mathrm{Bernoulli}(\sigma(a_z))$ produces a simultaneous intercept rectangle with probability at least $1 - \alpha$; a conditionally δ-correct rule then has total error at most $\alpha + \delta$. For the symmetric two-protocol class with $a \pm \zeta$, $r = 1$, $J_{\mathrm{eff}} \geq 3$, an antisymmetric direct contrast $|\theta| \leq \eta = O(\Delta^2)$, and *independent* protocol-specific response balls $\|h_\pm\|_{W_\pm} \leq \gamma = O(\Delta^2)$,

$$I_{\mathrm{rc}}^{\diamond}(\Delta, \zeta, \eta, \gamma) = \frac{1}{2}[\Gamma_\Delta(\zeta, \eta) - \gamma]_+^2 + O(\Delta^5), \qquad \diamond \in \{\mathrm{bal}, \mathrm{sel}\}, \tag{32}$$

uniformly for $|\zeta| = O(\Delta)$. The closed form uses this exchange-symmetric product class. Correlated response errors with the same marginal balls generally have a different coefficient and need not have $w = 1/2$ as the leading selectable optimizer.

Protocol-validity hierarchy. Compact joint $\mathcal{U}$: $\Phi_{\mathcal{U}}(w)$ is the fixed-design, fixed-boundary coefficient. Convex joint $\mathcal{U}$: $\Psi_{\mathcal{U}}$ is the Chebyshev maximin value. Rectangular $\mathcal{U}$: the interval formula and pair optimum hold. Reference family $\mathcal{W}_0$: d_{amb} and d_{opp} are distinct. Symmetric product class: Γ_Δ gives the explicit crossover.

Appendices F, G, H, and I give the proofs and counterexamples.

If h_n is the largest pilot logit half-width, then on simultaneous coverage clipping is eventually inactive and $h_n \asymp \sqrt{\log(2L/\alpha_n)/n_{\min,n}}$, up to constants determined by the compact response range. Hence

$$\frac{n_{\min,n}\zeta_n^2}{\log(2L/\alpha_n)} \longrightarrow \infty \tag{33}$$

is equivalent to these particular intervals satisfying $h_n = o(|\zeta_n|)$. This is not a general minimax necessity statement: the pilot identifies a_z, not the validity radii; those are external.

4.1 Product-structured critical window

The critical law is a separate two-protocol specialization of Theorem C. It uses an antisymmetric direct-effect contrast and independent response-error balls across the two protocols; it is not a corollary of the full-library rectangular pair theorem, nor a formula for arbitrary correlated validity sets.

Let $x_j = y_j - \bar{y}$, $\nu = \sum_j \omega_j x_j^2$, $\rho = \sigma(a)$, and $\lambda_0 = \rho(1-\rho)^2$. Define

$$A = \sqrt{\lambda_0 \nu}, \qquad N = \sqrt{2C_{\mathrm{sing}}}, \qquad C_{\mathrm{sing}} = C_{J,1}. \tag{34}$$

For $b_\pm = (1 - \sigma(a \pm \zeta))^{-1}$, write $\delta b = (b_+ - b_-)/2$, $\bar{b} = (b_+ + b_-)/2$, and

$$m_\Delta = [\Delta|\delta b| - \bar{b}\eta]_+, \qquad \Gamma_\Delta = \sqrt{A^2 m_\Delta^2 + N^2\Delta^4}. \tag{35}$$

The first term is the protocol contrast surviving direct-effect uncertainty; the second is the quartic curvature present even with one protocol.

There exist $\Delta_0, c_0, K > 0$ such that, uniformly for $0 < \Delta \le \Delta_0$, $|\zeta| \le c_0\Delta$, and $0 \le \eta, \gamma \le c_0\Delta^2$,

$$\left| I_{\mathrm{rc}}^{\diamond}(\Delta, \zeta, \eta, \gamma) - \frac{1}{2}[\Gamma_\Delta - \gamma]_+^2 \right| \le K\Delta^5, \qquad \diamond \in \{\mathrm{bal}, \mathrm{sel}\}. \tag{36}$$

The balanced compound and selectable experiments are different statistical experiments, although exchange symmetry and the product uncertainty class make their leading critical coefficients agree. The exact product-class response radius satisfies $\gamma_{\mathrm{alias}}^{\mathrm{rc}}(\Delta, \zeta, \eta) = \Gamma_\Delta + O(\Delta^3)$, with strict attainment determined by whether a radius minimizer has negative deployment-reference utility.

If

$$\zeta_\Delta = \vartheta\Delta + O(\Delta^2), \qquad \eta_\Delta = \kappa\Delta^2 + O(\Delta^3), \qquad \gamma_\Delta = \tau\Delta^2 + O(\Delta^3),$$

then

$$\frac{I_{\mathrm{rc}}^{\diamond}(\Delta, \zeta_\Delta, \eta_\Delta, \gamma_\Delta)}{\Delta^4} \longrightarrow \frac{1}{2}\left[\sqrt{A^2(e^a|\vartheta| - b\kappa)_+^2 + N^2} - \tau\right]_+^2, \qquad b = (1-\rho)^{-1}. \tag{37}$$

At exact validity this becomes $C_\lambda\vartheta^2 + C_{\mathrm{sing}}$, where $C_\lambda = \rho^3\nu/2$. If $\zeta = o(\Delta)$ and both validity budgets are $o(\Delta^2)$, the quartic coefficient survives. Every class-correct rule obeys

$$\mathbb{E}[\tau_n] \ge \frac{\mathrm{kl}(1-\delta_n, \delta_n)}{I_{\mathrm{rc}}^{\diamond}(\Delta_n, \zeta_n, \eta_n, \gamma_n)}. \tag{38}$$

Why the dependence assumption matters. In leading tangent–normal coordinates, write the two protocol residuals as (X, Y) and $(-X, Y)$. Independent response balls of radius γ give

$$I_{\mathrm{prod}}^{\mathrm{lead}} = \frac{1}{2}[\sqrt{X^2+Y^2} - \gamma]_+^2, \qquad I_{\mathrm{shared}}^{\mathrm{lead}} = \frac{1}{2}\{X^2 + [Y-\gamma]_+^2\}, \tag{39}$$

where the second formula corresponds to a perfectly shared response error $h_+ = h_-$. The two joint classes have identical marginal balls but different coefficients. Thus the radial law in equation 36 is a product-class statement; general correlated errors retain only the intrinsic product-space distance characterization of Theorem A.

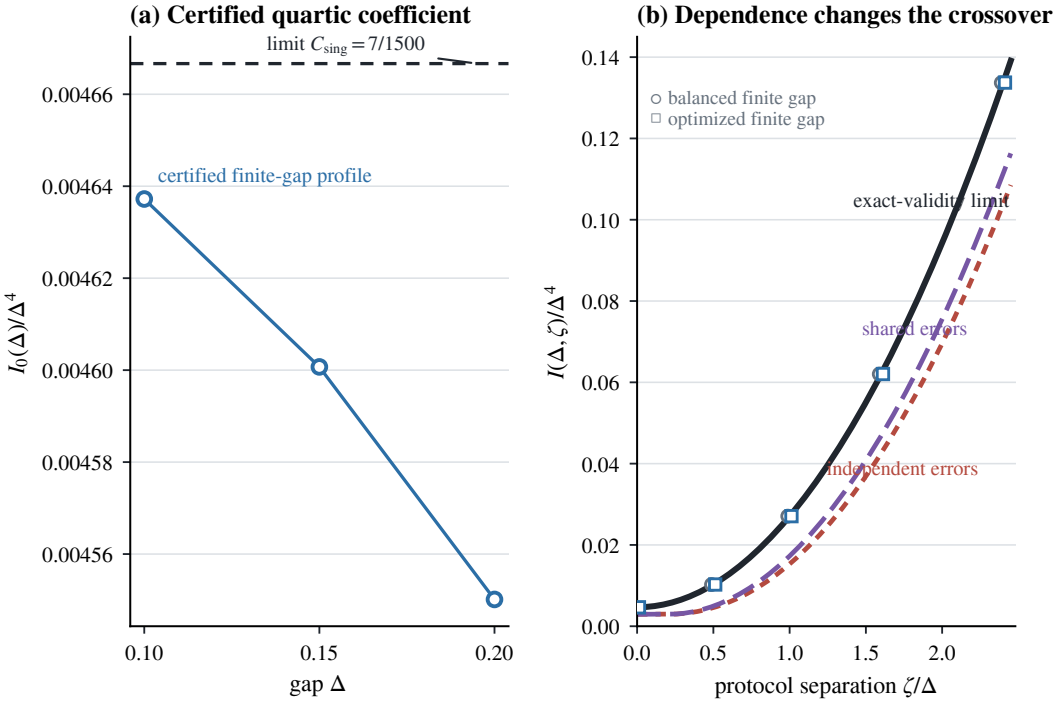


Figure 2: Certified finite-gap behavior. Left: complete-law profiles approach the quartic coefficient. Right: finite-gap crossover values approach the exact-validity limit; independent and shared errors with equal marginal radii give different curves. The analytic theorems establish the general claims.

Proof mechanism. A third-order target expansion in the protocol-specific KL metrics, followed by profiling the common affine nuisance, separates tangent and normal terms. Sampled residuals are $O(\Delta^2)$, so the cubic KL term is lower order; endpoint and boundary-localization arguments handle $w \in \{0, 1\}$ and the wrong utility. Appendix I gives the derivation and coupling counterexample.

5 What changes the information order?

The support threshold decides whether one protocol identifies the sign at all. Conditional on identification, the residual $D_{J,r}$ controls how quickly information vanishes. Protocols act differently: they create a first-order contrast across response mechanisms, provided validity errors do not absorb it.

Experiment	Profiled information	Fixed-confidence implication
$J_{\text{eff}} \le r + 1$, one protocol	exact zero by interpolation	no uniform error below $1/2$ on a class containing the alias
$J_{\text{eff}} \ge r + 2$, one protocol	$C_{J,r}\Delta^{2(r+1)} + O(\Delta^{2r+3})$	every class-correct rule obeys an $\Omega(\Delta^{-2(r+1)} \log(1/\delta))$ lower bound
Fixed deployment-reference sign with $\Psi_{\mathcal{U}} > 0$	$I^{\star}_{\text{joint}}(\Delta) = \frac{1}{2}\Psi_{\mathcal{U}}\Delta^2 + O(\Delta^3)$; rectangular budgets give $\Psi_{\mathcal{U}} = (\Theta^{\text{excl}}_{\star})^2$	every class-correct rule has a quadratic fixed-confidence lower-bound scale; the critical product class crosses to quartic or zero information

For a deployment-reference family, the strict-opposite coefficient is d_{opp}; the ambiguity coefficient d_{amb} can be smaller.

Table 1: Correctness and every denominator use the same row-wise class; all sample scales are lower bounds.

Resolution and conditioning. Crossing $J_{\text{eff}} = r + 1$ changes exact nonidentification into finite-order identification, but it need not give a well-conditioned experiment. A rare essential category or a near score collision can drive $C_{J,r}$ to zero; label count alone is therefore insufficient.

Reference and validity. The target is reference indexed: these are reference-indexed decisions under a fixed deployment reference $\mathcal{W}_0$, not a protocol-free ordering. The radius to lose a sign and the radius to attain the opposite sign can differ, and sampling cannot choose the estimand. For rectangular budgets, equation 31 trades protocol separation against direct-effect and sharing uncertainty. Correlated sets must instead be profiled jointly and can require three protocols. The pilot estimates a_z, not exclusion or the sensitivity radii.

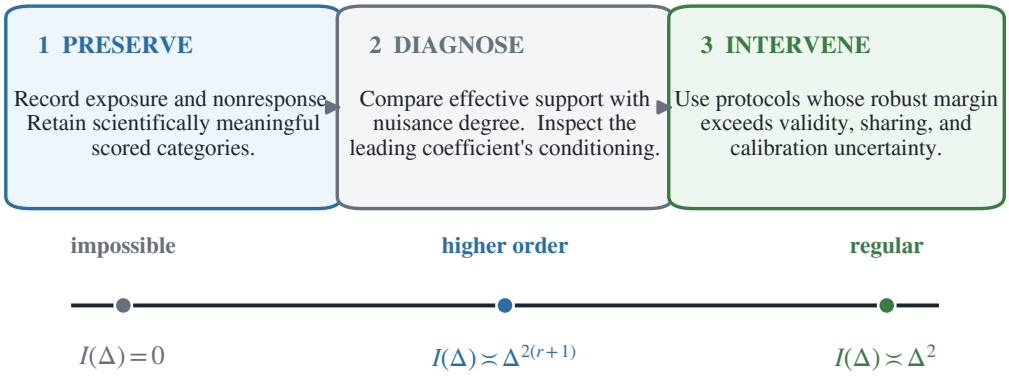
Zero profiles and uniformity. A zero tangent coefficient may leave a quartic profile, an unattained zero strict infimum, or an exact alias. A boundary-only zero still precludes uniform δ-correctness for $\delta < 1/2$. Quantitative Hausdorff convergence gives the displayed remainders; qualitative convergence gives only little-o laws. No attaining procedure is claimed.

5.1 Scope and design implications

The explicit $2(r+1)$ law is proved at a category-independent response baseline; other analytic baselines are classified by the contact theorem. The exact finite-gap scored robustness theorem is one-sided and local; fixed nonvanishing violations are outside scope, and no matching same-sample stopping rule is claimed.

The design rule is: **preserve** exposure, nonresponse, and meaningful scores; **diagnose** effective support, nuisance degree, and coefficient conditioning; and **intervene** only when calibrated protocol separation exceeds validity, sharing, and calibration uncertainty. Compact correlated sets must be profiled jointly; convex sets yield Chebyshev maximin design, while pair formulas require rectangular budgets; identical protocolwise marginal budgets do not determine the correlated crossover. Feedback resolution and protocol validity decide whether learning is ambiguous, nonidentifiable, higher-order, or regular.

From information geometry to feedback design



AI USE STATEMENT

Generative AI tools assisted with literature discovery, conceptual framing, formulation of mathematical claims, proof drafting and checking, code development, figure preparation, and manuscript editing. The authors reviewed all AI-assisted material. Every retained AI-assisted proof cluster, including the joint-validity, scientific-estimand, correlated-design, reference-robust decision, approximate-exclusion, and robust-crossover extensions, was independently reconstructed and line-checked by an author who did not prepare its initial draft. The cited novelty comparisons were checked against the original sources, and the symbolic, interval, derivative, KKT, and clean-build certificates were independently reproduced. The authors take responsibility for the final text, claims, proofs, code, and artifacts.

REPRODUCIBILITY STATEMENT

The main text states the assumptions and finite-gap theorem forms used by every headline claim. The appendices contain the full proof chain, proof-hardening lemmas, row-wise change-of-measure theorem, failure-mode matrix, and certified-validation specification. The anonymous archive includes exact symbolic identities, interval enclosures, four-way gradient checks, optimizer KKT ledgers, standalone reference-robust and correlated-design certificates, deterministic tests, two isolated clean-source reproduction transcripts, and SHA-256 manifests; `verify.sh` rebuilds and checks the complete package. The completed, anonymous, hash-bound proof-reconstruction record is summarized in `HUMAN_VERIFICATION_LEDGER.md`; identifying signed records are retained privately during double-blind review.

REFERENCES

Akhil Agnihotri, Rahul Jain, Deepak Ramachandran, and Zheng Wen. Best policy learning from trajectory preference feedback. In *Proceedings of The 29th International Conference on Artificial Intelligence and Statistics*, volume 300 of *Proceedings of Machine Learning Research*, pp. 2116–2124, 2026.

Viktor Bengs, Björn Haddenhorst, and Eyke Hüllermeier. Identifying copeland winners in dueling bandits with indifferences. In *Proceedings of The 27th International Conference on Artificial Intelligence and Statistics*, volume 238 of *Proceedings of Machine Learning Research*, pp. 226–234, 2024.

David Blackwell. Equivalent comparisons of experiments. *The Annals of Mathematical Statistics*, 24 (2):265–272, 1953.

Jian Chen, Zhehao Li, and Xiaojie Mao. Learning with selectively labeled data from multiple decision-makers. In *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, pp. 8480–8519. PMLR, 2025.

Herman Chernoff. Sequential design of experiments. *The Annals of Mathematical Statistics*, 30(3): 755–770, 1959.

Keertana Chidambaram, Karthik Vinay Seetharaman, and Vasilis Syrgkanis. Direct preference optimization with unobserved preference heterogeneity: The necessity of ternary preferences. In *Proceedings of the 29th International Conference on Artificial Intelligence and Statistics*, volume 300 of *Proceedings of Machine Learning Research*, pp. 4555–4563. PMLR, 2026.

Timothy G. Conley, Christian B. Hansen, and Peter E. Rossi. Plausibly exogenous. *The Review of Economics and Statistics*, 94(1):260–272, 2012. doi: 10.1162/REST_a_00139.

Regina Dittrich, Brian Francis, Reinhold Hatzinger, and Walter Katzenbeisser. Missing observations in paired comparison data. *Statistical Modelling*, 12(2):117–143, 2012. doi: 10.1177/1471082X1001200201.

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

Aurélien Garivier and Emilie Kaufmann. Optimal best arm identification with fixed confidence. In *Proceedings of the 29th Conference on Learning Theory*, volume 49 of *Proceedings of Machine Learning Research*, pp. 998–1027, 2016.

Nhat Ho and XuanLong Nguyen. Singularity structures and impacts on parameter estimation in finite mixtures of distributions. *SIAM Journal on Mathematics of Data Science*, 1(4):730–758, 2019. doi: 10.1137/18M122947X.

Emilie Kaufmann, Olivier Cappé, and Aurélien Garivier. On the complexity of best-arm identification in multi-armed bandit models. *Journal of Machine Learning Research*, 17(1):1–42, 2016.

Seok-Jin Kim. Nearly optimal best arm identification for semiparametric bandits. In *Proceedings of the 29th International Conference on Artificial Intelligence and Statistics*, volume 300 of *Proceedings of Machine Learning Research*, pp. 4402–4410. PMLR, 2026.

Seok-Jin Kim, Gi-Soo Kim, and Min-Hwan Oh. Experimental design for semiparametric bandits. In *Proceedings of the 38th Conference on Learning Theory*, volume 291 of *Proceedings of Machine Learning Research*, pp. 3215–3252. PMLR, 2025.

Yi Lin, Ryan Martin, and Min Yang. On optimal designs for nonregular models. *The Annals of Statistics*, 47(6):3335–3359, 2019. doi: 10.1214/18-AOS1780.

Shang Liu, Yu Pan, Guanting Chen, and Xiaocheng Li. Reward modeling with ordinal feedback: Wisdom of the crowd. In *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, pp. 39190–39218. PMLR, 2025.

Subhojyoti Mukherjee, Anusha Lalitha, Kousha Kalantari, Aniket Deshmukh, Ge Liu, Yifei Ma, and Branislav Kveton. Optimal design for human preference elicitation. In *Advances in Neural Information Processing Systems*, volume 37, 2024.

Helmut Rieder. Neighborhoods as nuisance parameters? robustness versus semiparametrics. *arXiv preprint arXiv:1412.1711*, 2014. URL `https://arxiv.org/abs/1412.1711`.

Guillaume Sagnol. Computing optimal designs of multiresponse experiments reduces to second-order cone programming. *arXiv preprint arXiv:0912.5467*, 2010. URL `https://arxiv.org/abs/0912.5467`.

Andreas Schlaginhaufen, Reda Ouhamma, and Maryam Kamgarpour. Efficient preference-based reinforcement learning: Randomized exploration meets experimental design. In *Advances in Neural Information Processing Systems*, volume 38, 2025. doi: 10.52202/085713-0407.

Rui Song, Michael R. Kosorok, and Jason P. Fine. On asymptotically optimal tests under loss of identifiability in semiparametric models. *The Annals of Statistics*, 37(5A):2409–2444, 2009. doi: 10.1214/08-AOS643.

BaoLuo Sun, Lan Liu, Wang Miao, Kathleen Wirth, James Robins, and Eric J. Tchetgen Tchetgen. Semiparametric estimation with data missing not at random using an instrumental variable. *Statistica Sinica*, 28(4):1965–1983, 2018. doi: 10.5705/ss.202016.0324.

Aad van der Vaart. Higher order tangent spaces and influence functions. *Statistical Science*, 29(4): 679–686, 2014.

A Literature positioning and claim boundary

This appendix records the evidence behind the main-text related-work paragraph. It is a *positioning record*, not a claim that no unpublished or concurrent work exists. The comparison is restricted to the cited public papers and to the exact theorem statements used in this manuscript. The governing rule is conservative: whenever an ingredient already exists, the manuscript says so explicitly and locates novelty only in the theorem-level conjunction that is actually proved. Relative to these sources, the combined theorem package appears distinct; this statement is a bounded comparison, not a priority certificate.

A.1 CLOSEST LINES OF WORK

Table 2: Selected closest public work and theorem-level distinction.

Literature line	Established result in the cited source	Position of this paper
Controlled sequential design (Chernoff, 1959; Kaufmann et al., 2016; Garivier & Kaufmann, 2016)	Chernoff-style max–min KL design, stopped-history change of measure, and asymptotically optimal Track-and-Stop allocation in regular controlled experiments.	We do not claim a new generic characteristic time or change-of-measure inequality. The new calculation is the exact *nuisance-profiled, one-sided decision* information order and coefficient at a singular selective-response boundary.
Nonregular optimal design (Lin et al., 2019)	Hellinger information extends Fisher-information design to certain nonregular models and yields local minimax lower bounds and nonregular design criteria.	We do not claim that Hellinger geometry or nonregular design is new. Theorems 1 and 2 compute a model-specific normal contact order, exact finite-alphabet coefficient, and decision-side localization after nonlinear nuisance profiling.
Robust tangent neighborhoods (Rieder, 2014)	Convex tangent neighborhoods may be treated as nuisance directions; in robust testing, least-favorable first-order directions arise from projection onto a difference of tangent sets.	We do not claim projection onto convex difference sets as a first-order idea. Theorem 1 proves its higher-order analogue for analytic normal jets, two-sided graphical uncertainty, unrestricted-fiber localization, and one-sided decision profiling.
MNAR identification with instruments (Sun et al., 2018)	Under an exclusion instrument affecting nonresponse but not the outcome, the full-data distribution can be identified under stated conditions; semiparametric estimators are developed for mean functionals.	We do not claim the first IV treatment of MNAR data. Our protocol variable is used to change the order of a controlled decision experiment, and the target is a wrong-decision KL profile rather than a population-mean estimator.
Approximate exclusion and plausibly exogenous instruments (Conley et al., 2012)	Sensitivity analysis is developed for instruments with bounded direct effects, relaxing exact exclusion while preserving informative inference in classical IV models.	We do not claim the first relaxation of exclusion. Theorem 3 instead gives the joint product-space coefficient for correlated validity errors. A closed-form optimal protocol pair is derived only for the protocol-separable rectangular subclass, together with a separate finite-gap alias theorem.
Selective labels with multiple decision-makers (Chen et al., 2025)	Heterogeneous decision-makers are used as instrumental variables to establish exact or partial identification of classification risk under selective labels, followed by robust learning procedures.	We do not claim that multiple decision-makers or selection instruments are new. Our theorem uses controlled protocol assignment, a shared response nuisance, and calibrated intercepts, but exact exclusion is the zero-radius case: bounded direct effects are allowed and receive an explicit cancellation threshold and finite-gap radius.

Literature line	Established result in the cited source	Position of this paper
Optimal human preference elicitation (Mukherjee et al., 2024)	Optimal-design policies are developed for eliciting absolute and ranking feedback from lists of items, with efficient algorithms and statistical analysis.	We do not claim the first optimal preference-query design. Our design criterion is the exact nuisance-profiled decision information of the complete response/nonresponse law and may change polynomial order, not only an information constant.
Preference-based RL with experimental design (Schlaginhaufen et al., 2025)	Randomized exploration and optimal experimental design select informative trajectory comparisons and reduce preference-query complexity in regular PbRL models.	We do not claim that experimental design for trajectory preferences is new. The present paper isolates a finite selective-feedback experiment and proves an impossible–singular–regular transition caused by endogenous response nuisance.
Best-policy learning from preferences (Agnihotri et al., 2026)	Best-policy learning from noisy binary trajectory preferences is studied with offline data plus online pure exploration and Bayesian simple-regret guarantees.	We do not claim the first best-policy objective under preference feedback. Our target is the local fixed-confidence information lower bound under outcome-dependent response and nuisance profiling; no attaining sequential upper algorithm is claimed.
Ordinal and tie feedback (Liu et al., 2025; Bengs et al., 2024)	Fine-grained ordinal and tied feedback is modeled for reward learning, and ternary indifference feedback is used for fixed-confidence winner identification in dueling bandits.	We do not claim that ties or ordinal labels are new or always beneficial. Our theorem identifies when numerical score support crosses an exact nuisance-interpolation threshold and computes the surviving orthogonal-polynomial coefficient.
Missing paired comparisons (Dittrich et al., 2012)	Paired-comparison outcomes and missing indicators are modeled jointly under MCAR and MNAR mechanisms through a sample-selection likelihood.	We do not claim the first MNAR paired-comparison model. Our target is the wrong-decision contact order, exact alias threshold, robustness radius, and controlled protocol geometry of the complete experiment.
Ternary identification under latent heterogeneity (Chidambaram et al., 2026)	Binary comparisons can fail to identify latent heterogeneous preferences, while rankings over at least three responses restore identification and support algorithms for heterogeneous preference learning.	Our obstruction is different: an endogenous response nuisance interpolates the complete law after one scored judgment. We derive a support–degree threshold, coefficient, and finite-gap alias radius rather than a latent-mixture ranking theorem.
Semiparametric experimental design and BAI (Kim et al., 2025; Kim, 2026)	Orthogonalized experimental design yields regret, PAC, and BAI guarantees, and later work supplies an attainable instance-dependent lower bound with an efficient near-optimal upper algorithm.	We do not claim nuisance-aware design itself. Here the nuisance enters the observation law nonlinearly and may absorb several likelihood derivatives, changing the information order and identification class.

Literature line	Established result in the cited source	Position of this paper
Loss of identifiability and algebraic singularity (Song et al., 2009; Ho & Nguyen, 2019)	Semiparametric tests under nuisance nonidentifiability and algebraic singularity indices for finite mixtures establish nonstandard testing or estimation behavior when ordinary regularity fails.	We do not claim higher-order nonidentifiability as a generic phenomenon. Our theorem is decision-specific, controlled, finite-alphabet, globally localized, and coupled to exact scored-feedback and protocol-design consequences.
Comparison of experiments (Blackwell, 1953)	Parameter-independent garbling cannot improve an experiment; equality is characterized by sufficiency of the processed observation for the model pair.	Blackwell monotonicity is used, not claimed. The paper derives the model-specific equality condition for duplicate-score merging and the strict collapse that occurs when coarsening crosses the effective-support threshold.
Higher-order tangent methods (van der Vaart, 2014)	Higher-order tangent spaces and influence functions explain why first-order semiparametric scores may be insufficient.	The paper specializes this principle to an analytic finite-alphabet observation manifold, proves an explicit KL coefficient and remainder, and separates local profiling from global wrong-decision localization.

A.2 Claim-language boundaries

The main paper follows the following language contract.

Do not claim	Admissible replacement
"the first use of instruments for MNAR feedback"	"valid calibrated protocols can change the decision-specific contact order under a shared response nuisance"
"the first optimal design for preference queries"	"the exact local design is evaluated after selective-response nuisance profiling and can change information order"
"the first best-policy or best-arm identification from preferences"	"the complete response/nonresponse profile has an exact support–nuisance phase transition"
"the first nuisance-aware BAI method"	"nonlinear response nuisance can absorb multiple likelihood derivatives; the first unabsorbed normal jet determines order"
"the first nonregular, Hellinger, or tangent-neighborhood design"	"higher-order robust normal contact yields a two-sided Minkowski-difference coefficient; the scored model adds an exact one-sided finite-gap radius"
"ties are informative" as a generic novelty claim	"a necessary scored category can move the experiment across an exact interpolation threshold; rare or colliding categories quantify conditioning"
"a new Blackwell theorem"	"classical data processing yields a model-specific equality condition and coarsening-collapse corollary"

Table 3: Claim-language contract used by the static literature audit.

A.3 The theorem-level distinction

Relative to the selected public papers, the submission claims only the following conjunction, with each part backed by a numbered theorem:

1. a general two-sided robust normal-contact theorem for a positive finite-alphabet signal curve, arbitrary nonempty compact normal cross-sections, global fiber localization, and an explicit Minkowski-difference coefficient;

2. an exactly solved scored selective-feedback model in which effective support versus nuisance degree determines exact aliasing or order $2(r+1)$, with an orthogonal-polynomial coefficient and conditioning law;
3. a one-sided norm-thickened nuisance class with an exact alias radius, a positive-part information law, and a strict-versus-closed alias classification;
4. response protocols whose correlated validity errors admit a joint product-space coefficient; whose rectangular subclass admits exact pairwise thresholds and a finite-gap radius; and whose symmetric critical specialization yields a quadratic–quartic crossover and class-indexed lower bound.

No selected source states this four-part theorem package. This is an evidence-based comparison to the cited public literature, not an exhaustive priority proof.

A.4 Main-text placement decision

The main text names the closest established contributions before stating the paper's narrower distinction. It therefore does not rely on the appendix to disclose collisions. The extended matrix above is retained so that every reviewer-facing novelty sentence can be traced to a concrete source and a deliberately narrower claim.

B Claim ledger and proof map

The abstract contains exactly three scientific claims. Table 4 records their formal support and scope.

Abstract claim	Main result	Proof obligations and scope
Two-sided robust normal contact determines profiled KL order	Theorem 1	Analytic tubular projection; upper/lower graphical convergence; Minkowski-difference distance; uniform KL–Hellinger conversion; chart/frame invariance; isolated-fiber localization; exact-zero topology.
Exact and robust support–nuisance phase diagram	Theorem 2 and its Sections 3–3.1 corollaries	Lagrange interpolation; Vandermonde positivity; complete-law metric; global sign identification; uniform finite-gap remainder; Hankel determinant identity; exact tube geometry; strict/closed alias topology.
Protocol restoration under imperfect validity and crossover	Theorem 3	Deployment-reference estimand; reference-robust positive/negative cones; exact ambiguity and opposite-decision radii; sharp support-function sensitivity; reference-profiled normal coefficient; gauge-fixed row-wise class; quantitative versus qualitative graphical convergence; general joint target-slope coefficient and compact-joint fixed-boundary profile; one-sided transfer to strict wrong-decision information (Corollary 51); exact finite-gap radius on the same class; convex-joint Chebyshev duality (Theorem 65); active-set/KKT characterization and the $d_{\mathcal{U}} + 2$ support bound (Theorem 66); exact rank-one and ellipsoidal reductions; rectangular interval/pair corollary; relaxation-price bound; pilot coverage; product-structured robust critical crossover; coupling counterexample showing that identical marginal response balls do not determine the coefficient; class-indexed lower bounds.

Table 4: Each abstract claim maps to one of the three primary main-text theorem clusters with visible assumptions.

The active manuscript makes no theorem-level claim about dynamic control, a tensor design, or an attaining sequential upper algorithm. Under imperfect exclusion, the nominal scalar is a pre-specified deployment-reference utility. The reference-robust extension permits a compact family of acceptable deployment mixtures, returns an ambiguous decision when their signs disagree, and profiles the same family in its normal coefficient. The sampling design does not choose the reference family, and no protocol-free common utility is claimed without an anchor. The lower bound is class relative: correctness and the alternative set use the same compact model class. The exact support exponent is restricted to the category-independent true-response stratum; general baselines are first classified by Theorem 1. The general contact theorem permits two-sided local uncertainty. The exact finite-gap scored thickening theorem is one-sided. The approximate-validity theorem retains an arbitrary compact joint validity set for its fixed-boundary coefficient and exact finite-gap radius. Corollary 51 supplies the additional localization needed before that coefficient is used as strict wrong-decision information. For compact convex joint sets, robust design is the attained weighted Chebyshev problem in Theorem 65; optimal weights are active-set multipliers and may require up to $\dim(\operatorname{aff}\mathcal{U}) + 2$ protocols. The interval and optimal-pair formulas are restricted to the protocol-separable rectangular subclass. Its cubic remainder requires $O(\Delta)$ Hausdorff convergence after normalization; qualitative convergence yields only $o(\Delta^2)$. The robust crossover uses an explicit exchange-symmetric two-protocol product class with an antisymmetric scalar direct effect, independent protocol-specific response perturbations, and order-Δ^2 validity radii. Arbitrary correlated critical uncertainty retains only the intrinsic product-space distance law; the radial formula, equality of balanced/selectable

leading coefficients, and $w = 1/2$ optimizer are not claimed there. Calibration uses an independent pilot and does not identify the validity radii.

C Proofs of the contact, exact phase, and conditioning results

C.1 Proof of Theorem 1

Let

$$L(t,\xi) = \frac{1}{2}\|F(t) - G(\xi)\|_2^2, \qquad \Phi(t,\xi) = \nabla_\xi L(t,\xi) = -DG(\xi)^\top (F(t) - G(\xi)).$$

At the base, $\Phi(0,0) = 0$ and $D_\xi \Phi(0,0) = B^\top B \succ 0$. The analytic implicit-function theorem therefore gives a unique analytic stationary branch $\xi_H(t)$ for small t. The exact Hessian of $L(t,\cdot)$ is

$$D_\xi^2 L(t,\xi) = DG(\xi)^\top DG(\xi) - \sum_o (F_o(t) - G_o(\xi)) D^2 G_o(\xi).$$

Continuity and $B^\top B \succ 0$ make this Hessian uniformly positive definite on a common compact chart, so the stationary branch is the unique chart-global nearest point. Its stationarity condition is

$$DG(\xi_H(t))^\top R(t) = 0. \tag{40}$$

Differentiating at zero and using $R(0) = 0$ gives

$$B^\top (F'(0) - B\xi_H'(0)) = 0, \qquad R'(0) = (I - BB^\dagger) F'(0).$$

More generally R is analytic. If its first nonzero derivative is $R^{(s)}(0)$, then

$$R(t) = \frac{t^s}{s!} R^{(s)}(0) + O(|t|^{s+1}), \qquad \frac{1}{2}\|R(t)\|_2^2 = \frac{\|R^{(s)}(0)\|_2^2}{2(s!)^2}|t|^{2s} + O(|t|^{2s+1}). \tag{41}$$

If all derivatives vanish, the real-analytic identity theorem gives $R(t) \equiv 0$ locally. Since the square-root map is injective on the positive simplex, P_t then equals the projected nuisance law exactly.

For positive probability vectors p, q in a common compact simplex interior, a coordinatewise Taylor expansion gives uniformly

$$D_{\mathrm{KL}}(p\|q) = 4H^2(p,q) + O(\|p-q\|_2^3), \qquad H^2(p,q) = \frac{1}{2}\|S(p) - S(q)\|_2^2. \tag{42}$$

Evaluating KL at the Hellinger projection gives the upper bound in equation 5. Conversely, the upper bound makes any KL minimizer converge to $P_\star$ by Pinsker's inequality. Local injectivity of the visible nuisance chart then places it in the same chart, where its Hellinger distance is at least the nearest-point distance in equation 41; applying equation 42 yields the matching lower coefficient. Finally, the isolated-fiber assumption supplies a positive KL gap outside the chart while the local candidate tends to zero, so every unrestricted minimizer is eventually local. The same argument also shows that every KL minimizer has square-root residual $O(|t|^s)$.

The construction is coordinate invariant. If $\widetilde{G} = G \circ \phi$ is another analytic visible chart with invertible $D\phi(0)$, then $\widetilde{G}$ and G have the same local image, the same nearest physical nuisance law, and the same residual $R(t)$. Equivalently, $\mathrm{range}(D\widetilde{G}(0)) = \mathrm{range}(B)$, so the normal projector $I - BB^\dagger$ is unchanged. On a compact regular stratum on which positivity, the minimum singular value of the visible Jacobian, and the outer-fiber gap are bounded uniformly, every constant in equation 5 is locally uniform in the base point. This proves Theorem 1. □

One-sided transfer. Let $\kappa(v+d, v)$ be the fixed-base nuisance profile when the true signal is $v + d$ and the boundary coordinate is $v \le 0$. Suppose uniformly for $v \uparrow 0$ and $d \downarrow 0$, $\kappa(v+d, v) = c(v)d^{2s} + o(d^{2s})$, with $c(v) \to c(0) > 0$, and suppose $\kappa(0, v) > 0$ for every fixed $v < 0$. For $I(\Delta) = \min_{v \le 0} \kappa(\Delta, v)$, the boundary candidate gives $I(\Delta) \le (c(0) + o(1))\Delta^{2s}$. Compactness forces every minimizer $v_\Delta \to 0$; otherwise a negative limiting v would have positive cost. Put $d_\Delta = \Delta - v_\Delta \ge \Delta$. Then

$$I(\Delta) = (c(v_\Delta) + o(1)) d_\Delta^{2s}.$$

Comparison with the boundary upper bound yields $d_\Delta/\Delta \to 1$, hence $I(\Delta) = c(0)\Delta^{2s} + o(\Delta^{2s})$. The quantitative version follows by retaining the uniform remainder.

C.2 Proof of Theorem 2

Exact matching equation. For a candidate wrong utility $u \le 0$, equality of every answered-category probability in equation 7 is equivalent to

$$V_r\eta = T_{\Delta,u}, \qquad T_{\Delta,u,j} = G_a(R_{\Delta,u,j}), \qquad R_{\Delta,u,j} = (\Delta - u)y_j + \psi(u) - \psi(\Delta). \tag{43}$$

Once the answered masses agree, the $\perp$ masses agree because both complete laws sum to one.

Subthreshold interpolation. If $J_{\text{eff}} \le r + 1$, the evaluation map from degree-r polynomials to functions on the distinct positive-mass support is surjective by Lagrange interpolation. Hence the finite target vector in equation 43 belongs to V_r for every sufficiently small (Δ, u) for which it is defined. Choosing, for example, $u = -\Delta < 0$ and reducing Δ if necessary keeps every implied response probability strictly inside $(0, 1)$ and the coefficient vector inside the nuisance neighborhood. This gives an exact strict wrong-decision alias.

Global sign identification above the threshold. Suppose $P_{\Delta,0} = P_{u,V_r\eta}$ for some $\Delta > 0$ and $u \le 0$. Then a polynomial q of degree at most r satisfies

$$q(y_j) = G_a(c + dy_j), \qquad d = \Delta - u > 0, \qquad c = \psi(u) - \psi(\Delta), \tag{44}$$

for every positive-mass score. Writing $\rho = \sigma(a)$,

$$G_a(x) = x + \log(1 - \rho) - \log(1 - \rho e^x)$$

yields, throughout its domain,

$$G_a^{(k)}(x) > 0 \quad (k \ge 1).$$

If equation 44 held at $r+2$ distinct scores, repeated Rolle's theorem would produce a point at which the $(r+1)$st derivative of $q(y) - G_a(c+dy)$ vanished. But this derivative equals $-d^{r+1}G_a^{(r+1)}(c+dy) < 0$, a contradiction. Thus no positive-gap opposite-decision alias exists when $J_{\text{eff}} \ge r + 2$.

First noncancelable term. Let $d = \Delta - u$. On a compact local rectangle, $R_{\Delta,u,j} = d(y_j - \psi'(0)) + O(d^2)$ uniformly. Expanding G_a to order $s = r + 1$ and collecting all score polynomials of degree at most r gives

$$T_{\Delta,u} = v_{\Delta,u} + \beta_s(a)d^s\mathbf{y}^{[s]} + O(d^{s+1}), \qquad v_{\Delta,u} \in V_r. \tag{45}$$

Therefore

$$\operatorname{dist}_W(T_{\Delta,u}, V_r) = A_{J,r}d^s + O(d^{s+1}), \qquad A_{J,r} = \beta_s(a)\|D_{J,r}\|_W. \tag{46}$$

The vector $D_{J,r}$ is nonzero: otherwise the degree-s monomial would agree with a degree-r polynomial at at least $s + 1$ distinct points, contradicting Vandermonde nonsingularity.

Complete-law quadratic metric. At exact matching, perturb the response-logit vector by h. Direct differentiation of the complete multinomial law gives

$$\nabla_h^2 D_{\text{KL}}(P_{0,0}\|P_{0,h})\big|_{h=0} = W = \rho(1-\rho)^2 \operatorname{diag}(\boldsymbol{\omega}) + \rho^2(1-\rho)\boldsymbol{\omega}\boldsymbol{\omega}^\top. \tag{47}$$

Uniform positivity and the integral Taylor formula imply, uniformly for localized parameters,

$$D_{\text{KL}}(P_{\Delta,0}\|P_{u,V_r\eta}) = \frac{1}{2}\|V_r\eta - T_{\Delta,u}\|_W^2 + O\big(d\|h\|_2^2 + \|h\|_2^3\big). \tag{48}$$

Projection onto V_r and equation 46 give the fixed-utility profile

$$K_\Delta(u) = C_{J,r}d^{2s} + O(d^{2s+1}), \qquad C_{J,r} = \frac{1}{2}A_{J,r}^2, \tag{49}$$

uniformly for $0 < \Delta \le d \le d_0$.

Global localization and least-favorable utility. The boundary candidate $u = 0$ gives $K_\Delta(0) \le C\Delta^{2s}$. By compactness and uniqueness of the base observation fiber, alternatives outside a fixed local neighborhood have a positive KL gap, so every global minimizer is local. Pinsker's inequality and inverse-Lipschitzness of the finite response-logit map then give residual size $O(d^s)$ and validate equation 48 at every exact minimizer. The lower bound in equation 49, compared with the boundary upper bound, yields $d \le M\Delta$. Write $d = \Delta + h$ with $h \ge 0$. Optimality relative to the boundary candidate and the uniform remainder imply

$$C_{J,r}\big((\Delta + h)^{2s} - \Delta^{2s}\big) \le K'\Delta^{2s+1}.$$

The mean-value theorem gives $(\Delta+h)^{2s}-\Delta^{2s} \ge 2s\Delta^{2s-1}h$, hence $h = -u = O(\Delta^2)$. Substitution proves equation 9. □

Lemma 4 (Strict and closed scalar profiles agree)**.** *Fix a sufficiently small $\Delta > 0$. Let*

$$I_{\mathrm{cl}}(\Delta) = \min_{u\in[-u_0,0]} K_\Delta(u), \qquad I_{\mathrm{str}}(\Delta) = \inf_{u\in[-u_0,0)} K_\Delta(u).$$

Then $I_{\mathrm{str}}(\Delta) = I_{\mathrm{cl}}(\Delta)$. Equality of the values does not imply that the strict infimum is attained.

Proof. The compact nuisance set is independent of u, and the complete-law KL objective is jointly continuous in (u, η). Berge's minimum theorem therefore makes $K_\Delta(u) = \min_\eta D_{\mathrm{KL}}(P_{\Delta,0}\|P_{u,\eta})$ continuous on $[-u_0, 0]$. Since $[-u_0, 0)$ is dense in $[-u_0, 0]$, the infimum of a continuous function over the former equals the minimum over the latter. If every minimizer of the closed problem is $u = 0$, the equal strict infimum is approached by $u_n \uparrow 0$ and is not attained in the strict domain. □

C.3 Proof of Theorem 2

Let X_s be the $(s + 1) \times J$ Vandermonde matrix with $(X_s)_{k,j} = y_j^k$, $0 \le k \le s$, and let $\Omega = \mathrm{diag}(\omega_1, \dots, \omega_J)$. Then $H_s = X_s\Omega X_s^\top$. Cauchy–Binet gives

$$\det H_s = \sum_{|S|=s+1} \det(X_{s,S})^2 \prod_{j\in S} \omega_j, \tag{50}$$

and $\det(X_{s,S}) = \prod_{i<j,\ i,j\in S}(y_j - y_i)$.

Let $D_{J,r}$ be the W-orthogonal residual of y^s after projection on V_r. Because $1 \in V_r$,

$$0 = \langle D_{J,r}, 1\rangle_W = \rho(1-\rho)\sum_j \omega_j D_{J,r}(j),$$

so the residual has ω-mean zero. The rank-one part of W therefore vanishes on $D_{J,r}$, and its W-orthogonality to every polynomial in V_r becomes ordinary $L^2(\omega)$ orthogonality. Hence $D_{J,r}$ is the monic degree-s orthogonal-polynomial residual and

$$\|D_{J,r}\|_W^2 = \rho(1-\rho)^2\frac{\det H_s}{\det H_{s-1}}, \qquad C_{J,r} = \frac{1}{2}\beta_s(a)^2\rho(1-\rho)^2\frac{\det H_s}{\det H_{s-1}}. \tag{51}$$

Under $y'_j = c + by_j$ with $b > 0$, the monic degree-s residual scales by b^s, so $C_{J,r}$ scales by b^{2s} after the reciprocal utility rescaling that leaves the exponential tilt unchanged. Since $\mathrm{Var}(Y')^s = b^{2s}\mathrm{Var}(Y)^s$, the ratio $\chi_{J,r}$ in equation 14 is positive-affine invariant. In the minimal-support case, every summand in $\det H_s$ contains every mass and every pairwise score difference, while $\det H_{s-1}$ has a positive finite limit if all unaffected masses and separations stay uniformly bounded away from zero. Hence a necessary mass ε gives a linear decay and a necessary collision h gives a quadratic decay. Simultaneous degeneracies are governed by the full determinant ratio and are not obtained by multiplying the one-parameter statements. If support is redundant, Cauchy–Binet leaves a positive limiting summand precisely when at least $s + 1$ distinct positive-mass scores remain. □

C.4 Location of the remaining repaired proofs

The compact-class thickening theorem, its exact endpoint topology, the nonaligned-norm and ellipsoidal corollaries, the approximate-sharing tangent coefficient and finite-gap alias radius, the independent-pilot calibration theorem, and the exact-validity coefficient-level crossover are proved

in Appendix E. The gauge-fixed approximate-exclusion class, quantitative/qualitative coefficient theorem, and finite-gap exclusion radius are proved in Appendix F. Appendix I then proves the joint imperfect-validity quadratic–quartic crossover and its exact critical alias radius. These appendices supersede earlier full-subspace tube arguments and any inference from tangent cancellation to finite-gap aliasing. The standalone coercivity and row-wise transportation lemmas in Appendix K remain supporting results.

D TWO-SIDED ROBUST NORMAL-CONTACT GEOMETRY

This appendix proves the general robust normal-contact theorem used in Theorem 1. The argument is finite-dimensional and works in square-root likelihood coordinates. Its substantive assumptions are visible: a regular embedded nuisance observation manifold, compact normal uncertainty cross-sections with graphical convergence, uniform positivity, and an isolated remote fiber for unrestricted profiling.

D.1 INTRINSIC NORMAL COORDINATES AND UNCERTAINTY GRAPHS

Let

$$\mathbb{S}_+^{M-1} = \{x \in (0,\infty)^M : \|x\|_2 = 1\}$$

be the positive square-root probability sphere. Let $G : U \subset \mathbb{R}^q \to \mathbb{S}_+^{M-1}$ be an analytic embedded chart, $x_0 = G(0)$, and write $\mathcal{N}_{\mathrm{loc}} = G(U)$. After shrinking U, the tubular-neighborhood theorem gives a unique nearest-point projection $\pi_{\mathcal{N}}$ on a neighborhood of x_0. For the analytic signal $F(t) = S(P_t)$, define

$$x_t = \pi_{\mathcal{N}}(F(t)) = G(\xi_t), \qquad R_t = F(t) - x_t.$$

For any chart point ξ, let $T_\xi = \operatorname{range} DG(\xi)$ and $\mathsf{N}_\xi = (T_\xi \oplus \operatorname{span}\{G(\xi)\})^\perp$. Along the projected signal write $T_t = T_{\xi_t}$ and

$$\mathsf{N}_t = (T_t \oplus \operatorname{span}\{x_t\})^\perp.$$

This is the normal space inside the probability sphere. A local analytic orthogonal trivialization is an analytic family of isometries $U_\xi : \mathsf{N}_0 \to \mathsf{N}_\xi$ with $U_0 = I$; along the projected signal we abbreviate $U_t = U_{\xi_t}$.

Lemma 5 (Analytic normal-bundle trivialization)**.** *After shrinking U, such a family U_ξ exists. It is unique only up to an analytic orthogonal change of frame, which leaves every distance appearing below invariant.*

Proof. Because G is an analytic immersion into the sphere, the orthogonal normal projector is the analytic constant-rank matrix

$$\Pi_{\mathsf{N}_\xi} = I - G(\xi)G(\xi)^\top - DG(\xi)\big(DG(\xi)^\top DG(\xi)\big)^{-1} DG(\xi)^\top.$$

Choose fixed ambient vectors whose projections form a basis of N_0. Their projected vectors remain linearly independent after shrinking U. Analytic Gram–Schmidt produces an analytic orthonormal frame of N_ξ; mapping the frame at zero to this frame defines U_ξ. Two orthonormal frames differ by an analytic orthogonal matrix. □

The radial component of R_t is higher order. Indeed, $\|x_t + R_t\|_2 = \|x_t\|_2 = 1$ gives

$$\langle x_t, R_t\rangle = -\frac{1}{2}\|R_t\|_2^2. \tag{52}$$

Hence the first nonzero term of R_t belongs to N_t. If the contact order is $s < \infty$, write intrinsically

$$U_t^{-1}\Pi_{\mathsf{N}_t} R_t = t^s r_s + O(t^{s+1}), \qquad r_s \in \mathsf{N}_0 \setminus \{0\}. \tag{53}$$

For the analytic residual of Section 2, $r_s = R^{(s)}(0)/s!$ after identifying the normal spaces.

Definition 6 (Order-s normal-graph uncertainty)**.** Let $\mathcal{B}_{\mathrm{T}}, \mathcal{B}_{\mathrm{A}} \subset \mathsf{N}_0$ be nonempty compact sets containing zero. A two-sided uncertainty family has an order-s normal graph with these limiting

cross-sections if there are compact sets $\mathcal{B}_{\mathrm{T}}(t), \mathcal{B}_{\mathrm{A}}(t) \subset \mathsf{N}_0$ and continuous parameterizations by actual positive laws $P^{\mathrm{T}}_{t,a}$ and $Q^{\mathrm{A}}_{t,\xi,b}$ such that

$$d_{\mathrm{H}}(\mathcal{B}_{\mathrm{T}}(t), \mathcal{B}_{\mathrm{T}}) + d_{\mathrm{H}}(\mathcal{B}_{\mathrm{A}}(t), \mathcal{B}_{\mathrm{A}}) \le Ct, \tag{54}$$

$$S(P^{\mathrm{T}}_{t,a}) = F(t) + t^s U_t a + E_{\mathrm{T}}(t,a), \qquad \sup_a \|E_{\mathrm{T}}(t,a)\|_2 \le Ct^{s+1}, \tag{55}$$

$$S(Q^{\mathrm{A}}_{t,\xi,b}) = G(\xi) + t^s U_\xi b + E_{\mathrm{A}}(t,\xi,b), \quad \sup_{\xi,b} \|E_{\mathrm{A}}(t,\xi,b)\|_2 \le Ct^{s+1}. \tag{56}$$

Here the suprema are over the displayed compact cross-sections and one fixed compact chart. The maps parameterize the whole local uncertainty sets, not only selected candidate paths. Replacing the $O(t)$ and $O(t^{s+1})$ bounds by $o(1)$ and $o(t^s)$ defines the qualitative graphical-convergence version.

The definition is intrinsic. Equivalently, the rescaled normal projections of the entire true and alternative uncertainty graphs converge, in the common normal fiber, in both the upper and lower Painlevé–Kuratowski senses. For compact sets this is equivalent to the Hausdorff convergence in equation 54. Thus the definition first quotients out the nuisance tangent space and then records only the surviving normal uncertainty. An order-s tangential perturbation does not change the leading cross-section: by the full rank of $DG(\xi_t)$ it is absorbed by an order-s shift of ξ, while the induced normal error is $O(t^{2s}) = O(t^{s+1})$. Under another orthogonal trivialization all limiting cross-sections are rotated by the same orthogonal map.

Define the local two-sided robust profile

$$\mathcal{I}_{\mathrm{loc}}(t) = \inf_{\substack{a \in \mathcal{B}_{\mathrm{T}}(t),\ b \in \mathcal{B}_{\mathrm{A}}(t) \\ \xi \in U_{\mathrm{c}}}} D_{\mathrm{KL}}(P^{\mathrm{T}}_{t,a} \| Q^{\mathrm{A}}_{t,\xi,b}), \tag{57}$$

where $U_{\mathrm{c}} \Subset U$ is a compact chart containing every ξ_t for small t.

D.2 Geometric lemmas

Lemma 7 (Uniform tubular estimates)**.** *After shrinking the chart there are constants $c, C > 0$ such that, uniformly in small t and $\xi \in U_{\mathrm{c}}$,*

$$\|G(\xi) - G(\xi_t)\|_2 \ge c\|\xi - \xi_t\|_2, \tag{58}$$

$$\|U_\xi - U_t\|_{\mathrm{op}} \le C\|\xi - \xi_t\|_2, \tag{59}$$

$$\|\Pi_{\mathsf{N}_t}(G(\xi) - G(\xi_t))\|_2 \le C\|\xi - \xi_t\|_2^2. \tag{60}$$

Proof. The first inequality follows from full column rank and compactness after shrinking the chart. The second follows from analyticity of the chosen normal frame. For the third, Taylor expand G at ξ_t. The linear term lies in T_t and is killed by Π_{N_t}; the second derivative is uniformly bounded on the compact chart. □

Lemma 8 (Hausdorff stability of the normal difference set)**.** *Let $r_t \to r$ and let compact sets $A_t \to A$, $B_t \to B$ in Hausdorff distance. Then*

$$\mathrm{dist}(r_t, B_t - A_t) \longrightarrow \mathrm{dist}(r, B - A).$$

If all three errors are $O(t)$, the difference of the two distances is $O(t)$.

Proof. For compact sets, $d_{\mathrm{H}}(B_t - A_t, B - A) \le d_{\mathrm{H}}(B_t, B) + d_{\mathrm{H}}(A_t, A)$. The distance to a nonempty closed set is one-Lipschitz in its point and is stable under Hausdorff perturbations of the set. Combining the two bounds proves the claim. □

Lemma 9 (Profiled square-root distance)**.** *Under Definition 6, let*

$$d_\star = \mathrm{dist}(r_s, \mathcal{B}_{\mathrm{A}} - \mathcal{B}_{\mathrm{T}}).$$

Then the minimum square-root distance between the two local uncertainty sets satisfies

$$\delta_t := \inf_{a,b,\xi} \|S(P^{\mathrm{T}}_{t,a}) - S(Q^{\mathrm{A}}_{t,\xi,b})\|_2 = t^s d_\star + O(t^{s+1}). \tag{61}$$

With qualitative graphical convergence, $\delta_t / t^s \to d_\star$. Moreover every $O(t^s)$-near-minimizing nuisance coordinate satisfies $\|\xi - \xi_t\|_2 = O(t^s)$.

Proof. Put $r_t = t^{-s}U_t^{-1}\Pi_{\mathsf{N}_t}R_t$. By equation 53, $r_t = r_s + O(t)$. Lemma 8 therefore gives

$$d_t := \operatorname{dist}(r_t, \mathcal{B}_{\mathrm{A}}(t) - \mathcal{B}_{\mathrm{T}}(t)) = d_\star + O(t). \tag{62}$$

For the upper bound, choose a_t, b_t within $O(t)$ of a minimizing pair and set $\xi = \xi_t$. Equations equation 55– equation 56 give

$$S(P^{\mathrm{T}}_{t,a_t}) - S(Q^{\mathrm{A}}_{t,\xi_t,b_t}) = R_t + t^s U_t(a_t - b_t) + O(t^{s+1}),$$

so $\delta_t \le t^s d_t + O(t^{s+1})$.

For the lower bound, take a minimizing sequence. The upper bound makes its square-root distance $O(t^s)$. Since the two uncertainty displacements and R_t are also $O(t^s)$,

$$\|G(\xi) - G(\xi_t)\|_2 = O(t^s).$$

The inverse-Lipschitz bound equation 58 yields $\|\xi - \xi_t\|_2 = O(t^s)$. Projecting the difference onto N_t and using Lemma 7 gives

$$\Pi_{\mathsf{N}_t}\{S(P^{\mathrm{T}}_{t,a}) - S(Q^{\mathrm{A}}_{t,\xi,b})\} = \Pi_{\mathsf{N}_t}R_t + t^s U_t(a - b) + O(t^{s+1}).$$

Indeed the normal part of $G(\xi) - G(\xi_t)$ is $O(t^{2s})$, and the frame change contributes $t^s O(\|\xi - \xi_t\|) = O(t^{2s})$; both are $O(t^{s+1})$ because $s \ge 1$. Taking norms and then the infimum gives $\delta_t \ge t^s d_t - O(t^{s+1})$. Together with equation 62, this proves equation 61. □

Lemma 10 (Uniform KL–square-root conversion)**.** *On a uniformly positive finite alphabet there are $c, C > 0$ such that, whenever p and q lie in one sufficiently small compact neighborhood,*

$$\left|D_{\mathrm{KL}}(p\|q) - 2\|S(p) - S(q)\|_2^2\right| \le C\|S(p) - S(q)\|_2^3, \tag{63}$$

and $c\|S(p) - S(q)\|_2^2 \le D_{\mathrm{KL}}(p\|q) \le C\|S(p) - S(q)\|_2^2$.

Proof. Write $q_i = p_i + h_i$ and expand the finite sum $\sum_i p_i \log(p_i/q_i)$ uniformly on a compact set bounded away from zero. The linear term vanishes because $\sum_i h_i = 0$. Expanding $\sqrt{p_i + h_i}$ gives $D_{\mathrm{KL}}(p\|q) = 4H^2(p,q) + O(\|p - q\|_2^3) = 2\|S(p) - S(q)\|_2^2 + O(\|S(p) - S(q)\|_2^3)$. Compact positivity supplies the two-sided quadratic coercivity. □

D.3 The robust normal-contact theorem

Theorem 11 (Two-sided robust normal-contact law)**.** *Assume equation 53 and the quantitative normal-graph conditions of Definition 6. With*

$$r_t = t^{-s}U_t^{-1}\Pi_{\mathsf{N}_t}R_t, \qquad d_t = \operatorname{dist}(r_t, \mathcal{B}_{\mathrm{A}}(t) - \mathcal{B}_{\mathrm{T}}(t)),$$

the finite-t normal form is

$$\mathcal{I}_{\mathrm{loc}}(t) = 2t^{2s}d_t^2 + O(t^{2s+1}). \tag{64}$$

Consequently,

$$\mathcal{I}_{\mathrm{loc}}(t) = 2t^{2s}\operatorname{dist}^2\big(r_s, \mathcal{B}_{\mathrm{A}} - \mathcal{B}_{\mathrm{T}}\big) + O(t^{2s+1}). \tag{65}$$

Under qualitative graphical convergence,

$$\frac{\mathcal{I}_{\mathrm{loc}}(t)}{t^{2s}} \longrightarrow 2\operatorname{dist}^2\big(r_s, \mathcal{B}_{\mathrm{A}} - \mathcal{B}_{\mathrm{T}}\big). \tag{66}$$

If the full alternative class satisfies the robust isolated-fiber condition

$$\liminf_{t\downarrow 0} \inf_{\substack{a\in\mathcal{B}_{\mathrm{T}}(t)\\ Q\in\mathcal{Q}_t^{\mathrm{glob}}\setminus\mathcal{Q}_t^{\mathrm{loc}}}} D_{\mathrm{KL}}(P^{\mathrm{T}}_{t,a}\|Q) > 0, \tag{67}$$

then the unrestricted global robust profile has the same expansion. The coefficient is invariant under analytic reparameterization of the nuisance chart and under the choice of orthogonal normal frame. No convexity is required: nonempty compact cross-sections suffice. Convexity is used only by later optimization corollaries when uniqueness or stability of a nearest point, a minimax interchange, or a measurable design selector is required.

Proof. Lemma 9 gives $\delta_t = t^s d_\star + O(t^{s+1})$. Evaluating KL at a pair achieving the square-root-distance upper bound and applying Lemma 10 yields

$$\mathcal{I}_{\mathrm{loc}}(t) \le 2\delta_t^2 + O(t^{3s}).$$

Conversely, this upper bound is $O(t^{2s})$. Quadratic coercivity in Lemma 10 forces every exact KL minimizer to have square-root residual $O(t^s)$. The uniform expansion equation 63 is therefore valid at the KL minimizer, whose square-root distance is at least δ_t. Hence

$$\mathcal{I}_{\mathrm{loc}}(t) \ge 2\delta_t^2 - O(t^{3s}).$$

Since $3s \ge 2s+1$, the two bounds first prove equation 64. The distance expansion equation 62 then proves equation 65. The same argument with little-o bounds proves equation 66.

Under equation 67, the local candidate tends to zero whereas every remote alternative has a fixed positive cost. Thus every unrestricted minimizer is eventually local, proving the global statement.

An analytic change of nuisance coordinates leaves the physical manifold, its tubular projection, and its normal bundle unchanged. If another normal frame is $\widetilde{U}_t = U_t O_t$ with O_t orthogonal, then $\widetilde{r}_s = O_0^\top r_s$ and both cross-sections are transformed by the same $O_0^\top$. Euclidean distance to their Minkowski difference is unchanged. □

Corollary 12 (One-sided and two-sided balls)**.** *If $\mathcal{B}_{\mathrm{T}}$ and $\mathcal{B}_{\mathrm{A}}$ are Euclidean balls in the normal space with radii κ_{T} and κ_{A}, then*

$$\mathcal{I}(t) = 2t^{2s}\big(\|r_s\|_2 - \kappa_{\mathrm{T}} - \kappa_{\mathrm{A}}\big)_+^2 + O(t^{2s+1}). \tag{68}$$

The one-sided law is the case $\kappa_{\mathrm{T}} = 0$.

Proof. The Minkowski difference of centered Euclidean balls of radii κ_{A} and κ_{T} is the centered ball of radius $\kappa_{\mathrm{A}} + \kappa_{\mathrm{T}}$. The distance of r_s to that ball is the positive part in equation 68. □

Corollary 13 (Ellipsoidal normal uncertainty)**.** *For a one-sided ellipsoid $\mathcal{B}_{\mathrm{A}} = \{b : b^\top M b \le \kappa^2\}$ with $M \succ 0$, the leading coefficient is $2\|r_s - b^\star\|_2^2$, where $b^\star = r_s$ if $r_s^\top M r_s \le \kappa^2$ and otherwise*

$$b^\star = (I + \lambda M)^{-1} r_s, \qquad (b^\star)^\top M b^\star = \kappa^2, \tag{69}$$

for the unique $\lambda > 0$ solving the displayed constraint.

Proof. This is the KKT system for Euclidean projection onto a compact strictly convex ellipsoid. The scalar constraint is strictly decreasing in λ, so the positive solution is unique. □

D.4 Decision transfer and exact-zero topology

Proposition 14 (Robust one-sided decision transfer)**.** *Let $v \le 0$ index the wrong side of a decision boundary and let $d > 0$ be the signal displacement from v. Suppose Theorem 11 holds uniformly for $v \uparrow 0$ and $d \downarrow 0$, with $d_\star(v)^2 = d_\star(0)^2 + O(|v|)$ and $d_\star(0) > 0$. Suppose every fixed remote $v < 0$ is positively separated. Then for a true gap $\Delta > 0$, every least-favorable wrong utility satisfies*

$$v_\Delta = O(\Delta^2), \qquad \frac{\Delta - v_\Delta}{\Delta} \to 1,$$

and the one-sided robust decision profile is

$$2d_\star(0)^2 \Delta^{2s} + O(\Delta^{2s+1}). \tag{70}$$

Proof. The boundary candidate gives an $O(\Delta^{2s})$ upper bound. Compactness and remote separation force $v_\Delta \to 0$. With $h_\Delta = -v_\Delta \ge 0$ and $d_\Delta = \Delta + h_\Delta$, the uniform theorem gives $I(\Delta) = 2d_\star(v_\Delta)^2 d_\Delta^{2s} + O(d_\Delta^{2s+1})$. Comparison with the boundary candidate first yields $d_\Delta = O(\Delta)$. Using $d_\star(v_\Delta)^2 = d_\star(0)^2 + O(h_\Delta)$ and the mean-value bound $(\Delta + h)^{2s} - \Delta^{2s} \ge 2s\Delta^{2s-1}h$, optimality yields $h_\Delta \le C\Delta^2$. Substitution proves equation 70. □

Proposition 15 (Closed zero, strict infimum, and strict alias)**.** *Fix $t > 0$. Let $\mathfrak{P}_t$ be compact, let $\mathfrak{Q}_t^{\mathrm{cl}}$ be a compact closed alternative set, and let $\mathfrak{Q}_t^{\circ}$ be dense in $\mathfrak{Q}_t^{\mathrm{cl}}$. Then*

$$\inf_{P\in\mathfrak{P}_t, Q\in\mathfrak{Q}_t^{\circ}} D_{\mathrm{KL}}(P\|Q) = \min_{P\in\mathfrak{P}_t, Q\in\mathfrak{Q}_t^{\mathrm{cl}}} D_{\mathrm{KL}}(P\|Q), \tag{71}$$

$$\min_{P\in\mathfrak{P}_t, Q\in\mathfrak{Q}_t^{\mathrm{cl}}} D_{\mathrm{KL}}(P\|Q) = 0 \iff \mathfrak{P}_t \cap \mathfrak{Q}_t^{\mathrm{cl}} \neq \varnothing. \tag{72}$$

A strict exact alias exists iff $\mathfrak{P}_t \cap \mathfrak{Q}_t^{\circ} \neq \varnothing$. Thus a zero strict infimum may be unattained. If strict wrong-decision kernels converge in total variation to a true kernel at which the rule terminates almost surely, then no procedure can be uniformly δ-correct for $\delta < 1/2$.

Proof. Joint continuity of finite-alphabet KL on the uniformly positive compact set and density of $\mathfrak{Q}_t^{\circ}$ prove equation 71. Nonnegativity of KL and equality iff the laws coincide prove equation 72 and the strict-alias criterion. For the final statement, apply Lemma 28: for every fixed horizon T, adaptive transcript total variation is at most T times the maximum one-step kernel total variation. Almost-sure termination at the true model supplies a finite T carrying probability greater than $1 - \delta - \eta$ for the true recommendation event. A sufficiently close strict wrong-decision alternative gives that same event probability greater than $1 - \delta - 2\eta > \delta$, contradicting correctness when $\eta < (1 - 2\delta)/4$. □

A zero distance in equation 65 is only a leading-order cancellation statement. It does not, by itself, assert finite-t equality. Exact finite-gap aliasing is governed by Proposition 15 and the model-specific alias radius.

D.5 Specializations used in the paper

Corollary 16 (Scored-model one- and two-sided thickening)**.** *In the scored response-logit model, let h denote a local logit perturbation and let L be the derivative of the square-root likelihood map at the matched law. Then*

$$L^{\top}L = \frac{1}{4}W, \qquad \|Lh\|_2 = \frac{1}{2}\|h\|_W. \tag{73}$$

Let the first unabsorbed logit-normal target be $A_{J,r}\Delta^s\widehat{r}$ with $\|\widehat{r}\|_W = 1$. If true and alternative logit uncertainty are W-balls of radii ϵ_{T} and ϵ_{A}, then

$$I_{\epsilon_{\mathrm{T}},\epsilon_{\mathrm{A}}}(\Delta) = \frac{1}{2}\big(A_{J,r}\Delta^s - \epsilon_{\mathrm{T}} - \epsilon_{\mathrm{A}}\big)_+^2 + O(\Delta^{2s+1}), \tag{74}$$

whenever the radii are $O(\Delta^s)$. The exact finite-gap theorem in Section 3.1 is the one-sided specialization $\epsilon_{\mathrm{T}} = 0$ with its model-specific compact-class and endpoint analysis.

Proof. The complete-law expansions give $D_{\mathrm{KL}}(P\|P_h) = \frac{1}{2}h^{\top}Wh + O(\|h\|^3)$, while Lemma 10 gives $D_{\mathrm{KL}}(P\|P_h) = 2\|Lh\|_2^2 + O(\|h\|^3)$. Equality of the quadratic forms proves equation 73. A W-ball of radius ϵ therefore maps to a Euclidean normal ball of radius $\epsilon/2$, and the nominal normal jet has Euclidean norm $A_{J,r}/2$. Corollary 12 gives equation 74. □

Corollary 17 (Approximate sharing as product-space robust contact)**.** *For a finite protocol library, form the product square-root space with the block complete-law metric. The shared nuisance evaluations form its tangent subspace; the protocol-specific deviations form the alternative normal cross-section. Theorem 11 gives the tangent coefficient as the squared distance of the first protocol-disagreement jet to that cross-section. Under the aligned per-protocol balls of Theorem 3, this distance reduces to the exact positive-part formula stated there. Allowing uncertainty on the true protocol mechanism as well replaces the alternative cross-section by the Minkowski difference of the alternative and true cross-sections.*

Proof. For fixed design w, draw the revealed protocol label from w and then observe its complete response/nonresponse outcome. This compound experiment has a square-root space equal to the weighted Hilbert direct sum of the protocol spaces, and its KL is the weighted sum of protocol KL divergences. Profiling the shared nuisance is orthogonal projection onto the shared tangent subspace. The remaining protocol-specific perturbations are normal directions. Apply Theorem 11; the scalar reduction in Theorem 33 evaluates the resulting distance for the aligned balls. □

D.6 WHY THE ASSUMPTIONS ARE VISIBLE

Example 18 (Remote escape without an outer-fiber gap)**.** Let the signal be $P_t = \mathrm{Bernoulli}(1/2 + t)$ and let the local nuisance germ consist only of $Q_0 = \mathrm{Bernoulli}(1/2)$. Add a disconnected noncompact nuisance branch $Q_\eta = \mathrm{Bernoulli}(1/2 + 1/\eta)$, $\eta \geq 2$. The local contact coefficient is positive, but $Q_{1/t} = P_t$, so the unrestricted profile is zero. The minimizing nuisance escapes to infinity. This shows why local normal geometry needs compactness or the explicit remote-fiber gap equation 67.

Example 19 (Nonconvex cross-sections and optimizer instability)**.** In $\mathbb{R}^2$, take $r = (0, 1)$, $\mathcal{B}_{\mathrm{T}} = \{0\}$, and $\mathcal{B}_{\mathrm{A}} = \{(-1, 0), (1, 0)\}$. The distance law remains valid, but there are two nearest pairs. Perturbing r to $(\epsilon, 1)$ selects one or the other discontinuously as the sign of ϵ changes. This nonconvex compact set still satisfies Theorem 11, because the coefficient proof uses only compactness and Hausdorff convergence. Convexity is useful for unique stable projections and measurable robust-design selectors; without it, optimizing pairs can be nonunique and discontinuous.

Example 20 (Failure of graphical convergence)**.** Let $r = 1$ in a one-dimensional normal space and let the alternative scaled cross-section alternate along a sequence between $\{0\}$ and $\{1\}$. The normalized robust information alternates between 2 and 0. Hence a single limiting coefficient need not exist without upper and lower graphical convergence of the scaled uncertainty sets.

Uniformity ledger. In Theorem 11, $t \downarrow 0$ while the finite alphabet, compact chart, cross-sections, contact order, and positivity bounds are fixed. The $O(t^{2s+1})$ constant is uniform over the displayed compact uncertainty sets and over compact regular strata on which the tubular radius, visible-Jacobian singular value, positivity constant, graph-remainder bounds, and remote-fiber gap are bounded away from zero. The theorem separates local from unrestricted profiling. Proposition 15 separates closed alternatives, strict infima, and attained strict aliases. Proposition 14 is the additional step needed for one-sided wrong-decision profiling.

E ACCEPTANCE-GATE REPAIRS: EXACT CLASSES, TOPOLOGY, CALIBRATION, AND CROSSOVER

This appendix records the repaired statements used by the submission-facing paper. All parameter sets below are compact, all complete output probabilities are bounded away from zero on the displayed neighborhoods, and all nuisance coefficient sets contain the true coefficient in their relative interior. Constants may depend on these fixed compact sets, on the finite score support, and on the response baseline, but not on the local variables that tend to zero.

E.1 CORRECTNESS, ALTERNATIVES, AND CLASS ALIGNMENT

Definition 21 (Almost-surely terminating fixed-confidence correctness)**.** Let $\mathcal{M}$ be a model class with decision map d, and let $\mathcal{M}_{\mathrm{uni}}$ be its unique-decision subset. A predictable experiment selection rule, stopping time τ, and recommendation $\widehat{d}$ are δ-correct on $\mathcal{M}$ when

$$\mathbb{P}_m(\tau < \infty) = 1, \qquad \mathbb{P}_m(\widehat{d} = d(m)) \geq 1 - \delta \quad \text{for every } m \in \mathcal{M}_{\mathrm{uni}}. \tag{75}$$

Nontermination is therefore counted as failure.

For a unique-decision model m, write

$$\mathrm{Alt}^{\circ}_{\mathcal{M}}(m) = \{m' \in \mathcal{M}_{\mathrm{uni}} : d(m') \neq d(m)\}, \tag{76}$$

$$\mathrm{Alt}_{\mathcal{M}}(m) = \mathrm{cl}_{\mathcal{M}}\, \mathrm{Alt}^{\circ}_{\mathcal{M}}(m). \tag{77}$$

The class-relative characteristic information is

$$\Gamma_{\mathcal{M}}(m) = \sup_{w \in \Delta(\mathcal{A})} \inf_{m' \in \mathrm{Alt}^{\circ}_{\mathcal{M}}(m)} \sum_{e \in \mathcal{A}} w_e D_{\mathrm{KL}}(P^e_m \| P^e_{m'}). \tag{78}$$

The infimum is unchanged if the strict set is replaced by its closure whenever the one-step KL objective is continuous and the strict set is dense in the closed set.

Theorem 22 (Row-wise class-aligned transportation bound)**.** *For each row n, let the procedure be δ_n-correct on the model class $\mathcal{M}_n$, and define $\Gamma_{\mathcal{M}_n}$ using the same class. Then, for every $m_n \in (\mathcal{M}_n)_{\mathrm{uni}}$,*

$$\mathbb{E}_{m_n}[\tau_n] \geq \frac{\mathrm{kl}(1-\delta_n, \delta_n)}{\Gamma_{\mathcal{M}_n}(m_n)}. \tag{79}$$

No fixed-class result is being applied to a moving family.

Proof. Fix one row and one strict wrong-decision alternative m'. The stopped-history chain rule and predictable experiment selection give

$$D_{\mathrm{KL}}(\mathcal{L}_m(H_\tau) \| \mathcal{L}_{m'}(H_\tau)) = \sum_e \mathbb{E}_m[N_e(\tau)] D_{\mathrm{KL}}(P_m^e \| P_{m'}^e).$$

Data processing through the event $\{\widehat{d} = d(m)\}$ and equation 75 lower-bound the left side by $\mathrm{kl}(1-\delta, \delta)$. Divide by $\mathbb{E}_m \tau$, normalize the expected counts to a design in $\Delta(\mathcal{A})$, infimize over strict alternatives, and maximize over designs. Apply this argument separately in every row. Almost-sure termination justifies the stopped-history law; the usual truncation argument covers unbounded integrable stopping times. □

E.2 The exact compact thickened class

Let $V_r \in \mathbb{R}^{J \times (r+1)}$ be the score-evaluation matrix and let $\mathcal{H} \subset \mathbb{R}^{r+1}$ be compact with $B_2(0, h_0) \subset \mathcal{H}$ for some $h_0 > 0$. Put

$$\mathcal{Q}_r^0 = V_r \mathcal{H}, \qquad \mathcal{Q}_r^W(\epsilon) = \{V_r \eta + g : \eta \in \mathcal{H},\ \|g\|_W \leq \epsilon\}, \tag{80}$$

and, for a fixed utility radius $u_0 > 0$,

$$\mathcal{K}_\epsilon = [-u_0, u_0] \times \mathcal{Q}_r^W(\epsilon), \qquad (\mathcal{K}_\epsilon)_{\mathrm{uni}} = \{(u, q) \in \mathcal{K}_\epsilon : u \neq 0\}. \tag{81}$$

The decision is the sign of u, the true model is $m_\Delta = (\Delta, 0)$, and its strict wrong-decision alternatives have $u < 0$. The closed alternatives additionally allow $u = 0$.

For the exact target $T_{\Delta,u}$, define

$$d_{\mathcal{H}}(\Delta, u) = \mathrm{dist}_W(T_{\Delta,u}, \mathcal{Q}_r^0), \tag{82}$$

$$\epsilon_{\mathrm{alias}}(\Delta) = \min_{-u_0 \leq u \leq 0} d_{\mathcal{H}}(\Delta, u). \tag{83}$$

The minimum exists, but it need not be unique. Let $\mathcal{U}_\Delta^\star$ denote its nonempty compact set of minimizing utilities.

Lemma 23 (Uniform local interiority of projected coefficients)**.** *There are $d_0 > 0$ and $h_1 \in (0, h_0)$ such that for every $0 \leq \Delta \leq d = \Delta - u \leq d_0$, the W-orthogonal projection of $T_{\Delta,u}$ onto $\mathrm{range}(V_r)$ has a unique coefficient $\widehat{\eta}(\Delta, u)$ satisfying*

$$\|\widehat{\eta}(\Delta, u)\|_2 \leq Cd \leq h_1. \tag{84}$$

Consequently, throughout this local region,

$$\mathrm{dist}_W(T_{\Delta,u}, \mathcal{Q}_r^0) = \mathrm{dist}_W(T_{\Delta,u}, \mathrm{range}(V_r)), \tag{85}$$

and

$$\mathrm{dist}_W(T_{\Delta,u}, \mathcal{Q}_r^W(\epsilon)) = \big(d_{\mathcal{H}}(\Delta, u) - \epsilon\big)_+. \tag{86}$$

Proof. The W-projection coefficient is

$$\widehat{\eta} = (V_r^\top W V_r)^{-1} V_r^\top W T_{\Delta,u}.$$

The Gram matrix is positive definite because the support contains at least $r+1$ distinct positive-mass scores. The exact target satisfies $\|T_{\Delta,u}\|_2 \leq Cd$ on a compact local rectangle, so equation 84 follows. Reduce d_0 until the coefficient lies in $B_2(0, h_1) \subset \mathcal{H}$. This proves equation 85. The nearest point in the linear subspace therefore belongs to the relative interior of the compact evaluation set. Radially truncate its normal residual in the W norm; the usual Hilbert space tube identity then gives equation 86 without replacing the statistical class by an unbounded subspace. □

Lemma 24 (Uniform complete-KL approximation by compact tube distance)**.** *For*

$$K_{\Delta,\epsilon}(u) = \min_{q \in \mathcal{Q}_r^W(\epsilon)} D_{\mathrm{KL}}(P_{\Delta,0} \| P_{u,q}), \qquad d = \Delta - u,$$

there are $d_0, \epsilon_0, C > 0$ such that, uniformly over

$$0 \le \Delta \le d \le d_0, \qquad 0 \le \epsilon \le \epsilon_0, \tag{87}$$

we have

$$\left| K_{\Delta,\epsilon}(u) - \frac{1}{2} \operatorname{dist}_W^2\big(T_{\Delta,u}, \mathcal{Q}_r^W(\epsilon)\big) \right| \le C d^{2s+1}, \qquad s = r + 1. \tag{88}$$

The same statement covers the positive-distance and exact-zero regions. If $T_{\Delta,u} \in \mathcal{Q}_r^W(\epsilon)$, both terms on the left of equation 88 are exactly zero. The constants depend only on the declared compact utility and nuisance neighborhoods, the positivity margin, ϵ_0, and the isolated-fiber constants.

Proof. Write $h = q - T_{\Delta,u}$ and let $\mathcal{K}_{\Delta,u}(h) = D_{\mathrm{KL}}(P_{\Delta,0} \| P_{u,T_{\Delta,u}+h})$. At $h = 0$ the complete laws match, so the value and gradient vanish. On one common compact positive neighborhood, the exact integral-Hessian identity and smoothness of the complete-law map give

$$\left| \mathcal{K}_{\Delta,u}(h) - \frac{1}{2} \|h\|_W^2 \right| \le C\big(d\|h\|_W^2 + \|h\|_W^3\big) \tag{89}$$

whenever $\|h\|_W$ is sufficiently small; the same neighborhood has a uniform quadratic lower Hessian bound.

Put $\delta = \operatorname{dist}_W(T_{\Delta,u}, \mathcal{Q}_r^W(\epsilon))$. By Lemma 23, $\delta = (d_{\mathcal{H}}(\Delta, u) - \epsilon)_+ \le C d^s$. A nearest compact-tube point therefore gives, by equation 89,

$$K_{\Delta,\epsilon}(u) \le \frac{1}{2}\delta^2 + C d^{2s+1},$$

where $3s \ge 2s + 1$. If $\delta = 0$, that point is the exact target and the profile equals zero, proving the claim in the exact-zero region.

Assume $\delta > 0$ and let $q^\star$ be an exact global minimizer. The preceding upper bound is $O(d^{2s})$. The uniform outer-fiber gap first places $q^\star$ in the common local chart. Pinsker's inequality and the local inverse-Lipschitz property of the complete response-logit map then imply $\|q^\star - T_{\Delta,u}\|_W = O(d^s)$. Applying equation 89 at this minimizer and using the definition of δ yields

$$K_{\Delta,\epsilon}(u) \ge \frac{1}{2}\delta^2 - C d^{2s+1}.$$

Combining the two bounds proves equation 88 uniformly on equation 87. □

Theorem 25 (One-sided alternative-class thickening, exact topology, and metric stability)**.** *Let $s = r + 1$, $J_{\mathrm{eff}} \ge r + 2$, and $A_{J,r} = \beta_s(a)\|(I - \Pi_r^W)\mathbf{y}^{[s]}\|_W$. There are $\Delta_0, \epsilon_0, K, L < \infty$ such that, for $0 < \Delta \le \Delta_0$ and $0 \le \epsilon \le \epsilon_0$,*

$$A_{J,r}\Delta^s - K\Delta^{s+1} \le \epsilon_{\mathrm{alias}}(\Delta) \le A_{J,r}\Delta^s + K\Delta^{s+1}, \tag{90}$$

$$\left| I_\epsilon(\Delta) - \frac{1}{2}(\epsilon_{\mathrm{alias}}(\Delta) - \epsilon)_+^2 \right| \le K\Delta^{2s+1}. \tag{91}$$

Here I_ϵ profiles only the wrong-decision alternative class in $\mathcal{K}_\epsilon$; the true law is not perturbed. For each fixed Δ, $\epsilon \mapsto I_\epsilon(\Delta)$ is nonincreasing and L-Lipschitz on $[0, \epsilon_0]$.

Furthermore:

- *(i) $I_\epsilon(\Delta) > 0$ iff $\epsilon < \epsilon_{\mathrm{alias}}(\Delta)$;*
- *(ii) at $\epsilon = \epsilon_{\mathrm{alias}}(\Delta)$, the closed profile and strict infimum are zero;*
- *(iii) a strict exact alias is attained at equality iff $\mathcal{U}_\Delta^\star \cap (-u_0, 0) \ne \varnothing$;*
- *(iv) for $\epsilon > \epsilon_{\mathrm{alias}}(\Delta)$, a strict exact alias exists.*

No uniqueness of the alias-radius minimizer is assumed or claimed.

Proof. The target expansion and Lemma 23 give, uniformly for $d = \Delta - u$ in the localized region,

$$d_{\mathcal{H}}(\Delta, u) = A_{J,r} d^s + O(d^{s+1}).$$

The boundary candidate and the positive outer-fiber KL gap localize every radius minimizer to $d = O(\Delta)$; comparison of the homogeneous leading term then gives $d = \Delta + O(\Delta^2)$. This proves equation 90.

By equation 86 and Lemma 24, uniformly at every fixed wrong utility,

$$K_{\Delta,\epsilon}(u) = \frac{1}{2}(d_{\mathcal{H}}(\Delta, u) - \epsilon)_+^2 + O(d^{2s+1}).$$

The lemma includes both the exact-zero and positive-distance regions and localizes every exact global minimizer. Minimization over $u \le 0$ proves equation 91.

Monotonicity follows from nested alternative classes. On the compact positive parameter rectangle, the KL objective has a uniformly bounded gradient in the logit vector, while the Hausdorff distance between $\mathcal{Q}_r^W(\epsilon_1)$ and $\mathcal{Q}_r^W(\epsilon_2)$ is at most $|\epsilon_1 - \epsilon_2|$ in the W norm. Norm equivalence and the elementary value function comparison therefore give the asserted Lipschitz bound.

KL is zero exactly at equality of complete laws, which is equivalent to the matching target belonging to the thickened class. Compactness gives closed attainment, continuity in u gives equality of strict and closed infima, and the characterization by $\mathcal{U}_\Delta^\star$ follows. If the radius is strictly larger, continuity of $T_{\Delta,u}$ as $u \uparrow 0$ moves a boundary-only target strictly inside the tube, producing a strict alias. □

Corollary 26 (Nonaligned-norm critical exponent)**.** *Let $\|\cdot\|_{\mathsf{X}}$ be any norm satisfying, on $\mathbb{R}^J$,*

$$c_-\|g\|_{\mathsf{X}} \le \|g\|_W \le c_+\|g\|_{\mathsf{X}} \qquad (0 < c_- \le c_+ < \infty). \tag{92}$$

Let $\epsilon_{\mathrm{alias}}^{\mathsf{X}}(\Delta)$ be the exact alias radius when the alternative perturbation obeys $\|g\|_{\mathsf{X}} \le \epsilon$. Then

$$\frac{A_{J,r}}{c_+}\Delta^s - K\Delta^{s+1} \le \epsilon_{\mathrm{alias}}^{\mathsf{X}}(\Delta) \le \frac{A_{J,r}}{c_-}\Delta^s + K\Delta^{s+1}. \tag{93}$$

Thus the critical exponent $\epsilon \asymp \Delta^{r+1}$ is invariant under every fixed equivalent norm, although the exact critical constant need not be.

Proof. The norm inequalities imply the ball inclusions $B_W(c_-\epsilon) \subseteq B_{\mathsf{X}}(\epsilon) \subseteq B_W(c_+\epsilon)$. Apply the exact W-radius characterization and solve the two inclusions for the smallest admissible X radius. □

Corollary 27 (Ellipsoidal normal-section formula)**.** *Let $M \succ 0$ and $E_M(\epsilon) = \{g : g^\top M g \le \epsilon^2\}$. Let $N_W = \mathrm{range}(V_r)^{\perp_W}$, P_N^W be the W-orthogonal projector, and define the quotient norm*

$$\|n\|_{M/V}^2 = \inf_{v \in \mathrm{range}(V_r)} (n+v)^\top M (n+v), \qquad n \in N_W. \tag{94}$$

Then, whenever the projected coefficient is interior as in Lemma 23,

$$\mathrm{dist}_W(T, \mathcal{Q}_r^0 + E_M(\epsilon)) = \mathrm{dist}_W\left(P_N^W T, \{n \in N_W : \|n\|_{M/V} \le \epsilon\}\right). \tag{95}$$

The information-aligned positive-part formula is the special case in which the quotient ellipsoid is a W ball in N_W.

Proof. Every g decomposes uniquely as $P_N^W g + v$ with $v \in \mathrm{range}(V_r)$. After the nuisance tangent is reoptimized, only the normal component remains. For a prescribed normal component n, the least M cost among all lifts is exactly equation 94. Taking the W-nearest feasible normal component proves equation 95. □

E.3 Adaptive transcript continuity at a boundary-only zero

Lemma 28 (Finite-horizon adaptive transcript TV bound)**.** *Consider two controlled finite-alphabet experiments with*

$$\eta = \sup_{e \in \mathcal{A}} \|P_m^e - P_{m'}^e\|_{\mathrm{TV}}. \tag{96}$$

Under any common predictable, possibly randomized experiment-selection rule, the laws of the first T actions and observations satisfy

$$\|\mathcal{L}_m(H_T) - \mathcal{L}_{m'}(H_T)\|_{\mathrm{TV}} \le T\eta. \tag{97}$$

Proof. Couple the internal random seed identically. Conditional on a common history, the algorithm uses the same experiment-selection kernel. Couple the next observation maximally; its conditional mismatch probability is at most η. A union bound over T rounds gives a coupling whose transcript mismatch probability is at most $T\eta$, which upper-bounds total variation. Equivalently, the same result follows by induction using contraction of total variation under common Markov kernels. □

Proposition 29 (Uniform correctness fails at an unattained boundary zero)**.** *Let m have a unique decision and let strict wrong-decision models m_k obey*

$$\sup_e \|P^e_{m_k} - P^e_m\|_{\mathrm{TV}} \longrightarrow 0. \tag{98}$$

No procedure can be δ-correct on a class containing $m, m_1, m_2, \ldots$ when $\delta < 1/2$.

Proof. Suppose otherwise. Almost-sure termination and correctness at m imply that, for every $\eta_0 > 0$, some finite T satisfies

$$\mathbb{P}_m(\tau \le T, \widehat{d} = d(m)) \ge 1 - \delta - \eta_0.$$

Choose $\eta_0 < (1 - 2\delta)/4$. Lemma 28 and equation 98 imply, for all sufficiently large k,

$$\mathbb{P}_{m_k}(\tau \le T, \widehat{d} = d(m)) \ge 1 - \delta - 2\eta_0 > \delta.$$

But $d(m_k) \neq d(m)$, so δ-correctness under m_k requires the last probability to be at most δ, a contradiction. □

E.4 Post-processing and the correctly scoped collapse theorem

Proposition 30 (Blackwell monotonicity and equality)**.** *For every parameter-independent Markov kernel A,*

$$\Gamma_{A,\mathcal{M}}(m) \le \Gamma_{\mathcal{M}}(m). \tag{99}$$

For a positive pair (P, Q), equality in $D_{\mathrm{KL}}(AP\|AQ) \le D_{\mathrm{KL}}(P\|Q)$ holds iff the likelihood ratio dP/dQ is measurable with respect to the processed output, equivalently iff the conditional law of the original output given the processed output is the same under P and Q. For deterministic coarsening this means the likelihood ratio is constant on every nonempty fiber.

Proof. Apply the log-sum inequality to each processed output and then sum. Its equality condition is proportionality of the two input mass vectors within every output fiber, which is exactly conditional-distribution sufficiency. Apply the pairwise inequality inside the infimum in equation 78 and then take the supremum over designs. □

Proposition 31 (Order-preserving deterministic low-bin collapse)**.** *Let a deterministic coarsening partition the reported scored categories into K nonempty contiguous bins in increasing score order, optionally sending additional categories to the observed null symbol. If $K \le r + 1$, then, for all sufficiently small $\Delta > 0$, the coarsened complete experiment has an exact strict opposite-decision alias.*

Proof. Let $B_1 < \cdots < B_K$ be the ordered bins and $q_\xi(t) = \sum_{\ell=0}^{K-1} \xi_\ell t^\ell$. The kth reported-bin matching equation at opposite utility $-\Delta$ is

$$F_k(\Delta, \xi) = \sum_{j \in B_k} p_{-\Delta}(j)\sigma(a + q_\xi(y_j)) - \rho \sum_{j \in B_k} p_\Delta(j) = 0.$$

At $(\Delta, \xi) = (0, 0)$, the nuisance Jacobian is

$$[D_\xi F(0,0)]_{k,\ell+1} = \rho(1-\rho) \sum_{j \in B_k} \omega_j y_j^\ell.$$

Its determinant is

$$\rho^K (1-\rho)^K \sum_{j_1 \in B_1, \ldots, j_K \in B_K} \left(\prod_{k=1}^K \omega_{j_k} \right) \prod_{a<b} (y_{j_b} - y_{j_a}),$$

which is strictly positive because every selected score from bin a is below every selected score from bin $b > a$. The analytic implicit-function theorem therefore supplies $\xi(\Delta) \to 0$ solving all K equations. Embed these coefficients in the degree-r nuisance. The reported-bin masses match, and total mass one forces the coarsened observed-null masses to match. □

E.5 APPROXIMATE SHARING: TANGENT COEFFICIENT AND FINITE-GAP RADIUS

For a protocol support $S \subseteq [L]$, let $T_{z,\Delta,u}$ be the exact response-logit target for protocol z and let W_z be its base complete-law metric. Define the closed finite-gap sharing radius

$$\gamma_{\mathrm{alias},S}(\Delta) = \min_{-u_0 \le u \le 0} \min_{q \in \mathcal{Q}_r^0} \max_{z \in S} \|T_{z,\Delta,u} - q\|_{W_z}. \tag{100}$$

The strict radius is the same numerical infimum with $u < 0$, although it need not be attained at equality.

Lemma 32 (Local admissibility of the common protocol center)**.** *Let S be any nonempty support in the finite protocol library. There are $d_0, C > 0$, uniform over S, such that for $-u_0 \le u \le 0$ and $d = \Delta - u \le d_0$, every minimizer*

$$q^\star \in \operatorname*{arg\,min}_{q \in \mathrm{range}(V_r)} \max_{z \in S} \|T_{z,\Delta,u} - q\|_{W_z} \tag{101}$$

satisfies

$$\|q^\star\|_2 \le Cd. \tag{102}$$

Consequently $q^\star = V_r \eta^\star$ with $\|\eta^\star\|_2 \le Cd$, so after reducing d_0 every full-space minimizer belongs to the fixed relative-interior set $V_r B_2(0, h_1) \subset \mathcal{Q}_r^0$. Thus the radius over the actual compact admissible nuisance class in equation 100 equals the full-evaluation-space radius throughout the local region.

Proof. The finite protocol library and compact local rectangle give $\|T_{z,\Delta,u}\|_2 \le C_T d$ uniformly in z. Every W_z is positive definite, and finiteness supplies uniform constants $m\|v\|_2 \le \|v\|_{W_z} \le M\|v\|_2$. The maximum-distance objective in equation 101 is continuous and coercive, so a minimizer exists. The candidate $q = 0$ yields objective at most $MC_T d$. For any fixed $z_0 \in S$, every minimizer therefore satisfies

$$\|q^\star\|_{W_{z_0}} \le \|q^\star - T_{z_0,\Delta,u}\|_{W_{z_0}} + \|T_{z_0,\Delta,u}\|_{W_{z_0}} \le 2MC_T d.$$

Norm equivalence proves equation 102. Because the score-evaluation matrix has full column rank on the declared effective support, $\eta^\star = V_r^\dagger q^\star$ and hence $\|\eta^\star\|_2 \le \|V_r^\dagger\|_{\mathrm{op}} Cd$. Nuisance interiority gives $B_2(0, h_1) \subset \mathcal{H}$ for a fixed $h_1 > 0$; reducing d_0 proves the compact admissibility and equality of the two local radii. □

Theorem 33 (Exact robust tangent coefficient under one-sided protocol sharing error)**.** *Let $x_j = y_j - \bar{y}$, $\nu = \sum_j \omega_j x_j^2$, $\alpha_z = \sqrt{\nu \lambda_z}$, and $b_z = (1-\rho_z)^{-1}$. Give each wrong-decision protocol an independent information-unit deviation e_z satisfying $\|e_z\|_{W_z} \le \tau$. For every design w,*

$$I_\tau(w) = \min_{c \in \mathbb{R}} \sum_z w_z (\alpha_z |c - b_z| - \tau)_+^2. \tag{103}$$

The scalar reduction is simultaneous for all protocols. The square root of the coefficient is one-Lipschitz and the coefficient is nonincreasing in τ.

For two protocols i, j with weights $p, 1-p$,

$$I_\tau(p; i, j) = \frac{p(1-p)\alpha_i^2 \alpha_j^2}{p\alpha_i^2 + (1-p)\alpha_j^2} \left[|b_i - b_j| - \tau \left(\alpha_i^{-1} + \alpha_j^{-1} \right) \right]_+^2. \tag{104}$$

Let $\Theta_{ij} = \alpha_i \alpha_j |b_i - b_j| / (\alpha_i + \alpha_j)$ and $\Theta_\star = \max_{i<j} \Theta_{ij}$. Then

$$\max_w I_\tau(w) = (\Theta_\star - \tau)_+^2. \tag{105}$$

Every pair in $\arg\max_{i<j} \Theta_{ij}$ attains the value with weights $w_i = \alpha_j/(\alpha_i + \alpha_j)$ and $w_j = \alpha_i/(\alpha_i + \alpha_j)$; uniqueness is not claimed when pairs tie.

Proof. Write a common nuisance vector as $q = m\mathbf{1} + cx + h$, where $\langle h, 1 \rangle_\omega = \langle h, x \rangle_\omega = 0$. Since $W_z = \lambda_z \operatorname{diag}(\omega) + \kappa_z \omega\omega^\top$,

$$\|q - b_z x\|_{W_z}^2 = c_{0,z} m^2 + \lambda_z \nu (c - b_z)^2 + \lambda_z \|h\|_\omega^2.$$

Setting $m = h = 0$ weakly decreases every protocol residual simultaneously, which proves the common scalar reduction. Radial projection onto each W_z ball gives equation 103; distance to an expanding ball is one-Lipschitz. Solving the two-point weighted least-squares problem gives equation 104.

For $0 \le \tau \le \bar{\tau}$, every minimizing center belongs to the compact interval

$$C_{\bar{\tau}} = \left[\min_z b_z - \bar{\tau}/\alpha_{\min} - 1, \max_z b_z + \bar{\tau}/\alpha_{\min} + 1\right].$$

On $\Delta([L]) \times C_{\bar{\tau}}$, the objective is affine, hence concave, in w and continuous convex in c. Sion's theorem therefore gives

$$\max_w \min_c \sum_z w_z f_z(c) = \min_c \max_z f_z(c).$$

The intervals $[b_z - t/\alpha_z, b_z + t/\alpha_z]$ have common intersection iff they intersect pairwise. Hence the smallest t is $\max_{i<j} \Theta_{ij} = \Theta_\star$, proving equation 105 and the tie statement. □

Proposition 34 (Finite-gap sharing alias radius)**.** *For every support S with $\Theta(S) > 0$, where*

$$\Theta(S) = \min_{c\in\mathbb{R}} \max_{z\in S} \alpha_z |c - b_z| = \max_{i<j\in S} \Theta_{ij}, \tag{106}$$

there are $\Delta_0, K > 0$ such that

$$|\gamma_{\mathrm{alias},S}(\Delta) - \Theta(S)\Delta| \le K\Delta^2, \qquad 0 < \Delta \le \Delta_0. \tag{107}$$

Below this exact radius the finite-gap profile is positive; above it a strict exact alias exists. At equality, strict attainment is equivalent to attainment of the radius at a strictly negative utility. Vanishing of the tangent coefficient alone does not imply finite-gap aliasing.

Proof. Uniformly over the finite protocol library,

$$T_{z,\Delta,u} = db_z x + O(d^2), \qquad d = \Delta - u.$$

By Lemma 32, the compact statistical radius equals the radius over $\mathrm{range}(V_r)$ throughout the local region. The weighted Chebyshev-radius functional $(T_z)_z \mapsto \inf_{q\in\mathrm{range}(V_r)} \max_z \|T_z - q\|_{W_z}$ is one-Lipschitz in the maximum product norm. Its value at the leading targets $(db_z x)_z$ is exactly $d\Theta(S)$. Thus the fixed-utility compact-class radius is $d\Theta(S) + O(d^2)$. The boundary candidate gives an $O(\Delta)$ upper bound; compact fiber isolation localizes every minimizer and the positive leading coefficient gives $d = \Delta + O(\Delta^2)$. This proves equation 107. Exact matching is equivalent to the target tuple belonging to the product of protocol-deviation balls, giving the stated positive/zero/strict-attainment classification exactly as in Theorem 25. □

The budgets in Theorem 33 are equal in local information units. Equal Euclidean or coefficient-unit budgets generally produce protocol-dependent information radii; the norm-equivalence argument of Corollary 26 gives the corresponding critical-scale bounds.

E.6 Calibration as an independent statistical experiment

Fix $0 < \underline{\rho} < \rho_{\min} < \rho_{\max} < \overline{\rho} < 1$ and assume $\rho_z = \sigma(a_z) \in [\rho_{\min}, \rho_{\max}]$. The wider interval $[\underline{\rho}, \overline{\rho}]$ is the known clipping range; the strict interior margin is used only for the two-sided width comparison below. Before the decision stage, protocol z receives an independent homogeneous pilot

$$S_{z,i} \overset{\mathrm{iid}}{\sim} \mathrm{Bernoulli}(\rho_z), \qquad i = 1, \dots, n_z, \tag{108}$$

constructed so that category-dependent response-nuisance terms are absent. Let $\widehat{\rho}_z = n_z^{-1} \sum_i S_{z,i}$ and

$$e_z(\alpha) = \sqrt{\frac{\log(2L/\alpha)}{2n_z}}. \tag{109}$$

Let

$$\widetilde{A}_z = [\widehat{\rho}_z - e_z, \widehat{\rho}_z + e_z] \cap [\underline{\rho}, \overline{\rho}].$$

If this intersection is empty, set $\widetilde{A}_z = [\underline{\rho}, \overline{\rho}]$ as a measurable full-range fallback. Transform its endpoints by logit to obtain A_z and put $A = \prod_z A_z$. The fallback is irrelevant on the simultaneous coverage event but makes the selected design well defined on every pilot outcome.

Theorem 35 (Finite-sample calibration certificate and total confidence)**.** *The pilot satisfies*

$$\mathbb{P}(a^\star \in A) \geq 1 - \alpha. \tag{110}$$

The robust criterion $I(w; A) = \inf_{a \in A} I_{\text{eff}}(w; a)$ is continuous and concave in w. The minimum-Euclidean-norm maximizer

$$w_{\text{rob}}(A) = \arg\min\{\|w\|_2^2 : w \in \arg\max_{v \in \Delta([L])} I(v; A)\} \tag{111}$$

is unique and Borel measurable as a function of the random interval endpoints. On the coverage event,

$$I_{\text{eff}}(w_{\text{rob}}(A); a^\star) \geq I^\star(A), \tag{112}$$

and therefore

$$\mathbb{P}\big(I_{\text{eff}}(w_{\text{rob}}(A); a^\star) \geq I^\star(A)\big) \geq 1 - \alpha. \tag{113}$$

If the conditionally selected decision procedure is almost surely terminating and has conditional error at most δ whenever $a^\star \in A$, then

$$\mathbb{P}(\textit{decision error}) \leq \alpha + \delta. \tag{114}$$

The reported total sample cost is $N_{\text{pilot}} + \tau_\delta$, where $N_{\text{pilot}} = \sum_z n_z$.

Proof. Hoeffding's inequality gives $\mathbb{P}(|\widehat{\rho}_z - \rho_z| > e_z) \leq \alpha/L$; a union bound and monotonicity of logit prove equation 110. Joint continuity of $I_{\text{eff}}(w; a)$ on the compact product implies continuity of its infimum over A, while the infimum of functions affine in the design representation is concave. The maximizer set is nonempty, compact, and convex. Strict convexity of $\|w\|_2^2$ gives the unique secondary selector. The closed-graph/measurable maximum theorem makes it Borel measurable. On $a^\star \in A$, set inclusion gives $I_{\text{eff}}(w_{\text{rob}}; a^\star) \geq \inf_{a \in A} I_{\text{eff}}(w_{\text{rob}}; a) = I^\star(A)$. Finally, split the error event according to coverage and use the conditional δ guarantee, proving equation 114. □

For row n, let $\ell_{z,n}$ and $u_{z,n}$ be the probability endpoints of $\widetilde{A}_{z,n}$ and define the reported logit half-widths

$$h_{z,n} = \frac{1}{2}\{\text{logit}(u_{z,n}) - \text{logit}(\ell_{z,n})\}, \qquad h_n = \max_z h_{z,n}. \tag{115}$$

Lemma 36 (Two-sided Hoeffding–logit width comparison)**.** *Let $n_{\min,n} = \min_z n_{z,n}$, assume $\log(2L/\alpha_n)/n_{\min,n} \to 0$, and let $\mathcal{E}_n = \bigcap_z \{|\widehat{\rho}_{z,n} - \rho_{z,n}| \leq e_{z,n}\}$. There are fixed $0 < c_h \leq C_h < \infty$ such that on $\mathcal{E}_n$, for all sufficiently large n,*

$$c_h \sqrt{\frac{\log(2L/\alpha_n)}{n_{\min,n}}} \leq h_n \leq C_h \sqrt{\frac{\log(2L/\alpha_n)}{n_{\min,n}}}. \tag{116}$$

Consequently, on the simultaneous coverage event, $h_n = o(|\zeta_n|)$ if and only if

$$\frac{n_{\min,n}\zeta_n^2}{\log(2L/\alpha_n)} \longrightarrow \infty. \tag{117}$$

This equivalence concerns these particular Hoeffding intervals followed by the logit transform; it is not a general minimax necessity statement.

Proof. Let $m_\rho = \min\{\rho_{\min} - \underline{\rho}, \overline{\rho} - \rho_{\max}\} > 0$. The shrinking-width assumption gives $\max_z e_{z,n} \to 0$, so eventually $2\max_z e_{z,n} < m_\rho$. On $\mathcal{E}_n$ this implies $[\widehat{\rho}_{z,n} - e_{z,n}, \widehat{\rho}_{z,n} + e_{z,n}] \subset [\underline{\rho}, \overline{\rho}]$, so clipping and the fallback are inactive. The mean-value theorem gives

$$h_{z,n} = e_{z,n}\ \text{logit}'(\xi_{z,n}) \quad \text{for some } \xi_{z,n} \in [\underline{\rho}, \overline{\rho}].$$

The derivative $\text{logit}'(p) = 1/[p(1-p)]$ has fixed positive lower and finite upper bounds on this compact interval. Since $\max_z e_{z,n} = \sqrt{\log(2L/\alpha_n)/(2n_{\min,n})}$, taking the maximum over z proves equation 116 and the equivalence. □

On $|\zeta_n| \asymp \Delta_n$, the same condition retains the logarithmic factor; the pilot cost is included explicitly in total cost.

E.7 COEFFICIENT-LEVEL QUADRATIC–QUARTIC CROSSOVER

Assume $r = 1$ and $J_{\text{eff}} \geq 3$. Normalize $\sum_j \omega_j = 1$, set

$$\bar{y} = \sum_j \omega_j y_j, \qquad x_j = y_j - \bar{y}, \qquad \nu = \sum_j \omega_j x_j^2, \qquad V_1 = \text{span}\{\mathbf{1}, x\}, \tag{118}$$

and let Π_1^ω be the $L^2(\omega)$ projection onto V_1. Put $q_2 = (I - \Pi_1^\omega)x^{\odot 2} \neq 0$. For $\varsigma \in \{-1, +1\}$ define

$$\rho_\varsigma(\zeta) = \sigma(a + \varsigma\zeta), \tag{119}$$

$$b_\varsigma(\zeta) = (1 - \rho_\varsigma(\zeta))^{-1}, \tag{120}$$

$$\lambda_\varsigma(\zeta) = \rho_\varsigma(\zeta)(1 - \rho_\varsigma(\zeta))^2, \tag{121}$$

$$W_\varsigma(\zeta) = \lambda_\varsigma(\zeta)\,\text{diag}(\omega) + \rho_\varsigma(\zeta)^2(1 - \rho_\varsigma(\zeta))\omega\omega^\top. \tag{122}$$

Let $\rho = \sigma(a)$, $\lambda_0 = \rho(1 - \rho)^2$, and

$$C_{\text{sing}} = \frac{1}{2}\lambda_0\beta_2(a)^2\|q_2\|_\omega^2 = C_{J,1}, \tag{123}$$

$$C_\lambda = \frac{1}{2}\rho^3\nu. \tag{124}$$

For a fixed design $w_+ = w, w_- = 1 - w$, define the exact quadratic coefficient

$$C_w^{(2)}(\zeta) = \frac{\nu}{2}\frac{w(1-w)\lambda_+(\zeta)\lambda_-(\zeta)}{w\lambda_+(\zeta) + (1-w)\lambda_-(\zeta)}(b_+(\zeta) - b_-(\zeta))^2, \tag{125}$$

with value zero at $w = 0, 1$. Put $C_{\text{bal}}(\zeta) = C_{1/2}^{(2)}(\zeta)$ and $C_{\text{opt}}^{(2)}(\zeta) = \max_w C_w^{(2)}(\zeta)$.

Lemma 37 (Exact local quadratic optimizer and even openings)**.** *For $\zeta \neq 0$, the unique quadratic-coefficient optimizer is*

$$w_{\text{quad}}^\star(\zeta) = \frac{\sqrt{\lambda_-(\zeta)}}{\sqrt{\lambda_+(\zeta)} + \sqrt{\lambda_-(\zeta)}}. \tag{126}$$

Moreover,

$$C_{\text{bal}}(\zeta) = C_\lambda\zeta^2 + O(\zeta^4), \tag{127}$$

$$C_{\text{opt}}^{(2)}(\zeta) = C_\lambda\zeta^2 + O(\zeta^4), \tag{128}$$

and $C_{\text{opt}}^{(2)}(\zeta) - C_{\text{bal}}(\zeta) \geq 0$ is $O(\zeta^4)$. The balanced and optimized exact global profiles are even functions of ζ.

Proof. Differentiate equation 125; the unique interior critical point is equation 126, and the endpoint values are zero. Substitution gives

$$C_{\text{opt}}^{(2)}(\zeta) = \frac{\nu}{2}\frac{\lambda_+\lambda_-}{(\sqrt{\lambda_+} + \sqrt{\lambda_-})^2}(b_+ - b_-)^2.$$

Since $b_+ - b_- = 2e^a\zeta + O(\zeta^3)$ and $\lambda_\pm = \lambda_0 + O(\zeta)$, both symmetric coefficients have leading constant $\nu\lambda_0 e^{2a}/2 = \rho^3\nu/2$. Direct subtraction yields

$$C_{\text{opt}}^{(2)} - C_{\text{bal}} = \frac{\nu\lambda_+\lambda_-(b_+ - b_-)^2(\sqrt{\lambda_+} - \sqrt{\lambda_-})^2}{4(\lambda_+ + \lambda_-)(\sqrt{\lambda_+} + \sqrt{\lambda_-})^2}.$$

Protocol exchange maps (ζ, w) to $(-\zeta, 1 - w)$. It leaves the balanced profile invariant and, after maximizing over w, leaves the selectable profile invariant. Thus the globally profiled balanced and optimized values are even; no odd $\zeta\Delta^3$ term can occur in either theorem. □

For $d = \Delta - u$, define the exact target

$$T_{\varsigma,j} = G_{a+\varsigma\zeta}\big(dy_j + \psi(u) - \psi(\Delta)\big), \tag{129}$$

and decompose $T_\varsigma = V_\varsigma + N_\varsigma$ using Π_1^ω and $I - \Pi_1^\omega$.

Lemma 38 (Weighted target expansion and exact orthogonality)**.** *Uniformly for* $0 < \Delta \le d \le d_0$ *and* $|\zeta| \le \zeta_0$*,*

$$\|V_+ - V_- - d(b_+ - b_-)x\|_2 \le K|\zeta|d^2, \tag{130}$$

$$N_\varsigma = \beta_2(a + \varsigma\zeta)d^2 q_2 + R_{N,\varsigma}, \qquad \|R_{N,\varsigma}\|_2 \le Kd^3. \tag{131}$$

For every $v \in V_1$*,* $n \in V_1^{\perp_\omega}$*, and both protocols,*

$$v^\top W_\varsigma(\zeta) n = 0. \tag{132}$$

Proof. Write $R_{\Delta,u} = dx + \kappa_{\Delta,u}\mathbf{1}$ with $|\kappa_{\Delta,u}| \le Kd^2$. Joint Taylor expansion gives $G_{a+\varsigma\zeta}(R) = b_\varsigma R + \beta_2(a+\varsigma\zeta)R^2 + O(|R|^3)$. The linear term is tangent, while the normal projection of R^2 is exactly $d^2 q_2$. Taking the protocol difference gives equation 130; normal projection gives equation 131. Finally, $\omega^\top n = \langle v, n\rangle_\omega = 0$, and the two terms of W_ς therefore both pair v and n to zero. □

Let

$$D_w^2(\Delta, u, \zeta) = \inf_{q \in V_1} \sum_\varsigma w_\varsigma \|q - T_\varsigma\|^2_{W_\varsigma(\zeta)}. \tag{133}$$

Proposition 39 (Explicit weighted target-distance remainder)**.** *Uniformly over* $w \in [0,1]$*,*

$$\left| D_w^2 - 2C_w^{(2)}(\zeta)d^2 - 2C_{\mathrm{sing}}d^4 \right| \le K\left(\zeta^2 d^3 + |\zeta|d^4 + d^5\right). \tag{134}$$

This includes the endpoint designs: there the quadratic coefficient is zero and the one-protocol quartic coefficient equals $C_{\mathrm{sing}} + O(|\zeta|)$*.*

Proof. By equation 132, the profile splits exactly into a tangent quadratic problem and the two normal squared norms. The tangent profile is a uniformly bounded quadratic form of $V_+ - V_-$. Perturbing its leading argument $d(b_+ - b_-)x = O(|\zeta|d)$ by the $O(|\zeta|d^2)$ error in equation 130 changes its value by $O(\zeta^2 d^3)$. Direct minimization of the leading scalar problem gives $2C_w^{(2)}(\zeta)d^2$.

For the normal part, equation 131 and $\|q_2\|^2_{W_\varsigma} = \lambda_\varsigma \|q_2\|^2_\omega$ give

$$d^4 \|q_2\|^2_\omega \sum_\varsigma w_\varsigma \lambda_\varsigma \beta_2 (a + \varsigma\zeta)^2 + O(d^5).$$

The bracket differs from $\lambda_0 \beta_2(a)^2$ by $O(|\zeta|)$ uniformly in w, proving equation 134. No interior lower bound on $w(1-w)$ is used. □

Lemma 40 (Individual control of every sampled crossover residual)**.** *Uniformly over* $w \in [0,1]$*,* $0 < \Delta \le d = \Delta - u \le d_0$*, and* $|\zeta| \le \zeta_0$*, every exact fixed-utility minimizer* $q^\star$ *satisfies*

$$\max_{\varsigma : w_\varsigma > 0} \|q^\star - T_\varsigma\|_2 \le Cd. \tag{135}$$

Consequently,

$$\sum_\varsigma w_\varsigma \|q^\star - T_\varsigma\|_2^3 \le Cd \sum_\varsigma w_\varsigma \|q^\star - T_\varsigma\|_2^2. \tag{136}$$

Proof. Every exact target obeys $\|T_\varsigma\|_2 \le C_T d$ uniformly. The projected candidate used in Proposition 39 has exact KL value at most $C(\zeta^2 d^2 + d^4) \le C' d^2$. Hence the exact minimizer has no larger value. Global fiber localization places it in the common residual chart. Choose a sampled protocol ς_0 with $w_{\varsigma_0} \ge 1/2$ (at an endpoint it is the only sampled protocol). Uniform single-protocol KL coercivity gives

$$c w_{\varsigma_0} \|q^\star - T_{\varsigma_0}\|_2^2 \le K_w(\Delta, u, \zeta) \le C' d^2,$$

so $\|q^\star - T_{\varsigma_0}\|_2 \le Cd$ and therefore $\|q^\star\|_2 \le Cd$. For every other sampled protocol, $\|q^\star - T_\varsigma\|_2 \le \|q^\star\|_2 + \|T_\varsigma\|_2 \le Cd$. This proves equation 135; multiplying the pointwise bound by the squared residual and summing gives equation 136. □

Theorem 41 (Explicit fixed-utility complete-KL expansion)**.** *Let*

$$K_w(\Delta, u, \zeta) = \inf_{q \in \mathcal{Q}_1^0} \sum_{\varsigma} w_\varsigma D_{\mathrm{KL}}(P_{\Delta,0}^{\varsigma,\zeta} \| P_{u,q}^{\varsigma,\zeta}). \tag{137}$$

There are $d_0, \zeta_0, K > 0$ such that, uniformly for $w \in [0,1]$, $0 < \Delta \le d = \Delta - u \le d_0$, and $|\zeta| \le \zeta_0$,

$$\left| K_w(\Delta, u, \zeta) - C_w^{(2)}(\zeta) d^2 - C_{\mathrm{sing}} d^4 \right| \le K \left(\zeta^2 d^3 + |\zeta| d^4 + d^5 \right). \tag{138}$$

Every exact minimizer $q^\star$ satisfies

$$\sum_{\varsigma} w_\varsigma \| q^\star - T_\varsigma \|_2^2 \le K(\zeta^2 d^2 + d^4). \tag{139}$$

Proof. At the exact protocol target, the residual KL Hessian is the positive matrix $W_\varsigma(u, d, \zeta)$ obtained by pulling back the complete multinomial information metric. Uniform positivity gives

$$\| W_\varsigma(u, d, \zeta) - W_\varsigma(\zeta) \|_{\mathrm{op}} \le Kd$$

and a uniform Hessian lower bound on one common residual neighborhood. The integral Taylor formula gives

$$D_{\mathrm{KL}}(P_{\Delta,0}^{\varsigma,\zeta} \| P_{u,T_\varsigma + h}^{\varsigma,\zeta}) = \tfrac{1}{2} h^\top W_\varsigma(u, d, \zeta) h + O(\|h\|_2^3).$$

The minimizer of D_w^2 is an admissible nuisance for small d and has residual $O(|\zeta| d + d^2)$, giving the upper bound. Conversely, that candidate makes the exact profiled KL $O(\zeta^2 d^2 + d^4)$. The isolated-fiber gap forces every exact KL minimizer into the common local chart; uniform Hessian coercivity then proves equation 139. Lemma 40 gives $\max_{w_\varsigma > 0} \|h_\varsigma\|_2 \le Cd$ and therefore the explicit weighted cubic bound equation 136. The metric perturbation has the same $d \sum_\varsigma w_\varsigma \|h_\varsigma\|_2^2$ bound. Hence both are $O(\zeta^2 d^3 + d^5)$. Combine with Proposition 39, whose normal coefficient variation supplies the $|\zeta| d^4$ term. □

Theorem 42 (Complete balanced and selectable crossover)**.** *Let I^{bal} be the balanced compound profile and let $I^{\mathrm{sel}} = \max_{w \in [0,1]} \min_{u \le 0} K_w(\Delta, u, \zeta)$. There are $\Delta_0, \zeta_0, K > 0$ such that*

$$\left| I^{\mathrm{bal}}(\Delta, \zeta) - C_{\mathrm{bal}}(\zeta) \Delta^2 - C_{\mathrm{sing}} \Delta^4 \right| \le K \left(\zeta^2 \Delta^3 + |\zeta| \Delta^4 + \Delta^5 \right), \tag{140}$$

$$\left| I^{\mathrm{sel}}(\Delta, \zeta) - C_{\mathrm{opt}}^{(2)}(\zeta) \Delta^2 - C_{\mathrm{sing}} \Delta^4 \right| \le K \left(\zeta^2 \Delta^3 + |\zeta| \Delta^4 + \Delta^5 \right). \tag{141}$$

Consequently, for $\diamond \in \{\mathrm{bal}, \mathrm{sel}\}$,

$$\left| I^{\diamond}(\Delta, \zeta) - C_\lambda \zeta^2 \Delta^2 - C_{\mathrm{sing}} \Delta^4 \right| \le K \left(\zeta^4 \Delta^2 + \zeta^2 \Delta^3 + |\zeta| \Delta^4 + \Delta^5 \right). \tag{142}$$

For every least-favorable utility of the balanced profile, and for the witness $w_{\mathrm{quad}}^\star(\zeta)$ used in the selectable lower bound,

$$d = \Delta - u = \Delta + O(\Delta^2) \tag{143}$$

uniformly on the joint rectangle. If $\liminf |\zeta_n| / \Delta_n > 0$, every exact full-profile maximizing selectable design converges to $w_{\mathrm{quad}}^\star(\zeta_n)$.

Proof. For the balanced profile, Theorem 41 at $w = 1/2$ and the boundary candidate give an upper bound of order $\zeta^2 \Delta^2 + \Delta^4$. Its lower counterpart and the positive quartic coefficient localize every minimizing d to $d \le M\Delta$. Let $h = d - \Delta \ge 0$, $A = C_{\mathrm{bal}}(\zeta)$, and $B = C_{\mathrm{sing}}$. Comparing the exact minimizer with the boundary candidate gives

$$h(2A\Delta + 4B\Delta^3) \le K(\zeta^2 \Delta^3 + |\zeta| \Delta^4 + \Delta^5).$$

Since $A \asymp \zeta^2$ for small ζ,

$$h \le K\Delta^2 \frac{(\zeta/\Delta)^2 + |\zeta|/\Delta + 1}{(\zeta/\Delta)^2 + 1} \le K' \Delta^2.$$

This proves equation 143; substitution in equation 138 gives equation 140.

For the selectable upper bound, evaluate every fixed-design profile at the boundary and maximize the resulting right side. For the lower bound, use the exact quadratic optimizer $w_{\mathrm{quad}}^\star(\zeta)$, which lies in a

fixed interior interval for small ζ, and repeat the preceding global localization argument. This proves equation 141. Apply equation 128 to obtain equation 142.

For design localization, divide the fixed-design expansion by the two-scale normalizer. Uniformly in w, $C_w^{(2)}(\zeta)/\zeta^2 \to 4C_\lambda w(1-w)$, whose unique maximizer is $1/2$, while $w^\star_{\mathrm{quad}}(\zeta) \to 1/2$. When the regular scale is a nonvanishing fraction of the total scale, any design separated from the exact quadratic optimizer loses a fixed normalized amount, contradicting exact full-profile optimality. □

The singular coefficients in the balanced and selectable theorems are exactly the same at the coincidence base $\zeta = 0$ and agree to leading order in the joint limit. They are not asserted to be identical at fixed nonzero ζ; the $|\zeta|\Delta^4$ term records that variation.

E.8 Validated computation and its logical scope

The numerical supplement records the entire compact nuisance box, every accepted Hessian-subdivision box and its interval Gershgorin lower bound, the Krawczyk images, exact algebraic inputs, derivative cross-checks, and hashes of all generated outputs. The global-optimality argument is not Krawczyk inclusion alone: interval Hessian bounds establish strict convexity on the whole convex box; Krawczyk inclusion establishes a unique interior stationary point; strict convexity then excludes every other interior point and every boundary face. The release includes a clean-run transcript and an executable audit that checks that the reported subdivision boxes cover the declared domain and that every reported lower bound is positive.

This is a validated numerical audit, not a formal proof-assistant development. The analytic theorems above remain the source of the general claims.

F Approximate exclusion with joint protocol uncertainty

This appendix separates two logically different results. For an arbitrary compact *joint* validity set, the local quadratic coefficient is a product-space distance. The interval-center reduction and the closed-form optimal-pair law require the stronger assumption that validity errors are protocol-separable after quotienting the common-shift direction. Keeping this distinction explicit is essential: coordinatewise marginal bounds do not characterize a correlated joint uncertainty set.

F.1 Scientific estimand, reference gauge, and sensitivity

Under imperfect exclusion, protocol z may change the latent judgment itself. The primitive latent object is therefore the vector of protocol-potential utilities $V = (V_1, \dots, V_L)$, not an unqualified scalar "common utility." A scalar decision target becomes well defined only after the scientific population of protocol variants has been specified.

Definition 43 (Deployment-reference utility)**.** Fix a reference distribution $\bar{w} \in \Delta([L])$ with strictly positive coordinates. This fixed scientific reference is chosen before data collection and independently of the experimental design and of the adaptive sampling design. For $V \in \mathbb{R}^L$, define

$$u_{\bar{w}}(V) = \bar{w}^\top V, \qquad \xi_{\bar{w}}(V) = V - u_{\bar{w}}(V)\mathbf{1}. \tag{144}$$

Here $u_{\bar{w}}$ is the deployment-reference utility. The decision studied under imperfect exclusion is $d_{\bar{w}}(V) = \mathrm{sgn}(u_{\bar{w}}(V))$ whenever $u_{\bar{w}}(V) \neq 0$.

Proposition 44 (Reference-gauge representation and observational invariance)**.** *For every $V \in \mathbb{R}^L$ and every fixed $\bar{w} \in \Delta([L])$:*

(i) there is a unique decomposition

$$V = u\mathbf{1} + \xi, \qquad \bar{w}^\top \xi = 0, \tag{145}$$

namely $u = u_{\bar{w}}(V)$ and $\xi = \xi_{\bar{w}}(V)$;

(ii) for another reference $\bar{w}' \in \Delta([L])$,

$$u_{\bar{w}'}(V) = u_{\bar{w}}(V) + (\bar{w}' - \bar{w})^\top \xi_{\bar{w}}(V), \tag{146}$$

$$\xi_{\bar{w}'}(V) = \xi_{\bar{w}}(V) - \big((\bar{w}' - \bar{w})^\top \xi_{\bar{w}}(V)\big)\mathbf{1}; \tag{147}$$

therefore $\xi_{\bar{w},i} - \xi_{\bar{w},j} = V_i - V_j$ is independent of the reference, whereas $u_{\bar{w}}$ generally is not;

(iii) the raw parameterization $V_z = u + \xi_z$ is invariant under

$$(u, \xi) \longmapsto (u + c, \xi - c\mathbf{1}). \tag{148}$$

Consequently every complete protocol law is unchanged by equation 148;

(iv) without a fixed reference condition or an anchor protocol, the sign of an auxiliary scalar u is not a functional of V and is not identified even if the entire vector V were known.

Proof. Set $u = \bar{w}^\top V$ and $\xi = V - u\mathbf{1}$. Since $\bar{w}^\top \mathbf{1} = 1$, one has $\bar{w}^\top \xi = 0$, proving existence. If $V = u'\mathbf{1} + \xi'$ and $\bar{w}^\top \xi' = 0$, applying $\bar{w}^\top$ gives $u' = \bar{w}^\top V = u$, hence $\xi' = \xi$, proving uniqueness.

For a second reference,

$$u_{\bar{w}'}(V) = \bar{w}'^\top (u_{\bar{w}}\mathbf{1} + \xi_{\bar{w}}) = u_{\bar{w}} + (\bar{w}' - \bar{w})^\top \xi_{\bar{w}},$$

because $\bar{w}^\top \xi_{\bar{w}} = 0$; subtracting this scalar multiple of $\mathbf{1}$ from V proves equation 147. Pairwise contrasts are unaffected by adding a common multiple of $\mathbf{1}$.

Every protocol law below depends on the latent part only through $V_z = u + \xi_z$. The transformation equation 148 preserves all V_z and therefore preserves every answered-category and nonresponse probability. Finally, if no gauge or anchor is imposed, any $c \in \mathbb{R}$ gives the same V but replaces u by $u + c$; choosing c on either side of $-u$ can reverse its sign. Thus a protocol-free scalar decision is not defined by the experiment. □

Theorem 45 (Sharp reference sensitivity and classwise sign stability)**.** *Fix $\bar{w} \in \Delta([L])$, a nonempty compact reference family $\mathcal{W}_0 \subseteq \Delta([L])$ containing $\bar{w}$, and a nonempty compact contrast class $\Xi \subseteq \mathsf{G}_{\bar{w}} := \{\xi : \bar{w}^\top \xi = 0\}$. Let*

$$h_\Xi(s) = \max_{\xi \in \Xi} s^\top \xi$$

be its support function and define

$$R_+(\bar{w}; \mathcal{W}_0, \Xi) = \max_{\bar{w}' \in \mathcal{W}_0} h_\Xi(\bar{w} - \bar{w}'), \qquad R_-(\bar{w}; \mathcal{W}_0, \Xi) = \max_{\bar{w}' \in \mathcal{W}_0} h_\Xi(\bar{w}' - \bar{w}). \tag{149}$$

For the potential-utility class $\mathcal{V}(u, \Xi) = \{u\mathbf{1} + \xi : \xi \in \Xi\}$,

$$\min_{\substack{V \in \mathcal{V}(u,\Xi) \\ \bar{w}' \in \mathcal{W}_0}} \bar{w}'^\top V = u - R_+(\bar{w}; \mathcal{W}_0, \Xi), \tag{150}$$

$$\max_{\substack{V \in \mathcal{V}(u,\Xi) \\ \bar{w}' \in \mathcal{W}_0}} \bar{w}'^\top V = u + R_-(\bar{w}; \mathcal{W}_0, \Xi). \tag{151}$$

Consequently:

(i) every $V \in \mathcal{V}(u, \Xi)$ has positive utility under every $\bar{w}' \in \mathcal{W}_0$ if and only if $u > R_+$;

(ii) every $V \in \mathcal{V}(u, \Xi)$ has negative utility under every $\bar{w}' \in \mathcal{W}_0$ if and only if $u < -R_-$;

(iii) at $u = R_+$ or $u = -R_-$, a zero-utility boundary is attained; beyond the corresponding threshold, an opposite or ambiguous reference-indexed decision is attained;

(iv) if $\Xi = -\Xi$, then $R_+ = R_- =: R_{\mathrm{ref}}$, and the sign is uniform over $\mathcal{W}_0$ and Ξ exactly when $|u| > R_{\mathrm{ref}}$.

Proof. For $V = u\mathbf{1} + \xi$ and $\bar{w}' \in \mathcal{W}_0$,

$$\bar{w}'^\top V = u + (\bar{w}' - \bar{w})^\top \xi,$$

using $\bar{w}^\top \xi = 0$. For fixed $\bar{w}'$, minimization over Ξ gives

$$u + \min_{\xi \in \Xi} (\bar{w}' - \bar{w})^\top \xi = u - h_\Xi(\bar{w} - \bar{w}'),$$

and maximization gives $u + h_\Xi(\bar{w}' - \bar{w})$. Taking the remaining minimum or maximum over $\mathcal{W}_0$ proves equation 150–equation 151. Compactness guarantees attainment.

The strict sign statements are equivalent to the minimum being positive or the maximum being negative. At equality the attained extremizer has utility zero; if the inequality is reversed, the same extremizer has the opposite weak sign. If $\Xi = -\Xi$, then $h_\Xi(-s) = h_\Xi(s)$, so the two radii coincide and the final statement follows. □

Corollary 46 (Fixed-model sensitivity to the deployment reference)**.** *For a fixed $V = u_{\bar{w}}\mathbf{1} + \xi_{\bar{w}}$ and any $\bar{w}' \in \Delta([L])$,*

$$u_{\bar{w}'} - u_{\bar{w}} = (\bar{w}' - \bar{w})^\top \xi_{\bar{w}}. \tag{152}$$

If $\xi_{\bar{w}} \in \Xi$, then

$$|u_{\bar{w}'} - u_{\bar{w}}| \le h_\Xi^{\mathrm{sym}}(\bar{w}' - \bar{w}), \qquad h_\Xi^{\mathrm{sym}}(s) = \max\{h_\Xi(s), h_\Xi(-s)\}. \tag{153}$$

For a compact reference family $\mathcal{W}_0$, the sign of this fixed model is positive for every reference exactly when

$$u_{\bar{w}} > \max_{\bar{w}' \in \mathcal{W}_0} (\bar{w} - \bar{w}')^\top \xi_{\bar{w}},$$

and negative for every reference exactly when

$$u_{\bar{w}} < -\max_{\bar{w}' \in \mathcal{W}_0} (\bar{w}' - \bar{w})^\top \xi_{\bar{w}}.$$

Proof. equation 152 is equation 146. The support-function bound follows by applying the definition to both signs. Taking the minimum and maximum over $\mathcal{W}_0$ for the fixed contrast gives the two exact sign conditions. □

Example 47 (Why a reference or anchor is necessary)**.** For two protocols, take $V = (-1, 1)$. The first-protocol reference e_1 gives $u_{e_1} = -1$, while the second-protocol reference e_2 gives $u_{e_2} = 1$. The potential-utility vector and every observation law are unchanged; only the scientific aggregation rule differs. Thus, without a pre-specified reference or an anchor protocol, the sign of a purported protocol-free common utility is not an identified decision target.

F.2 Gauge-fixed row-wise class and induced target-slope uncertainty

With the deployment reference fixed, write $u = u_{\bar{w}}$ and $\xi = \xi_{\bar{w}}$ below. For triangular-array row n, let Ξ_n be a nonempty compact subset of the gauge hyperplane $\{\xi : \bar{w}^\top \xi = 0\}$ and let $\mathcal{H}_n$ be a nonempty compact joint set of protocol-specific response deviations. Define

$$\mathcal{M}_n^{\mathrm{excl}} = \{(u, q, \xi, h) : u \in [-u_0, u_0],\ q \in \mathcal{Q}_r^0,\ \xi \in \Xi_n,\ h \in \mathcal{H}_n\}, \tag{154}$$

with complete protocol law

$$P^z_{u,q,\xi,h}(j) = p_{u+\xi_z}(j)\sigma(a_z + q_j + h_{z,j}), \tag{155}$$

$$P^z_{u,q,\xi,h}(\perp) = 1 - \sum_j P^z_{u,q,\xi,h}(j). \tag{156}$$

The unique-decision subset is

$$\mathcal{M}_{n,\mathrm{uni}}^{\mathrm{excl}} = \{m = (u, q, \xi, h) \in \mathcal{M}_n^{\mathrm{excl}} : u \neq 0\}, \qquad d_n(m) = \mathrm{sgn}(u).$$

For a positive true model m, define

$$\mathrm{Alt}_n^\circ(m) = \{m' \in \mathcal{M}_{n,\mathrm{uni}}^{\mathrm{excl}} : u(m') < 0\}, \tag{157}$$

$$\mathrm{Alt}_n^{\mathrm{cl}}(m) = \{m' \in \mathcal{M}_n^{\mathrm{excl}} : u(m') \le 0\}. \tag{158}$$

A row-n rule is δ_n-correct on this class when

$$\Pr_m(\tau_n < \infty) = 1, \qquad \Pr_m(\widehat{d}_n = d_n(m)) \ge 1 - \delta_n \quad \forall m \in \mathcal{M}_{n,\mathrm{uni}}^{\mathrm{excl}}. \tag{159}$$

Its characteristic information is

$$\Gamma_n(m) = \sup_{w \in \Delta([L])} \inf_{m' \in \mathrm{Alt}_n^\circ(m)} \sum_{z=1}^{L} w_z D_{\mathrm{KL}}(P_m^z \| P_{m'}^z)\,. \tag{160}$$

Thus correctness, strict and closed alternatives, and every information denominator below use the same compact row-wise class and action set.

Throughout this appendix $r \ge 1$, so the shared nuisance contains the affine score direction. Put $x_j = y_j - \bar{y}$, $\nu = \sum_j \omega_j x_j^2$, and

$$\rho_z = \sigma(a_z), \qquad b_z = (1 - \rho_z)^{-1}, \qquad \lambda_z = \rho_z(1 - \rho_z)^2, \qquad \alpha_z = \sqrt{\nu\lambda_z}. \tag{161}$$

The quotient target-slope space is $\mathsf{Q} = \mathbb{R}^L/\operatorname{span}\{\mathbf{1}\}$. We identify it with the fixed gauge section $\mathsf{G}_{\bar{w}} = \{v \in \mathbb{R}^L : \bar{w}^\top v = 0\}$. After converting direct-utility and response-logit perturbations to the common slope coordinate along x, let $\mathcal{U}_\Delta \subset \mathsf{G}_{\bar{w}}$ be the compact *joint* set of normalized target-slope perturbations. The local model assumption is the uniform target representation

$$T_{z,\Delta,v} = \Delta(b_z + v_z)x + R_{z,\Delta,v}, \qquad \max_{z,v\in\mathcal{U}_\Delta} \|R_{z,\Delta,v}\|_{W_z} \le C\Delta^2. \tag{162}$$

The set $\mathcal{U}_\Delta$ is the image of the actual joint raw uncertainty class; it is not replaced by the product of its coordinate projections. The quantitative theorem assumes

$$d_H(\mathcal{U}_\Delta, \mathcal{U}) \le C\Delta \tag{163}$$

for a nonempty compact $\mathcal{U} \subset \mathsf{G}_{\bar{w}}$. The qualitative form assumes only $d_H(\mathcal{U}_\Delta, \mathcal{U}) = o(1)$.

A direct latent-utility contrast of normalized size κ_z shifts the slope by

$$\chi_z = b_z\kappa_z. \tag{164}$$

A response-logit perturbation $g_z = s_z x$ has local information norm $\|g_z\|_{W_z} = \alpha_z|s_z|$, so a radius τ_z in information units corresponds to a slope radius τ_z/α_z. These conversions are used only in the rectangular corollary below.

F.3 General joint-uncertainty coefficient

For a design $w \in \Delta([L])$, define

$$\Phi_{\mathcal{U}}(w) = \inf_{\substack{c\in\mathbb{R}\\ v\in\mathcal{U}}} \sum_{z=1}^{L} w_z\alpha_z^2(c - b_z - v_z)^2. \tag{165}$$

The common center c profiles the shared score direction. Because adding $s\mathbf{1}$ to a representative v and adding s to c leaves the objective unchanged, equation 165 is well defined on the quotient Q.

Lemma 48 (Quotient invariance, compact centers, and set stability)**.** *Let $\mathcal{U}, \mathcal{V} \subset \mathsf{G}_{\bar{w}}$ be nonempty compact sets. Then:*

(i) the value in equation 165 is independent of the chosen quotient section;
(ii) every minimizing center may be chosen in the compact interval

$$C_{\mathcal{U}} = \left[\min_{z,v\in\mathcal{U}}(b_z + v_z), \max_{z,v\in\mathcal{U}}(b_z + v_z)\right];$$

(iii) $w \mapsto \Phi_{\mathcal{U}}(w)$ is continuous and concave on $\Delta([L])$, so a maximizing design exists;
(iv) with $\alpha_{\max} = \max_z \alpha_z$,

$$\left|\sqrt{\Phi_{\mathcal{U}}(w)} - \sqrt{\Phi_{\mathcal{V}}(w)}\right| \le \alpha_{\max} d_H(\mathcal{U}, \mathcal{V}) \tag{166}$$

for every w.

Proof. If $v' = v + s\mathbf{1}$, the change of variable $c' = c + s$ preserves every residual, proving (i). If c is above the upper endpoint of $C_{\mathcal{U}}$, replacing c by that endpoint weakly decreases every squared residual for every $v \in \mathcal{U}$; the lower side is identical. This proves (ii) and gives a common compact domain. For fixed (c, v) the objective is affine in w; its infimum is concave. Joint continuity on the compact domain and the maximum theorem give continuity and attainment, proving (iii).

For (iv), use the weighted seminorm $\|r\|_w^2 = \sum_z w_z\alpha_z^2 r_z^2$. If $v \in \mathcal{U}$ and $v' \in \mathcal{V}$ satisfy $\|v - v'\|_2 \le d_H(\mathcal{U}, \mathcal{V}) + \epsilon$, then $\|v - v'\|_w \le \alpha_{\max}(d_H(\mathcal{U}, \mathcal{V}) + \epsilon)$. Apply the triangle inequality to an ϵ-minimizer of the distance from b to $\{c\mathbf{1} - v : c \in \mathbb{R}, v \in \mathcal{U}\}$, and then exchange $\mathcal{U}$ and $\mathcal{V}$. Letting $\epsilon \downarrow 0$ proves equation 166. □

Proposition 49 (Product-space coefficient under joint approximate validity)**.** *In the weighted product square-root-likelihood experiment, let r_w^{prot} be the first protocol-disagreement normal jet after profiling the shared nuisance tangent. If $\mathcal{B}_{\mathrm{T}}, \mathcal{B}_{\mathrm{A}}$ and $\mathcal{D}_{\mathrm{T}}, \mathcal{D}_{\mathrm{A}}$ are the compact true- and alternative-side response and direct-effect cross-sections, then*

$$K_w^{\mathrm{excl}}(t) = 2t^2 \operatorname{dist}^2\big(r_w^{\mathrm{prot}}, (\mathcal{B}_{\mathrm{A}} + \mathcal{D}_{\mathrm{A}}) - (\mathcal{B}_{\mathrm{T}} + \mathcal{D}_{\mathrm{T}})\big) + O(t^3). \tag{167}$$

No convexity is needed for this coefficient law.

Proof. Square-root likelihoods of the revealed-protocol compound experiment lie in $\bigoplus_z \sqrt{w_z}\,\mathcal{H}_z$, and KL is the weighted sum of protocol KL divergences. The shared response nuisance is a tangent subspace. True- and alternative-side perturbations combine through the Minkowski difference shown above. Theorem 11 with contact order one gives equation 167; its frame-invariance statement makes the coefficient intrinsic. □

Theorem 50 (Joint validity coefficient)**.** *Fix a finite protocol library, a compact intercept range, and the row-wise class equation 154. Assume complete-law positivity, local interiority of the shared nuisance coefficients, the isolated-fiber gap, and the joint target expansion equation 162. Then, uniformly in $w \in \Delta([L])$:*

(i) under equation 163,

$$K_w^{\mathrm{joint}}(\Delta, 0) = \frac{1}{2}\Delta^2 \Phi_{\mathcal{U}}(w) + O(\Delta^3); \tag{168}$$

(ii) if only $d_H(\mathcal{U}_\Delta, \mathcal{U}) = o(1)$ holds, then

$$K_w^{\mathrm{joint}}(\Delta, 0) = \frac{1}{2}\Delta^2 \Phi_{\mathcal{U}}(w) + o(\Delta^2). \tag{169}$$

The coefficient uses the actual joint set $\mathcal{U}$, not the product of its protocolwise marginal intervals.

Proof. For a common degree-r nuisance evaluation, write $q = m\mathbf{1} + sx + h$, where $\langle h, \mathbf{1}\rangle_\omega = \langle h, x\rangle_\omega = 0$. For every protocol,

$$W_z = \lambda_z \operatorname{diag}(\omega) + \kappa_z^{(W)} \omega\omega^\top,$$

and therefore, for a target slope β_z,

$$\|q - \beta_z x\|_{W_z}^2 = c_{0,z} m^2 + \lambda_z \nu (s - \beta_z)^2 + \lambda_z \|h\|_\omega^2. \tag{170}$$

Setting $m = h = 0$ weakly decreases all protocol residuals simultaneously. With $s = \Delta c$ and the target representation equation 162, the infimum of the leading quadratic KL term is exactly

$$\frac{1}{2}\Delta^2 \inf_{c, v \in \mathcal{U}_\Delta} \sum_z w_z \alpha_z^2 (c - b_z - v_z)^2 = \frac{1}{2}\Delta^2 \Phi_{\mathcal{U}_\Delta}(w).$$

The zero nuisance and any fixed $v \in \mathcal{U}_\Delta$ give an $O(\Delta^2)$ upper bound. The outer-fiber gap localizes every exact minimizer to the common regular chart. Uniform Hessian coercivity then gives $O(\Delta)$ sampled response-logit residuals and an $O(\Delta)$ common nuisance evaluation. Consequently the cubic KL remainder and the target error in equation 162 contribute $O(\Delta^3)$ uniformly in w, including endpoint designs (where only the sampled protocols matter). Thus

$$K_w^{\mathrm{joint}}(\Delta, 0) = \frac{1}{2}\Delta^2 \Phi_{\mathcal{U}_\Delta}(w) + O(\Delta^3). \tag{171}$$

Lemma 48 and $d_H(\mathcal{U}_\Delta, \mathcal{U}) = O(\Delta)$ imply $\Phi_{\mathcal{U}_\Delta}(w) = \Phi_{\mathcal{U}}(w) + O(\Delta)$ uniformly, which proves (i). Under qualitative convergence the same lemma gives $\Phi_{\mathcal{U}_\Delta}(w) = \Phi_{\mathcal{U}}(w) + o(1)$, proving (ii). □

Corollary 51 (Joint-validity one-sided decision transfer)**.** *For $u \le 0$, let $K_w^{\mathrm{joint}}(\Delta, u)$ be the fixed-utility profile on the same row-wise joint-validity class as Theorem 50, and define the strict wrong-decision profile*

$$I_w^{\mathrm{joint}}(\Delta) = \inf_{u<0} K_w^{\mathrm{joint}}(\Delta, u). \tag{172}$$

Assume that, uniformly for $w \in \Delta([L])$ and localized $u \le 0$, there are continuous normal distances $d_w(u)$ such that, with $d = \Delta - u$,

$$K_w^{\mathrm{joint}}(\Delta, u) = 2 d_w(u)^2 d^2 + O(d^3), \tag{173}$$

$$d_w(u)^2 = d_w(0)^2 + O(|u|), \qquad 4 d_w(0)^2 = \Phi_{\mathcal{U}}(w). \tag{174}$$

Assume also that strict alternatives are dense in the closed boundary slice and that every remote strict wrong-decision fiber is uniformly KL-separated. If

$$\Psi_{\mathcal{U}} := \max_{w \in \Delta([L])} \Phi_{\mathcal{U}}(w) > 0, \tag{175}$$

then, uniformly for w with $\Phi_{\mathcal{U}}(w) \geq \Psi_{\mathcal{U}}/2$, every minimizer of the closed wrong-decision profile and every $O(\Delta^3)$-near-minimizer of the strict profile satisfies

$$-u_{\Delta,w} = O(\Delta^2). \tag{176}$$

Consequently,

$$I_w^{\mathrm{joint}}(\Delta) = \frac{1}{2}\Phi_{\mathcal{U}}(w)\Delta^2 + O(\Delta^3), \tag{177}$$

$$I^{\star}_{\mathrm{joint}}(\Delta) := \max_{w \in \Delta([L])} I_w^{\mathrm{joint}}(\Delta) = \frac{1}{2}\Psi_{\mathcal{U}}\Delta^2 + O(\Delta^3). \tag{178}$$

Every row-wise procedure that is δ-correct on the same joint class therefore obeys

$$\mathbb{E}[\tau_\delta] \geq \frac{\mathrm{kl}(1-\delta,\delta)}{\frac{1}{2}\Psi_{\mathcal{U}}\Delta^2 + O(\Delta^3)}. \tag{179}$$

Thus the Chebyshev value becomes strict wrong-decision information only after this one-sided transfer; without the displayed assumptions it remains an optimized fixed-boundary coefficient.

Proof. Let $\mathcal{W}_+ = \{w : \Phi_{\mathcal{U}}(w) \geq \Psi_{\mathcal{U}}/2\}$. It is compact, and equation 174 gives $d_w(0)^2 \geq \Psi_{\mathcal{U}}/8$ on $\mathcal{W}_+$. The hypotheses of Proposition 14 therefore hold uniformly on $\mathcal{W}_+$ with contact order one. In particular, every closed-profile minimizer is localized at $-u = O(\Delta^2)$.

For completeness, the same localization is uniform for strict near-minimizers. Write $h = -u \geq 0$. The strict candidate $h = \Delta^2$ gives

$$I_w^{\mathrm{joint}}(\Delta) \leq K_w^{\mathrm{joint}}(\Delta, -\Delta^2) = 2d_w(0)^2\Delta^2 + O(\Delta^3) \tag{180}$$

uniformly in w. Remote separation first confines every $O(\Delta^3)$-near-minimizer to the local chart. Comparing the normal form with equation 180 then gives $\Delta + h = O(\Delta)$. On $\mathcal{W}_+$, equation 174 and the identity $(\Delta+h)^2 - \Delta^2 = 2\Delta h + h^2$ yield

$$K_w^{\mathrm{joint}}(\Delta, -h) - 2d_w(0)^2\Delta^2 \geq c\Delta h - C(\Delta^3 + h\Delta^2)$$

for constants independent of w. For sufficiently small Δ, any $O(\Delta^3)$-near-minimizer therefore has $h = O(\Delta^2)$, proving equation 176. Strict density and continuity identify the strict infimum with the closed boundary value. Substitution in equation 173 proves equation 177.

It remains to optimize over designs. If $w \notin \mathcal{W}_+$, equation 180 is at most $\Psi_{\mathcal{U}}\Delta^2/4 + O(\Delta^3)$. For $w^\star \in \arg\max_w \Phi_{\mathcal{U}}(w)$, equation 177 equals $\Psi_{\mathcal{U}}\Delta^2/2 + O(\Delta^3)$. Hence every maximizing design lies in $\mathcal{W}_+$ for small Δ, and maximizing the uniform fixed-design expansion proves equation 178. Finally, Theorem 22 applied to the identical row-wise class gives equation 179. □

F.4 Exact finite-gap radius for the same joint class

The tangent set $\mathcal{U}$ records the first-order image of the joint raw validity class. The exact finite-gap statement must retain the same coupling. For a protocol support $S \subseteq [L]$, let $\mathfrak{T}_{S,\Delta,u}(\gamma)$ be the compact set of vectors of exact response-logit matching targets induced by the joint direct-effect and response perturbations in the same row-wise class at budget γ. Let

$$\mathfrak{D}_S^0 = \{(q)_{z\in S} : q \in \mathcal{Q}_r^0\}$$

be the diagonal shared-nuisance set. We assume that the exact target correspondence is Hausdorff-continuous in (u,γ), is nested in γ, and has strict interior nesting in its common affine hull:

$$0 \leq \gamma' < \gamma \quad \Longrightarrow \quad \mathfrak{T}_{S,\Delta,u}(\gamma') \subset \operatorname{relint} \mathfrak{T}_{S,\Delta,u}(\gamma). \tag{181}$$

Norm balls, ellipsoids, polytopes with the origin in relative interior, and the rectangular protocol budgets below satisfy this condition. The normalized first-order graphical limit of this exact correspondence is the same joint set $\mathcal{U}_S$ used in Theorem 50.

Define the closed joint alias radius

$$\gamma^{\mathcal{U}}_{\mathrm{alias},S}(\Delta) = \inf\left\{\gamma \geq 0 : \begin{array}{l}\exists\, u \in [-u_0, 0],\ \mathbf{T} \in \mathfrak{T}_{S,\Delta,u}(\gamma),\\ \exists\, \mathbf{q} \in \mathfrak{D}_S^0 \text{ with } \mathbf{T} = \mathbf{q}\end{array}\right\}. \tag{182}$$

The strict radius is the same infimum with $u < 0$. Let $\mathcal{M}^{\mathrm{joint}}_{S,\Delta}(\gamma)$ denote the compact slice of equation 154 whose jointly admissible raw validity perturbations generate $\mathfrak{T}_{S,\Delta,u}(\gamma)$, and set

$$\mathrm{Alt}^{\mathrm{cl}}_{S,\gamma}(m_\Delta) = \{m' \in \mathcal{M}^{\mathrm{joint}}_{S,\Delta}(\gamma) : u(m') \le 0\},$$
$$\mathrm{Alt}^{\circ}_{S,\gamma}(m_\Delta) = \{m' \in \mathcal{M}^{\mathrm{joint}}_{S,\Delta}(\gamma) : u(m') < 0\}.$$

For the fixed true model m_Δ, define the support-level closed and strict profiles on this same joint class by

$$\mathcal{I}^{\mathrm{cl}}_{\gamma,S}(\Delta) = \inf_{m' \in \mathrm{Alt}^{\mathrm{cl}}_{S,\gamma}(m_\Delta)} \max_{z \in S} D_{\mathrm{KL}}(P^z_{m_\Delta} \| P^z_{m'}), \tag{183}$$

$$\mathcal{I}^{\circ}_{\gamma,S}(\Delta) = \inf_{m' \in \mathrm{Alt}^{\circ}_{S,\gamma}(m_\Delta)} \max_{z \in S} D_{\mathrm{KL}}(P^z_{m_\Delta} \| P^z_{m'}). \tag{184}$$

Here the alternative sets are the $u \le 0$ and $u < 0$ slices of the same compact raw class whose exact target image is $\mathfrak{T}_{S,\Delta,u}(\gamma)$.

Theorem 52 (Joint finite-gap radius and strict/closed topology)**.** *Assume complete-law positivity, compactness of the row-wise class, continuity of the exact likelihood and target maps, and equation 181. The radius in equation 182 is an extended nonnegative number. If it is finite, the infimum is attained and the strict and closed alias radii have the same numerical value. If it is infinite, the closed and strict profiles are positive for every finite budget. In the finite case:*

(i) if $\gamma < \gamma^{\mathcal{U}}_{\mathrm{alias},S}(\Delta)$, then $\mathcal{I}^{\mathrm{cl}}_{\gamma,S}(\Delta) > 0$ and hence also $\mathcal{I}^{\circ}_{\gamma,S}(\Delta) > 0$;
(ii) if $\gamma > \gamma^{\mathcal{U}}_{\mathrm{alias},S}(\Delta)$, an attained strict opposite-decision exact alias exists;
(iii) at equality, the closed profile is zero and the strict profile has infimum zero. A strict exact alias is attained at equality if and only if some radius-minimizing configuration has $u < 0$.

At a boundary-only equality point, no almost-surely terminating rule can be uniformly δ-correct for any $\delta < 1/2$ on the same row-wise class. No pairwise or interval representation is asserted for a general correlated joint uncertainty set.

Proof. For a finite candidate radius, restrict budgets to a compact interval containing a minimizing sequence. Compactness and Hausdorff continuity make the feasibility correspondence in equation 182 closed and give attainment. If no finite budget is feasible, the radius is infinite and the compact true-law and wrong-decision law sets remain disjoint for every finite budget, so their KL separation is positive. Hence it remains to treat the finite case. Exact inverse-response matching says that all complete protocol laws agree if and only if an exact target vector belongs to the diagonal set $\mathfrak{D}^0_S$. Thus the closed profile is zero exactly at or above the closed radius.

Let $\gamma_\star$ be the closed radius. If a minimizing configuration has $u < 0$, it is already strict. Otherwise take $u_k \uparrow 0$ with $u_k < 0$. Continuity of the exact target correspondence moves the minimizing target by $o(1)$. For every $\epsilon > 0$, strict interior nesting places a relative neighborhood of the radius-$\gamma_\star$ target inside the budget $\gamma_\star + \epsilon$. Hence, for all sufficiently large k, a strict configuration at u_k is feasible with budget $\gamma_\star + \epsilon$. The strict radius is therefore at most $\gamma_\star + \epsilon$ for every $\epsilon > 0$ and at least the closed radius, so the two numerical radii agree.

If $\gamma < \gamma_\star$, the compact true-law set and compact closed wrong-decision law set are disjoint. Positivity and continuity of KL then give a strictly positive minimum, proving (i). If $\gamma > \gamma_\star$, equality of strict and closed radius infima supplies a strict feasible configuration at some budget below γ; monotonicity in the budget gives an attained strict alias, proving (ii).

At equality, compactness gives a closed alias $\mathbf{T}_\star = \mathbf{q}_\star$ at some radius-minimizing utility $u_\star \le 0$. If $u_\star < 0$, the alias is strict. Suppose instead that every radius minimizer has $u_\star = 0$. Choose $u_k < 0$ with $u_k \uparrow 0$. Hausdorff continuity of $\mathfrak{T}_{S,\Delta,u}(\gamma_\star)$ at the *fixed budget* $\gamma_\star$ gives $\mathbf{T}_k \in \mathfrak{T}_{S,\Delta,u_k}(\gamma_\star)$ with $\mathbf{T}_k \to \mathbf{T}_\star = \mathbf{q}_\star$. Combine the strict raw model inducing $\mathbf{T}_k$ with the admissible shared nuisance $\mathbf{q}_\star \in \mathfrak{D}^0_S$. The resulting protocolwise response-logit residuals converge to zero. Uniform positivity and continuity of the complete likelihood therefore imply

$$\max_{z \in S} D_{\mathrm{KL}}(P^z_{m_\Delta} \| P^z_{m_k}) \longrightarrow 0.$$

Thus the strict profile has infimum zero at the exact budget $\gamma_\star$, without passing to $\gamma_\star + \epsilon$. Strict attainment at equality is exactly the existence of a radius minimizer with $u < 0$, proving (iii).

Finally, at a boundary-only zero, strict wrong-decision laws approach the true law uniformly over S. Lemma 28, almost-sure termination, and the usual finite-horizon truncation argument contradict uniform δ-correctness whenever $\delta < 1/2$. □

F.5 Rectangular specialization and the exact pair law

Suppose now that, after quotienting the common shift, validity errors are protocol-separable. Put

$$r_z = \chi_z + \frac{\tau_z}{\alpha_z}, \qquad \mathcal{U}_{\mathrm{rect}} = \prod_{z=1}^{L} [-r_z, r_z] \mod \mathrm{span}\{\mathbf{1}\}, \tag{185}$$

and define

$$\mathcal{J}_z = [b_z - r_z, b_z + r_z] = \left[b_z - \chi_z - \frac{\tau_z}{\alpha_z}, b_z + \chi_z + \frac{\tau_z}{\alpha_z} \right]. \tag{186}$$

Corollary 53 (Rectangular interval reduction and optimal-pair design)**.** *For the quotient-rectangular set equation 185,*

$$\Phi_{\mathcal{U}_{\mathrm{rect}}}(w) = \min_{c \in \mathbb{R}} \sum_{z=1}^{L} w_z \alpha_z^2 \operatorname{dist}^2(c, \mathcal{J}_z). \tag{187}$$

For two protocols i, j with weights $p, 1-p$, let

$$g_{ij} = [|b_i - b_j| - r_i - r_j]_+ . \tag{188}$$

Then

$$\Phi_{\mathcal{U}_{\mathrm{rect}}}(p; i, j) = \frac{p(1-p)\alpha_i^2 \alpha_j^2}{p\alpha_i^2 + (1-p)\alpha_j^2} \, g_{ij}^2. \tag{189}$$

The rectangular maximizing-pair allocation is

$$p_{ij}^\star = \frac{\alpha_j}{\alpha_i + \alpha_j}, \qquad 1 - p_{ij}^\star = \frac{\alpha_i}{\alpha_i + \alpha_j}, \tag{190}$$

and its value is $(\Theta_{ij}^{\mathrm{excl}})^2$, where

$$\Theta_{ij}^{\mathrm{excl}} = \frac{\alpha_i \alpha_j}{\alpha_i + \alpha_j} g_{ij}. \tag{191}$$

For the full rectangular library,

$$\max_{w \in \Delta([L])} \Phi_{\mathcal{U}_{\mathrm{rect}}}(w) = \left(\Theta_\star^{\mathrm{excl}}\right)^2, \qquad \Theta_\star^{\mathrm{excl}} = \max_{i<j} \Theta_{ij}^{\mathrm{excl}}. \tag{192}$$

Every rectangular maximizing pair with the weights equation 190 is optimal; ties are allowed. In particular,

$$\Theta_{ij}^{\mathrm{excl}} = \left[\underbrace{\frac{\alpha_i \alpha_j}{\alpha_i + \alpha_j} |b_i - b_j|}_{\textit{nominal separation}} - \underbrace{\frac{\alpha_i \alpha_j}{\alpha_i + \alpha_j} (\chi_i + \chi_j)}_{\textit{direct-effect loss}} - \underbrace{\frac{\alpha_j \tau_i + \alpha_i \tau_j}{\alpha_i + \alpha_j}}_{\textit{response-sharing loss}} \right]_+ . \tag{193}$$

No pair-support conclusion is asserted for a general correlated $\mathcal{U}$.

Proof. For fixed c, minimization over the product box is coordinatewise and gives equation 187. For two disjoint intervals, let g_{ij} be their distance and let $s \in [0, g_{ij}]$ be the center's distance from the nearer endpoint of $\mathcal{J}_i$. Minimizing

$$p\alpha_i^2 s^2 + (1-p)\alpha_j^2 (g_{ij} - s)^2$$

over s proves equation 189; calculus in p proves equation 190 and equation 191. Overlapping intervals give zero on both sides.

Restrict c to the compact interval containing all $\mathcal{J}_z$. Sion's theorem applies because the objective is affine in w and convex in c:

$$\max_w \min_c \sum_z w_z \alpha_z^2 \operatorname{dist}^2(c, \mathcal{J}_z) = \min_c \max_z \alpha_z^2 \operatorname{dist}^2(c, \mathcal{J}_z).$$

The right side is at most t^2 exactly when the expanded intervals $\mathcal{J}_z + [-t/\alpha_z, t/\alpha_z]$ have a common point. Intervals on the line have a common point iff every pair intersects. Pair (i, j) intersects exactly when

$$t \geq \frac{\alpha_i \alpha_j}{\alpha_i + \alpha_j} \operatorname{dist}(\mathcal{J}_i, \mathcal{J}_j).$$

Taking the largest threshold proves equation 192; the pairwise construction proves attainment and the tie statement. Expanding $r_i + r_j$ gives equation 193. □

Example 54 (Marginals do not determine the joint coefficient)**.** Take three protocols with $b = (2, 2, 3)$ and $\alpha_z = 1$. Begin with

$$\widetilde{S}_{12} = \{(t, t, t+1) : t \in [-1, 0]\},$$
$$\widetilde{S}_{13} = \{(t, t+1, t) : t \in [-1, 0]\},$$
$$\widetilde{S}_{23} = \{(t+1, t, t) : t \in [-1, 0]\},$$

put $v_0 = (0, 0, 1)$, and define $S_{ij} = \widetilde{S}_{ij} - v_0$ and $\mathcal{U}_{\mathrm{corr}} = S_{12} \cup S_{13} \cup S_{23}$ in the quotient. The nominal perturbation 0 belongs to $\mathcal{U}_{\mathrm{corr}}$. Its coordinate projections are $[-1, 1], [-1, 1], [-2, 0]$, exactly those of the rectangular envelope $\mathcal{U}_{\mathrm{box}} = [-1, 1] \times [-1, 1] \times [-2, 0]$. The rectangular relaxation reports zero for every design: $v = -v_0$ and $c = 2$ give $b + v = 2\mathbf{1}$.

Because $b + v = 2\mathbf{1} + \widetilde{v}$ for $v = \widetilde{v} - v_0$, common shifts cancel and direct minimization over the correlated set gives

$$\Phi_{\mathcal{U}_{\mathrm{corr}}}(w) = \min_{z \in \{1,2,3\}} w_z(1 - w_z). \tag{194}$$

Indeed, on S_{23} only coordinate one is displaced from the common level, and minimizing over (c, t) gives $w_1(1 - w_1)$; the other two segments are cyclic permutations. Every design supported on at most two protocols has value zero. In contrast, $w = (1/3, 1/3, 1/3)$ has value $2/9$, and this is the maximum because $\min_z w_z \leq 1/3$ and $s(1-s)$ is increasing on $[0, 1/2]$. Thus two sets with identical protocolwise marginal intervals have different robust coefficients, and the optimal-pair conclusion fails for the correlated set.

Corollary 55 (Rectangular quadratic restoration)**.** *Assume the quotient-rectangular class and $\Theta_\star^{\mathrm{excl}} > 0$. Under quantitative graphical convergence, every least-favorable wrong utility satisfies $\Delta - u = \Delta + O(\Delta^2)$ and*

$$I^\star_{\mathrm{excl}}(\Delta) = \frac{1}{2}(\Theta_\star^{\mathrm{excl}})^2 \Delta^2 + O(\Delta^3). \tag{195}$$

Under qualitative convergence, $(\Delta - u)/\Delta \to 1$ and the remainder is $o(\Delta^2)$. Every row-wise procedure correct on the same rectangular subclass of $\mathcal{M}_n^{\mathrm{excl}}$ obeys

$$\mathbb{E}[\tau_\delta] \geq \frac{\mathrm{kl}(1 - \delta, \delta)}{\frac{1}{2}(\Theta_\star^{\mathrm{excl}})^2 \Delta^2 + O(\Delta^3)} \tag{196}$$

in the quantitative case. A zero tangent coefficient is not by itself a finite-gap nonidentification statement.

Proof. For a pair (i, j), write

$$A_{ij} = \frac{\alpha_i \alpha_j}{\alpha_i + \alpha_j} |b_i - b_j|, \qquad B_{ij} = \frac{\alpha_i \alpha_j}{\alpha_i + \alpha_j} \left(\chi_i + \chi_j + \frac{\tau_i}{\alpha_i} + \frac{\tau_j}{\alpha_j} \right).$$

At fixed $d = \Delta - u$, the leading square root is $[dA_{ij} - \Delta B_{ij}]_+$. Choose a pair attaining the positive boundary margin $A_{ij} - B_{ij} > 0$. If $d = \Delta + s_d$, then

$$[dA_{ij} - \Delta B_{ij}]^2 - [\Delta(A_{ij} - B_{ij})]^2 = A_{ij} s_d \{2\Delta(A_{ij} - B_{ij}) + A_{ij} s_d\}.$$

The boundary candidate gives an $O(\Delta^2)$ upper bound. Remote-fiber localization and coercivity first give $d = O(\Delta)$; the quantitative fixed- utility remainder is then $O(\Delta^3)$. Optimality relative to the boundary candidate yields $s_d = O(\Delta^2)$. Substitution and Corollary 53 prove equation 195. The qualitative argument replaces $O(d^3)$ by $o(d^2)$ and gives $s_d = o(\Delta)$. The lower bound is Theorem 22 on the identical row-wise subclass. □

F.6 RECTANGULAR FINITE-GAP EXCLUSION RADIUS

For the exact scored experiment, let $\psi(v) = \log \sum_j \omega_j e^{v y_j}$. At true gap $\Delta > 0$ and wrong utility $u \le 0$, allow independently bounded protocol-specific direct perturbations $\xi_z^{\mathrm{T}}, \xi_z^{\mathrm{A}}$ with

$$|\xi_z^{\mathrm{T}}| \le \eta_z^{\mathrm{T}}(\Delta), \qquad |\xi_z^{\mathrm{A}}| \le \eta_z^{\mathrm{A}}(\Delta), \qquad \frac{\eta_z^{\mathrm{T}}(\Delta) + \eta_z^{\mathrm{A}}(\Delta)}{\Delta} = \kappa_z + O(\Delta). \tag{197}$$

These protocolwise bounds are the finite-gap rectangular specialization. The effective latent utilities are $v_z^{\mathrm{T}} = \Delta + \xi_z^{\mathrm{T}}$ and $v_z^{\mathrm{A}} = u + \xi_z^{\mathrm{A}}$. Equality of answered-category probabilities requires the alternative response-logit target

$$T^{\mathrm{excl}}_{z,\Delta,u,\xi^{\mathrm{T}},\xi^{\mathrm{A}},j} = G_{a_z}(R_{z,j}), \tag{198}$$

$$R_{z,j} = (v_z^{\mathrm{T}} - v_z^{\mathrm{A}}) y_j + \psi(v_z^{\mathrm{A}}) - \psi(v_z^{\mathrm{T}}). \tag{199}$$

For support S, define the minimum common alternative-side response radius needed for exact matching:

$$\gamma^{\mathrm{excl}}_{\mathrm{alias},S}(\Delta) = \min_{-u_0 \le u \le 0} \min_{\xi^{\mathrm{T}},\xi^{\mathrm{A}}} \min_{q \in \mathcal{Q}_r^0} \max_{z \in S} \|T^{\mathrm{excl}}_{z,\Delta,u,\xi^{\mathrm{T}},\xi^{\mathrm{A}}} - q\|_{W_z}. \tag{200}$$

The strict radius uses $u < 0$.

Lemma 56 (Uniform exact-target expansion and compact center)**.** *Uniformly over the rectangular direct-effect sets and localized utilities $d = \Delta - u = O(\Delta)$,*

$$T^{\mathrm{excl}}_{z,\Delta,u,\xi^{\mathrm{T}},\xi^{\mathrm{A}}} = b_z(d + \xi_z^{\mathrm{T}} - \xi_z^{\mathrm{A}})x + O(\Delta^2). \tag{201}$$

Every full-evaluation-space minimizer of the common-center problem in equation 200 has Euclidean norm $O(d + \Delta)$ and, for small Δ, lies in one fixed relative-interior subset of $\mathcal{Q}_r^0$.

Proof. Taylor expansion of ψ gives

$$R_{z,j} = (d + \xi_z^{\mathrm{T}} - \xi_z^{\mathrm{A}})(y_j - \bar{y}) + O((d + |\xi_z^{\mathrm{T}}| + |\xi_z^{\mathrm{A}}|)^2).$$

Since $G'_{a_z}(0) = b_z$, equation 201 follows. The zero center is a candidate of value $O(d+\Delta)$. Uniform norm equivalence and the triangle inequality bound every minimizing center by the same order. Full column rank of the visible evaluation map converts this to a coefficient bound, and nuisance interiority places the minimizer in the compact statistical class. □

Define the rectangular direct-effect separation constant

$$\Theta_S^{\mathrm{dir}}(\kappa) = \max_{i<j \in S} \frac{\alpha_i \alpha_j}{\alpha_i + \alpha_j} \left[|b_i - b_j| - b_i\kappa_i - b_j\kappa_j \right]_+ . \tag{202}$$

Theorem 57 (Rectangular finite-gap exclusion radius and topology)**.** *There are $\Delta_0, K > 0$ such that*

$$\left| \gamma^{\mathrm{excl}}_{\mathrm{alias},S}(\Delta) - \Theta_S^{\mathrm{dir}}(\kappa)\Delta \right| \le K\Delta^2, \qquad 0 < \Delta \le \Delta_0. \tag{203}$$

The closed and strict radius values have the same numerical infimum. For a common one-sided alternative response budget γ:

- *(i) if $\gamma < \gamma^{\mathrm{excl}}_{\mathrm{alias},S}(\Delta)$, the compact finite-gap profile is positive;*
- *(ii) if $\gamma > \gamma^{\mathrm{excl}}_{\mathrm{alias},S}(\Delta)$, an attained strict opposite-decision alias exists;*
- *(iii) at equality, the closed profile is zero; the strict infimum is attained exactly when some radius minimizer has $u < 0$.*

At a boundary-only equality point, uniform δ-correctness is impossible for every $\delta < 1/2$ by adaptive transcript continuity. If $\Theta_S^{\mathrm{dir}}(\kappa) = 0$, then the exact radius is $O(\Delta^2)$.

Proof. Let $d = \Delta - u$ and $\lambda = \Delta/d$. Lemma 56 and the rectangular pair formula give the fixed-utility leading radius

$$d \max_{i<j \in S} \frac{\alpha_i \alpha_j}{\alpha_i + \alpha_j} \left[|b_i - b_j| - \lambda(b_i\kappa_i + b_j\kappa_j) \right]_+ ,$$

with uniform $O(\Delta^2)$ error. The boundary candidate is $O(\Delta)$. Remote-fiber localization gives $d = O(\Delta)$. If the boundary constant is positive, an attaining pair increases its leading radius by

$A_{ij}(d-\Delta)$, so comparison with the boundary candidate yields $d = \Delta + O(\Delta^2)$. If the boundary constant vanishes, the boundary candidate is already $O(\Delta^2)$. This proves equation 203.

Exact matching is equivalent to the existence of a common $q \in \mathcal{Q}_r^0$ and protocol deviations of W_z-norm at most γ. Compactness gives closed attainment. Strict and closed numerical infima agree by continuity and density of $u < 0$ in $u \le 0$. With strict slack, a boundary minimizer can be perturbed to $u < 0$ while retaining feasibility. At equality, strict attainment is therefore equivalent to a negative minimizing utility. A boundary-only zero implies impossibility by Lemma 28 and almost-sure termination. □

Example 58 (Distinct protocols can lose rectangular restoration)**.** For two protocols i, j, if $|b_i - b_j| \le b_i\kappa_i + b_j\kappa_j$, then their direct-effect target intervals overlap and $\Theta_{ij}^{\mathrm{excl}} = 0$ even when $a_i \ne a_j$. Theorem 57 then gives $\gamma_{\mathrm{alias},\{i,j\}}^{\mathrm{excl}}(\Delta) = O(\Delta^2)$; any positive linear-scale alternative response error eventually creates a strict exact alias. Randomized assignment alone therefore does not preserve the protocol contrast.

F.7 Calibration uncertainty

Corollary 59 (Confidence-set design for a joint validity set)**.** *Let an independent pilot produce a simultaneous intercept rectangle A with $\Pr(a^\star \in A) \ge 1-\alpha$. Suppose $a \mapsto \mathcal{U}(a)$ is a compact-valued continuous joint target-slope correspondence. Define*

$$\Phi(w; A, \mathcal{U}) = \inf_{a \in A} \Phi_{\mathcal{U}(a)}(w; a), \qquad \Phi^\star(A, \mathcal{U}) = \max_w \Phi(w; A, \mathcal{U}). \tag{204}$$

The criterion is continuous and concave in w, a Borel-measurable minimum-norm maximizer exists, and on simultaneous coverage

$$\Phi_{\mathcal{U}(a^\star)}(w_{\mathrm{rob}}; a^\star) \ge \Phi^\star(A, \mathcal{U}). \tag{205}$$

Consequently a conditionally δ-correct rule on the covered row-wise class has total error at most $\alpha + \delta$. For rectangular budgets, $\mathcal{U}(a) = \mathcal{U}_{\mathrm{rect}}(a)$ and the interval and pair formulas of Corollary 53 apply scenario by scenario.

Proof. For fixed a, $\Phi_{\mathcal{U}(a)}(w; a)$ is the infimum of functions affine in w, hence concave. Lemma 48, compactness, and continuity of $a \mapsto (b(a), \alpha(a), \mathcal{U}(a))$ give joint continuity. Infimizing over compact A preserves concavity and continuity. The compact argmax correspondence admits the same Borel measurable minimum-norm selector used in Theorem 35. On $a^\star \in A$, the true coefficient dominates the infimum over A, proving equation 205; a union bound gives $\alpha + \delta$. □

Uniformity and scope. The joint coefficient theorem sends $\Delta \downarrow 0$ with the finite protocol library, score support, compact intercept range, positivity bounds, shared-nuisance interiority, and isolated-fiber constants fixed. The induced joint set $\mathcal{U}_\Delta$ may be correlated across protocols. The rectangular interval, pair-design, decision-profile, and finite-gap formulas require the protocol-separable subclass explicitly displayed above. A zero joint tangent coefficient is not called exact nonidentification without a finite-gap theorem on the same joint class. The validity radii are local sensitivity parameters; fixed nonvanishing direct or response violations are outside this local result.

G Reference-robust decision identification

This appendix replaces a single deployment reference by a compact family of scientifically acceptable references. It separates three questions that are not equivalent: whether the sign is invariant across the reference family, how far a model is from losing that invariant sign, and how much information is needed to reach an opposite unique decision. The first two admit sharp support-function formulas. The third is a joint cone-distance problem and, in general, does not reduce to a minimum over one reference.

G.1 Decision regions and exact alternative sets

Let $\mathcal{W}_0 \subseteq \Delta([L])$ be nonempty and compact. For a vector $V \in \mathbb{R}^L$ of protocol-potential utilities, define

$$\underline{u}_{\mathcal{W}_0}(V) = \min_{\omega \in \mathcal{W}_0} \omega^\top V, \qquad \overline{u}_{\mathcal{W}_0}(V) = \max_{\omega \in \mathcal{W}_0} \omega^\top V. \tag{206}$$

The reference-robust decision map is

$$D_{\mathcal{W}_0}(V) = \begin{cases} +, & \underline{u}_{\mathcal{W}_0}(V) > 0, \\ -, & \overline{u}_{\mathcal{W}_0}(V) < 0, \\ ?, & \underline{u}_{\mathcal{W}_0}(V) \le 0 \le \overline{u}_{\mathcal{W}_0}(V). \end{cases} \tag{207}$$

The symbol ? denotes partial identification of the sign, not failure of an estimator.

Proposition 60 (Reference-robust unique-decision geometry)**.** *Define*

$$\begin{aligned} \mathcal{P}_{\mathcal{W}_0} &= \{V : \underline{u}_{\mathcal{W}_0}(V) > 0\}, \\ \mathcal{N}_{\mathcal{W}_0} &= \{V : \overline{u}_{\mathcal{W}_0}(V) < 0\}, \\ \mathcal{A}_{\mathcal{W}_0} &= \{V : \underline{u}_{\mathcal{W}_0}(V) \le 0 \le \overline{u}_{\mathcal{W}_0}(V)\}. \end{aligned}$$

Then $\mathcal{P}_{\mathcal{W}_0}$ and $\mathcal{N}_{\mathcal{W}_0}$ are disjoint open convex cones, $\mathcal{A}_{\mathcal{W}_0}$ is closed, and these three sets partition $\mathbb{R}^L$. Their closures are

$$\overline{\mathcal{P}}_{\mathcal{W}_0} = \{V : \underline{u}_{\mathcal{W}_0}(V) \ge 0\}, \qquad \overline{\mathcal{N}}_{\mathcal{W}_0} = \{V : \overline{u}_{\mathcal{W}_0}(V) \le 0\}. \tag{208}$$

For a true positive model, the strict wrong-decision alternatives are exactly $\mathcal{N}_{\mathcal{W}_0}$ and their closure is $\overline{\mathcal{N}}_{\mathcal{W}_0}$; the negative case is symmetric.

Proof. The lower envelope is the pointwise minimum of linear functions and the upper envelope is the pointwise maximum. Hence $\mathcal{P}_{\mathcal{W}_0} = \bigcap_{\omega \in \mathcal{W}_0} \{V : \omega^\top V > 0\}$ and $\mathcal{N}_{\mathcal{W}_0} = \bigcap_{\omega \in \mathcal{W}_0} \{V : \omega^\top V < 0\}$ are open convex cones. Continuity of the two envelopes follows from compactness of $\mathcal{W}_0$, and therefore

$$\overline{\mathcal{P}}_{\mathcal{W}_0} \subseteq \{V : \underline{u}_{\mathcal{W}_0}(V) \ge 0\}, \qquad \overline{\mathcal{N}}_{\mathcal{W}_0} \subseteq \{V : \overline{u}_{\mathcal{W}_0}(V) \le 0\}.$$

For the reverse positive inclusion, take V with $\underline{u}_{\mathcal{W}_0}(V) \ge 0$. Since every $\omega \in \mathcal{W}_0$ satisfies $\omega^\top \mathbf{1} = 1$,

$$\underline{u}_{\mathcal{W}_0}(V + \epsilon \mathbf{1}) = \underline{u}_{\mathcal{W}_0}(V) + \epsilon > 0 \qquad (\epsilon > 0).$$

Thus $V + \epsilon \mathbf{1} \in \mathcal{P}_{\mathcal{W}_0}$ and converges to V. The negative inclusion follows identically from $V - \epsilon \mathbf{1} \in \mathcal{N}_{\mathcal{W}_0}$ whenever $\overline{u}_{\mathcal{W}_0}(V) \le 0$. This proves equation 208. Since the lower envelope never exceeds the upper envelope, every vector lies in exactly one of the three displayed regions. The alternative-set statement follows directly from the decision map. □

For a row-wise statistical class $\mathcal{M}_n$ with potential-utility map $V(m)$, define

$$\mathcal{M}_{n,+}^{\mathcal{W}_0} = \{m \in \mathcal{M}_n : V(m) \in \mathcal{P}_{\mathcal{W}_0}\}, \qquad \mathcal{M}_{n,-}^{\mathcal{W}_0} = \{m \in \mathcal{M}_n : V(m) \in \mathcal{N}_{\mathcal{W}_0}\}, \tag{209}$$

$$\mathrm{Alt}^{\circ}_{n,\mathcal{W}_0}(m) = \mathcal{M}_{n,-}^{\mathcal{W}_0}, \qquad \mathrm{Alt}^{\mathrm{cl}}_{n,\mathcal{W}_0}(m) = \{m' \in \mathcal{M}_n : \overline{u}_{\mathcal{W}_0}(V(m')) \le 0\} \tag{210}$$

for $m \in \mathcal{M}_{n,+}^{\mathcal{W}_0}$, with the symmetric definition for a negative true model. Correctness, the strict alternative set, and every information denominator below use the same row-wise class $\mathcal{M}_n$.

G.2 Exact ambiguity and opposite-decision radii

Fix any norm $\|\cdot\|$ on $\mathbb{R}^L$ and its dual norm $\|\cdot\|_*$. Because every $\omega \in \Delta([L])$ is nonzero, the dual norms in the next theorem are strictly positive.

Theorem 61 (Exact reference-ambiguity and opposite-decision radii)**.** *For $V \in \mathcal{P}_{\mathcal{W}_0}$, define*

$$\begin{aligned} \rho^+_{\mathrm{amb}}(V) &= \inf\{\|h\| : V + h \notin \mathcal{P}_{\mathcal{W}_0}\}, \\ \rho^+_{\mathrm{opp}}(V) &= \inf\{\|h\| : V + h \in \overline{\mathcal{N}}_{\mathcal{W}_0}\}. \end{aligned}$$

Then

$$\boxed{\rho^{+}_{\mathrm{amb}}(V) = \min_{\omega \in \mathcal{W}_0} \frac{\omega^\top V}{\|\omega\|_*}.} \tag{211}$$

Let $K_{\mathcal{W}_0} = \mathrm{cone}(\mathcal{W}_0)$. *The exact radius to an opposite closed decision is*

$$\boxed{\rho^{+}_{\mathrm{opp}}(V) = \max_{\substack{y \in K_{\mathcal{W}_0} \\ \|y\|_* \le 1}} y^\top V.} \tag{212}$$

For $V \in \mathcal{N}_{\mathcal{W}_0}$,

$$\rho^{-}_{\mathrm{amb}}(V) = \min_{\omega \in \mathcal{W}_0} \frac{-\omega^\top V}{\|\omega\|_*}, \tag{213}$$

$$\rho^{-}_{\mathrm{opp}}(V) = \max_{\substack{y \in K_{\mathcal{W}_0} \\ \|y\|_* \le 1}} (-y^\top V). \tag{214}$$

Always $\rho^{\pm}_{\mathrm{amb}} \le \rho^{\pm}_{\mathrm{opp}}$. *If* $\mathcal{W}_0$ *is a singleton, the two radii coincide.*

Proof. For a positive model,

$$\mathcal{P}^{c}_{\mathcal{W}_0} = \bigcup_{\omega \in \mathcal{W}_0} \{x : \omega^\top x \le 0\}.$$

The distance from V to the halfspace $\{x : \omega^\top x \le 0\}$ equals $\omega^\top V / \|\omega\|_*$. Taking the distance to the union and using compactness of $\mathcal{W}_0$ proves equation 211.

The closed negative region is

$$\overline{\mathcal{N}}_{\mathcal{W}_0} = \{x : \omega^\top x \le 0 \ \forall \omega \in \mathcal{W}_0\} = K^{\circ}_{\mathcal{W}_0},$$

where the polar convention is $K^\circ = \{x : y^\top x \le 0 \ \forall y \in K\}$. The cone $K_{\mathcal{W}_0}$ is closed because $\mathcal{W}_0$ is compact and lies in the affine hyperplane $\mathbf{1}^\top \omega = 1$.

For any $y \in K_{\mathcal{W}_0}$ with $\|y\|_* \le 1$ and any $x \in K^{\circ}_{\mathcal{W}_0}$,

$$y^\top V \le y^\top (V - x) \le \|V - x\|,$$

so the right-hand side of equation 212 is at most the distance. For the reverse inequality, let x^* minimize the distance to the closed convex cone $K^{\circ}_{\mathcal{W}_0}$. Finite-dimensional convex optimality gives a norming functional y^* with $\|y^*\|_* = 1$, $y^{*\top}(V - x^*) = \|V - x^*\|$, and $y^* \in N_{K^{\circ}_{\mathcal{W}_0}}(x^*)$. For a convex cone, this normal condition implies $y^* \in (K^{\circ}_{\mathcal{W}_0})^\circ = K_{\mathcal{W}_0}$ and $y^{*\top} x^* = 0$. Hence $y^{*\top} V = \|V - x^*\|$, proving the dual formula.

The negative formulas follow by applying the positive result to $-V$. The ambiguity complement contains the opposite closed region, so the first radius cannot exceed the second. For one reference both sets have the same boundary hyperplane, and equality follows. □

Example 62 (Ambiguity and opposite decision are different)**.** Let $\mathcal{W}_0 = \{e_1, e_2\}$, use the Euclidean norm, and take $V = (2, 1)$. Then the unique positive region is the positive orthant,

$$\rho^{+}_{\mathrm{amb}}(V) = 1, \qquad \rho^{+}_{\mathrm{opp}}(V) = \sqrt{5}.$$

The first number is the distance to losing positivity under one reference; the second is the distance to making both reference utilities nonpositive. Thus a reference-profiled ambiguity coefficient is not, by itself, the opposite-decision coefficient used by a binary fixed-confidence lower bound.

G.3 Normal information after profiling the reference

The next result is a direct decision-set specialization of the robust normal-contact theorem. It records explicitly which coefficient is obtained by profiling a reference and why the strict opposite-decision coefficient is generally different.

Let $\mathcal{G} \subset \mathbb{R}^L$ be a nonempty compact set of feasible local potential-utility directions. Let $\mathcal{R}_s : \mathcal{G} \to \mathsf{N}_0$ be the continuous order-s normal likelihood-jet map after profiling the ordinary nuisance manifold, and let $\mathcal{B} \subset \mathsf{N}_0$ be the compact limiting normal uncertainty difference set. For a robustly positive true direction $g_+ \in \mathcal{G}$, define

$$\begin{aligned}\mathcal{G}_{\neg+} &= \{g \in \mathcal{G} : \underline{u}_{\mathcal{W}_0}(g) \le 0\},\\ \mathcal{G}_{-} &= \{g \in \mathcal{G} : \overline{u}_{\mathcal{W}_0}(g) \le 0\},\\ \mathcal{G}_{\omega} &= \{g \in \mathcal{G} : \omega^\top g \le 0\}.\end{aligned}$$

Theorem 63 (Reference-profiled and opposite-decision normal coefficients)**.** *Assume the quantitative graphical-convergence, positivity, and outer-fiber conditions of Theorem 11, uniformly over the compact direction set $\mathcal{G}$. Then the closed profile for loss of a robust positive decision satisfies*

$$I_{\mathrm{amb}}(t) = 2t^{2s} d_{\mathrm{amb}}^2 + O(t^{2s+1}), \tag{215}$$

where

$$\begin{aligned} d_{\mathrm{amb}} &= \operatorname{dist}(\mathcal{R}_s(g_+), \mathcal{R}_s(\mathcal{G}_{\neg+}) + \mathcal{B}) && (216)\\ &= \min_{\omega \in \mathcal{W}_0} \operatorname{dist}(\mathcal{R}_s(g_+), \mathcal{R}_s(\mathcal{G}_{\omega}) + \mathcal{B})\,. && (217)\end{aligned}$$

The closed strict-opposite profile satisfies

$$I_{\mathrm{opp}}(t) = 2t^{2s} d_{\mathrm{opp}}^2 + O(t^{2s+1}), \qquad d_{\mathrm{opp}} = \operatorname{dist}(\mathcal{R}_s(g_+), \mathcal{R}_s(\mathcal{G}_{-}) + \mathcal{B})\,. \tag{218}$$

Because $\mathcal{G}_- \subseteq \mathcal{G}_{\neg+}$, $d_{\mathrm{amb}} \le d_{\mathrm{opp}}$. In general the second distance cannot be written as a minimum over one reference. Under qualitative graphical convergence, the normalized profiles converge to the same coefficients with little-o remainders.

Proof. The robust normal-contact theorem applies to any compact alternative normal cross-section. The loss-of-positive alternative directions are exactly $\mathcal{G}_{\neg+}$, while the opposite closed alternatives are exactly $\mathcal{G}_-$. This gives equation 215 and equation 218.

Moreover,

$$\mathcal{G}_{\neg+} = \bigcup_{\omega \in \mathcal{W}_0} \mathcal{G}_{\omega},$$

because the minimum reference utility is nonpositive exactly when at least one reference utility is nonpositive. Continuity of $(\omega, g) \mapsto \omega^\top g$ and compactness of $\mathcal{W}_0 \times \mathcal{G}$ make the indexed union compact. Distance to a compact union is the minimum of the component distances, proving equation 217. By contrast, $\mathcal{G}_- = \bigcap_\omega \mathcal{G}_\omega$, so only the inequality follows in general. □

A zero value in equation 217 is a leading-order cancellation. It does not, without the finite-gap intersection theorem, imply an attained ambiguous model or a strict opposite-decision alias.

G.4 Protocol specialization and contraction of the usable margin

For the local protocol experiment, retain the notation α_z, b_z from Section 4. Let $g \in \mathbb{R}^L$ be a local potential-utility direction and put

$$m_\omega(g) = \omega^\top g, \qquad \underline{m}_{\mathcal{W}_0}(g) = \min_{\omega \in \mathcal{W}_0} m_\omega(g). \tag{219}$$

Assume $\underline{m}_{\mathcal{W}_0}(g) > 0$. For a compact joint target-slope set $\mathcal{U}$, define the coefficient at scalar margin $m \ge 0$ by

$$\Phi_{\mathcal{U}}(w; m) = \inf_{\substack{c \in \mathbb{R}\\ v \in \mathcal{U}}} \sum_z w_z \alpha_z^2 (c - m b_z - v_z)^2. \tag{220}$$

Corollary 64 (Reference-profiled protocol coefficient)**.** *Suppose the inverse-response targets have, uniformly over $\omega \in \mathcal{W}_0$ and the compact validity class, the expansion*

$$T_{z,t,\omega,v} = t\{m_\omega(g) b_z + v_z\} x + O(t^2),$$

and the normalized joint validity sets converge at Hausdorff rate $O(t)$. Then, for every fixed protocol design w,

$$\boxed{K^{\mathrm{amb}}_{w,\mathcal{W}_0}(t;g) = \frac{1}{2}t^2 \min_{\omega\in\mathcal{W}_0} \Phi_{\mathcal{U}}(w; m_\omega(g)) + O(t^3).} \tag{221}$$

Under qualitative normalized-set convergence the remainder is $o(t^2)$. If $\mathcal{U}$ is convex and $0 \in \mathcal{U}$, then $m \mapsto \Phi_{\mathcal{U}}(w;m)$ is nondecreasing on $[0,\infty)$, and

$$\min_{\omega\in\mathcal{W}_0} \Phi_{\mathcal{U}}(w; m_\omega(g)) = \Phi_{\mathcal{U}}\big(w; \underline{m}_{\mathcal{W}_0}(g)\big). \tag{222}$$

Proof. For each fixed reference, the scalar-center reduction used in Theorem 3 gives the coefficient $\Phi_{\mathcal{U}}(w; m_\omega(g))$. The expansion is uniform on the compact reference family, so taking the infimum over ω preserves the $O(t^3)$ remainder and proves equation 221.

For monotonicity, put $C = \mathrm{span}\{\mathbf{1}\} - \mathcal{U}$ and use the weighted norm induced by $w_z\alpha_z^2$. The definition equation 220 is $\mathrm{dist}^2_w(mb, C)$ because $c\mathbf{1} - mb - v = -(mb - (c\mathbf{1} - v))$. Convexity of $\mathcal{U}$ and $0 \in \mathcal{U}$ make C convex and ensure $0 \in C$. If $0 \le m_1 \le m_2$ and $y_2 \in C$ is a nearest point to $m_2 b$, convexity and $0 \in C$ imply $(m_1/m_2)y_2 \in C$. Hence

$$\mathrm{dist}_w(m_1 b, C) \le \frac{m_1}{m_2}\,\mathrm{dist}_w(m_2 b, C) \le \mathrm{dist}_w(m_2 b, C).$$

This proves monotonicity and the worst-margin reduction. □

For protocol-separable rectangular budgets, let $r_z = \chi_z + \tau_z/\alpha_z$ and define

$$\mathcal{J}_z(m) = [mb_z - r_z, mb_z + r_z].$$

Then

$$\Phi_{\mathcal{U}_{\mathrm{rect}}}(w; m) = \min_c \sum_z w_z\alpha_z^2\,\mathrm{dist}^2(c, \mathcal{J}_z(m)). \tag{223}$$

For a pair i, j, the reference-robust surviving margin is

$$\boxed{\Theta^{\mathrm{ref}}_{ij} = \frac{\alpha_i\alpha_j}{\alpha_i+\alpha_j}\left[\underline{m}_{\mathcal{W}_0}(g)|b_i - b_j| - r_i - r_j\right]_+.} \tag{224}$$

The square-root allocation remains $p^*_{ij} = \alpha_j/(\alpha_i+\alpha_j)$. If the nominal reference is normalized by $\bar{w}^\top g = 1$, define

$$A_{ij} = \frac{\alpha_i\alpha_j}{\alpha_i+\alpha_j}|b_i - b_j|.$$

Then

$$\Theta^{\mathrm{ref}}_{ij} = \Bigg[A_{ij} - \underbrace{(1 - \underline{m}_{\mathcal{W}_0}(g))A_{ij}}_{\text{reference uncertainty}} - \underbrace{\frac{\alpha_i\alpha_j}{\alpha_i+\alpha_j}(\chi_i+\chi_j)}_{\text{direct-effect uncertainty}} - \underbrace{\frac{\alpha_j\tau_i + \alpha_i\tau_j}{\alpha_i+\alpha_j}}_{\text{response-sharing uncertainty}}\Bigg]_+. \tag{225}$$

Thus reference uncertainty contracts the usable protocol margin before the direct-effect and response-sharing losses are paid. This formula concerns loss of a robust unique sign. The strict opposite-decision information is the joint cone distance in equation 218 and may be strictly larger.

Proof of equation 224. For fixed scalar margin m, the two interval centers are mb_i and mb_j. The interval gap is $[m|b_i - b_j| - r_i - r_j]_+$. The two-protocol weighted interval calculation from Theorem 3 multiplies this gap by $\alpha_i\alpha_j/(\alpha_i+\alpha_j)$ and is maximized by the same square-root allocation. Corollary 64 sets $m = \underline{m}_{\mathcal{W}_0}(g)$. Expanding $r_z = \chi_z + \tau_z/\alpha_z$ gives equation 225. □

G.5 CLASS-RELATIVE FIXED-CONFIDENCE CONTRACT

For a compact row-wise class $\mathcal{M}_n$ equipped with $D_{\mathcal{W}_0}$, let $\mathcal{M}_{n,\mathrm{uni}}^{\mathcal{W}_0} = \mathcal{M}_{n,+}^{\mathcal{W}_0} \cup \mathcal{M}_{n,-}^{\mathcal{W}_0}$. A rule is δ_n-correct when it terminates almost surely and returns the correct sign for every model in this unique-decision set. For a positive true model, define

$$\Gamma_{n,\mathcal{W}_0}(m) = \sup_{w\in\Delta([L])} \inf_{m'\in\mathcal{M}_{n,-}^{\mathcal{W}_0}} \sum_z w_z D_{\mathrm{KL}}(P_m^z \| P_{m'}^z). \tag{226}$$

The row-wise transportation theorem gives

$$\mathbb{E}_m[\tau_n] \geq \frac{\mathrm{kl}(1-\delta_n, \delta_n)}{\Gamma_{n,\mathcal{W}_0}(m)}. \tag{227}$$

The denominator uses the strict opposite region, not the ambiguity boundary. At a zero strict infimum attained only on the closed boundary, the adaptive transcript-continuity argument from Appendix E rules out uniform $\delta_n < 1/2$ correctness.

H ROBUST PROTOCOL DESIGN UNDER CORRELATED VALIDITY UNCERTAINTY

This appendix turns the joint-validity coefficient into a complete robust-design theorem. The main result is a minimax identity: optimizing the sampling frequencies is exactly a weighted Chebyshev-center problem against the *joint* validity set. The interval and two-protocol formulas are recovered only for the protocol-separable rectangular specialization.

The minimax exchange, active-set characterization, sparse-support argument, and conic formulations below use classical convex optimal-design ideas related to Elfving geometry and multiresponse SOCP formulations (Elfving, 1952; Sagnol, 2010). The contribution here is the joint target-slope validity set induced by selective response, its decision-information interpretation, and the resulting failure of pair support under correlated validity errors with unchanged coordinate marginals.

H.1 SETUP

Fix a finite protocol library $[L] = \{1, \dots, L\}$, positive information weights $\alpha_z > 0$, and nominal target slopes $b_z \in \mathbb{R}$. A fixed reference gauge identifies the quotient $\mathbb{R}^L/\operatorname{span}\{\mathbf{1}\}$ with a linear section $\mathsf{G} \subset \mathbb{R}^L$. Let $\mathcal{U} \subset \mathsf{G}$ be nonempty, compact, and convex. For $w \in \Delta([L])$, $c \in \mathbb{R}$, and $v \in \mathcal{U}$, define

$$F(w, c, v) = \sum_{z=1}^{L} w_z \alpha_z^2 (c - b_z - v_z)^2, \tag{228}$$

$$\Phi_{\mathcal{U}}(w) = \min_{c\in\mathbb{R},\, v\in\mathcal{U}} F(w, c, v), \tag{229}$$

$$\Psi_{\mathcal{U}} = \max_{w\in\Delta([L])} \Phi_{\mathcal{U}}(w). \tag{230}$$

The objective is invariant under replacing a representative v by $v + s\mathbf{1}$ and c by $c + s$, so it is intrinsic to the quotient.

Set

$$C_{\mathcal{U}} = \left[\min_{z,\, v\in\mathcal{U}} (b_z + v_z),\ \max_{z,\, v\in\mathcal{U}} (b_z + v_z) \right] \tag{231}$$

and $K_{\mathcal{U}} = C_{\mathcal{U}} \times \mathcal{U}$. For $y = (c, v)$, write

$$f_z(y) = \alpha_z^2 (c - b_z - v_z)^2, \qquad G(y) = \max_{z\in[L]} f_z(y). \tag{232}$$

H.2 ROBUST CHEBYSHEV DUALITY

Theorem 65 (Correlated-validity design duality)**.** *Under the setup above,*

$$\boxed{\max_{w\in\Delta([L])} \Phi_{\mathcal{U}}(w) = \min_{(c,v)\in K_{\mathcal{U}}} \max_{z\in[L]} \alpha_z^2 (c - b_z - v_z)^2.} \tag{233}$$

Both extrema are attained. Equivalently, the robust protocol-design value is the squared radius of a weighted Chebyshev center of the jointly perturbed slope vector $b + v$.

If the statistical profile satisfies, uniformly in w,

$$K_w^{\mathrm{joint}}(\Delta, 0) = \frac{1}{2}\Delta^2 \Phi_{\mathcal{U}}(w) + O(\Delta^3), \tag{234}$$

then the optimized experiment obeys

$$\sup_{w \in \Delta([L])} K_w^{\mathrm{joint}}(\Delta, 0) = \frac{1}{2}\Delta^2 \Psi_{\mathcal{U}} + O(\Delta^3). \tag{235}$$

The corresponding statement with $o(\Delta^2)$ holds under qualitative graphical convergence.

Proof. First restrict the center. If c exceeds the upper endpoint of $C_{\mathcal{U}}$, replacing it by that endpoint weakly decreases $|c - b_z - v_z|$ simultaneously for every z and every $v \in \mathcal{U}$. The lower endpoint is identical. Hence every inner minimum in equation 229 may be taken over the compact convex set $K_{\mathcal{U}}$.

On $\Delta([L]) \times K_{\mathcal{U}}$, the map F is continuous, affine in w, and convex in (c, v). Sion's minimax theorem therefore gives

$$\max_{w \in \Delta([L])} \min_{(c,v) \in K_{\mathcal{U}}} F(w, c, v) = \min_{(c,v) \in K_{\mathcal{U}}} \max_{w \in \Delta([L])} F(w, c, v).$$

For fixed (c, v), maximizing an affine function over the simplex selects a largest coordinate:

$$\max_{w \in \Delta([L])} F(w, c, v) = \max_z f_z(c, v).$$

This proves equation 233. Compactness gives attainment on both sides.

Finally, let the uniform remainder in equation 234 be bounded in absolute value by $C\Delta^3$. Taking suprema in w changes either side by at most that same amount, proving equation 235. The little-o statement follows by the same uniform argument. □

H.3 Active protocols, KKT conditions, and support size

For a closed convex set K and $y \in K$, write $N_K(y)$ for its normal cone. All normal cones below are relative to the affine hull of K.

Theorem 66 (Active-set characterization and sparse optimal designs)**.** *Let $y^\star = (c^\star, v^\star)$ solve the right-hand side of equation 233, and set*

$$t^\star = G(y^\star), \qquad A^\star = \{z : f_z(y^\star) = t^\star\}. \tag{236}$$

Then there exists an optimal design $w^\star \in \Delta([L])$ such that

$$\mathrm{supp}(w^\star) \subseteq A^\star, \tag{237}$$

$$0 \in \sum_{z \in A^\star} w_z^\star \nabla f_z(y^\star) + N_{K_{\mathcal{U}}}(y^\star). \tag{238}$$

For this design, $y^\star$ minimizes $F(w^\star, \cdot)$ and

$$\Phi_{\mathcal{U}}(w^\star) = F(w^\star, y^\star) = t^\star = \Psi_{\mathcal{U}}. \tag{239}$$

Conversely, every saddle pair $(w^\star, y^\star)$ satisfying equation 239 assigns positive mass only to active protocols.

Let $d_{\mathcal{U}} = \dim \mathrm{aff}(\mathcal{U})$. An optimal design can always be chosen with

$$\boxed{|\,\mathrm{supp}(w^\star)| \le d_{\mathcal{U}} + 2.} \tag{240}$$

This bound includes boundary solutions; no relative-interior assumption is needed.

Proof. Because $G = \max_z f_z$ is a finite maximum of differentiable convex functions,

$$\partial G(y^\star) = \operatorname{conv}\{\nabla f_z(y^\star) : z \in A^\star\}. \tag{241}$$

The convex first-order condition for minimizing G over $K_{\mathcal{U}}$ is

$$0 \in \partial G(y^\star) + N_{K_{\mathcal{U}}}(y^\star).$$

Combining this with equation 241 yields coefficients $w_z^\star \geq 0$, supported on $A^\star$ and summing to one, for which equation 238 holds.

The function $F(w^\star, \cdot) = \sum_z w_z^\star f_z$ is convex. Condition equation 238 is therefore sufficient for $y^\star$ to minimize it. Since every positively weighted protocol is active,

$$F(w^\star, y^\star) = \sum_z w_z^\star t^\star = t^\star.$$

For any design w,

$$\min_y F(w, y) \leq F(w, y^\star) \leq \max_z f_z(y^\star) = t^\star.$$

Thus $w^\star$ attains the maximin value and equation 239 follows. Conversely, if a saddle design placed positive mass on a protocol with $f_z(y^\star) < t^\star$, then $F(w^\star, y^\star) < t^\star$, contradicting equation 239.

For the support bound, let $A = \operatorname{lin}(\operatorname{aff} K_{\mathcal{U}} - y^\star)$ and let P_A denote orthogonal projection onto A. Its dimension is at most $d_{\mathcal{U}} + 1$. The relative first-order condition can be written

$$0 \in \operatorname{conv}\{P_A \nabla f_z(y^\star) : z \in A^\star\} + N_{K_{\mathcal{U}}}^A(y^\star).$$

Choose g in the displayed convex hull and n in the relative normal cone with $g + n = 0$. By Caratheodory's theorem in the vector space A, the same vector g is a convex combination of at most $\dim(A) + 1 \leq d_{\mathcal{U}} + 2$ projected active gradients. Retaining n shows that the reduced combination still satisfies the relative KKT condition. The preceding argument then turns those coefficients into an optimal design. Since the normal cone was taken relative to the affine hull, this argument applies equally to interior and boundary minimizers. □

The gradients in equation 238 are explicit. If $r_z^\star = c^\star - b_z - v_z^\star$, then

$$\nabla_c f_z(y^\star) = 2\alpha_z^2 r_z^\star, \qquad \nabla_v f_z(y^\star) = -2\alpha_z^2 r_z^\star e_z, \tag{242}$$

with the v-gradient projected onto the fixed quotient section. Thus the optimal sampling frequencies are precisely multipliers balancing the active protocol residuals against the normal cone of the correlated validity set.

Proposition 67 (Measurable choice under random calibration sets)**.** *Let Θ be a standard Borel parameter space. Suppose $\theta \mapsto b(\theta)$ and $\theta \mapsto \alpha(\theta)$ are Borel, $\inf_{\theta,z} \alpha_z(\theta) > 0$, and $\theta \mapsto \mathcal{U}(\theta)$ is a measurable compact convex correspondence contained in one common compact set. Then the saddle correspondence associated with equation 233 has a measurable graph and nonempty compact values. Consequently, a Borel-measurable optimal center, validity vector, and design can be selected. One canonical choice is the unique minimum-Euclidean-norm point of the compact convex saddle set.*

Proof. The common compact bound gives one compact center interval containing every $C_{\mathcal{U}(\theta)}$. The objective is a Caratheodory integrand: Borel in θ and continuous in (w, c, v). The measurable maximum theorem implies that the maximin value and the inner minimum values are measurable and that their argmin and argmax correspondences have measurable graphs and nonempty compact values. The saddle set is the intersection of these graphs with the closed equal-value relation, hence is also measurable, compact-valued, and nonempty by Theorem 65. A measurable selection follows from the measurable selection theorem. The saddle set is convex, so minimizing the strictly convex squared Euclidean norm over it gives a unique point; applying the measurable maximum theorem once more makes this canonical selection Borel measurable. □

H.4 RANK-ONE CORRELATED VALIDITY

Theorem 68 (Exact rank-one reduction)**.** *Let*

$$\mathcal{U}_{\beta,\kappa} = \{\beta t : |t| \le \kappa\}, \qquad \beta \in \mathsf{G}. \tag{243}$$

Define

$$q_{ij} = \frac{\alpha_i \alpha_j}{\alpha_i + \alpha_j}, \qquad R(t) = \max_{i<j} q_{ij} \left| (b_i - b_j) + (\beta_i - \beta_j) t \right|. \tag{244}$$

Then

$$\boxed{\Psi_{\mathcal{U}_{\beta,\kappa}} = \left(\min_{|t| \le \kappa} R(t) \right)^2.} \tag{245}$$

The scalar function R is convex and piecewise linear. Its minimum is found exactly by evaluating the endpoints $\pm\kappa$ and the intersection points, inside $[-\kappa, \kappa]$, of the finitely many affine functions

$$\ell_{ij,s}(t) = s q_{ij} \big((b_i - b_j) + (\beta_i - \beta_j) t \big), \qquad s \in \{-1, +1\}. \tag{246}$$

An optimal design exists on at most three protocols, and three protocols can be necessary.

Proof. Fix t and put $d_z(t) = b_z + \beta_z t$. A radius $R \ge 0$ admits a center c with $\alpha_z |c - d_z(t)| \le R$ for every z if and only if the intervals

$$I_z(R, t) = \left[d_z(t) - \frac{R}{\alpha_z}, d_z(t) + \frac{R}{\alpha_z} \right]$$

have a common point. Intervals on the real line have a common intersection if and only if every pair intersects. The pair (i, j) intersects exactly when

$$|d_i(t) - d_j(t)| \le R \left(\frac{1}{\alpha_i} + \frac{1}{\alpha_j} \right),$$

or equivalently $R \ge q_{ij} |d_i(t) - d_j(t)|$. Hence the smallest radius at fixed t is $R(t)$. Minimizing over the jointly feasible scalar t and using Theorem 65 proves equation 245.

equation 246 writes R as a finite maximum of affine functions, so it is convex and piecewise linear. If a minimizer lies in the interior of one linearity cell, either the active slope is zero, in which case a cell endpoint is also a minimizer, or moving in one direction strictly decreases the function. Thus a minimizer occurs at an interval endpoint or at an intersection of two affine pieces. Finally, $d_{\mathcal{U}} = 1$ in Theorem 66, giving support at most three. The following example proves necessity. □

Example 69 (A rank-one class that requires three protocols)**.** Take

$$\alpha = (1, 1, 1), \qquad b = (2, 3, 2), \qquad \beta = (-1, 0, 1), \qquad \kappa = 1. \tag{247}$$

At scalar t, the three perturbed slopes are $(2 - t, 3, 2 + t)$. Their smallest equal-weight Chebyshev radius is

$$R(t) = \frac{1}{2} \big(1 + |t| \big), \qquad |t| \le 1. \tag{248}$$

Thus $t^\star = 0$, $c^\star = 5/2$, and

$$\Psi_{\mathcal{U}_{\beta,1}} = \frac{1}{4}. \tag{249}$$

The three residuals are $(1/2, -1/2, 1/2)$. The unique KKT design is

$$w^\star = \left(\frac{1}{4}, \frac{1}{2}, \frac{1}{4} \right). \tag{250}$$

Indeed, the two KKT balance equations are $w_1 - w_2 + w_3 = 0$ and $w_1 - w_3 = 0$.

Every two-protocol design has value zero: protocols $(1, 2)$ are fitted by $(c, t) = (3, -1)$, protocols $(2, 3)$ by $(3, 1)$, and protocols $(1, 3)$ by $(2, 0)$. Hence no design supported on two protocols is optimal. The rectangular envelope with the same marginals, $[-1, 1] \times \{0\} \times [-1, 1]$, also has value zero because $v = (1, 0, 1)$ and $c = 3$ equalize all three slopes. This single convex example shows that identical marginal intervals can hide positive correlated information and that the rectangular optimal-pair law is not generic.

H.5 Ellipsoidal validity uncertainty

Let G have dimension d and let $B \in \mathbb{R}^{L\times d}$ have columns forming a basis of G. Consider

$$\mathcal{U}_{M,\kappa} = \{Bx : x^\top Mx \le \kappa^2\}, \qquad M \succ 0. \tag{251}$$

Theorem 70 (Ellipsoidal fixed-design solution)**.** *Fix a full-support design w. Set*

$$d_z = w_z\alpha_z^2, \quad D = \operatorname{diag}(d_1,\dots,d_L), \quad s = \mathbf{1}^\top d, \quad H = D - \frac{dd^\top}{s}. \tag{252}$$

Let

$$Q = B^\top HB, \qquad h = B^\top Hb, \qquad q_0 = b^\top Hb. \tag{253}$$

Then $Q \succ 0$, and

$$\Phi_{\mathcal{U}_{M,\kappa}}(w) = \min_{x^\top Mx\le\kappa^2} \left(q_0 + 2h^\top x + x^\top Qx\right). \tag{254}$$

If $x_0 = -Q^{-1}h$ satisfies $x_0^\top Mx_0 \le \kappa^2$, then $x^\star = x_0$ and the multiplier is $\lambda = 0$. Otherwise there is a unique $\lambda > 0$ such that

$$x_\lambda = -(Q+\lambda M)^{-1}h, \tag{255}$$

$$x_\lambda^\top Mx_\lambda = \kappa^2, \tag{256}$$

and $x^\star = x_\lambda$. The profiled center and value are

$$c^\star = \frac{d^\top(b+Bx^\star)}{s}, \tag{257}$$

$$\Phi_{\mathcal{U}_{M,\kappa}}(w) = q_0 - h^\top(Q+\lambda M)^{-1}h - \lambda\kappa^2. \tag{258}$$

After whitening M, equation 256 is a monotone one-dimensional secular equation. Thus the fixed-design coefficient is computable by one scalar root search.

Proof. For fixed x, minimizing $\sum_z d_z(c - b_z - (Bx)_z)^2$ over c gives equation 257. Substitution yields $(b+Bx)^\top H(b+Bx)$ and therefore equation 254. Since w has full support, the nullspace of H is $\operatorname{span}\{\mathbf{1}\}$. The fixed quotient section satisfies $\mathsf{G}\cap\operatorname{span}\{\mathbf{1}\} = \{0\}$, so $x^\top Qx = (Bx)^\top H(Bx) > 0$ for every nonzero x, proving $Q \succ 0$.

The problem in equation 254 is a convex quadratic over an ellipsoid. Its KKT conditions are

$$(Q+\lambda M)x = -h, \quad \lambda \ge 0, \quad x^\top Mx \le \kappa^2, \quad \lambda(x^\top Mx - \kappa^2) = 0. \tag{259}$$

They are necessary and sufficient. If the unconstrained minimizer is feasible, $\lambda = 0$. Otherwise the constraint is active and $Q + \lambda M \succ 0$, giving equation 255.

Let $M^{-1/2}QM^{-1/2} = U\operatorname{diag}(\sigma_1,\dots,\sigma_d)U^\top$ and $g = U^\top M^{-1/2}h$. Then

$$x_\lambda^\top Mx_\lambda = \sum_{k=1}^{d} \frac{g_k^2}{(\sigma_k+\lambda)^2}. \tag{260}$$

When the constraint is active, the right side is continuous and strictly decreasing from a value larger than κ^2 to zero, so the positive root is unique. Finally, substituting $(Q+\lambda M)x^\star = -h$ into the quadratic gives equation 258. □

Equivalently, the optimized ellipsoidal design is the second-order cone program

$$\min_{c,x,R} R \quad \text{s.t.} \quad \alpha_z|c - b_z - (Bx)_z| \le R\,\forall z, \qquad x^\top Mx \le \kappa^2. \tag{261}$$

The multipliers of the active protocol inequalities yield the optimal design weights through Theorem 66.

H.6 Price of a rectangular relaxation

Proposition 71 (Monotonicity and Hausdorff sensitivity)**.** *Let $\mathcal{U}, \mathcal{V} \subset \mathsf{G}$ be nonempty compact sets and let $\alpha_{\max} = \max_z \alpha_z$.*

(i) If $\mathcal{U} \subseteq \mathcal{V}$, then for every design w,

$$\Phi_{\mathcal{U}}(w) \geq \Phi_{\mathcal{V}}(w), \qquad \Psi_{\mathcal{U}} \geq \Psi_{\mathcal{V}}. \tag{262}$$

(ii) For every w,

$$\left|\sqrt{\Phi_{\mathcal{U}}(w)} - \sqrt{\Phi_{\mathcal{V}}(w)}\right| \leq \alpha_{\max} d_H(\mathcal{U}, \mathcal{V}), \tag{263}$$

and

$$\left|\sqrt{\Psi_{\mathcal{U}}} - \sqrt{\Psi_{\mathcal{V}}}\right| \leq \alpha_{\max} d_H(\mathcal{U}, \mathcal{V}). \tag{264}$$

(iii) If $\mathcal{U} \subseteq \mathcal{V}$ and $w_{\mathcal{V}} \in \arg\max_w \Phi_{\mathcal{V}}(w)$, then

$$0 \leq \sqrt{\Psi_{\mathcal{U}}} - \sqrt{\Phi_{\mathcal{U}}(w_{\mathcal{V}})} \leq \alpha_{\max} d_H(\mathcal{U}, \mathcal{V}). \tag{265}$$

Thus a rectangular envelope gives a conservative information certificate, and its optimal design is quantitatively near-optimal whenever the envelope is close to the true correlated set.

Proof. Part (i) follows because the infimum in equation 229 is taken over a larger set under $\mathcal{V}$. For part (ii), use the weighted norm $\|r\|_w^2 = \sum_z w_z \alpha_z^2 r_z^2$. If $v \in \mathcal{U}$ is matched by $v' \in \mathcal{V}$ within Euclidean distance $d_H + \epsilon$, then $\|v - v'\|_w \leq \alpha_{\max}(d_H + \epsilon)$. The triangle inequality for distance to the sets $\{c\mathbf{1} - b - v : c \in \mathbb{R}, v \in \mathcal{U}\}$ and $\{c\mathbf{1} - b - v : c \in \mathbb{R}, v \in \mathcal{V}\}$ gives equation 263 after exchanging the sets and sending $\epsilon \downarrow 0$. Taking maxima in w and using $\sqrt{\max_w \Phi(w)} = \max_w \sqrt{\Phi(w)}$ proves equation 264.

For part (iii), monotonicity gives $\sqrt{\Phi_{\mathcal{U}}(w_{\mathcal{V}})} \geq \sqrt{\Phi_{\mathcal{V}}(w_{\mathcal{V}})} = \sqrt{\Psi_{\mathcal{V}}}$, while equation 264 gives $\sqrt{\Psi_{\mathcal{U}}} \leq \sqrt{\Psi_{\mathcal{V}}} + \alpha_{\max} d_H(\mathcal{U}, \mathcal{V})$. Combining the two inequalities proves equation 265. □

Interpretation. The general joint theorem is a correlated robust-design problem; the rectangular pair-selection theorem is only a special case. The number of protocols needed by an optimal design is controlled by the affine dimension of the scientific validity mechanism. One-dimensional correlated mechanisms can require three protocols, while the special interval structure of rectangular uncertainty collapses the design to a pair. Rectangular sensitivity analysis remains useful because it is conservative and comes with an explicit approximation bound, but its pair law must not be transferred to a correlated class.

I Robust quadratic–quartic crossover under imperfect validity

This appendix unifies the exact-validity crossover of Appendix E.7 with the approximate-exclusion geometry of Appendix F. The result is deliberately stated for a symmetric two-protocol benchmark with an explicit quartic coefficient and a specified dependence structure. Direct-effect uncertainty is a single antisymmetric scalar contrast, and the two protocol-specific response perturbations range independently over their local information balls. Under this product structure, response uncertainty contracts the joint tangent–normal radius and yields a radial positive-part law. The same closed form is generally false for correlated response perturbations, even when their protocolwise marginal balls are identical; Proposition 74 gives an exact counterexample.

I.1 Gauge-fixed row-wise model class

For a raw decomposition $(\widetilde{u}, \widetilde{\xi}_+, \widetilde{\xi}_-)$, the protocol utilities are $\widetilde{u} + \widetilde{\xi}_\varsigma$. The transformation

$$(\widetilde{u}, \widetilde{\xi}_+, \widetilde{\xi}_-) \mapsto (\widetilde{u} + c, \widetilde{\xi}_+ - c, \widetilde{\xi}_- - c)$$

leaves both protocol utilities unchanged. The quotient-invariant common coordinate and contrast are

$$u = \widetilde{u} + \frac{\widetilde{\xi}_+ + \widetilde{\xi}_-}{2}, \qquad \theta = \frac{\widetilde{\xi}_+ - \widetilde{\xi}_-}{2}. \tag{266}$$

Equivalently, the unique uniform-mean gauge has $\xi_\varsigma = \varsigma\theta$ and protocol utilities $u + \varsigma\theta$. The decision map is the quotient-invariant sign of u; pairwise protocol contrasts are unchanged by every raw reparameterization. The reference gauge is fixed independently of the experimental design and before every design optimization; it therefore does not depend on the optimized allocation.

Let $r = 1$, $J_{\mathrm{eff}} \geq 3$, and let the two response intercepts be $a + \varsigma\zeta$, $\varsigma \in \{-1, +1\}$. For nonnegative radii η and γ, define the compact model class

$$\mathcal{M}^{\mathrm{rc}}_{\zeta,\eta,\gamma} = \left\{ (u, q, \theta, h_+, h_-) : \begin{array}{c} u \in [-u_0, u_0],\ q \in \mathcal{Q}^0_1,\ |\theta| \leq \eta, \\ \|h_\varsigma\|_{W_\varsigma(\zeta)} \leq \gamma, \quad \varsigma \in \{-1, +1\} \end{array} \right\}. \tag{267}$$

The complete law under protocol ς is

$$P^{\varsigma,\zeta}_{u,q,\theta,h}(j) = p_{u+\varsigma\theta}(j)\sigma(a + \varsigma\zeta + q_j + h_{\varsigma,j}), \tag{268}$$

$$P^{\varsigma,\zeta}_{u,q,\theta,h}(\perp) = 1 - \sum_j P^{\varsigma,\zeta}_{u,q,\theta,h}(j). \tag{269}$$

Here θ is a protocol-specific direct effect and h_ς is a protocol-specific response deviation. The nominal true model is $m_\Delta = (\Delta, 0, 0, 0, 0)$. Strict wrong-decision alternatives have $u < 0$; closed alternatives have $u \leq 0$. A row-wise procedure is required to be correct on the same class equation 267 used in its characteristic-information denominator.

The radii are local sensitivity budgets. In the critical theorem below, $\eta = O(\Delta^2)$ and $\gamma = O(\Delta^2)$; a fixed nonvanishing exclusion or response violation is outside this shrinking local class. The intercepts $a \pm \zeta$ may be calibrated by the independent pilot, whereas η and γ are externally supplied sensitivity bounds rather than identified pilot parameters.

The dependence assumption is part of the theorem. The exact class in equation 267 is the product

$$\mathcal{C}^{\mathrm{prod}}_\Delta = [-\eta, \eta] \times B_{W_+(\zeta)}(\gamma) \times B_{W_-(\zeta)}(\gamma), \tag{270}$$

so the scalar direct contrast is chosen independently of the two response perturbations, and the response perturbations may be chosen independently across protocols. More generally, we allow only product sets $\Theta_\Delta \times \mathcal{H}_{+,\Delta} \times \mathcal{H}_{-,\Delta}$ with $\Theta_\Delta = -\Theta_\Delta$ and protocol-exchange symmetry. The quantitative theorem requires

$$\begin{aligned} d_H(\Theta_\Delta, [-\eta, \eta]) &\leq K d^3, \\ d_H^{W_\varsigma(\zeta)}(\mathcal{H}_{\varsigma,\Delta}, B_{W_\varsigma(\zeta)}(\gamma)) &\leq K d^3, \qquad \varsigma \in \{-1, +1\}, \end{aligned} \tag{271}$$

uniformly on the critical rectangle, where $d = \Delta - u$. On the decision-profile critical rectangle this is rate-$O(\Delta)$ convergence after division by the critical Δ^2 scale; before utility profiling it is equivalently rate-$O(d)$ relative to the d^2 scale. Replacing the right sides by $o(d^2)$ gives the qualitative version. The response variables may be interpreted as net alternative-minus-true perturbations; for symmetric true- and alternative-side balls, γ is the sum of their radii.

For an arbitrary correlated joint set in $\Theta \times \mathcal{H}_+ \times \mathcal{H}_-$, the general robust normal-contact theorem still gives an intrinsic distance coefficient in the weighted product space. It does not, without additional structure, imply the radial contraction $[\Gamma_{\mathrm{crit}} - \gamma]_+$, equality of the balanced and selectable leading coefficients, or the optimizer $w = 1/2$.

For $d = \Delta - u$ and a direct contrast θ, exact equality of answered category probabilities requires the response-logit target

$$T_{\varsigma,\theta,j} = G_{a+\varsigma\zeta}(R_{\varsigma,\theta,j}), \tag{272}$$

$$R_{\varsigma,\theta,j} = (d - \varsigma\theta)y_j + \psi(u + \varsigma\theta) - \psi(\Delta). \tag{273}$$

Thus $q + h_\varsigma = T_{\varsigma,\theta}$ gives exact equality of the complete laws, including nonresponse.

For design $w_+ = w$, $w_- = 1 - w$, define the fixed-utility robust profile

$$K^{\mathrm{rc}}_w(\Delta, u, \zeta, \eta, \gamma) = \inf_{\substack{q \in \mathcal{Q}^0_1,\ |\theta| \leq \eta \\ \|h_\varsigma\|_{W_\varsigma(\zeta)} \leq \gamma}} \sum_\varsigma w_\varsigma D_{\mathrm{KL}}\left(P^{\varsigma,\zeta}_{\Delta,0,0,0} \middle\| P^{\varsigma,\zeta}_{u,q,\theta,h} \right). \tag{274}$$

The balanced and selectable decision profiles are

$$I^{\mathrm{bal}}_{\mathrm{rc}}(\Delta,\zeta,\eta,\gamma) = \min_{-u_0\le u\le 0} K^{\mathrm{rc}}_{1/2}(\Delta,u,\zeta,\eta,\gamma),$$
$$I^{\mathrm{sel}}_{\mathrm{rc}}(\Delta,\zeta,\eta,\gamma) = \max_{w\in[0,1]}\min_{-u_0\le u\le 0} K^{\mathrm{rc}}_{w}(\Delta,u,\zeta,\eta,\gamma). \tag{275}$$

The balanced compound action randomizes internally and reveals the protocol; the selectable experiment lets the learner choose the asymptotic frequency. They are distinct statistical experiments.

I.2 Critical tangent–normal geometry

Retain the notation of Appendix E.7: $x_j = y_j - \bar{y}$, $\nu = \sum_j \omega_j x_j^2$, $q_2 = (I - \Pi_1^{\omega})x^{\odot 2}$, $\rho = \sigma(a)$, $\lambda_0 = \rho(1-\rho)^2$, and $C_{\mathrm{sing}} = \frac{1}{2}\lambda_0\beta_2(a)^2\|q_2\|_\omega^2$. Put

$$A = \sqrt{\lambda_0\nu}, \qquad N = \sqrt{2C_{\mathrm{sing}}} = \sqrt{\lambda_0}\,\beta_2(a)\|q_2\|_\omega. \tag{276}$$

For $b_\varsigma(\zeta) = (1-\sigma(a+\varsigma\zeta))^{-1}$ define

$$\delta b(\zeta) = \frac{b_+(\zeta) - b_-(\zeta)}{2}, \qquad \bar{b}(\zeta) = \frac{b_+(\zeta) + b_-(\zeta)}{2}, \tag{277}$$

and the direct-effect-reduced tangent half-gap

$$m(d,\zeta,\eta) = \left[d\,|\delta b(\zeta)| - \bar{b}(\zeta)\eta\right]_+. \tag{278}$$

Finally define the joint tangent–normal radius

$$\Gamma_{\mathrm{crit}}(d,\zeta,\eta) = \sqrt{A^2 m(d,\zeta,\eta)^2 + N^2 d^4}. \tag{279}$$

The first term is the surviving first-order protocol contrast; the second is the one-protocol quartic normal curvature. A response-deviation ball acts on their joint information norm and therefore subtracts from Γ_{crit}, not from the two components separately.

Lemma 72 (Critical exact-target expansion)**.** *There are $d_0, c_0, K > 0$ such that, uniformly for*

$$0 < \Delta \le d = \Delta - u \le d_0, \quad |\zeta| \le c_0 d, \quad 0 \le \eta \le c_0 d^2, \quad |\theta| \le \eta, \tag{280}$$

there exist $v_\theta \in V_1$ and scalars g_θ satisfying

$$T_{\varsigma,\theta} = v_\theta + \varsigma g_\theta x + \beta_2(a)d^2 q_2 + E_{\varsigma,\theta}, \tag{281}$$
$$|g_\theta - (d\delta b(\zeta) - \bar{b}(\zeta)\theta)| \le Kd^3, \qquad \|E_{\varsigma,\theta}\|_2 \le Kd^3. \tag{282}$$

Moreover, $\|W_\varsigma(u,d,\zeta) - W_0\|_{\mathrm{op}} \le Kd$ and the $W_\varsigma(\zeta)$-balls of radius $\gamma = O(d^2)$ are at Hausdorff distance $O(d^3)$ from the corresponding W_0-balls.

Proof. Put $d_\varsigma = d - \varsigma\theta$. Taylor expansion of the log partition function gives

$$R_{\varsigma,\theta} = d_\varsigma x + \kappa_\varsigma \mathbf{1}, \qquad |\kappa_\varsigma| \le Kd^2.$$

Uniformly on the critical rectangle,

$$G_{a+\varsigma\zeta}(R) = b_\varsigma R + \beta_2(a+\varsigma\zeta)R^{\odot 2} + O(\|R\|_\infty^3).$$

The order-d and projected order-d^2 terms lie in V_1. Their protocol average is absorbed into v_θ. The half-difference of the linear score coefficients is exactly $d\delta b - \bar{b}\theta$. Every protocol difference generated by the projected quadratic term is $O(|\zeta|d^2 + d|\theta|) = O(d^3)$. The common normal projection of $R^{\odot 2}$ is $d^2 q_2$; replacing $\beta_2(a+\varsigma\zeta)d_\varsigma^2$ by $\beta_2(a)d^2$ costs $O(|\zeta|d^2 + d|\theta|) = O(d^3)$. This proves equation 281– equation 282. Analyticity and uniform positivity give the metric estimate. Equivalent positive-definite norms whose matrices differ by $O(d)$ have radius-$O(d^2)$ balls at Hausdorff distance $O(d^3)$. □

For $B = Nd^2$ and $\gamma \ge 0$ define the even function

$$\phi_{d,\gamma}(t) = \left[\sqrt{A^2t^2 + B^2} - \gamma\right]_+^2. \tag{283}$$

Lemma 73 (Balanced and selectable leading product geometry)**.** *The function $\phi_{d,\gamma}$ is convex and nondecreasing in $|t|$. For a fixed tangent half-gap $g \geq 0$, define*

$$F_w(g) = \frac{1}{2} \min_{c \in \mathbb{R}} \{ w\phi_{d,\gamma}(c-g) + (1-w)\phi_{d,\gamma}(c+g) \}. \tag{284}$$

Then $F_w(g)$ is concave and symmetric in w, so $\max_{w\in[0,1]} F_w(g) = F_{1/2}(g)$. Moreover,

$$F_{1/2}(g) = \frac{1}{2}\phi_{d,\gamma}(g). \tag{285}$$

Allowing $|\theta| \leq \eta$ replaces g by $m(d,\zeta,\eta)$ in equation 285.

Proof. On the region $\sqrt{A^2t^2+B^2} > \gamma$,

$$\phi''_{d,\gamma}(t) = 2A^2 \left(1 - \frac{\gamma B^2}{(A^2t^2+B^2)^{3/2}}\right) \geq 0,$$

because $\gamma < \sqrt{A^2t^2+B^2}$ there. On the complementary interval the function is zero, and the one-sided derivative at the joining points is zero; thus the whole function is convex. It is even and nondecreasing in $|t|$. For fixed c, the objective in equation 284 is affine in w, hence its infimum is concave. The change $c \mapsto -c$ shows $F_w = F_{1-w}$, so a concave symmetric function is maximized at $w = 1/2$. For $w = 1/2$, convexity and symmetry make $c = 0$ a minimizer, yielding equation 285.

The direct contrast changes the leading tangent half-gap from $d\delta b$ to $d\delta b - \bar{b}\theta$. The balanced value is nondecreasing in the absolute gap, so minimizing over $|\theta| \leq \eta$ gives exactly $m(d,\zeta,\eta)$. Since the full leading selectable objective is itself an infimum of functions affine in w and is invariant under protocol exchange, its maximum is likewise attained at $w = 1/2$. □

Proposition 74 (Dependence of the critical law on response-error coupling)**.** *The radial positive-part law is a consequence of the product response class. At the leading balanced geometry, write the two protocol residuals in an orthonormal tangent–normal coordinate system as*

$$r_+ = (X, Y), \qquad r_- = (-X, Y), \qquad X = Ag, \quad Y = Nd^2,$$

with $X, Y \geq 0$. Consider two joint response-error sets having the same protocolwise marginal ball $B_2(\gamma)$.

Independent product errors. *If $\mathcal{H}_{\mathrm{prod}} = B_2(\gamma) \times B_2(\gamma)$, then the leading KL profile is*

$$\mathcal{I}_{\mathrm{prod}}(X, Y, \gamma) = \frac{1}{2}\left[\sqrt{X^2+Y^2} - \gamma\right]_+^2. \tag{286}$$

Perfectly shared errors. *If $\mathcal{H}_{\mathrm{sh}} = \{(h,h) : \|h\|_2 \leq \gamma\}$, then the leading KL profile is*

$$\mathcal{I}_{\mathrm{sh}}(X, Y, \gamma) = \frac{1}{2}\left\{X^2 + [Y-\gamma]_+^2\right\}. \tag{287}$$

Hence identical protocolwise marginal balls do not determine the crossover coefficient. In particular, the product formula equation 286 cannot be asserted for an arbitrary correlated joint response-error set.

Proof. For the product class the two response perturbations are optimized separately. The Euclidean distance from either r_ς to $B_2(\gamma)$ is $[\sqrt{X^2+Y^2} - \gamma]_+$. The balanced product squared distance is the average of two identical squared distances, and the complete-law KL contributes one half of that product distance, proving equation 286.

For the shared class, the balanced product squared distance is

$$\min_{\|h\|_2 \leq \gamma} \frac{1}{2}\left(\|r_+ - h\|_2^2 + \|r_- - h\|_2^2\right).$$

Writing $h = (h_1, h_2)$ and expanding gives

$$\frac{1}{2}\left(\|r_+ - h\|_2^2 + \|r_- - h\|_2^2\right) = X^2 + h_1^2 + (Y - h_2)^2.$$

The minimum over the ball is attained at $h_1 = 0$ and $h_2 = \min\{Y, \gamma\}$, so the product squared distance is $X^2 + [Y-\gamma]_+^2$. Multiplication by the local KL factor $1/2$ proves equation 287. The two joint classes have the same coordinate projections, but the formulas differ whenever both the tangent component and the coupling constraint matter. □

Lemma 75 (Critical maximum-radius geometry)**.** *Let $B = Nd^2$ and let the direct-effect-adjusted tangent half-gap range over $g_\theta = d\delta b(\zeta) - \bar{b}(\zeta)\theta$ with $|\theta| \le \eta$. Then*

$$\min_{|\theta|\le\eta} \min_{c\in\mathbb{R}} \max_{\varsigma\in\{-1,+1\}} \sqrt{A^2(c-\varsigma g_\theta)^2 + B^2} = \Gamma_{\mathrm{crit}}(d,\zeta,\eta). \tag{288}$$

The same identity holds after replacing the Euclidean scalar coordinates by the corresponding orthogonal W_0 tangent and normal vectors.

Proof. For fixed g, the elementary inequality $\max\{|c-g|, |c+g|\} \ge |g|$ is an equality at $c = 0$. Because $r \mapsto \sqrt{A^2r^2 + B^2}$ is nondecreasing on $r \ge 0$, the inner minimum therefore equals $\sqrt{A^2g^2 + B^2}$. Minimizing over $|\theta| \le \eta$ replaces $|g|$ by $[d|\delta b| - \bar{b}\eta]_+ = m(d,\zeta,\eta)$, proving the identity. Exact tangent–normal orthogonality identifies the scalar norm with the W_0 product norm. □

Lemma 76 (Stability of critical squared distances)**.** *Let H_d and H_0 be positive-definite matrices on a fixed finite-dimensional space with eigenvalues in one compact interval and $\|H_d - H_0\|_{\mathrm{op}} \le Kd$. Let x_d, x_0 and nonempty compact sets S_d, S_0 satisfy*

$$\|x_d - x_0\|_2 \le Kd^3, \qquad d_H(S_d, S_0) \le Kd^3, \qquad \mathrm{dist}_{H_0}(x_0, S_0) \le Kd^2.$$

Assume that all nearest points involved lie in one fixed compact chart. Then

$$\left|\mathrm{dist}^2_{H_d}(x_d, S_d) - \mathrm{dist}^2_{H_0}(x_0, S_0)\right| \le K'd^5. \tag{289}$$

For a finite product with coordinate metrics satisfying the same bounds, the conclusion also holds uniformly after weighting the coordinate squared distances by any w in the simplex, provided every coordinate displacement is $O(d^2)$.

Proof. Uniform eigenvalue bounds make the Euclidean and H_0, H_d norms uniformly equivalent. Distance in a fixed norm is one-Lipschitz in its point argument and changes by at most the Hausdorff distance when the feasible set changes. Consequently,

$$|\mathrm{dist}_{H_0}(x_d, S_d) - \mathrm{dist}_{H_0}(x_0, S_0)| \le Kd^3.$$

Both distances are $O(d^2)$, so $|a^2 - b^2| \le (a+b)|a-b|$ gives an $O(d^5)$ change in squared distance. On the localized chart, every relevant displacement is $O(d^2)$; hence replacing H_0 by H_d changes its squared norm by at most $\|H_d - H_0\|_{\mathrm{op}} O(d^4) = O(d^5)$. Applying the two comparisons in both directions proves equation 289. For a finite weighted product, apply the coordinatewise estimate before multiplying by the weights. The weighted error is at most the maximum coordinate error because the weights sum to one. Hence no positive lower bound on a nonzero weight is required; zero-weight coordinates contribute nothing. □

Lemma 77 (Uniform robust KL reduction on the critical rectangle)**.** *There are $d_0, c_0, K > 0$ such that, uniformly on equation 280 with $0 \le \gamma \le c_0 d^2$,*

$$\left| K^{\mathrm{rc}}_{1/2}(\Delta, u, \zeta, \eta, \gamma) - \frac{1}{2}\left[\Gamma_{\mathrm{crit}}(d,\zeta,\eta) - \gamma\right]^2_+ \right| \le Kd^5, \tag{290}$$

$$\left| \max_{w\in[0,1]} K^{\mathrm{rc}}_w(\Delta, u, \zeta, \eta, \gamma) - \frac{1}{2}\left[\Gamma_{\mathrm{crit}}(d,\zeta,\eta) - \gamma\right]^2_+ \right| \le Kd^5. \tag{291}$$

Every exact minimizer lies in the common local chart. Every sampled residual has Euclidean norm $O(d^2)$, uniformly over all $w \in [0,1]$.

Proof. The candidate obtained from Lemmas 72 and 73 has product residual $O(d^2)$ and exact KL $O(d^4)$. The isolated-fiber gap therefore localizes every exact minimizer to the common target chart. Choose a sampled protocol ς_0 with $w_{\varsigma_0} \ge 1/2$; at an endpoint it is the only sampled protocol. Uniform Hessian coercivity gives $\|q^\star + h^\star_{\varsigma_0} - T_{\varsigma_0,\theta^\star}\|_2 = O(d^2)$. Because $\|h^\star_{\varsigma_0}\|_2 = O(d^2)$ and the two exact targets differ by $O(d^2)$ on the critical rectangle, $q^\star$ lies within $O(d^2)$ of both targets. Every sampled response deviation is itself $O(d^2)$, so every sampled residual is $O(d^2)$ uniformly in w, including weights tending to zero.

At the exact target, the complete-law KL expansion is

$$D_{\mathrm{KL}}(P\|Q_{T+h}) = \frac{1}{2}h^\top W_\varsigma(u,d,\zeta)h + O(\|h\|_2^3).$$

The cubic term is $O(d^6)$. Replacing the exact metrics and their balls by W_0 changes the weighted squared distance by $O(d^5)$. The target-set Hausdorff error in Lemma 72 is $O(d^3)$. The graphical uncertainty error in equation 271 contributes the same order. Since all leading residuals are $O(d^2)$, either perturbation changes the squared distance by $O(d^5)$. The compact nuisance set and its full tangent subspace agree on these local minimizers by nuisance interiority. Lemma 76 now gives matching upper and lower comparisons between the exact squared distance and the leading product-space squared distance. The uniform comparison also commutes with $\max_{w\in[0,1]}$ because $|\max_w f_w - \max_w g_w| \le \sup_w |f_w - g_w|$. Lemma 73 evaluates the leading distance for the balanced design and for the selectable maximum, proving equation 290– equation 291. □

I.3 Decision profile, exact radius, and the robust crossover law

Define the balanced and selectable decision profiles by minimizing the fixed-utility quantities over $u \le 0$. Define also the exact alternative-side response radius needed for complete matching,

$$\gamma^{\mathrm{rc}}_{\mathrm{alias}}(\Delta,\zeta,\eta) = \min_{-u_0\le u\le 0}\ \min_{|\theta|\le\eta}\ \min_{q\in\mathcal{Q}_1^0}\ \max_{\varsigma\in\{-1,+1\}} \|T_{\varsigma,\theta} - q\|_{W_\varsigma(\zeta)}. \tag{292}$$

The strict radius uses $u < 0$.

Theorem 78 (Product-structured robust quadratic–quartic crossover)**.** *There are $\Delta_0, c_0, K > 0$ such that, uniformly for*

$$0 < \Delta \le \Delta_0, \qquad |\zeta| \le c_0\Delta, \qquad 0 \le \eta \le c_0\Delta^2, \qquad 0 \le \gamma \le c_0\Delta^2, \tag{293}$$

and for $\diamond \in \{\mathrm{bal}, \mathrm{sel}\}$, with the exact product class equation 270 or compact product sets satisfying equation 271,

$$\left| I^\diamond_{\mathrm{rc}}(\Delta,\zeta,\eta,\gamma) - \frac{1}{2}\left[\Gamma_{\mathrm{crit}}(\Delta,\zeta,\eta) - \gamma\right]_+^2 \right| \le K\Delta^5. \tag{294}$$

The balanced compound and selectable experiments are distinct statistical experiments, but under the stated exchange-symmetric product class they have the same leading critical coefficient; $w = 1/2$ is a leading-order selectable optimizer. No such conclusion is asserted for an arbitrary correlated joint validity set. The exact alias radius satisfies

$$|\gamma^{\mathrm{rc}}_{\mathrm{alias}}(\Delta,\zeta,\eta) - \Gamma_{\mathrm{crit}}(\Delta,\zeta,\eta)| \le K\Delta^3. \tag{295}$$

Consequently, below the exact radius the closed profile is positive. Above it, an attained strict wrong-decision alias exists. At equality there is a zero closed profile; the strict infimum is either an unattained strict infimum at a boundary-only minimizer or an attained strict alias when some radius minimizer has $u < 0$. The boundary-only case rules out uniform correctness by the adaptive transcript continuity lemma of Appendix E.

Suppose, more specifically, that for a constant $\vartheta \in \mathbb{R}$ and $\kappa, \tau \ge 0$,

$$\zeta_\Delta = \vartheta\Delta + O(\Delta^2), \qquad \eta_\Delta = \kappa\Delta^2 + O(\Delta^3), \qquad \gamma_\Delta = \tau\Delta^2 + O(\Delta^3). \tag{296}$$

Then

$$\frac{I^\diamond_{\mathrm{rc}}(\Delta,\zeta_\Delta,\eta_\Delta,\gamma_\Delta)}{\Delta^4} \longrightarrow \frac{1}{2}\left[\sqrt{A^2\left(e^a|\vartheta| - b\kappa\right)_+^2 + N^2} - \tau\right]_+^2, \qquad b = (1-\rho)^{-1}. \tag{297}$$

Under only $o(1)$ convergence of the normalized uncertainty sets and scaled parameters, the same limit holds with an $o(\Delta^4)$ rather than an $O(\Delta^5)$ remainder.

Proof. The upper bound in equation 294 evaluates the fixed-utility profile at the boundary $u = 0$. For the lower bound, compact remote-fiber separation first restricts every near minimizer to $d \le M\Delta$. The function $d \mapsto \Gamma_{\mathrm{crit}}(d,\zeta,\eta)$ is nondecreasing for $d \ge \Delta$: both $d \mapsto [d|\delta b| - \bar{b}\eta]_+$ and $d \mapsto Nd^2$ are nondecreasing. Lemma 77 therefore bounds every fixed-utility value below by the boundary leading value minus $K\Delta^5$. This proves equation 294 without assuming that the leading coefficient is

positive. If the normalized limit in equation 297 is strictly positive, comparison with the boundary candidate further gives $d = \Delta + O(\Delta^2)$.

Lemma 75 evaluates the leading maximum-norm problem. The target expansion, metric perturbation, compact-class interiority, and graphical rate in equation 271 perturb that maximum radius by at most $O(d^3)$, so

$$\inf_{|\theta|\le\eta, q\in\mathcal{Q}_1^0} \max_{\varsigma} \|T_{\varsigma,\theta} - q\|_{W_\varsigma(\zeta)} = \Gamma_{\mathrm{crit}}(d,\zeta,\eta) + O(d^3).$$

The leading radius is nondecreasing in d, so minimization over $u \le 0$ gives equation 295. Exact complete-law matching is possible precisely when the response radius reaches equation 292. Compactness, density of $u < 0$ in $u \le 0$, slack above the radius, and adaptive transcript continuity give the strict/closed classification exactly as in Theorem 57.

Finally, $\delta b(\zeta) = e^a\zeta + O(\zeta^3)$ and $\bar{b}(\zeta) = b + O(\zeta^2)$. Substituting equation 296 into equation 278– equation 279 and then equation 294 proves equation 297. Qualitative graphical convergence changes the $O(d^3)$ target-set approximation to $o(d^2)$ and hence the $O(\Delta^5)$ squared-distance error to $o(\Delta^4)$. □

Corollary 79 (Regular, critical, singular, and zero-information regimes)**.** *Under the product dependence assumptions of Theorem 78, the theorem contains the following regimes.*

- *(i) If calibrated protocol separation stays bounded away from zero and the gap-scale direct-effect and response-sharing losses leave a positive margin, Theorem 50 gives regular quadratic information. This outer regime is not obtained by extrapolating the critical remainder in Theorem 78.*
- *(ii) If $|\zeta| \asymp \Delta$, $\eta \asymp \Delta^2$, and $\gamma \asymp \Delta^2$, the critical coefficient is the scaling function in equation 297.*
- *(iii) If $\zeta = o(\Delta)$, $\eta = o(\Delta^2)$, and $\gamma = o(\Delta^2)$, then $I^{\diamond}_{\mathrm{rc}} = C_{\mathrm{sing}}\Delta^4 + o(\Delta^4)$.*
- *(iv) If the response radius exceeds the exact radius in equation 295, an attained strict alias exists; at the exact radius the endpoint topology determines strict attainment.*

At exact validity $(\kappa,\tau) = (0,0)$, equation 297 reduces to $C_{\mathrm{reg}}\vartheta^2 + C_{\mathrm{sing}}$, recovering the exact-validity crossover. Direct-effect uncertainty shrinks only the tangent component at leading order, whereas response uncertainty can also consume the quartic normal component.

Corollary 80 (Row-wise fixed-confidence lower bound)**.** *Let $(\Delta_n,\zeta_n,\eta_n,\gamma_n)$ satisfy the quantitative or qualitative scalings above. Require almost-sure termination and δ_n-correctness on the same class $\mathcal{M}^{\mathrm{rc}}_{\zeta_n,\eta_n,\gamma_n}$ used in the information denominator. Then, for the balanced compound or selectable action set,*

$$\mathbb{E}_{m_n}[\tau_n] \ge \frac{\mathrm{kl}(1-\delta_n,\delta_n)}{I^{\diamond}_{\mathrm{rc}}(\Delta_n,\zeta_n,\eta_n,\gamma_n)}. \tag{298}$$

If the limit coefficient in equation 297 is positive and $\delta_n \downarrow 0$, this yields the corresponding Δ_n^{-4} critical lower-bound constant. If the exact radius is reached or exceeded, the strict/closed topology gives direct impossibility rather than a formal infinite-denominator statement.

Proof. This is Theorem 22 applied row by row to the explicit class equation 267. Almost-sure termination and the same strict alternative set are used in correctness and in the information denominator. □

J Complete assumption and failure-mode map

The phase diagram depends on structural, rate-critical, and scope assumptions that should not be conflated. Table 5 records the exact conclusion that fails when each is removed.

Assumption removed	Concrete failure	Surviving conclusion
Observe $\perp$ and exposure counts	Conditioning on response can make distinct complete laws identical	No answer-only version of the theorem
Sufficient effective score support	Degree-r interpolation creates an exact opposite-decision alias	Impossibility, not a slower rate
Isolated base observation fiber	A remote nuisance branch can dominate the local contact expansion	Only the local-chart profile is certified
Nuisance shared across protocols	Each protocol can cancel its own reward score independently	Multiple protocols need not restore order two
Fixed deployment reference or anchor protocol	The decomposition into a scalar utility and a common protocol effect is not unique; the sign can change with the reference	Only protocol-potential utilities, pairwise contrasts, or explicitly reference-indexed decisions are defined
Declared compact reference family	The sign may be positive for one acceptable deployment mixture and negative for another	The reference-robust decision is ambiguous; the exact ambiguity radius remains defined
Protocol-separable validity errors	Coordinatewise marginal intervals do not characterize a correlated joint set	The product-space coefficient remains valid, but the interval and pair laws may fail
Independent protocol-specific response errors in the critical crossover	Perfectly shared or otherwise correlated response errors can have the same marginal balls but a different tangent–normal contraction	The general normal-contact distance law survives, but the radial $[\Gamma_{\mathrm{crit}} - \gamma]_+$ formula and $w = 1/2$ conclusion may fail
Separated calibration intervals	The adversary chooses coincident intercepts	The finite-sample robust coefficient is exactly zero
Independent calibration pilot	Plug-in likelihood reuses data and ignores intercept uncertainty	Joint profiling or sample splitting is required
Analytic nuisance geometry	A flat nonanalytic contact can vanish to every finite order without local equality	Finite jet order need not characterize the profile
Compact graphical uncertainty and outer gap	Scaled cross-sections can oscillate or a minimizing nuisance can escape to a remote branch	No single robust coefficient or global transfer is certified
Row-indexed correctness class	The denominator changes with n but correctness is asserted on another class	The triangular-array lower bound is invalid

Table 5: The assumptions have identifiable mathematical roles; none is hidden as a generic "regularity condition."

Two distinctions are especially important. First, complete-law positivity is not only technical: it supplies a uniformly coercive KL Hessian and an inverse-Lipschitz response map. If category or response probabilities approach zero, the constants in Theorems 2–3 need not remain uniform, and the apparent support count can overstate usable information. The conditioning theorem quantifies the corresponding rare-mass and score-collision limits rather than claiming uniformity there.

Second, local contact and global identification are separate obligations. The normal derivative computes the best nuisance cancellation inside one embedded germ. It cannot exclude a distant parameter with the same complete law. That is why every unrestricted profile theorem includes an isolated-fiber gap, while Theorem 2 separately proves global sign identification by a repeated-Rolle argument. The same separation governs strict aliases: a zero infimum over an open wrong-decision set may be unattained, yet still rule out uniform correctness through continuity of adaptive transcript laws.

K CROSSOVER MODEL AND STANDING ASSUMPTIONS

Let the scored category alphabet be $\{1, \dots, J\}$, with distinct scores $y_1 < \cdots < y_J$ and positive base weights $\omega_j > 0$, normalized so that $\sum_j \omega_j = 1$. Define the finite exponential-tilt family

$$p_u(j) = \omega_j \exp\{uy_j - \psi(u)\}, \qquad \psi(u) = \log \sum_{k=1}^{J} \omega_k e^{uy_k}. \tag{299}$$

Let $\sigma(t) = (1 + e^{-t})^{-1}$. There are two protocol labels $\varsigma \in \{-1, +1\}$, with response intercepts

$$a_\varsigma(\zeta) = a + \varsigma\zeta. \tag{300}$$

The shared degree-one response nuisance is represented by its evaluation vector $q = (q_1, \dots, q_J)$ in

$$V_1 = \operatorname{span}\{\mathbf{1}, \boldsymbol{y}\} \subset \mathbb{R}^J, \qquad \boldsymbol{y} = (y_1, \dots, y_J)^\top. \tag{301}$$

Fix a compact admissible nuisance set $\mathcal{Q}_1 \subset V_1$ that contains a relative Euclidean neighborhood of zero. Every profiled nuisance optimization below is over $\mathcal{Q}_1$; whenever an unconstrained projection onto V_1 is used as a candidate, the local rectangle is reduced so that the candidate belongs to $\mathcal{Q}_1$. The complete law, including nonresponse, is

$$P^{\varsigma,\zeta}_{u,q}(j) = p_u(j)\sigma\big(a + \varsigma\zeta + q_j\big), \qquad j = 1, \dots, J, \tag{302}$$

$$P^{\varsigma,\zeta}_{u,q}(\perp) = 1 - \sum_{j=1}^{J} P^{\varsigma,\zeta}_{u,q}(j). \tag{303}$$

The true law is $P^{\varsigma,\zeta}_{\Delta,0}$, with $\Delta > 0$, and a wrong-decision utility satisfies $u \le 0$. Put

$$d = \Delta - u \ge \Delta. \tag{304}$$

For fixed $(\Delta, u, \zeta, \varsigma)$, define the exact single-protocol matching target

$$T_{\varsigma,j}(\Delta, u, \zeta) = \operatorname{logit}\left(\frac{p_\Delta(j)}{p_u(j)}\sigma(a + \varsigma\zeta)\right) - (a + \varsigma\zeta). \tag{305}$$

Whenever the argument of logit lies in $(0, 1)$, this target satisfies

$$P^{\varsigma,\zeta}_{u,T_\varsigma} = P^{\varsigma,\zeta}_{\Delta,0}. \tag{306}$$

For a residual $h \in \mathbb{R}^J$, define

$$F_\varsigma(\Delta, u, \zeta; h) := D_{\mathrm{KL}}\left(P^{\varsigma,\zeta}_{\Delta,0} \,\middle\|\, P^{\varsigma,\zeta}_{u,T_\varsigma + h}\right). \tag{307}$$

Assumption 81 (Local compact positivity and isolated fiber)**.** There exist compact intervals containing 0 in their interiors for the utility and nuisance coordinates such that:

(A1) all complete probabilities in the relevant local parameter neighborhood are bounded below by a common constant $p_* > 0$;
(A2) the nuisance set contains a Euclidean neighborhood of $0 \in V_1$;
(A3) at $(\Delta, u, \zeta) = (0, 0, 0)$, the only $q \in V_1$ in the chosen compact nuisance set satisfying $P^{\varsigma,0}_{0,q} = P^{\varsigma,0}_{0,0}$ is $q = 0$;
(A4) the same zero-law fiber is separated from the complement of a fixed local nuisance neighborhood by a positive KL gap.

Assumption 81 is exactly the finite-alphabet positivity, nuisance-interiority, and isolated-fiber condition used by the existing global profiling and explicit-remainder theorems. It is not an additional asymptotic assumption.

L UNIFORM CROSSOVER COERCIVITY

L.1 EXISTENCE OF A COMMON LOCAL RECTANGLE

Lemma 82 (Uniform existence and boundedness of the matching targets)**.** *Under Assumption 81, there exist $\epsilon_0 > 0$, $B_T < \infty$, and $\eta \in (0, 1/2)$ such that, whenever*

$$0 \le \Delta \le \epsilon_0, \qquad -\epsilon_0 \le u \le 0, \qquad d = \Delta - u \le \epsilon_0, \qquad |\zeta| \le \epsilon_0, \tag{308}$$

then the targets in equation 305 are well defined and satisfy

$$\|T_\varsigma(\Delta, u, \zeta)\|_2 \le B_T d, \tag{309}$$

for both ς. Moreover, the matched category-wise response probabilities

$$\ell_{\varsigma,j} := \sigma(a + \varsigma\zeta + T_{\varsigma,j}) = \frac{p_\Delta(j)}{p_u(j)}\sigma(a + \varsigma\zeta) \tag{310}$$

lie in $[\eta, 1 - \eta]$, uniformly over the rectangle.

Proof. At $(\Delta, u, \zeta) = (0, 0, 0)$, the argument of every logit in equation 305 equals $\sigma(a) \in (0, 1)$. The maps $u \mapsto p_u(j)$ are analytic and strictly positive. Since the alphabet is finite, continuity gives a common rectangle on which all arguments remain in a compact subinterval $[\eta, 1 - \eta] \subset (0, 1)$. This proves existence and the second assertion.

Write

$$R_j(\Delta, u) = \log \frac{p_\Delta(j)}{p_u(j)} = dy_j + \psi(u) - \psi(u + d). \tag{311}$$

On a compact utility interval, Taylor's theorem gives $\max_j |R_j(\Delta, u)| \le C_R d$. For

$$G_b(x) := \operatorname{logit}\big(\sigma(b)e^x\big) - b, \tag{312}$$

we have $T_{\varsigma,j} = G_{a+\varsigma\zeta}(R_j)$. The derivative of G_b is uniformly bounded on the compact set selected above, and $G_b(0) = 0$. The mean-value theorem therefore yields equation 309. □

L.2 EXACT RESIDUAL HESSIAN

Lemma 83 (Residual information matrix at exact matching)**.** *Fix a point in the rectangle of Lemma 82. Let*

$$t_j = p_\Delta(j)\rho_\varsigma, \qquad \rho_\varsigma = \sigma(a + \varsigma\zeta), \qquad t_\perp = 1 - \rho_\varsigma, \tag{313}$$

and let $b_j = 1 - \ell_{\varsigma,j}$, where $\ell_{\varsigma,j}$ is given by equation 310. Then

$$\nabla_h F_\varsigma(\Delta, u, \zeta; 0) = 0 \tag{314}$$

and

$$H^0_\varsigma := \nabla^2_h F_\varsigma(\Delta, u, \zeta; 0) = \operatorname{diag}(t_j b_j^2)_{j=1}^J + \frac{\boldsymbol{s}\boldsymbol{s}^\top}{t_\perp}, \qquad s_j = t_j b_j. \tag{315}$$

Consequently, there are constants $0 < m_0 \le M_0 < \infty$, independent of the local parameters and ς, such that

$$m_0 I_J \preceq H^0_\varsigma \preceq M_0 I_J. \tag{316}$$

Proof. For brevity write

$$q_j(h) = p_u(j)\sigma(a + \varsigma\zeta + T_{\varsigma,j} + h_j), \qquad q_\perp(h) = 1 - \sum_j q_j(h). \tag{317}$$

By equation 306, $q(0) = t$. Hence $F_\varsigma(0) = 0$, and the gradient vanishes because KL is minimized at equality.

A direct differentiation gives

$$\partial_k q_j(0) = \mathbf{1}\{j = k\}t_j b_j, \qquad \partial_k q_\perp(0) = -t_k b_k. \tag{318}$$

At equality, the Hessian of $D_{\mathrm{KL}}(t\|q)$ with respect to a smooth local coordinate equals the pullback of the multinomial information metric. Thus

$$v^\top H_\varsigma^0 v = \sum_{j=1}^{J} \frac{(t_j b_j v_j)^2}{t_j} + \frac{\left(\sum_j t_j b_j v_j\right)^2}{t_\perp} \tag{319}$$

$$= \sum_{j=1}^{J} t_j b_j^2 v_j^2 + \frac{\left(\sum_j t_j b_j v_j\right)^2}{t_\perp}, \tag{320}$$

which proves equation 315. Uniform positivity gives lower bounds on t_j, $b_j = 1 - \ell_{\varsigma,j}$, and $t_\perp$, and compactness gives upper bounds. In particular,

$$v^\top H_\varsigma^0 v \geq \left(\min_j t_j b_j^2\right) \|v\|_2^2, \tag{321}$$

while the operator norm is uniformly bounded above. This proves equation 316. □

Theorem 84 (Uniform residual coercivity)**.** *Under Assumption 81, there exist $\epsilon \in (0, \epsilon_0]$ and constants $0 < c_0 \leq C_0 < \infty$ such that, simultaneously for both protocols, all parameters satisfying equation 308 with ϵ in place of ϵ_0, and all h with $\|h\|_2 \leq \epsilon$,*

$$\nabla_h^2 F_\varsigma(\Delta, u, \zeta; h) \succeq c_0 I_J, \tag{322}$$

moreover

$$c_0\|h\|_2^2 \leq F_\varsigma(\Delta, u, \zeta; h) \leq C_0\|h\|_2^2. \tag{323}$$

The constants are uniform over the entire two-scale rectangle; in particular, they do not deteriorate when $|\zeta|/d$ tends to zero, a finite constant, or infinity.

Proof. On the compact set supplied by Lemma 82, enlarged by $\|h\|_2 \leq \epsilon_0$ and then reduced if needed, every complete alternative probability remains bounded below. Therefore F_ς is analytic in all arguments, and its Hessian is jointly continuous.

By Lemma 83, the minimum eigenvalue of the Hessian at $h = 0$ is uniformly at least m_0. Uniform continuity of the Hessian on the compact parameter set gives an $\epsilon > 0$ such that

$$\left\|\nabla_h^2 F_\varsigma(\Delta, u, \zeta; h) - H_\varsigma^0\right\|_{\mathrm{op}} \leq \frac{m_0}{2} \tag{324}$$

whenever $\|h\|_2 \leq \epsilon$. Hence

$$\nabla_h^2 F_\varsigma(\Delta, u, \zeta; h) \succeq \frac{m_0}{2} I_J. \tag{325}$$

A compact upper bound $M_1 I_J$ follows in the same way.

Since $F_\varsigma(0) = 0$ and $\nabla_h F_\varsigma(0) = 0$, the exact integral Taylor formula yields

$$F_\varsigma(h) = \int_0^1 (1-t) h^\top \nabla_h^2 F_\varsigma(th) h \, dt. \tag{326}$$

The Hessian bounds in equation 326 imply

$$\frac{m_0}{4}\|h\|_2^2 \leq F_\varsigma(h) \leq \frac{M_1}{2}\|h\|_2^2. \tag{327}$$

Reducing c_0 if necessary makes the same symbol valid in both equation 322 and equation 323. □

L.3 Target geometry and the missing residual estimate

Let Π_1 be the Euclidean orthogonal projector onto V_1, and let $\Pi_1^\perp = I - \Pi_1$.

Lemma 85 (Two-scale target geometry)**.** *Assume $J \geq 3$. There are nonzero vectors $r \in V_1$ and $n \in V_1^\perp$, and a constant $C < \infty$, such that uniformly over the local rectangle,*

$$\Pi_1 T_+ - \Pi_1 T_- = 2\zeta d\, r + e_T, \qquad \|e_T\|_2 \leq C\big(|\zeta| d^2 + |\zeta|^3 d\big), \tag{328}$$

$$\Pi_1^\perp T_\varsigma = d^2 n + e_{N,\varsigma}, \qquad \|e_{N,\varsigma}\|_2 \leq C\big(|\zeta| d^2 + d^3\big). \tag{329}$$

In the logistic model the leading vectors can be taken exactly as

$$r = e^a\big(\boldsymbol{y} - \psi'(0)\mathbf{1}\big), \qquad n = \frac{G_a''(0)}{2}\Pi_1^{\perp}(\boldsymbol{y}^{\odot 2}), \tag{330}$$

where $\boldsymbol{y}^{\odot 2} = (y_1^2, \dots, y_J^2)^\top$. *Both are nonzero: the scores are not all equal,* $e^a > 0$, $G_a''(0) > 0$, *and three distinct scores prevent* y^2 *from belonging to* $\operatorname{span}\{1, y\}$.

Proof. Write $\Delta = u + d$. The log-density ratio entering the target is

$$R_j = dy_j + \psi(u) - \psi(u+d) = d\big(y_j - \psi'(u)\big) - \frac{d^2}{2}\psi''(u) + O(d^3), \tag{331}$$

with a uniform remainder on a compact utility interval. Recall

$$T_{\varsigma,j} = G_{a+\varsigma\zeta}(R_j), \qquad G_b(x) = \operatorname{logit}(\sigma(b)e^x) - b. \tag{332}$$

At zero,

$$G_b'(0) = \frac{1}{1-\sigma(b)} = 1 + e^b, \qquad \frac{1}{2}G_b''(0) = \frac{\sigma(b)}{2(1-\sigma(b))^2} > 0. \tag{333}$$

Taylor expansion of G_b combined with equation 331 shows that the order-d term is $G_b'(0)d(\boldsymbol{y} - \psi'(u)\mathbf{1})$, hence lies in V_1. The order-d^2 term contains $\frac{1}{2}G_b''(0)d^2(\boldsymbol{y} - \psi'(u)\mathbf{1})^{\odot 2}$; after projection onto $V_1^{\perp}$, its leading nonzero part at $(u, \zeta) = (0, 0)$ is the vector n in equation 330. All other order-d^2 components lie in V_1, and the dependence on u contributes $O(d^3)$ because $|u| \le d$.

The difference between the two order-d coefficients is

$$G_{a+\zeta}'(0) - G_{a-\zeta}'(0) = 2\zeta e^a + O(\zeta^3), \tag{334}$$

which gives equation 328. Smooth dependence of all second-order coefficients on the intercept gives equation 329. Finally, $r \neq 0$ because the scores are nonconstant. If $y^2 = c_0 + c_1 y$ at three distinct score values, the quadratic polynomial $t^2 - c_1 t - c_0$ would have at least three roots and hence vanish identically, a contradiction. Thus $n \neq 0$. □

For weights $w_+ = w$, $w_- = 1 - w$, define the Euclidean target distance

$$D_w(\Delta, u, \zeta) = \inf_{q \in V_1} \sum_{\varsigma \in \{-1,+1\}} w_\varsigma \|q - T_\varsigma\|_2^2. \tag{335}$$

Corollary 86 (Uniform two-scale target distance)**.** *For every* $\eta_w \in (0, 1/2)$, *there are constants* $0 < c_D \le C_D < \infty$ *and a sufficiently small common rectangle such that*

$$c_D\big(\zeta^2 d^2 + d^4\big) \le D_w(\Delta, u, \zeta) \le C_D\big(\zeta^2 d^2 + d^4\big) \tag{336}$$

for all $w \in [\eta_w, 1 - \eta_w]$. *For arbitrary* $w \in [0, 1]$, *the upper bound remains valid.*

Proof. Write $V_\varsigma = \Pi_1 T_\varsigma$ and $N_\varsigma = \Pi_1^{\perp} T_\varsigma$. Orthogonality gives the exact identity

$$\begin{aligned} D_w &= w\|N_+\|_2^2 + (1-w)\|N_-\|_2^2 + \inf_{q \in V_1}\{w\|q - V_+\|_2^2 + (1-w)\|q - V_-\|_2^2\} \\ &= w\|N_+\|_2^2 + (1-w)\|N_-\|_2^2 + w(1-w)\|V_+ - V_-\|_2^2, \end{aligned} \tag{337}$$

where the minimizing tangent vector is $q = wV_+ + (1-w)V_-$. Substitute equation 328–equation 329. Because $r \neq 0$ and $n \neq 0$, the leading terms in equation 337 are respectively a positive multiple of $w(1-w)\zeta^2 d^2$ and a positive multiple of d^4. Every error term is bounded by a sufficiently small constant times $\zeta^2 d^2 + d^4$, uniformly on a reduced rectangle; this follows from Cauchy–Schwarz and $2|\zeta| d^3 \le \zeta^2 d^2 + d^4$. For $w \in [\eta_w, 1 - \eta_w]$, both leading coefficients are bounded below, proving the two-sided bound. Dropping the lower bound on $w(1-w)$ leaves the uniform upper bound for all $w \in [0, 1]$. □

Define the fixed-utility profiled KL

$$K_w(\Delta, u, \zeta) = \inf_{q \in \mathcal{Q}_1} \sum_{\varsigma} w_\varsigma F_\varsigma\big(\Delta, u, \zeta; q - T_\varsigma\big). \tag{338}$$

Theorem 87 (Profiled residual control)**.** *Continuity and compactness ensure that the infimum in equation 338 is attained. Let $q^* = q^*(\Delta, u, \zeta, w)$ be any exact minimizer, and set*

$$h_\varsigma^* = q^* - T_\varsigma. \tag{339}$$

Then, uniformly for $w \in [0,1]$,

$$\sum_\varsigma w_\varsigma \|h_\varsigma^*\|_2^2 \le C_H\big(\zeta^2 d^2 + d^4\big) \tag{340}$$

for a finite constant C_H. For every $\eta_w \in (0, 1/2)$, there are constants $0 < c_K \le C_K < \infty$ such that, uniformly over $w \in [\eta_w, 1-\eta_w]$,

$$c_K\big(\zeta^2 d^2 + d^4\big) \le K_w(\Delta, u, \zeta) \le C_K\big(\zeta^2 d^2 + d^4\big). \tag{341}$$

Proof. Let q^{cand} minimize the Euclidean target distance equation 335. By Corollary 86,

$$\sum_\varsigma w_\varsigma \|q^{\mathrm{cand}} - T_\varsigma\|_2^2 \le C_D(\zeta^2 d^2 + d^4). \tag{342}$$

For sufficiently small local parameters, the candidate lies in the nuisance interior. The upper half of Theorem 84 gives

$$K_w(\Delta, u, \zeta) \le C_0 C_D(\zeta^2 d^2 + d^4). \tag{343}$$

It remains to justify applying local coercivity to an arbitrary global minimizer. The right side of equation 343 tends to zero. Assumption 81(A3)–(A4) gives a positive KL gap outside a fixed local nuisance neighborhood at the base law. By continuity, the gap remains positive on a smaller parameter rectangle. Hence every exact minimizer is local for sufficiently small parameters. The lower half of Theorem 84 then yields

$$c_0 \sum_\varsigma w_\varsigma \|h_\varsigma^*\|_2^2 \le K_w(\Delta, u, \zeta). \tag{344}$$

Combining this with equation 343 proves equation 340.

For interior weights, apply Theorem 84 and minimize over $q \in \mathcal{Q}_1$. The lower bound may then be relaxed to the infimum over all of V_1, while the upper projection candidate lies in $\mathcal{Q}_1$ after reducing the rectangle:

$$c_0 D_w \le K_w \le C_0 D_w. \tag{345}$$

Corollary 86 now gives equation 341. □

Remark 88 (What the lemma closes). The estimate equation 340 is precisely the load-bearing statement needed to control the exact KL metric perturbation, the cubic Taylor remainder, and every cross term uniformly on the two-scale rectangle. It is stronger than an optimizer-specific heuristic: it follows from a uniform Hessian lower bound and applies to every exact profiled minimizer.

M TRIANGULAR-ARRAY FIXED-CONFIDENCE LOWER BOUNDS

The only safe way to let $\zeta = \zeta_n$ vary is to index the entire experiment class, correctness requirement, alternative set, and characteristic information by the same row n.

M.1 GENERIC ROW-WISE THEOREM

For each n, let $\mathcal{A}_n$ be a finite action set, $\mathcal{M}_n$ a model class, $d_n : \mathcal{M}_n \to \mathcal{D}_n$ a decision map, and $P_{m,a}^{(n)}$ the one-step observation law under model m and action a, on a finite alphabet $\mathcal{O}_{n,a}$. Assume that for each fixed row there is a constant $p_{*,n} > 0$ such that $P_{m,a}^{(n)}(o) \ge p_{*,n}$ for all m, a, o. No lower bound uniform in n is required. A row-n procedure consists of a predictable action rule, a stopping time τ_n, and an $\mathcal{F}_{\tau_n}$-measurable recommendation $\widehat{d}_n$.

Definition 89 (Row-wise correctness)**.** The row-n procedure is δ_n-correct on $\mathcal{M}_n$ if, for every model $m \in \mathcal{M}_n$ with a unique decision,

$$\mathbb{P}_m^{(n)}(\tau_n < \infty) = 1, \qquad \mathbb{P}_m^{(n)}\big(\widehat{d}_n = d_n(m)\big) \ge 1 - \delta_n. \tag{346}$$

For $x, y \in (0,1)$, write

$$\mathrm{kl}(x,y) = x \log \frac{x}{y} + (1-x) \log \frac{1-x}{1-y} \tag{347}$$

for the binary relative entropy.

For a true model $m \in \mathcal{M}_n$, define the strict wrong-decision set

$$\mathrm{Alt}_n^\circ(m) = \{m' \in \mathcal{M}_n : \ d_n(m') \neq d_n(m), \ d_n(m') \text{ unique}\} \tag{348}$$

and the row-wise characteristic information

$$\Gamma_n(m) = \sup_{w \in \Delta(\mathcal{A}_n)} \inf_{m' \in \mathrm{Alt}_n^\circ(m)} \sum_{a \in \mathcal{A}_n} w_a D_{\mathrm{KL}}\left(P_{m,a}^{(n)} \middle\| P_{m',a}^{(n)}\right). \tag{349}$$

Theorem 90 (Triangular-array transportation lower bound)**.** *For every n, every $\delta_n \in (0, 1/2)$, every row-n procedure that is δ_n-correct on $\mathcal{M}_n$, and every true model $m_n \in \mathcal{M}_n$ with a unique decision,*

$$\mathbb{E}_{m_n}^{(n)}[\tau_n] \geq \frac{\mathrm{kl}(1-\delta_n, \delta_n)}{\Gamma_n(m_n)}, \tag{350}$$

with the convention that the right side is $+\infty$ when $\Gamma_n(m_n) = 0$. No common ambient model class across rows is required.

Proof. Fix a row n, suppress the superscript (n), and fix a strict alternative $m' \in \mathrm{Alt}_n^\circ(m_n)$. Let

$$E_n = \{\widehat{d}_n = d_n(m_n)\} \in \mathcal{F}_{\tau_n}. \tag{351}$$

Correctness on the same row-wise class gives

$$\mathbb{P}_{m_n}(E_n) \geq 1 - \delta_n, \qquad \mathbb{P}_{m'}(E_n) \leq \delta_n. \tag{352}$$

If $\mathbb{E}_{m_n} \tau_n = \infty$, the claim is immediate. Otherwise, row-wise positivity and finite alphabets bound every one-step log-likelihood ratio by a finite row-dependent constant B_n. For $\tau_n \wedge T$, the ordinary chain rule gives the expected-count identity. The stopped log-likelihood ratios converge almost surely as $T \to \infty$ and are dominated in absolute value by $B_n \tau_n$, which is integrable. Dominated convergence therefore yields the stopped-history identity

$$D_{\mathrm{KL}}(\mathcal{L}_{m_n}(H_{\tau_n}) \,\|\, \mathcal{L}_{m'}(H_{\tau_n})) = \sum_{a \in \mathcal{A}_n} \mathbb{E}_{m_n}[N_{n,a}(\tau_n)] D_{\mathrm{KL}}(P_{m_n,a} \| P_{m',a}). \tag{353}$$

The action-selection factors cancel from the likelihood ratio because the procedure uses the same predictable kernel under both models.

Data processing from the stopped transcript to the binary event E_n, followed by equation 352, yields the next display: for $\delta_n < 1/2$, binary relative entropy is increasing in its first argument on $[1-\delta_n, 1]$ and decreasing in its second argument on $[0, \delta_n]$. Thus

$$\sum_a \mathbb{E}_{m_n}[N_{n,a}(\tau_n)] D_{\mathrm{KL}}(P_{m_n,a} \| P_{m',a}) \geq \mathrm{kl}(1-\delta_n, \delta_n). \tag{354}$$

Let

$$\bar{w}_{n,a} = \frac{\mathbb{E}_{m_n}[N_{n,a}(\tau_n)]}{\mathbb{E}_{m_n}[\tau_n]}. \tag{355}$$

Taking the infimum of equation 354 over strict alternatives gives

$$\mathbb{E}_{m_n}[\tau_n] \inf_{m' \in \mathrm{Alt}_n^\circ(m_n)} \sum_a \bar{w}_{n,a} D_{\mathrm{KL}}(P_{m_n,a} \| P_{m',a}) \geq \mathrm{kl}(1-\delta_n, \delta_n). \tag{356}$$

The infimum at the realized expected design is at most the supremum over all designs, namely $\Gamma_n(m_n)$. This proves equation 350. □

Remark 91 (Closed versus strict alternatives). If the paper defines the displayed profile over the closure of the wrong-decision set, the lower bound still uses strict alternatives. In the scalar crossover model, strict utilities $u < 0$ are dense in the closed set $u \leq 0$, and the finite-alphabet KL objective is continuous. Therefore the strict and closed infima are equal, although the closed infimum may be attained at $u = 0$ while the strict infimum is not.

M.2 SELECTABLE AND BALANCED CROSSOVER ARRAYS

Fix a compact utility interval $\mathcal{U} = [-u_0, u_0]$. For each ζ, let $\mathfrak{R}_\zeta^{\mathrm{sel}}$ be the model class $\{(u, q) : u \in \mathcal{U},\ q \in \mathcal{Q}_1\}$, with decision $\mathrm{sgn}(u)$ for $u \neq 0$ and two separately selectable actions $\varsigma \in \{-1, +1\}$. Let $\mathfrak{R}_\zeta^{\mathrm{bal}}$ have the same parameter set and decision map, but one compound action in which ς is drawn uniformly, revealed, and followed by the corresponding observation. Correctness is required only at the unique-decision models $u \neq 0$. For the true model $(\Delta, 0)$, define

$$\Gamma_\zeta^{\mathrm{sel}}(\Delta) = \sup_{w\in[0,1]} \inf_{-u_0\le u<0,\ q\in\mathcal{Q}_1} \sum_\varsigma w_\varsigma D_{\mathrm{KL}}\left(P_{\Delta,0}^{\varsigma,\zeta} \middle\| P_{u,q}^{\varsigma,\zeta}\right), \tag{357}$$

$$\Gamma_\zeta^{\mathrm{bal}}(\Delta) = \inf_{-u_0\le u<0,\ q\in\mathcal{Q}_1} \frac{1}{2}\sum_\varsigma D_{\mathrm{KL}}\left(P_{\Delta,0}^{\varsigma,\zeta} \middle\| P_{u,q}^{\varsigma,\zeta}\right). \tag{358}$$

The identity equation 358 follows from the KL chain rule because the compound protocol label has the same model-independent $1/2$ law under both models.

Corollary 92 (Class-aligned crossover lower bounds)**.** *Let $(\Delta_n, \zeta_n, \delta_n)$ be any sequence with $\Delta_n > 0$, $\delta_n \in (0, 1/2)$, and let $\diamond \in \{\mathrm{sel}, \mathrm{bal}\}$. Suppose that for every n, the row-n procedure is δ_n-correct on the same class $\mathfrak{R}_{\zeta_n}^\diamond$ used to define $\Gamma_{\zeta_n}^\diamond$. Then*

$$\mathbb{E}_{\Delta_n,\zeta_n}^\diamond[\tau_n] \ge \frac{\mathrm{kl}(1-\delta_n, \delta_n)}{\Gamma_{\zeta_n}^\diamond(\Delta_n)}. \tag{359}$$

Proof. Apply Theorem 90 separately at each row. For the selectable class, the row-wise design simplex is $[0, 1]$, giving equation 357. For the balanced class, there is one compound action, whose one-step KL is the average in equation 358. No theorem for a fixed ζ is applied across rows. □

The coefficient-level crossover expansion proved in Theorem 42 has the following uniform form on one fixed local rectangle: for each $\diamond \in \{\mathrm{sel}, \mathrm{bal}\}$,

$$\left|\Gamma_\zeta^\diamond(\Delta) - C_\lambda\zeta^2\Delta^2 - C_{\mathrm{sing}}\Delta^4\right| \le K_\diamond\left(\zeta^4\Delta^2 + \zeta^2\Delta^3 + |\zeta|\Delta^4 + \Delta^5\right), \tag{360}$$

where $C_\lambda, C_{\mathrm{sing}} > 0$.

Corollary 93 (Finite-row and asymptotic phase lower bounds)**.** *Assume $(\Delta_n, \zeta_n) \to (0, 0)$, and let the procedures satisfy the class alignment in Corollary 92. For all sufficiently large n,*

$$\mathbb{E}[\tau_n] \ge \frac{\mathrm{kl}(1-\delta_n, \delta_n)}{C_\lambda\zeta_n^2\Delta_n^2 + C_{\mathrm{sing}}\Delta_n^4 + K_\diamond R_n}, \tag{361}$$

where

$$R_n = \zeta_n^4\Delta_n^2 + \zeta_n^2\Delta_n^3 + |\zeta_n|\Delta_n^4 + \Delta_n^5. \tag{362}$$

If additionally $\delta_n \to 0$, then:

(i) if $|\zeta_n|/\Delta_n \to \infty$,

$$\liminf_{n\to\infty} \frac{\zeta_n^2\Delta_n^2\,\mathbb{E}[\tau_n]}{\log(1/\delta_n)} \ge \frac{1}{C_\lambda}; \tag{363}$$

(ii) if $\zeta_n/\Delta_n \to \kappa \in \mathbb{R}$,

$$\liminf_{n\to\infty} \frac{\Delta_n^4\,\mathbb{E}[\tau_n]}{\log(1/\delta_n)} \ge \frac{1}{C_\lambda\kappa^2 + C_{\mathrm{sing}}}; \tag{364}$$

(iii) if $|\zeta_n|/\Delta_n \to 0$,

$$\liminf_{n\to\infty} \frac{\Delta_n^4\,\mathbb{E}[\tau_n]}{\log(1/\delta_n)} \ge \frac{1}{C_{\mathrm{sing}}}. \tag{365}$$

Proof. The finite-row inequality follows from equation 359 and the upper half of equation 360. Also,

$$\frac{\mathrm{kl}(1-\delta_n, \delta_n)}{\log(1/\delta_n)} \longrightarrow 1. \tag{366}$$

In regime (i), every remainder term divided by $\zeta_n^2\Delta_n^2$ tends to zero:

$$\zeta_n^2, \quad \Delta_n, \quad \frac{\Delta_n^2}{|\zeta_n|}, \quad \frac{\Delta_n^3}{\zeta_n^2} \longrightarrow 0. \tag{367}$$

In regime (ii), divide by Δ_n^4; every remainder ratio tends to zero and the leading coefficient tends to $C_\lambda\kappa^2 + C_{\text{sing}}$. In regime (iii), division by Δ_n^4 gives remainder ratios ζ_n^4/Δ_n^2, ζ_n^2/Δ_n, $|\zeta_n|$, and Δ_n, all tending to zero. The stated liminf bounds follow. □

N CORRECT UNIVERSAL-MASKING QUANTIFIERS

Let $\mathcal{E}$ be a finite experiment set. Experiment e has a finite latent category set $\mathcal{J}_e$. Under latent model $i \in \{0,1\}$, let

$$p_i^e(j) > 0, \qquad \sum_{j\in\mathcal{J}_e} p_i^e(j) = 1. \tag{368}$$

An unrestricted selective-response mechanism consists of propensities $\rho_i^e(j) \in (0,1)$, and the complete observed kernel is

$$P_i^e(j) = p_i^e(j)\rho_i^e(j), \tag{369}$$

$$P_i^e(\perp) = 1 - \sum_j p_i^e(j)\rho_i^e(j). \tag{370}$$

Define the overlap ratio

$$\kappa(p_0,p_1) = \min_{e,j} \frac{\min\{p_0^e(j), p_1^e(j)\}}{\max\{p_0^e(j), p_1^e(j)\}} \in (0,1]. \tag{371}$$

Theorem 94 (Exact fixed-pair masking and its sharp response floor)**.** *For every fixed pair of strictly positive latent experiment families and every $\gamma \in (0,1)$, define*

$$q_e(j) = (1-\gamma)\min\{p_0^e(j), p_1^e(j)\}, \tag{372}$$

$$\rho_i^e(j) = \frac{q_e(j)}{p_i^e(j)}. \tag{373}$$

Then

$$(1-\gamma)\kappa(p_0,p_1) \le \rho_i^e(j) \le 1-\gamma < 1, \tag{374}$$

all complete kernels are identical,

$$P_0^e = P_1^e \qquad \forall e, \tag{375}$$

and their common nonresponse probability is at least γ.

Conversely, for any exact mask of this fixed latent pair, if

$$\alpha = \min_{i,e,j} \rho_i^e(j), \tag{376}$$

then

$$\alpha \le \kappa(p_0,p_1). \tag{377}$$

Thus the supremum achievable common response floor for the fixed pair is exactly $\kappa(p_0,p_1)$; it need not be close to one.

Proof. The bounds in equation 374 follow directly from equation 372–equation 373. The answered mass under either latent model is $q_e(j)$, so all answered probabilities agree. The nonresponse probabilities also agree because they are one minus the same sum. Moreover,

$$P_i^e(\perp) = 1 - (1-\gamma)\sum_j \min\{p_0^e(j), p_1^e(j)\} \ge \gamma. \tag{378}$$

This proves the constructive half.

For the converse, exact equality of the answered mass at (e, j) gives a common value

$$q_e(j) = p_0^e(j)\rho_0^e(j) = p_1^e(j)\rho_1^e(j). \tag{379}$$

If every propensity is at least α, then

$$q_e(j) \geq \alpha \max\{p_0^e(j), p_1^e(j)\}. \tag{380}$$

Since every propensity is at most one,

$$q_e(j) \leq \min\{p_0^e(j), p_1^e(j)\}. \tag{381}$$

Therefore α is at most the ratio in equation 371 for every (e, j), proving equation 377. Letting $\gamma \downarrow 0$ in the construction proves sharpness of the supremum. □

Corollary 95 (High-response masking requires close latent laws)**.** *Suppose a sequence of latent pairs admits exact masks with*

$$\alpha_n := \min_{i,e,j} \rho_{i,n}^e(j) \longrightarrow 1. \tag{382}$$

Then $\kappa(p_{0,n}, p_{1,n}) \to 1$, and hence, uniformly over the finite experiment-category set,

$$\frac{p_{0,n}^e(j)}{p_{1,n}^e(j)} \longrightarrow 1. \tag{383}$$

In particular,

$$\max_e \|p_{0,n}^e - p_{1,n}^e\|_{\mathrm{TV}} \longrightarrow 0. \tag{384}$$

Proof. The converse bound equation 377 gives $1 \geq \kappa_n \geq \alpha_n \to 1$. For positive numbers, the minimum-to-maximum ratio tending to one is equivalent to their ratio tending to one. Finiteness makes this convergence uniform. Total variation convergence follows componentwise. □

Theorem 96 (Arbitrarily rare censoring along an opposite-decision boundary path)**.** *Let $\theta \mapsto p_\theta^e(j)$ be continuous and strictly positive for every member of a finite experiment family. Suppose θ_0 lies on a decision boundary and there is a direction v such that, for every sufficiently small $\epsilon > 0$,*

$$\theta_+(\epsilon) = \theta_0 + \epsilon v, \qquad \theta_-(\epsilon) = \theta_0 - \epsilon v \tag{385}$$

produce opposite unique decisions. Then, for every prescribed response floor $\alpha \in (0, 1)$, there exists $\epsilon_\alpha > 0$ and exact masking propensities for the pair $(\theta_+(\epsilon), \theta_-(\epsilon))$, $0 < \epsilon \leq \epsilon_\alpha$, satisfying

$$\rho_i^e(j) \geq \alpha \qquad \forall i, e, j. \tag{386}$$

Thus exact nonidentification can persist with response probability arbitrarily close to one, but the latent pair changes with α and approaches the decision boundary.

Proof. By continuity and finiteness,

$$\kappa\big(p_{\theta_+(\epsilon)}, p_{\theta_-(\epsilon)}\big) \longrightarrow 1 \qquad (\epsilon \downarrow 0). \tag{387}$$

Fix $\alpha \in (0, 1)$, and choose

$$\gamma = \frac{1-\alpha}{2}, \qquad 1 - \gamma = \frac{1+\alpha}{2}. \tag{388}$$

Since $2\alpha/(1+\alpha) < 1$, choose ϵ_α so that

$$\kappa\big(p_{\theta_+(\epsilon)}, p_{\theta_-(\epsilon)}\big) \geq \frac{2\alpha}{1+\alpha} \tag{389}$$

for $0 < \epsilon \leq \epsilon_\alpha$. The construction of Theorem 94 then gives

$$\rho_i^e(j) \geq (1-\gamma)\kappa \geq \alpha. \tag{390}$$

The two decisions are opposite by assumption, and the complete laws are identical by Theorem 94. □

Proposition 97 (Converse in the scalar exponential-tilt decision model)**.** *In the scalar model equation 299, suppose the decision boundary is $u = 0$, with opposite decisions for $u_+ > 0 > u_-$. On any compact utility interval, if exact masks for pairs $(u_{+,n}, u_{-,n})$ have minimum response probability tending to one, then*

$$u_{+,n} \to 0, \qquad u_{-,n} \to 0. \tag{391}$$

Proof. Corollary 95 gives $\|p_{u_{+,n}} - p_{u_{-,n}}\|_{\mathrm{TV}} \to 0$. The mean-score map is

$$\mu(u) = \sum_j y_j p_u(j) = \psi'(u), \tag{392}$$

and

$$\mu'(u) = \psi''(u) = \mathrm{Var}_{p_u}(Y) > 0 \tag{393}$$

because the support contains at least two distinct scores. Hence μ is strictly increasing, and on a compact interval its inverse is uniformly continuous. Total variation convergence implies $|\mu(u_{+,n}) - \mu(u_{-,n})| \to 0$, so $|u_{+,n} - u_{-,n}| \to 0$. Since $u_{-,n} < 0 < u_{+,n}$, both are bounded in absolute value by $u_{+,n} - u_{-,n}$, proving the claim. □

Corollary 98 (Adaptive transcript invariance and decision impossibility)**.** *If two models have identical complete one-step kernels for every admissible experiment, then every adaptive transcript, stopped transcript, and randomized recommendation has the same law under the two models. If their unique correct decisions differ, every procedure has error probability at least $1/2$ on one of them. Consequently, no procedure can be uniformly δ-correct on a class containing both models for any $\delta < 1/2$.*

Proof. Couple the algorithm's internal random seed to be identical under both models. At time one, the experiment-selection kernel is the same. Conditional on the selected experiment, the observation kernels are equal, so the joint law of the first action and observation is equal. Inductively, conditional on any common past transcript, the algorithm uses the same next-action kernel, and the models use the same next-observation kernel. Thus every finite transcript law is equal. Equality extends to the infinite transcript and to every measurable stopped-history statistic.

Let A_i be the event that the recommendation equals model i's unique correct decision. These events are disjoint, while their probabilities are evaluated under one common recommendation law. Therefore

$$\mathbb{P}_0(A_0) + \mathbb{P}_1(A_1) \le 1, \tag{394}$$

so at least one success probability is at most $1/2$. □

Remark 99 (Correct quantifier statement). The valid pair of claims is:

(1) for every fixed strictly positive latent pair, there exists a strictly positive exact mask;
(2) for every target response floor $\alpha < 1$, there exists a sufficiently near-boundary opposite-decision pair admitting an exact mask with all responses at least α.

It is generally false that one fixed nonidentical latent pair can be masked with response probabilities arbitrarily close to one. The sharp obstruction is $\kappa(p_0, p_1) < 1$.

O CERTIFIED NUMERICAL VALIDATION

The computations are used as proof-audit certificates rather than as substitutes for the analytic theorems. We require four independent gates. First, rational and symbolic arithmetic verifies exact equality of every answered-category probability and the observed nonresponse probability for the alias and coarsening constructions. Second, interval Hessian subdivision and a Krawczyk operator certify finite-gap profile minimizers and enclose their values. Third, every analytic gradient used by optimization is checked against automatic differentiation, central finite differences, and arbitrary-precision directional differentiation. Fourth, every reported protocol optimizer is accompanied by primal feasibility, active constraints, stationarity, complementarity, and dual-feasibility residuals.

Table 6: Interval-certified fixed-boundary profile values. The interval Hessian cover proves global strong convexity on the declared nuisance box; Krawczyk inclusion proves existence and uniqueness of the stationary point.

Δ	lower bound	upper bound	relative gap	cover boxes
0.2	7.280208153322e-06	7.280208153322e-06	1.95e-95	8
0.15	2.329105707886e-06	2.329105707886e-06	6.59e-95	10
0.1	4.637217352930e-07	4.637217352930e-07	2.68e-94	10

Table 7: Certified-validation gates.

Gate	Status
Exact symbolic alias identities	PASS
Tie-deletion exact coarsening alias	PASS
Interval strong-convexity cover	PASS
Krawczyk root inclusion	PASS
One-sided utility KKT sign	PASS
Four-way gradient checks	PASS
Exact simplex KKT and dual feasibility	PASS

For the symmetric ternary degree-one benchmark with scores $(-1, 0, 1)$, weights $(7/20, 3/10, 7/20)$, and response probability $2/5$, define the fixed-boundary profile

$$K_\Delta^\partial = \min_{\eta \in [-0.05, 0.05] \times [0.05, 0.45]} D_{\mathrm{KL}}(P_{\Delta,0} \| P_{0,\eta}).$$

The interval cover proves global strong convexity on the full displayed closed box. An exact area-and-disjointness audit verifies that every reported subdivision box lies in the domain, that box interiors do not overlap, and that their union has the full domain area, including every boundary face. Strict Krawczyk inclusion then proves an interior stationary point; global strong convexity makes it the unique global minimizer. The resulting lower and upper bounds are shown in Table 6. The upper interval endpoint of the utility derivative at $u = 0$ is negative in every root box, yielding the active-bound KKT multiplier for the constraint $u \le 0$. The analytic localization theorem remains the proof that the unrestricted global profile enters this local chart.

The four symbolic alias cases and the tie-deletion coarsening certificate have exact zero discrepancy in the complete law. Across the scalar and two-protocol complete-KL objectives, the largest analytic–automatic-differentiation discrepancy is 3.65×10^{-17}, the largest analytic–central-difference discrepancy is 3.15×10^{-11}, and the largest analytic–90-digit directional discrepancy is 8.53×10^{-17}. For the four-protocol local coefficient with rational baselines $(1/8, 3/8, 5/8, 7/8)$, exact algebra gives

$$w^\star = \left(\frac{\sqrt{7} - 1}{6}, 0, 0, \frac{7 - \sqrt{7}}{6} \right), \qquad \Phi(w^\star) = 1 - \frac{\sqrt{7}}{4}.$$

The active KKT residuals and complementarity residual are exactly zero; the two inactive dual slacks are $-5/7 + 2\sqrt{7}/7 > 0$ and $-3/7 + 2\sqrt{7}/7 > 0$.

The anonymous supplement contains the complete certificate, machine-readable interval boxes, interval Jacobians, Krawczyk preconditioners and images, the exact cover audit, a clean-run transcript, the KKT ledger, deterministic tests, pinned dependencies, and SHA-256 manifests. Numerical output is described only as validation; all headline claims continue to rest on the analytic proofs.