LEARNING RECOURSE FROM DATA: SHARP COST RATES AND LIMITS OF OPTIMAL-ACTION RECOVERY

Anonymous authors
Paper under double-blind review

ABSTRACT

Counterfactual recourse identifies admissible changes that lead to a favorable prediction. It is usually computed for a fitted predictor treated as fixed, yet that predictor is itself estimated from finite data. We study recovery of the minimum recourse cost and the complete optimal-action set defined by an unknown population decision score. In smooth nonparametric binary regression with Euclidean cost and known translated convex action sets, we give design and boundary–action conditions under which score estimation error controls cost error uniformly over queries. We establish matching minimax bounds on explicit full-action and binding-box families, including cases where constraints change the minimizing action and its cost. The two targets nevertheless have different statistical limits: on a fixed regular family with full observational coverage, minimum cost is uniformly consistently estimable, but the complete exact optimal-action set is not. Arbitrarily small score perturbations can break ties between distant optima while barely changing cost. Uniqueness and uniform quadratic growth provide a positive counterpart through an upper bound for action-location recovery. Controlled experiments demonstrate finite-query Euclidean recovery and show that changing action constraints can amplify cost error for the same fitted score. These results distinguish estimating the minimum required change from recovering every minimum-cost choice.

1 INTRODUCTION

As machine-learning models increasingly inform consequential decisions, counterfactual explanations have become an important approach to explaining model behavior (Rudin, 2019; Wachter et al., 2018). They ask a question that ordinary prediction cannot answer: *for an individual receiving an unfavorable prediction, which changes would lead the model to make a favorable one?* Classical formulations seek nearby inputs that alter the model's prediction (Mothilal et al., 2020), while algorithmic recourse further requires changes that respect the individual's action constraints (Ustun et al., 2019; Karimi et al., 2021). Recent work further emphasizes the importance of actionability: a prediction may be fixed over the available actions, and features that are important for prediction need not be responsive for recourse (Kothari et al., 2024; Cheon et al., 2025). The minimum-cost formulation studied here seeks the nearest admissible input in the model's favorable region, as illustrated in Figure 1 (left).

A standard view of counterfactual explanation is post hoc: once a predictive model has been trained, it is treated as fixed, and recourse asks which feasible changes would alter its decision (Ustun et al., 2019). In many data-driven decision settings, however, the model being explained is itself learned from finite historical data. Credit assessment provides a concrete example: a lender may train a predictive model from past applications and repayment outcomes, and then use that fitted model to explain a rejection or identify changes that could alter its prediction (Cheon et al., 2025). The demand for explanations of credit decisions is reflected in U.S. Regulation B, which requires adverse-action notices to provide specific reasons or disclose the applicant's right to obtain them (Consumer Financial Protection Bureau, n.d.). The fitted predictor is therefore both an object to be explained and an estimate constructed from a finite sample. Its training process supplies statistical structure that can be used to study the recourse it induces. Figure 1 (right) illustrates this perspective: a learned score $\widehat{g}$ can induce a different favorable region and different optimal actions from those of the population score it estimates. The statistical question is how these induced quantities approach their population counterparts as data accumulate. Related work has already established that recourse can be sensitive to model retraining and model uncertainty (Rawal et al., 2020; Turbal et al., 2025), and statistical

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.

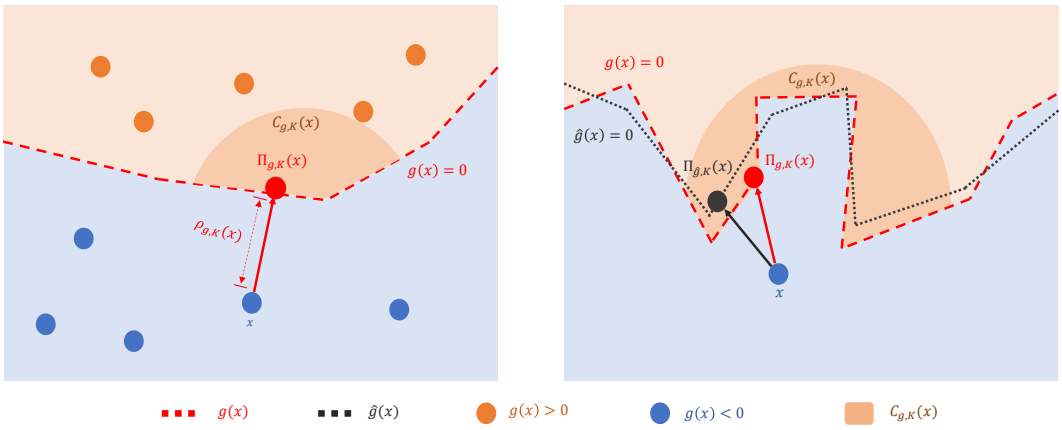


Figure 1: Counterfactual recourse and an estimated predictor: a schematic. Left: light orange and blue indicate favorable and unfavorable regions; darker orange denotes the feasible counterfactual set $C_{g,K}(x)$. Its minimum distance from x is $\rho_{g,K}(x)$, and its minimum-cost solutions form $\Pi_{g,K}(x)$. Right: the red dashed curve represents the zero-level boundary of the population score g; the black dotted curve represents that of the score $\widehat{g}$ learned from data. Background colors refer to g. The marked solutions illustrate the corresponding change in recourse.

validity guarantees have been developed for counterfactuals constructed from empirical generalized linear models (Carrizosa & Navas-Orozco, 2026). Complementing these validity and robustness guarantees, we ask: *which population recourse quantities can be recovered from the finite labeled data used to fit the predictor, and at what rate?* This question treats recourse as statistical recovery as well as post-hoc optimization.

To formalize this perspective, let $g(z) = \Pr(Y = 1 \mid X = z) - 1/2$ denote the *population decision score*, and view the fitted score $\widehat{g}$ as its finite-sample estimate. We study how well the minimum recourse cost and the complete optimal-action set induced by $\widehat{g}$ recover their population counterparts, uniformly over a fixed query set. This formulation lets us study identification, statistical rates, and the stability of optimal actions within a common framework. Specifically, we ask three questions. **First**, when can the minimum recourse cost be learned reliably? We establish sufficient conditions linking score accuracy to recourse recovery. Under regular design and boundary–action geometry, uniform score error controls feasible-set and minimum-cost error across the query set. Missing observational support can prevent identification, while tangential intersections can destabilize the feasible set even when the boundary converges. These are distinct observational and geometric obstacles. **Second**, how quickly can the minimum cost be learned? For smooth nonparametric binary regression, we obtain a query-uniform cost estimation rate and matching minimax lower bounds on explicit full-action and binding-box families. The latter retain constraints that change the minimizing action and its cost. **Third**, does reliable cost estimation imply reliable recovery of the optimal actions? We construct a fixed family satisfying the same cost-estimation conditions on which minimum cost is uniformly consistently estimable, whereas the complete exact optimal-action set admits no uniformly consistent estimator. The obstruction is a tie between distant optima that a small score error can break with little change in cost. Under uniqueness and uniform quadratic growth, we obtain a complementary plug-in upper bound for action location.

Model examples make the assumptions concrete for logistic, smooth feature-based, and smooth neural scores under explicit design and action conditions. Controlled experiments measure Euclidean cost and action-location recovery and isolate constraint effects at a fixed fitted score. On public datasets, we examine when geometric diagnostics explain refit-induced action displacement beyond score shift alone.

Contributions. We make three main contributions.

- *Statistical formulation and geometric error transfer.* We formulate recourse induced by a learned decision score as a population recovery problem. We establish explicit conditions under which uniform score estimation error controls feasible counterfactual sets and minimum costs across a fixed query set. Complementary constructions distinguish non-identification caused by missing data support from feasible-set instability at tangential boundary–action intersections.
- *Sharp learning rates under binding action constraints.* For smooth nonparametric binary regression, we establish a query-uniform cost estimation rate and matching minimax lower bounds on explicit full-action and binding-box families. In the latter, the constraints change the population optimum and remain binding under the statistically hard alternatives, establishing sharpness in genuinely constrained problems.
- *Statistical separation of cost and optimal-action recovery.* We construct a fixed family with full observational coverage and common geometric regularity constants on which minimum cost is uniformly consistently estimable, whereas the complete exact optimal-action set has nonvanishing minimax Hausdorff risk. Under uniqueness and uniform quadratic growth, we obtain a class-uniform plug-in upper bound for action location.

2 Related Work

Recourse and action constraints. Beyond minimum-cost prediction changes, recourse has causal, distributional, and global formulations (Karimi et al., 2021; Dominguez-Olmedo et al., 2022; Nguyen et al., 2022; Bewley et al., 2024). Diversity and preference elicitation address the value of alternative actions (Mothilal et al., 2020; De Toni et al., 2024). We study the minimum cost and complete optimal-action set defined by the population score under known constraints.

Robustness, multiplicity, and compatibility. Input sensitivity (Slack et al., 2021) differs from validity under model changes (Rawal et al., 2020; Black et al., 2022) and robust optimization against specified perturbations (Upadhyay et al., 2021; Jiang et al., 2024; Turbal et al., 2025). Multiple-output robustness studies guarantees for a collection of recommendations (Leofante & Potyka, 2024), and statistical validity for empirical generalized linear models is studied by Carrizosa & Navas-Orozco (2026). Robust Compatibility (RC) and Modular Robust Compatibility (MRC) distinguish location and distance sensitivity to nearby inputs, with empirical estimation and asymptotic learning guarantees (Xu et al., 2026). Our estimation targets are instead the population minimum-cost function and complete optimal-action set, with losses uniform over a fixed query set.

Boundary and level-set estimation. Level-set and smooth-discrimination theory supply statistical rates for boundary recovery (Tsybakov, 1997; Rigollet & Vert, 2009; Singh et al., 2009; Mammen & Tsybakov, 1999). Recourse additionally intersects the favorable region with query-dependent actions. Our tangency example shows why score and boundary convergence alone do not control these intersections.

Variational stability. Set convergence, intersection regularity, and optimizer growth are classical perturbation tools (Rockafellar & Wets, 1998; Bonnans & Shapiro, 2000; Dontchev & Rockafellar, 2014; Bauschke et al., 1999; Kruger, 2006). Our geometric bound verifies recourse-specific conditions relative to a regular population score, allowing arbitrary continuous score perturbations. Appendix H compares its quantifiers and residual bounds with classical implicit-map stability (Cuong & Kruger, 2021, Proposition 5.11).

3 Statistical Formulation

We now formalize recourse as a statistical recovery problem. The population decision score is unknown, while the action model and query set are fixed and known. From a finite labeled sample, we estimate the score and study how well the induced minimum recourse cost and complete optimal-action set recover their population counterparts, uniformly over queries.

Data and population targets. Let $\mathcal{X} = [0,1]^d$. We observe n independent pairs (X_i, Y_i) with $X_i \sim \mu$ and $Y_i \mid X_i = z \sim \text{Bernoulli}(1/2 + g(z))$, where $g(z) = \Pr(Y = 1 \mid X = z) - 1/2$ is the population decision score. A known compact convex template K defines the admissible destination set $A_K(x) = (x + K) \cap \mathcal{X}$. Thus z is an admissible destination and $z - x$ is the associated change. For an unfavorable query x, define

$$\begin{aligned} C_{g,K}(x) &= A_K(x) \cap \{z : g(z) \geq 0\}, \\ \rho_{g,K}(x) &= \min_{z \in C_{g,K}(x)} \|z - x\|_2, \qquad \Pi_{g,K}(x) = \operatorname*{argmin}_{z \in C_{g,K}(x)} \|z - x\|_2. \end{aligned} \tag{3.1}$$

The minimum recourse cost $\rho_{g,K}(x)$ is the value of this optimization problem. Its complete optimizer set $\Pi_{g,K}(x)$ is the *optimal-action menu*, indexed by destinations rather than changes. Full actions mean $A_{K_{\text{full}}}(x) = \mathcal{X}$. The query set $Q \subseteq \text{int}(\mathcal{X})$ is compact and fixed before sampling. Population feasible sets are nonempty. Each lower-bound construction fixes its own admissible (Q, K) before the sample size is chosen. The score and action assumptions are specified in Appendix A.

What constitutes recovery? For nonempty compact sets, use the symmetric Hausdorff distance

$$d_H(E, F) = \max\{\sup_{z \in E} \text{dist}(z, F), \sup_{z \in F} \text{dist}(z, E)\}.$$

We measure error simultaneously over the query set:

$$\begin{aligned} L_{\rho,Q}(\widetilde{\rho}, \rho) &= \sup_{x \in Q} |\widetilde{\rho}(x) - \rho(x)|, \\ L_{\Pi,Q}(\widetilde{\Pi}, \Pi) &= \sup_{x \in Q} d_H(\widetilde{\Pi}(x), \Pi(x)). \end{aligned} \tag{3.2}$$

We omit Q when it is fixed. In the second target, *exact* refers to population optimality; recovery is measured by Hausdorff error, not by exact finite-sample equality of sets. Section 5 explains why preserving the complete menu is more demanding than selecting one action. Write $\theta = (K, \mu, g)$ for a mechanism and P_θ^n for its sample law.

Statistical regimes. Three classes separate the requirements for different targets. $\Theta_{\text{out}}(Q)$ imposes the score-range, action, query, and feasibility conditions but permits missing observational support. $\Theta_{\text{val}}(\beta; Q)$ adds a design density bounded above and away from zero, Hölder smoothness $\beta \geq 2$, and regular boundary–action geometry. $\Theta_{\text{loc}}(\beta; Q)$ further requires a unique optimizer and uniform quadratic growth. The class constants and Q are fixed across members. Query-uniform losses control all queries simultaneously; class-uniform guarantees additionally take a supremum over mechanisms. Appendix A.4 gives complete definitions; Section 4.1 explains the geometry.

Plug-in estimation. We estimate g by boundary-corrected local-polynomial regression and solve (3.1) with g replaced by $\widehat{g}_n$. Appendix C specifies the estimator, its dependence on the confidence level, and the convention on singular samples. For the regular classes and sufficiently large n, it provides a continuous estimate and a uniform score bound $\|\widehat{g}_n - g\|_{\infty,\mathcal{X}} \leq e_{n,\delta}$ with probability at least $1 - \delta$, where

$$r_{n,\delta} = \left(\frac{\log(C_0 n/\delta)}{n}\right)^{\beta/(2\beta+d)}, \qquad e_{n,\delta} = C_e r_{n,\delta}. \tag{3.3}$$

We suppress the confidence index on the estimator. If an estimated feasible set is empty, define $\widehat{\rho}_n(x) = \text{diam}(\mathcal{X})$ and $\widehat{\Pi}_n(x) = \{x\}$ to keep the losses finite. Our high-probability value guarantee makes this convention inactive. The upper bounds analyze this specified estimator; the minimax lower bounds apply to all estimators of the stated target.

Identification comes before estimation. The observed law must contain information about the regions available for recourse. The following construction isolates failure of this requirement.

Theorem 1 (An off-support identification obstruction)**.** *There are a fixed admissible Q_{out} and two mechanisms $\theta_0, \theta_1 \in \Theta_{\text{out}}(Q_{\text{out}})$ with the same known K such that $P_{\theta_0}^n = P_{\theta_1}^n$ for every $n \geq 1$, but*

$$L_{\rho,Q_{\text{out}}}(\rho_{\theta_0}, \rho_{\theta_1}) \geq \Delta_\rho, \qquad L_{\Pi,Q_{\text{out}}}(\Pi_{\theta_0}, \Pi_{\theta_1}) \geq \Delta_\Pi$$

for fixed positive constants. The scores differ only on an open set of zero design probability containing an additional actionable threshold point.

No amount of sampling distinguishes these mechanisms. The construction and proof are in Appendix B.1; the value results below instead assume full observational coverage.

4 Learning the Minimum Recourse Cost

This section answers the first two questions in the introduction: when score accuracy controls recourse cost, and how rapidly that cost can be estimated. The two steps require different arguments. Geometry gives an upper bound; statistically indistinguishable alternatives establish its sharpness.

4.1 From score accuracy to recourse cost accuracy

Recourse depends on the intersection of a favorable region with $A_K(x)$. Regularity of the decision boundary alone does not control this intersection. We use two additional conditions, with constants shared over $x \in Q$. *Boundary–action coverage* requires

$$\operatorname{dist}(b, B_g \cap A_K(x)) \le \kappa_{\mathrm{cov}} \operatorname{dist}(b, A_K(x))$$

for boundary points b sufficiently close to the action set. *Radial crossing* requires, at each $c \in B_g \cap A_K(x)$, a unit vector v_c such that $c + [0, s_A]v_c \subset A_K(x)$ and $n_g(c)^\top v_c \ge \tau > 0$. It provides an actual feasible segment into the favorable side. The constant τ lower-bounds the feasible direction's alignment with the favorable normal. These geometric conditions are distinct from observational coverage. Together with a regular normal tube, a positive crossing slope a, and feasible depth, they yield the following deterministic result.

Theorem 2 (Query-uniform geometric transfer)**.** *Fix admissible (Q, K). Suppose g satisfies the gradient modulus, feasible-depth, boundary, and action conditions in Assumptions A.1–A.4. Let $\widetilde{g}$ be continuous with $\|\widetilde{g} - g\|_{\infty,\mathcal{X}} \le \varepsilon \le \varepsilon_{\mathrm{int}}$, where $\varepsilon_{\mathrm{int}}$ is defined in* (A.4)*. Then every perturbed feasible set is nonempty and*

$$\begin{aligned} &\sup_{x \in Q} d_H(C_{\widetilde{g},K}(x), C_{g,K}(x)) \le C_{\mathrm{int}}\varepsilon, \\ &\qquad L_\rho(\rho_{\widetilde{g},K}, \rho_{g,K}) \le C_{\mathrm{int}}\varepsilon, \qquad C_{\mathrm{int}} = \frac{1+\kappa_{\mathrm{cov}}}{a} + \frac{2}{a\tau}. \end{aligned} \tag{4.1}$$

This implication is deterministic and depends on the stated score and geometric conditions.

Why the transfer holds. A score bound sandwiches the estimated favorable set between $\{g \ge \varepsilon\}$ and $\{g \ge -\varepsilon\}$, so the two feasible sets can differ only near the population boundary. For an estimated-feasible point in this boundary band, the normal slope and boundary–action coverage move it to a genuine population-feasible intersection, at cost at most $(1+\kappa_{\mathrm{cov}})\varepsilon/a$. In the reverse direction, radial crossing moves a population-feasible boundary point a further $2\varepsilon/(a\tau)$ into the strictly favorable region, where it is also feasible for $\widetilde{g}$. These two corrections yield C_{int}. Distance to a set is Lipschitz in Hausdorff distance, giving the cost bound. The argument uses the geometry of the population score; the estimate $\widetilde{g}$ needs only continuity and a uniform error bound. Appendix B.2 supplies the full argument.

The next construction isolates the role of radial crossing: uniformly close scores and convergent decision boundaries can still induce feasible sets that remain a fixed distance apart.

Proposition 3 (Boundary convergence need not stabilize intersections)**.** *There exist fixed admissible Q, K and smooth scores g_0, g_ε with $\|g_\varepsilon - g_0\|_{\infty,\mathcal{X}} = O(\varepsilon)$ and $d_H(B_{g_\varepsilon}, B_{g_0}) = \varepsilon$, while all feasible sets remain nonempty and*

$$\liminf_{\varepsilon \downarrow 0} \sup_{x \in Q} d_H(C_{g_\varepsilon,K}(x), C_{g_0,K}(x)) > 0.$$

The baseline has regular boundary geometry and boundary–action coverage, but radial crossing fails at one query.

The perturbation removes a tangential feasible component while preserving recourse elsewhere. Full feasible-set recovery and scalar cost estimation therefore require different guarantees: the loss of a nonoptimal component can leave the minimum cost unchanged.

A concrete interpretation. For a logistic score $g(z) = \sigma(w^\top z + b) - 1/2$, the boundary slope is $\|w\|_2/4$. With sufficient domain slack, full actions admit the normal direction. A box can block that direction even though the score is unchanged; an available signed coordinate direction instead has alignment $|w_j|/\|w\|_2$. Appendix F gives complete model and domain conditions. This separates score accuracy from the directions available to a query.

4.2 Sharp cost rates under binding constraints

Combining Theorem 2 with a uniform score bound yields a plug-in estimator for the minimum recourse cost. Define $r_n = (\log n/n)^{\beta/(2\beta+d)}$ for $n \geq 2$. Could estimating the cost directly improve this rate? The next theorem rules out a uniform improvement on two explicit families, under full actions and under box constraints that change the minimizing action.

Theorem 4 (Uniform cost estimation and matching lower bounds)**.** *Let $d \geq 2$ and $\beta \geq 2$.* (i) Upper bound. *For each fixed admissible query problem, $C_{\mathrm{val}} = C_{\mathrm{int}}$ and, for $\delta \in (0, 1/2)$ and $n \geq n_{\mathrm{val}}(\delta)$,*

$$\sup_{\theta \in \Theta_{\mathrm{val}}(\beta;Q)} P_\theta^n\{L_\rho(\widehat{\rho}_n, \rho_{g,K}) > C_{\mathrm{val}} C_e r_{n,\delta}\} \leq \delta.$$

(ii) Full-action lower bound. *There exist a fixed query strip Q_{lb}, a full-action template K_{full}, and a nonempty $\Theta_{\mathrm{bump}}(\beta; Q_{\mathrm{lb}}) \subset \Theta_{\mathrm{val}}(\beta; Q_{\mathrm{lb}})$ such that, for all sufficiently large n,*

$$\inf_{\widetilde{\rho}_n} \sup_{\theta \in \Theta_{\mathrm{bump}}(\beta;Q_{\mathrm{lb}})} \mathbb{E}_\theta L_{\rho,Q_{\mathrm{lb}}}(\widetilde{\rho}_n, \rho_{g,K_{\mathrm{full}}}) \geq c_{\mathrm{val}} r_n.$$

(iii) Binding-box lower bound. *There also exist a fixed $K_\square = [-r, r]^d$, a compact $Q_\square$, and a nonempty $\Theta^\square_{\mathrm{bump}}(\beta; Q_\square) \subset \Theta_{\mathrm{val}}(\beta; Q_\square)$ such that*

$$\inf_{\widetilde{\rho}_n} \sup_{\theta \in \Theta^\square_{\mathrm{bump}}(\beta;Q_\square)} \mathbb{E}_\theta L_{\rho,Q_\square}(\widetilde{\rho}_n, \rho_{g,K_\square}) \geq c_\square r_n$$

for all sufficiently large n. At every hard query the baseline optimizer has a uniformly positive face multiplier; removing that face changes the optimizer and minimum cost, and every retained alternative keeps its optimizers on that face. Each infimum ranges over all measurable estimators given the sample and the corresponding known (Q, K).

Although cost is scalar at each query, estimating its function across queries retains nonparametric difficulty. The two lower bounds establish sharpness for their specified (Q, K) regimes, including a family with binding constraints and common regularity constants. The lower constructions use strict slope room $0 < a < \kappa < b$, whereas the upper bound permits $a = b$. Corollary C.10 places both bounds on the same expected-loss scale, completing the statistical rate comparison.

Proof idea. The central construction problem is to preserve common geometry and an active face while keeping the alternatives statistically indistinguishable. Localized bumps of width h and height h^β give

$$|\rho_j(x_j) - \rho_0(x_j)| \gtrsim h^\beta, \qquad D_{\mathrm{KL}}(P_j^n \| P_0^n) \lesssim n h^{2\beta+d}.$$

The proof first localizes each perturbed optimum; only then does a positive baseline multiplier rule out an interior solution. Packing $M_h \gtrsim h^{-(d-1)}$ alternatives and balancing their divergence against $\log M_h$ yields the rate. Appendices D.1 and D.2 give the full constructions.

5 Recovering Optimal Actions

We next ask whether recovering the minimum cost also preserves all choices attaining it. This is a stronger target than recommending a single action.

Why retain the full menu? Consider a two-stage decision: first minimize the primary cost, then select among the minimizers using a preference revealed later. Selecting one action at the first stage can discard the option preferred at the second. The complete optimal-action menu preserves all choices for this lexicographic objective. Its recovery loss has a precise secondary-utility interpretation. Let $\mathrm{Lip}_1(\mathcal{X})$ denote the real-valued 1-Lipschitz functions.

Proposition 5 (Secondary-utility interpretation). *For nonempty compact $E, F \subset \mathcal{X}$,*

$$d_H(E,F) = \sup_{u \in \mathrm{Lip}_1(\mathcal{X})} \left| \max_{z \in E} u(z) - \max_{z \in F} u(z) \right|.$$

The identity also holds with minima. Applying it querywise identifies $L_{\Pi,Q}$ with the worst distortion of the best attainable secondary utility over the two menus.

This interpretation connects menu recovery to downstream preferences. The next result asks whether the data needed to estimate the primary cost also suffice to preserve every exactly optimal option.

Theorem 6 (Statistical separation of cost and exact-menu recovery). *There exist a fixed admissible Q_{sep} and a smooth one-parameter family $\{g_t : 0 \le t \le t_0\}$ with common uniform design and K_{full} such that $\Theta_{\mathrm{sep}} \subset \Theta_{\mathrm{val}}(\beta; Q_{\mathrm{sep}})$ and*

$$\|g_t - g_0\|_\infty \le C_{\mathrm{sep}} t, \qquad L_{\rho, Q_{\mathrm{sep}}}(\rho_{g_t, K_{\mathrm{full}}}, \rho_{g_0, K_{\mathrm{full}}}) \le t.$$

Minimum cost is uniformly consistently estimable on this family, but there is a fixed $c_* > 0$ such that, for every $n \ge 1$,

$$\inf_{\widehat{\Pi}_n} \sup_{\theta \in \Theta_{\mathrm{sep}}} \mathbb{E}_\theta L_{\Pi, Q_{\mathrm{sep}}}(\widehat{\Pi}_n, \Pi_\theta) \ge c_*. \tag{5.1}$$

The infimum is over measurable estimators with nonempty compact values and measurable loss, given the sample and the known $(Q_{\mathrm{sep}}, K_{\mathrm{full}})$.

What the separation means. Full observational coverage and the value class's common regularity constants hold throughout this family. Consistent cost estimation can therefore coexist with irreducible worst-case error in preserving all exactly optimal options. The separation concerns uniform recovery of the *complete exact* menu over a fixed class. Approximate choices and preference-specific selectors define different statistical targets, as discussed in Appendix H.3.

Proof idea. Two separated favorable balls initially provide tied optima. Expanding one ball by t changes cost by at most t, while the other choice remains feasible but ceases to be exactly optimal. The menu distance is $2R_{\mathrm{sep}} - t$. Taking $t_n = c/\sqrt{n}$ keeps the sample laws difficult to distinguish while this gap remains positive. A two-point test then proves (5.1). The alternatives vary with n, but their parameter family does not (Appendix E).

Recovery under additional optimizer structure. A unique optimum and uniform quadratic growth provide sufficient additional structure for location recovery. Under Assumption A.5, excess cost controls squared displacement from the optimizer; the following result transfers this control to the estimator.

Theorem 7 (Conditional recovery of action location). *Fix Q and $\theta \in \Theta_{\mathrm{loc}}(\beta; Q)$. Let $\widetilde{g}$ be continuous and $\|\widetilde{g} - g\|_\infty \le \varepsilon$ with $\varepsilon \le \min\{\varepsilon_{\mathrm{int}}, m_Q/2\}$, $C_{\mathrm{int}}\varepsilon \le r_*/2$, and $4C_{\mathrm{int}}\varepsilon < \lambda r_*^2$. Then*

$$\sup_{x \in Q} d_H(\Pi_{\widetilde{g},K}(x), \{z_g(x)\}) \le C_{\mathrm{int}}\varepsilon + 2\sqrt{C_{\mathrm{int}}\varepsilon/\lambda} \le C_{\mathrm{loc}}\sqrt{\varepsilon},$$

where $C_{\mathrm{loc}} = C_{\mathrm{int}} + 2\sqrt{C_{\mathrm{int}}/\lambda}$. Consequently, for $n \ge n_{\mathrm{loc}}(\delta)$,

$$\sup_{\theta \in \Theta_{\mathrm{loc}}(\beta;Q)} P_\theta^n \{ L_\Pi(\widehat{\Pi}_n, \Pi_{g,K}) > C_{\mathrm{loc}} \sqrt{C_e r_{n,\delta}} \} \le \delta.$$

The feasible-set bound matches each estimated optimizer to a population-feasible point with $O(\varepsilon)$ excess cost. Quadratic growth then converts this objective gap into an $O(\sqrt{\varepsilon})$ displacement bound (Appendix E.4).

6 Experiments

We examine action-dependent error, finite-sample recovery, and geometric diagnostics after refitting. In the controlled experiments, errors are maximized over fixed queries within each fit and then averaged over independent fits. Figure 2 shows individual runs and aggregate behavior. Appendix G gives the populations, solvers, seeds, and complete records. Our code is available at `https://anonymous.4open.science/r/cf_rubust-3C33/`

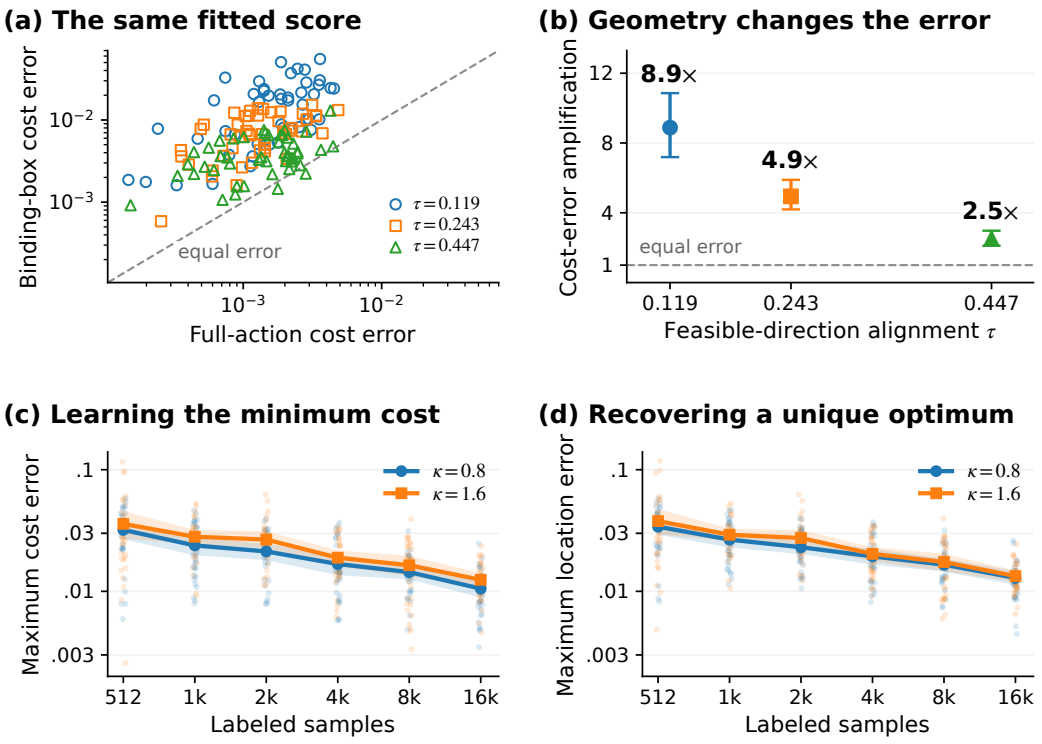


Figure 2: **Action constraints and statistical recovery.** (a) Each point pairs full-action and binding-box maximum cost errors from the same fit at $n = 32{,}768$: 48 fits at each of three alignments. (b) Ratios of mean binding-box to full-action errors, with paired-bootstrap 95% intervals; the dashed line is equality. (c,d) Euclidean cost and action-location errors over seven queries in two smooth populations. Faint points are individual fits; solid curves and bands are means and bootstrap 95% intervals over 30 fits per setting. Both recovery panels use the same axis scales. All displayed fitted problems are feasible.

Action constraints amplify error at a fixed score. We fit logistic regression to uniform-design Bernoulli samples and reuse each score under full, loose-box, binding-box, and tight-box actions. Across three alignments, seven sample sizes, and 48 replicates, this gives 1,008 fits. Within each alignment, all four action templates use the same 81 queries. Full and loose-box errors coincide. At $n = 32{,}768$, binding-box error exceeds full-action error in 142 of the 144 paired fits (Figure 2a). The ratios of mean errors are 8.9, 4.9, and 2.5 as feasible-direction alignment increases from 0.119 to 0.447 (panel b). At the smallest alignment, mean maximum cost error increases from 0.00191 to 0.01693, with every fitted binding-box problem feasible. The comparison isolates the error transfer through the available actions: the predictor and its score error are identical within each pair. Appendix G.1 includes the full sample-size sweep, feasibility measurements, and independent checks of the projection solver.

Euclidean cost and action locations improve with more data. We use two smooth binary-regression populations with parabolic favorable boundaries, indexed by curvature $\kappa \in \{0.8, 1.6\}$, and unique Euclidean projections. Local-linear estimates on a 33×33 grid define a continuous, piecewise-affine score. Enumerating its zero segments yields global Euclidean recourse for that interpolant; the population reference is computed independently from the analytic boundary. The 360 fits span six sample sizes and 30 replicates per population. As n increases from 512 to 16,384, the mean maximum cost error falls by factors of 3.0 and 2.9; location error falls by 2.6 and 2.8 (panels c,d). These are finite-query cost and Hausdorff errors for the interpolated score. Their decrease illustrates recovery in populations with the unique optima and quadratic growth used by Theorem 7.

Numerical resolution. At the largest sample size, fitting a 65×65 score grid to the same 60 training datasets changes mean maximum cost error by less than 0.6% and location error by less than 1.4%. Evaluating 25 instead of seven queries increases these errors by at most 0.9% and 5.3%, respectively. These checks separate score-grid and query-grid effects; Appendix G.2 reports the complete paired checks, including individual-query sensitivity.

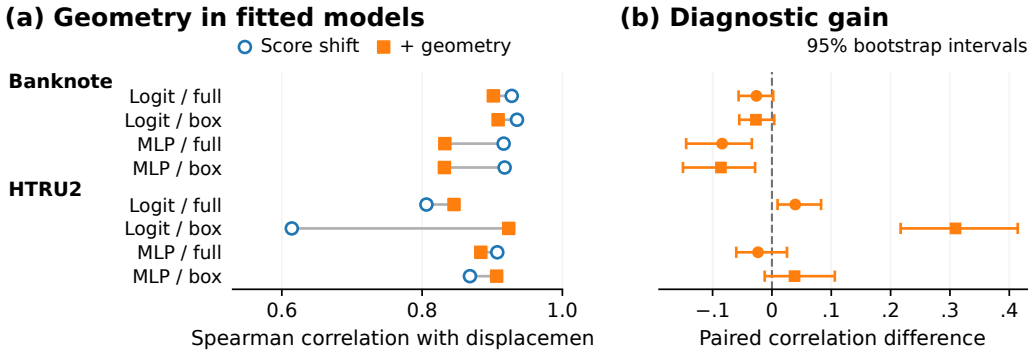


Figure 3: **Geometric diagnostics across model–action configurations.** (a) Spearman correlations between refit-induced optimizer displacement and score shift alone or a geometry-based index, for all eight dataset–model–action configurations. (b) Paired differences with 95% intervals from 2,000 crossed query/refit resamples, conditional on the three training splits. The fixed row order is shared across panels.

Geometry complements score-based diagnostics. On Banknote and HTRU2, we compare score shift with an index that additionally uses boundary slope, feasible-direction alignment, and local growth. Three splits and ten refits per split provide 7,600 eligible query–refit records. Figure 3 reports all eight configurations. For HTRU2 logistic–box, correlation rises from 0.61 to 0.92; adding alignment alone already yields 0.89, and the full-index gain is positive in each split. Among records with an active face, the same comparison is 0.71 versus 0.83; among records without one, it is 0.77 versus 0.82. Part of the combined gain thus reflects differences between action regimes. For Banknote MLP, score shift has the higher correlation: 0.92 versus 0.83. Appendix G.6 gives the complete fitting protocol, component analysis, and dependence-aware uncertainty procedure.

7 Conclusion

We studied the statistical recovery of population recourse from finite labeled data. Missing observational support can prevent identification, while regular design and boundary–action geometry yield uniform cost guarantees with matching minimax bounds on explicit full-action and binding-box families. Crucially, we construct a fixed regular family on which minimum cost is uniformly consistently estimable, whereas the complete exact optimal-action set admits no uniformly consistent estimator. Uniqueness and uniform quadratic growth provide a positive counterpart through an upper bound for action-location recovery. Controlled experiments illustrate action-dependent cost error at a fixed fitted score and finite-query Euclidean recovery as data accumulate.

These results distinguish learning how much change is required from preserving every minimum-cost choice. This distinction motivates studying approximate optimal-action sets with an explicit cost tolerance. Further questions concern weaker geometric conditions sufficient for cost estimation and sharp rates for action-location recovery. Extending the framework to discrete, sparse, or causal action models would broaden its scope beyond continuous features and known translated convex actions.

AI USE STATEMENT

AI-assisted tools were used for language polishing, grammar checking, and literature retrieval and discovery (e.g., identifying relevant related work). The authors take full responsibility for the content and integrity of the manuscript.

REFERENCES

Charalambos D. Aliprantis and Kim C. Border. *Infinite Dimensional Analysis: A Hitchhiker's Guide*. Springer, 3 edition, 2006. doi: 10.1007/3-540-29587-9.

Heinz H. Bauschke, Jonathan M. Borwein, and Wu Li. Strong conical hull intersection property, bounded linear regularity, Jameson's property (G), and error bounds in convex optimization. *Mathematical Programming*, 86(1):135–160, 1999. doi: 10.1007/s101070050083.

Tom Bewley, Salim I. Amoukou, Saumitra Mishra, Daniele Magazzeni, and Manuela Veloso. Counterfactual metarules for local and global recourse. In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pp. 3707–3724, 2024.

Emily Black, Zifan Wang, Matt Fredrikson, and Anupam Datta. Consistent counterfactuals for deep models. In *International Conference on Learning Representations*, 2022.

J. Frédéric Bonnans and Alexander Shapiro. *Perturbation Analysis of Optimization Problems*. Springer, 2000. doi: 10.1007/978-1-4612-1394-9.

Stéphane Boucheron, Gábor Lugosi, and Pascal Massart. *Concentration Inequalities: A Nonasymptotic Theory of Independence*. Oxford University Press, 2013. doi: 10.1093/acprof:oso/9780199535255.001.0001.

Emilio Carrizosa and Antonio Navas-Orozco. Building robust counterfactual explanations with statistical guarantees. *EURO Journal on Computational Optimization*, 14:100135, 2026. doi: 10.1016/j.ejco.2026.100135. URL https://doi.org/10.1016/j.ejco.2026.100135.

Seung Hyun Cheon, Anneke Wernerfelt, Sorelle Friedler, and Berk Ustun. Feature responsiveness scores: Model-agnostic explanations for recourse. In *International Conference on Learning Representations*, 2025. URL https://arxiv.org/abs/2410.22598.

Consumer Financial Protection Bureau. Regulation B: 12 CFR 1002.9, notifications. Electronic regulation and official interpretations, n.d. URL https://www.consumerfinance.gov/rules-policy/regulations/1002/9/. Current version accessed September 15, 2026.

Nguyen Duy Cuong and Alexander Y. Kruger. Uniform regularity of set-valued mappings and stability of implicit multifunctions. *Journal of Nonsmooth Analysis and Optimization*, 2:6599, 2021. doi: 10.46298/jnsao-2021-6599.

Giovanni De Toni, Paolo Viappiani, Stefano Teso, Bruno Lepri, and Andrea Passerini. Personalized algorithmic recourse with preference elicitation. *Transactions on Machine Learning Research*, 2024. URL https://openreview.net/forum?id=8sg2I9zXgO.

Ricardo Dominguez-Olmedo, Amir H. Karimi, and Bernhard Schölkopf. On the adversarial robustness of causal algorithmic recourse. In *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pp. 5324–5342, 2022.

Asen L. Dontchev and R. Tyrrell Rockafellar. *Implicit Functions and Solution Mappings: A View from Variational Analysis*. Springer, 2 edition, 2014. doi: 10.1007/978-1-4939-1037-3.

Junqi Jiang, Jianglin Lan, Francesco Leofante, Antonio Rago, and Francesca Toni. Provably robust and plausible counterfactual explanations for neural networks via robust optimisation. In *Proceedings of the 15th Asian Conference on Machine Learning*, volume 222 of *Proceedings of Machine Learning Research*, pp. 582–597, 2024.

Amir-Hossein Karimi, Bernhard Schölkopf, and Isabel Valera. Algorithmic recourse: From counterfactual explanations to interventions. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, pp. 353–362, 2021. doi: 10.1145/3442188.3445899.

Avni Kothari, Bogdan Kulynych, Tsui-Wei Weng, and Berk Ustun. Prediction without preclusion: Recourse verification with reachable sets. In *International Conference on Learning Representations*, 2024. URL https://arxiv.org/abs/2308.12820.

Steven G. Krantz and Harold R. Parks. *The Implicit Function Theorem: History, Theory, and Applications*. Birkhäuser, 2013. doi: 10.1007/978-1-4614-5981-1.

Alexander Y. Kruger. About regularity of collections of sets. *Set-Valued Analysis*, 14(2):187–206, 2006. doi: 10.1007/s11228-006-0014-8.

Francesco Leofante and Nico Potyka. Promoting counterfactual robustness through diversity. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pp. 21322–21330, 2024. doi: 10.1609/aaai.v38i19.30127.

Volker Lohweg. Banknote authentication. UCI Machine Learning Repository, 2012. CC BY 4.0.

Robert J. Lyon. HTRU2. UCI Machine Learning Repository, 2015. CC BY 4.0.

Enno Mammen and Alexandre B. Tsybakov. Smooth discrimination analysis. *The Annals of Statistics*, 27(6):1808–1829, 1999. doi: 10.1214/aos/1017939240.

Ramaravind K. Mothilal, Amit Sharma, and Chenhao Tan. Explaining machine learning classifiers through diverse counterfactual explanations. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, pp. 607–617, 2020. doi: 10.1145/3351095.3372850.

Tuan-Duy H. Nguyen, Ngoc Bui, Duy Nguyen, Man-Chung Yue, and Viet Anh Nguyen. Robust bayesian recourse. In *Proceedings of the 38th Conference on Uncertainty in Artificial Intelligence*, volume 180 of *Proceedings of Machine Learning Research*, pp. 1498–1508, 2022.

Kaivalya Rawal, Ece Kamar, and Himabindu Lakkaraju. Algorithmic recourse in the wild: Understanding the impact of data and model shifts. *arXiv preprint arXiv:2012.11788*, 2020.

Philippe Rigollet and Régis Vert. Optimal rates for plug-in estimators of density level sets. *Bernoulli*, 15(4):1154–1178, 2009. doi: 10.3150/09-BEJ184.

R. Tyrrell Rockafellar and Roger J.-B. Wets. *Variational Analysis*. Springer, 1998. doi: 10.1007/978-3-642-02431-3.

Cynthia Rudin. Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nature Machine Intelligence*, 1(5):206–215, 2019. doi: 10.1038/s42256-019-0048-x.

Saurabh Singh, Clayton Scott, and Robert Nowak. Adaptive hausdorff estimation of density level sets. *The Annals of Statistics*, 37(5B):2760–2782, 2009. doi: 10.1214/08-AOS661.

Dylan Slack, Anna Hilgard, Himabindu Lakkaraju, and Sameer Singh. Counterfactual explanations can be manipulated. In *Advances in Neural Information Processing Systems*, volume 34, pp. 62–75, 2021. URL https://proceedings.neurips.cc/paper_files/paper/2021/file/009c434cab57de48a31f6b669e7ba266-Paper.pdf.

Joel A. Tropp. User-friendly tail bounds for sums of random matrices. *Foundations of Computational Mathematics*, 12:389–434, 2012. doi: 10.1007/s10208-011-9099-z.

Alexandre B. Tsybakov. On nonparametric estimation of density level sets. *The Annals of Statistics*, 25(3):948–969, 1997. doi: 10.1214/aos/1069362732.

Alexandre B. Tsybakov. *Introduction to Nonparametric Estimation*. Springer, 2009. doi: 10.1007/b13794.

Bohdan Turbal, Iryna Voitsitska, and Lesia Semenova. ElliCE: Efficient and provably robust algorithmic recourse via the Rashomon sets. In *Advances in Neural Information Processing Systems*, volume 38, 2025.

Sohini Upadhyay, Shalmali Joshi, and Himabindu Lakkaraju. Towards robust and reliable algorithmic recourse. In *Advances in Neural Information Processing Systems*, volume 34, 2021.

Berk Ustun, Alexander Spangher, and Yang Liu. Actionable recourse in linear classification. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*, pp. 10–19, 2019. doi: 10.1145/3287560.3287566.

Sandra Wachter, Brent Mittelstadt, and Chris Russell. Counterfactual explanations without opening the black box: Automated decisions and the GDPR. *Harvard Journal of Law & Technology*, 31(2): 841–887, 2018.

Ao Xu, Zihao Li, Tieru Wu, and Yang Cao. Enhancing the robustness of counterfactual explanations via modular robust compatibility. *Expert Systems with Applications*, 304:130746, 2026. doi: 10.1016/j.eswa.2025.130746.

Appendix

The appendix is organized by the logical dependencies of the main results, followed by model examples, complete experimental records, and discussion. The contents and proof guide provide direct links to each result and its supporting argument.

Contents

A Population Assumptions and Statistical Classes **15**
A.1 Observation model and action constraints 15
A.2 Boundary geometry and feasible corrections 15
A.3 Additional structure for action location 16
A.4 Classes and quantifiers 16

B Identification and Geometric Transfer **16**
B.1 Identification without observational coverage 17
B.2 Query-uniform geometric transfer 19
B.3 Tangency and failure of feasible-set recovery 21

C Estimator Construction and Statistical Upper Bounds **21**
C.1 Estimator and admissible sample sizes 22
C.2 Uniform conditioning of the local fit 22
C.3 Bias and stochastic fluctuations 24
C.4 A continuous uniform score band 25
C.5 Cost recovery on the uniform score event 26
C.6 Expected risk and uniform consistency 26

D Sharp Lower Bounds for Recourse Cost **27**
D.1 Full-action minimax cost lower bound 27
D.2 Minimax sharpness with a binding box constraint 35

Reading order. Appendix A fixes the population classes. Appendix B establishes the deterministic and identification ingredients before Appendix C derives the statistical upper bounds. Appendix D verifies the separate hard families needed for minimax sharpness.

Contents (continued)

E Optimal-Action Recovery **42**
E.1 Secondary-utility interpretation of the menu loss 42
E.2 A fixed regular two-ball family 43
E.3 Statistical non-recovery of the exact menu 45
E.4 Conditional recovery under quadratic growth 45

F Model Examples and Action Geometry **46**
F.1 Logistic scores with full actions 47
F.2 Quadratic growth from Euclidean projection 48
F.3 Action constraints: full, box, and chartwise sets 48
F.4 Smooth feature-based scores 49
F.5 Neural scores and nonsmooth boundaries 50

G Experimental Protocols and Additional Results **51**
G.1 Changing actions at a fixed fitted score 51
G.2 Euclidean recovery with per-query records 53
G.3 Boundary-height tracking across dimensions 54
G.4 Selected graph-projection diagnostic 56
G.5 Boundary-face contact under score perturbations 57
G.6 Geometric diagnostics on public datasets 58
G.7 Exact constructions and the scope of empirical evidence 60

H Positioning and Scope **60**
H.1 From an estimated score to a constrained decision 60
H.2 Comparison with implicit multifunction stability 60
H.3 What exact-menu recovery asks of the data 61
H.4 Scope of the statistical model and empirical evidence 61

Main-result proof guide

Main-text result	Where to start
Theorem 1: identification	Section B.1, p. 17
Theorem 2: geometric transfer	Section B.2, p. 19
Proposition 3: tangency	Section B.3, p. 21
Theorem 4: value rates	Upper: Section C.5; lower: Sections D.1–D.2
Proposition 5: secondary utility	Section E.1, p. 42
Theorem 6: exact menus	Sections E.2–E.3, p. 43
Theorem 7: action location	Section E.4, p. 45

A Population Assumptions and Statistical Classes

The main results use three nested classes of mechanisms. This appendix fixes their common constants and separates population restrictions from estimator choices and from the reserves needed by particular lower-bound constructions. All conditions involving a query hold for every $x \in Q$ with the same constants. Within each statistical class, Q and these constants do not vary with the sample size or the member of the class.

A.1 Observation model and action constraints

Assumption A.1 (Design and smooth score)**.** The population decision score is continuous on $\mathcal{X} = [0,1]^d$ and takes values in $[-1/2+q, 1/2-q]$, where $q \in (0, 1/2)$ is fixed. In the regular classes, the design measure μ has a density p satisfying $0 < p_- \le p(z) \le p_+ < \infty$ almost everywhere on $\mathcal{X}$. The score belongs to the Hölder ball $\Sigma_d(\beta, L)$, with $\beta \ge 2$. For $\alpha_1 = \min\{1, \beta - 1\}$, let L_1 satisfy

$$\|\nabla g(z) - \nabla g(z')\|_2 \le L_1 \|z - z'\|_2^{\alpha_1}.$$

The score-range restriction keeps conditional Bernoulli probabilities away from zero and one. The design and smoothness restrictions support uniform score estimation; the outer identification class below retains only the continuity and score-range part of this assumption.

Assumption A.2 (Actions, queries, and feasible depth)**.** The known template K is nonempty, compact, and convex, contains zero, and satisfies $B(0, r_K) \subset K \subset B(0, R_K)$ for fixed positive radii. The query set is a nonempty compact set $Q \subset [s_Q, 1 - s_Q]^d$. Any additional strip or designated-query geometry is specified in the corresponding construction. For fixed $m_Q > 0$ and $0 < \underline{\rho}_Q \le \overline{\rho}_Q < \operatorname{diam}(\mathcal{X})$,

$$g(x) \le -m_Q, \qquad \underline{\rho}_Q \le \rho_{g,K}(x) \le \overline{\rho}_Q \quad (x \in Q).$$

Every population feasible set is nonempty, and

$$\inf_{x \in Q} \sup_{z \in A_K(x)} g(z) \ge \kappa_{\text{feas}} > 0.$$

The positive favorable-score reserve ensures that sufficiently accurate score estimation cannot remove every feasible point. It is stronger than bare nonemptiness and is used precisely when activating the plug-in problems.

A.2 Boundary geometry and feasible corrections

Assumption A.3 (Regular favorable boundary)**.** On $V_K = \bigcup_{x \in Q} A_K(x)$, the zero set $B_g = \{z : g(z) = 0\}$ is a nonempty compact C^1 hypersurface with favorable unit normal n_g. The normal map $T_g(b,t) = b + t n_g(b)$ is injective for $|t| \le r_0$. For fixed $0 < a \le b < \infty$,

$$a \le \nabla g(T_g(b_0, t))^\top n_g(b_0) \le b \quad (b_0 \in B_g,\ |t| \le r_0).$$

Outside the half-radius normal tube, $|g| \ge \gamma_0 > 0$ on V_K.

The normal slope converts small score residuals to distances from B_g. Tube injectivity identifies the relevant boundary point, while off-tube separation prevents additional near-zero components. Domain and chart qualifications for the model examples are made explicit in Appendix F.

Assumption A.4 (Regular boundary–action intersection)**.** There are $r_{\text{cov}} \le r_0/4$ and $\kappa_{\text{cov}} \ge 1$ such that

$$\operatorname{dist}(b, B_g \cap A_K(x)) \le \kappa_{\text{cov}} \operatorname{dist}(b, A_K(x)) \tag{A.1}$$

whenever $b \in B_g$ and $\operatorname{dist}(b, A_K(x)) < r_{\text{cov}}$. There are also $s_A > 0$ and $\tau \in (0, 1]$ such that every $c \in B_g \cap A_K(x)$ admits a unit vector v_c with

$$c + [0, s_A] v_c \subset A_K(x), \qquad n_g(c)^\top v_c \ge \tau. \tag{A.2}$$

Equation (A.1) corrects a nearby boundary point to an actual intersection. Equation (A.2) supplies a feasible segment of common positive length directed into the favorable side. A closed tangent-cone direction is not sufficient, and convexity alone does not supply these two properties. Appendix F.3 gives checks for full, box, and chartwise actions.

Uniform small-error threshold. The geometric proof uses the following constants:

$$s_{\mathrm{cross}} = \min\left\{ s_A, \frac{r_0}{2}, \left(\frac{a\tau}{2L_1}\right)^{1/\alpha_1} \right\}. \tag{A.3}$$

The last term is $+\infty$ when $L_1 = 0$. Define

$$\varepsilon_{\mathrm{int}} = \min\left\{ \frac{\gamma_0}{2}, \frac{ar_0}{8}, \frac{ar_{\mathrm{cov}}}{8}, \frac{\kappa_{\mathrm{feas}}}{2}, \frac{a\tau s_{\mathrm{cross}}}{2} \right\}. \tag{A.4}$$

This threshold collects the tube, boundary–action coverage, feasible-depth, and radial-crossing requirements. The transfer constant is $C_{\mathrm{int}} = (1 + \kappa_{\mathrm{cov}})/a + 2/(a\tau)$, and $C_{\mathrm{val}} = C_{\mathrm{int}}$ throughout.

A.3 Additional structure for action location

Assumption A.5 (Unique optimizer and uniform quadratic growth)**.** For every $x \in Q$, $\Pi_{g,K}(x) = \{z_g(x)\}$ and

$$\|z - x\|_2 - \rho_{g,K}(x) \geq \lambda \min\{\|z - z_g(x)\|_2^2, r_*^2\} \quad (z \in C_{g,K}(x))$$

for fixed $\lambda, r_* > 0$.

The inequality applies to every feasible point, with a truncated lower bound away from the optimum. Its local quadratic part converts objective error to displacement; its positive outer gap excludes distant approximate minimizers. The value class permits multiple optimal actions; the location class adds this growth condition and uniqueness.

A.4 Classes and quantifiers

For each fixed admissible Q, $\Theta_{\mathrm{out}}(Q)$ consists of mechanisms satisfying the continuity, score-range, action, query, and feasibility clauses of Assumptions A.1 and A.2. It may have incomplete design support and need not have Hölder smoothness. The value class $\Theta_{\mathrm{val}}(\beta; Q)$ additionally satisfies all of Assumptions A.1–A.4. The location class $\Theta_{\mathrm{loc}}(\beta; Q)$ adds Assumption A.5. For compatible common constants,

$$\Theta_{\mathrm{loc}}(\beta; Q) \subseteq \Theta_{\mathrm{val}}(\beta; Q) \subseteq \Theta_{\mathrm{out}}(Q).$$

The outer class supports the identification obstruction. Both value estimation and the exact-menu lower bound are studied in the value class; location recovery uses the smaller location class. Appendix C specifies the estimator separately. The positive margins in Appendices B.1, D, and E ensure that the constructed families satisfy these population conditions with common constants.

Table A.1: Which conditions support which conclusions?

Conditions	Role
Design and smoothness (Assumption A.1)	Uniform estimation of the population score
Actions and boundary (Assumptions A.2–A.4)	Nonempty perturbed problems and uniform feasible-set/value transfer
Optimizer growth (Assumption A.5)	Conversion of objective error to location error
Estimator specification (Assumption C.1)	Continuous local-polynomial estimator and admissible sample sizes
Construction reserves	Hard alternatives sharing the same class constants

B Identification and Geometric Transfer

This section separates two questions that precede rate calculations. Section B.1 proves Theorem 1 by constructing indistinguishable data laws. Section B.2 proves Theorem 2 without a sampling argument, and Section B.3 proves the feasible-set obstruction in Proposition 3. Only the population geometry in Appendix A is needed for the deterministic bridge.

B.1 IDENTIFICATION WITHOUT OBSERVATIONAL COVERAGE

We first specify a compatible outer construction and then realize it with two identical labeled-data laws. The added favorable set consists of a threshold point, which is included by the convention $g \geq 0$.

Assumption B.1 (Strict outer value room)**.** Fixed $r_{\mathrm{out}}, \Delta_{\mathrm{out}} > 0$ obey $\underline{\rho}_Q + 2\Delta_{\mathrm{out}} < r_{\mathrm{out}} < \overline{\rho}_Q - 2\Delta_{\mathrm{out}}$.

Assumption B.2 (Off-support outer closure)**.** $\Theta_{\mathrm{out}}(Q)$ is the maximal class satisfying the outer assumptions. If an outer member has an open μ-null set H and a continuous perturbation v supported in H such that $g + v$ preserves the score range and every outer query, value-range, and feasibility clause, then $(K, \mu, g + v) \in \Theta_{\mathrm{out}}(Q)$.

Construction B.1 (Off-support witness specification)**.** The following specification defines the pair used for the off-support obstruction. It consists of K_{out}, μ_{out}, continuous scores g^0_{out} and $g^1_{\mathrm{out}} = g^0_{\mathrm{out}} + v_{\mathrm{out}}$, an open μ_{out}-null set $H_{\mathrm{out}} \Subset \mathrm{int}(\mathcal{X})$ disjoint from Q, a continuous v_{out} supported in H_{out}, a compact favorable set $W_{\mathrm{out}} \subset H_{\mathrm{out}}$, and $x_{\mathrm{out}} \in Q$ such that $g^0_{\mathrm{out}} < 0$ on H_{out}, $D^+_{g^0_{\mathrm{out}}} \cap H_{\mathrm{out}} = \varnothing$ and $D^+_{g^1_{\mathrm{out}}} \cap H_{\mathrm{out}} = W_{\mathrm{out}}$. Both scores satisfy the range and query-margin conditions and have nonempty feasible sections. At the designated query,

$$\rho_{g^0_{\mathrm{out}}, K_{\mathrm{out}}}(x_{\mathrm{out}}) = r_{\mathrm{out}}, \quad \mathrm{dist}(x_{\mathrm{out}}, W_{\mathrm{out}}) = r_{\mathrm{out}} - \Delta_{\mathrm{out}}, \quad \mathrm{dist}(W_{\mathrm{out}}, \Pi_{g^0_{\mathrm{out}}, K_{\mathrm{out}}}(x_{\mathrm{out}})) > 0,$$

and $\inf_{x \in Q, w \in W_{\mathrm{out}}} \|w - x\|_2 \geq \underline{\rho}_Q$. Proposition B.3 constructs a perturbation satisfying these conditions.

Construction B.2 (Feasible outer witness)**.** The added favorable set is admissible at the designated query: $W_{\mathrm{out}} \subset A_{K_{\mathrm{out}}}(x_{\mathrm{out}})$.

Proposition B.3 (Realizability of the off-support witness chart)**.** *The outer construction specifications are nonempty. More precisely, for any $d \geq 2$ and $q \in (0, 1/2)$, the compatible specialization*

$$Q = Q_{\mathrm{out}} := U_Q \times \{1/5\}, \quad U_Q = [1/2 - a_Q, 1/2 + a_Q]^{d-1}, \quad a_Q = \frac{1}{100\sqrt{d-1}},$$

with

$$r_{\mathrm{out}} = \frac{3}{10}, \quad \Delta_{\mathrm{out}} = \frac{1}{25}, \quad \underline{\rho}_Q = \frac{1}{5}, \quad \overline{\rho}_Q = \frac{9}{20},$$

admits an explicit pair satisfying Constructions B.1 and B.2. The pair belongs to $\Theta_{\mathrm{out}}(Q_{\mathrm{out}})$ and has

$$P^n_{\theta^{\mathrm{out}}_0} = P^n_{\theta^{\mathrm{out}}_1} \quad (n \geq 1), \qquad \rho_0(x_\dagger) = \frac{3}{10}, \quad \rho_1(x_\dagger) = \frac{13}{50},$$

and $\Pi_0(x_\dagger) = \{b_{\mathrm{out}}\}$, $\Pi_1(x_\dagger) = \{w_{\mathrm{out}}\}$, where

$$\|b_{\mathrm{out}} - w_{\mathrm{out}}\|_2 = \frac{1}{25}.$$

Proof. Put $\lambda_{\mathrm{out}} = 1 - 2q$ and take the full action template $K_{\mathrm{out}} = [-1, 1]^d$. Let

$$x_\dagger = (1/2, \ldots, 1/2, 1/5), \quad b_{\mathrm{out}} = x_\dagger + (3/10)e_d, \quad w_{\mathrm{out}} = x_\dagger + (13/50)e_d,$$

and let $H_{\mathrm{out}} = B(w_{\mathrm{out}}, 3/100)$. Define

$$g^0_{\mathrm{out}}(z) = \lambda_{\mathrm{out}}(z_d - 1/2), \qquad \mu_{\mathrm{out}}(E) = \frac{\mathcal{L}^d(E \cap (\mathcal{X} \setminus H_{\mathrm{out}}))}{\mathcal{L}^d(\mathcal{X} \setminus H_{\mathrm{out}})}.$$

The last coordinate of $\overline{H}_{\mathrm{out}}$ lies in $[0.43, 0.49]$, while the query has last coordinate $1/5$. Hence $H_{\mathrm{out}} \Subset \mathrm{int}(\mathcal{X}) \setminus Q$ and the displayed measure is a probability with $\mu_{\mathrm{out}}(H_{\mathrm{out}}) = 0$. The full template gives $A_{K_{\mathrm{out}}}(x) = \mathcal{X}$. The score range is exactly

$$-1/2 + q \leq g^0_{\mathrm{out}}(z) \leq 1/2 - q,$$

and the baseline favorable set is the halfspace $\{z : z_d \geq 1/2\}$. For $x = (u, 1/5) \in Q$, its unique closest favorable point is $(u, 1/2)$, so the value is $3/10$ and the query score margin may be taken as $m_Q = \lambda_{\mathrm{out}}/5$. The value lies in $[1/5, 9/20]$, and the strict room inequality is

$$\frac{1}{5} + \frac{2}{25} = \frac{7}{25} < \frac{3}{10} < \frac{9}{20} - \frac{2}{25} = \frac{37}{100}.$$

Both scores retain $g(z) = \lambda_{\text{out}}/2$ on $z_d = 1$, so one may take $\kappa_{\text{feas}} = \lambda_{\text{out}}/4$. Thus the baseline satisfies all defining outer conditions. Moreover, $g^0_{\text{out}} < 0$ on H_{out}.

Let

$$\chi(r) = \left(1 - \frac{2r}{\Delta_{\text{out}}}\right)_+, \qquad v_{\text{out}}(z) = \lambda_{\text{out}}\Delta_{\text{out}}\chi(\|z - w_{\text{out}}\|_2), \qquad g^1_{\text{out}} = g^0_{\text{out}} + v_{\text{out}}.$$

The perturbation is continuous and supported in $\overline{B}(w_{\text{out}}, 1/50) \Subset H_{\text{out}}$. Write $z = w_{\text{out}} + h$ and $r = \|h\|_2$. Since $w_{\text{out},d} - 1/2 = -\Delta_{\text{out}}$, for $0 < r < 1/50$,

$$g^1_{\text{out}}(z) = \lambda_{\text{out}}(h_d - \Delta_{\text{out}}) + \lambda_{\text{out}}\Delta_{\text{out}}\left(1 - \frac{2r}{\Delta_{\text{out}}}\right) = \lambda_{\text{out}}(h_d - 2r) < 0,$$

because $h_d \le r$; at $r = 0$ the value is zero. At the cutoff edge and through the remainder of H_{out} the score is strictly negative, because $r_H = 3/100 < \Delta_{\text{out}}$. Consequently

$$D^+_{g^1_{\text{out}}} = D^+_{g^0_{\text{out}}} \cup \{w_{\text{out}}\}, \qquad W_{\text{out}} = \{w_{\text{out}}\}.$$

The perturbation is nonnegative but never creates a positive score on its support, so the score range is preserved. Queries are outside the support, and for $x = (u, 1/5) \in Q$,

$$\|x - w_{\text{out}}\|_2^2 = \|u - (1/2, \dots, 1/2)\|_2^2 + (13/50)^2 \in \left[(13/50)^2, \frac{1}{10000} + (13/50)^2\right].$$

Thus every perturbed query value belongs to $[13/50, \sqrt{677}/100] \subset [1/5, 9/20]$ and recourse remains nonempty. The threshold point belongs to $\mathcal{X} = A_{K_{\text{out}}}(x_\dagger)$, so it is an admissible destination. The off-support closure construction therefore puts both mechanisms in $\Theta_{\text{out}}(Q_{\text{out}})$.

The two mechanisms share $(K_{\text{out}}, \mu_{\text{out}})$ and their scores agree μ_{out}-almost everywhere, hence their one-sample labeled laws and all finite product laws coincide. Finally, w_{out} is the unique perturbed minimizer at $x_\dagger$, while b_{out} is the unique baseline minimizer. Their values differ by $3/10 - 13/50 = 1/25$, and their optimizer sets have Hausdorff distance $1/25$. This proves the claimed realization and all boundary cases of the cutoff. □

Proposition B.4 (Equal-law outer pair and exact gaps)**.** *Under the defining outer conditions and the explicit realization in Proposition B.3, there are two mechanisms $\theta_i = (K_{\text{out}}, \mu_{\text{out}}, g^i_{\text{out}})$ in $\Theta_{\text{out}}(Q_{\text{out}})$ such that, for every $n \ge 1$,*

$$P^n_{\theta_0} = P^n_{\theta_1}, \quad L_{\rho, Q_{\text{out}}}(\rho_{g^0_{\text{out}}, K_{\text{out}}}, \rho_{g^1_{\text{out}}, K_{\text{out}}}) \ge \Delta_{\text{out}},$$

and

$$L_{\Pi, Q_{\text{out}}}(\Pi_{g^0_{\text{out}}, K_{\text{out}}}, \Pi_{g^1_{\text{out}}, K_{\text{out}}}) \ge \operatorname{dist}(W_{\text{out}}, \Pi_{g^0_{\text{out}}, K_{\text{out}}}(x_{\text{out}})) > 0.$$

The cost functions are measurable, and the optimizer correspondences have measurable graphs and nonempty compact values on Q_{out}.

Proof. Write $D^+_i = D^+_{g^i_{\text{out}}}$ and $C_i(x) = C_{g^i_{\text{out}}, K_{\text{out}}}(x)$. The witness chart gives a common admissible template and design measure, continuous scores in the prescribed range, the query margin, and nonempty feasible sections for both scores. Proposition B.3 has proved these facts constructively, so the base member is in the outer class. The perturbation is continuous, supported in the open μ_{out}-null set H_{out}, and its chart checks preserve the outer conditions. The off-support closure assumption consequently places the perturbed member in the same class.

Each D^+_i is closed and each $A_{K_{\text{out}}}(x)$ is compact. Hence every $C_i(x)$ is compact and nonempty. Outside H_{out} the scores agree; inside it, the chart gives $D^+_0 \cap H_{\text{out}} = \varnothing$ and $D^+_1 \cap H_{\text{out}} = W_{\text{out}}$. Thus $D^+_1 = D^+_0 \cup W_{\text{out}}$. The reachability construction is used exactly here:

$$C_1(x_{\text{out}}) = (D^+_0 \cup W_{\text{out}}) \cap A_{K_{\text{out}}}(x_{\text{out}}) = C_0(x_{\text{out}}) \cup W_{\text{out}}.$$

This inclusion makes the off-support change relevant to the recourse target.

For any bounded measurable φ,

$$\mathbb{E}_{\theta_i}\varphi(X, Y) = \int_{\mathcal{X}} \big[\varphi(z, 1)(1/2 + g^i_{\text{out}}(z)) + \varphi(z, 0)(1/2 - g^i_{\text{out}}(z))\big]\, \mu_{\text{out}}(dz).$$

Because $v_{\mathrm{out}} = 0$ μ_{out}-almost everywhere, the integrands agree almost everywhere. Therefore $P_{\theta_0} = P_{\theta_1}$ and independence gives equality of every product law.

The chart states $\rho_0(x_{\mathrm{out}}) = r_{\mathrm{out}}$. Compactness of the added favorable set gives $w_* \in W_{\mathrm{out}}$ with $\|w_* - x_{\mathrm{out}}\|_2 = r_{\mathrm{out}} - \Delta_{\mathrm{out}}$. Since reachability puts w_* in $C_1(x_{\mathrm{out}})$,

$$\rho_1(x_{\mathrm{out}}) \le r_{\mathrm{out}} - \Delta_{\mathrm{out}}, \qquad L_\rho(\rho_0, \rho_1) \ge \Delta_{\mathrm{out}}.$$

Moreover, if $z \in \Pi_1(x_{\mathrm{out}})$ were in $C_0(x_{\mathrm{out}})$, then $\|z - x_{\mathrm{out}}\|_2 \ge r_{\mathrm{out}}$, contradicting the feasible point w_* at distance $r_{\mathrm{out}} - \Delta_{\mathrm{out}}$. Thus $\Pi_1(x_{\mathrm{out}}) \subset W_{\mathrm{out}}$. Put $\eta_{\mathrm{out}} = \mathrm{dist}(W_{\mathrm{out}}, \Pi_0(x_{\mathrm{out}})) > 0$. Then

$$d_H(\Pi_0(x_{\mathrm{out}}), \Pi_1(x_{\mathrm{out}})) \ge \sup_{z \in \Pi_1(x_{\mathrm{out}})} \mathrm{dist}(z, \Pi_0(x_{\mathrm{out}})) \ge \eta_{\mathrm{out}}.$$

Finally, the graph of $A_{K_{\mathrm{out}}}$ is $\{(x, z) : z - x \in K_{\mathrm{out}}\}$, a closed set. Intersecting it with the closed condition $g^i_{\mathrm{out}}(z) \ge 0$ gives a Borel compact-valued correspondence. We use the following compact-valued measurable maximum fact in this notation: if F has a Borel graph and nonempty compact sections in a compact metric space, and f is continuous, then

$$x \longmapsto \min_{z \in F(x)} f(x, z)$$

is Borel, while the argmin correspondence has nonempty compact sections and a measurable graph; see (Aliprantis & Border, 2006). The hypotheses here are exactly the closed graph, compact-section, and continuity conditions, so the result applies to $F = C_i$ and $f(x, z) = \|z - x\|_2$. Thus the values are Borel and the argmin graphs are measurable. The constants $\Delta_\rho = \Delta_{\mathrm{out}}$ and $\Delta_\Pi = \eta_{\mathrm{out}}$ are fixed by the fixed construction. □

B.2 Query-uniform geometric transfer

The deterministic argument uses three ingredients in order: localization near a regular zero set, feasible corrections between nearby score levels, and a strictly favorable witness that keeps the perturbed problem nonempty. All corrections are uniform over the query set.

Lemma B.5 (Signed level sandwich and tube localization)**.** *Under Assumption A.3, including its crossing clause, if $\|\widetilde{g} - g\|_{\infty,\mathcal{X}} \le \varepsilon$ and $\varepsilon \le \gamma_0/2$, then*

$$\{g \ge \varepsilon\} \subseteq \{\widetilde{g} \ge 0\} \subseteq \{g \ge -\varepsilon\}.$$

Every point in V_K with $-\varepsilon \le g < 0$ lies in the normal tube, and if $z = T_g(b, t)$ there, then

$$a|t| \le |g(z)| \le b|t|.$$

Proof. The inclusions follow from $\widetilde{g} \ge g - \varepsilon$ and $g \ge \widetilde{g} - \varepsilon$. If a point in V_K is at least $r_0/2$ from B_g, off-tube separation gives $|g| \ge \gamma_0$, so it cannot lie in $\{-\varepsilon \le g < 0\}$. The normal map is therefore valid. Along a normal fiber, $\varphi_b(t) = g(T_g(b, t))$ has derivative in $[a, b]$ and $\varphi_b(0) = 0$; integration gives the last display. □

Lemma B.6 (Feasible corrections between score levels)**.** *Let $\widetilde{g}$ be continuous with $\|\widetilde{g} - g\|_{\infty,\mathcal{X}} \le \varepsilon$. Assume the gradient modulus in Assumption A.1 and Assumptions A.3 and A.4. Suppose*

$$\varepsilon \le \min\left\{\frac{\gamma_0}{2}, \frac{a r_0}{8}, \frac{a r_{\mathrm{cov}}}{8}, \frac{a\tau s_{\mathrm{cross}}}{2}\right\}.$$

Then, uniformly in $x \in Q$,

$$\sup_{z \in A_K(x) \cap \{g \ge -\varepsilon\}} \mathrm{dist}(z, C_{g,K}(x)) \le \frac{1 + \kappa_{\mathrm{cov}}}{a}\varepsilon,$$

and every $z \in C_{g,K}(x)$ lies within

$$C_{\mathrm{int}}\varepsilon, \qquad C_{\mathrm{int}} = \frac{1 + \kappa_{\mathrm{cov}}}{a} + \frac{2}{a\tau},$$

of a point in $A_K(x) \cap \{g \ge \varepsilon\}$.

Proof. Fix $x \in Q$ and write $A = A_K(x)$ and $C = A \cap \{g \geq 0\}$. Take $z \in A \cap \{g \geq -\varepsilon\}$ with $g(z) < 0$. Lemma B.5 gives $z = T_g(b,t)$ for a point $b \in B_g$ and $|t| \leq \varepsilon/a$. Since $z \in A$,

$$\operatorname{dist}(b, A) \leq |t| < r_{\mathrm{cov}}.$$

The coverage inequality supplies $c \in B_g \cap A$ with $\|b - c\|_2 \leq \kappa_{\mathrm{cov}}|t|$. Hence

$$\operatorname{dist}(z, C) \leq \|z - c\|_2 \leq \frac{1 + \kappa_{\mathrm{cov}}}{a}\varepsilon. \tag{B.1}$$

The distance is zero when $g(z) \geq 0$.

For the reverse direction, take $z \in C$ with $g(z) < \varepsilon$. Tube localization gives $z = T_g(b,t)$ with $0 \leq t \leq \varepsilon/a$. Coverage again yields an actual point $c \in B_g \cap A$ and

$$\|z - c\|_2 \leq \frac{1 + \kappa_{\mathrm{cov}}}{a}\varepsilon. \tag{B.2}$$

Let v_c be the radial direction from Assumption A.4 and set $s = 2\varepsilon/(a\tau)$. The activation gives $s \leq s_{\mathrm{cross}} \leq s_A$, so $y = c + sv_c \in A$. Since $\nabla g(c)^\top v_c \geq a\tau$ and $L_1 s_{\mathrm{cross}}^{\alpha_1} \leq a\tau/2$,

$$\begin{aligned} g(y) - g(c) &= \int_0^s \nabla g(c + uv_c)^\top v_c \, du \\ &\geq \int_0^s (a\tau - L_1 u^{\alpha_1}) \, du \geq \frac{a\tau}{2} s = \varepsilon. \end{aligned}$$

Thus $y \in A \cap \{g \geq \varepsilon\}$ and

$$\|z - y\|_2 \leq \left(\frac{1 + \kappa_{\mathrm{cov}}}{a} + \frac{2}{a\tau}\right)\varepsilon.$$

The three displayed correction bounds prove both claims. At $\varepsilon = 0$ the sets coincide. The radial condition (A.2) makes each segment feasible for its full required length. □

Lemma B.7 (Nonempty estimated feasible sets)**.** *Let g be continuous and $\|\widetilde{g} - g\|_{\infty,\mathcal{X}} \leq \varepsilon$. Under Assumption A.2, if $\varepsilon \leq \kappa_{\mathrm{feas}}/2$, then for every $x \in Q$ there is $z_x \in A_K(x)$ with $g(z_x) \geq \kappa_{\mathrm{feas}}$, hence $z_x \in A_K(x) \cap \{g \geq 2\varepsilon\} \cap \{\widetilde{g} \geq 0\}$.*

Proof. Continuity and compactness of $A_K(x)$ make its score supremum attainable; feasible depth therefore supplies z_x. The score perturbation gives $\widetilde{g}(z_x) \geq \kappa_{\mathrm{feas}} - \varepsilon \geq \kappa_{\mathrm{feas}}/2$. Thus both the deep population set and the plug-in feasible set are nonempty. □

Proposition B.8 (Feasible-set and cost error transfer)**.** *Under the hypotheses of Lemmas B.5–B.7, with C_{int} as defined in* (4.1)*,*

$$\sup_{x \in Q} d_H(C_{\widetilde{g},K}(x), C_{g,K}(x)) \leq C_{\mathrm{int}}\varepsilon, \qquad \sup_{x \in Q} |\rho_{\widetilde{g},K}(x) - \rho_{g,K}(x)| \leq C_{\mathrm{int}}\varepsilon.$$

Proof. The level sandwich gives

$$A_K(x) \cap \{g \geq \varepsilon\} \subseteq C_{\widetilde{g},K}(x) \subseteq A_K(x) \cap \{g \geq -\varepsilon\}.$$

The first directed bound follows from Lemma B.6. For the reverse direction, take $y \in C_{g,K}(x)$. The second correction in that lemma gives $y' \in A_K(x) \cap \{g \geq \varepsilon\}$ with $\|y - y'\|_2 \leq C_{\mathrm{int}}\varepsilon$. Since $\widetilde{g}(y') \geq 0$, y' belongs to the plug-in set. Thus both directed distances are at most $C_{\mathrm{int}}\varepsilon$, proving the Hausdorff bound. Lemma B.7 gives nonemptiness. Compactness follows because $A_K(x)$ is compact and the superlevel sets of the continuous scores g and $\widetilde{g}$ are closed.

For nonempty compact A, B, the triangle inequality gives

$$|\operatorname{dist}(x, A) - \operatorname{dist}(x, B)| \leq d_H(A, B).$$

Apply this pointwise and take the supremum over Q. □

Proof of Theorem 2. Apply Proposition B.8 with the displayed activation in Theorem 2. Lemma B.7 makes both compact feasible sets nonempty. Lemma B.6 bounds the two directed set distances by $C_{\mathrm{int}}\varepsilon$, so their maximum is the claimed Hausdorff distance. Finally, $|\operatorname{dist}(x, E) - \operatorname{dist}(x, F)| \leq d_H(E, F)$ for nonempty compact E, F, which gives the value statement uniformly over Q. □

B.3 TANGENCY AND FAILURE OF FEASIBLE-SET RECOVERY

The following construction preserves boundary convergence and nonempty recourse but removes one tangential feasible component. It proves the feasible-set claim in Proposition 3.

Proof of Proposition 3. Work in $X = [0,1]^2$. Let

$$o = (1/2, 1/2), \quad R = 1/5, \quad \delta = 1/50, \quad K = \overline{B}(0, R),$$

$$Q = \{o, o \pm \delta e_1, o \pm \delta e_2\}, \qquad d = o + (2/25, 2/25), \quad r = 1/100,$$

and put $D = \overline{B}(d, r)$. There is no cube clipping, so $A_K(x) = \overline{B}(x, R)$. The displayed clearances imply that D lies strictly inside every $A_K(x)$ while no query lies in D.

For $0 \le \varepsilon \le 1/100$, define

$$g_\varepsilon(z) = -\kappa \left(\|z - o\|_2^2 - (R + \varepsilon)^2 \right) \left(r^2 - \|z - d\|_2^2 \right), \tag{B.3}$$

where $\kappa > 0$ is small enough for any fixed score-range and finite Hölder budget required here. The two spheres are disjoint and nested, and direct sign inspection gives

$$\{g_\varepsilon \ge 0\} = D \cup \{z : \|z - o\|_2 \ge R + \varepsilon\}. \tag{B.4}$$

Moreover,

$$\|g_\varepsilon - g_0\|_{\infty, \mathcal{X}} \le \kappa(2R\varepsilon + \varepsilon^2) \sup_{z \in \mathcal{X}} |r^2 - \|z - d\|_2^2| = O(\varepsilon). \tag{B.5}$$

We next identify the complete zero set in V_K. For every unit u, choose $v \in \{\pm e_1, \pm e_2\}$ with $u^\top v \ge 1/\sqrt{2}$. Then

$$\|(R + \varepsilon)u - \delta v\|_2^2 \le (R + 1/100)^2 + \delta^2 - \sqrt{2}\,\delta(R + 1/100) < R^2.$$

Thus $V_K = \bigcup_{x \in Q} A_K(x)$ contains the whole circle $S_{R+\varepsilon}(o)$, as well as $S_r(d)$, and hence

$$B_{g_\varepsilon} = S_{R+\varepsilon}(o) \cup S_r(d), \qquad d_H(B_{g_\varepsilon}, B_{g_0}) = \varepsilon. \tag{B.6}$$

At the designated query,

$$C_{g_0, K}(o) = D \cup S_R(o), \qquad C_{g_\varepsilon, K}(o) = D \quad (\varepsilon > 0).$$

Both are nonempty, but

$$d_H(C_{g_\varepsilon, K}(o), C_{g_0, K}(o)) = R + \|d - o\|_2 - r > 0. \tag{B.7}$$

The zero set is a disjoint union of smooth circles with a common positive tube radius and nonzero normal derivative. At $x = o$, both components meet the action. For the four shifted queries, $S_R(o)$ meets the action-circle transversally and $S_r(d)$ remains in its interior. Compactness and finiteness of Q therefore give boundary–action coverage with one finite constant. Radial crossing nevertheless fails at $x = o$. On $S_R(o)$ the favorable normal is $n(c) = (c - o)/R$. If a unit v generated a feasible segment $c + [0, s]v \subset \overline{B}(o, R)$, then

$$R^2 + 2Rs\, n(c)^\top v + s^2 \le R^2,$$

so $n(c)^\top v \le -s/(2R) < 0$. No positive crossing constant is possible. Equations (B.5)–(B.7) prove Proposition 3.

□

C ESTIMATOR CONSTRUCTION AND STATISTICAL UPPER BOUNDS

This appendix gives the complete sample-to-score argument. We first specify a continuous estimator on every sample, then prove uniform concentration, and finally combine that event with the deterministic geometric theorem proved in Appendix B.2. The confidence index is suppressed on $\widehat{g}_n$, as in the main text.

C.1 Estimator and admissible sample sizes

Assumption C.1 (Kernel, bandwidth, and finite-sample thresholds)**.** The kernel $\mathsf{K} : \mathbb{R}^d \to [0, \infty)$ is bounded, Lipschitz, supported on $[-1, 1]^d$, and strictly positive near zero. Its polynomial moment matrices through degree $\lfloor \beta \rfloor$ have a common positive lower bound on the cube truncations induced by $\mathcal{X}$ and $0 < h \le 1$. Fix $c_h > 0$ and $C_0 \ge e$, and set

$$h_{n,\delta} = c_h \left(\frac{\log(C_0 n/\delta)}{n} \right)^{1/(2\beta+d)},$$

$$e_{n,\delta} = C_{\mathrm{LP}} \left[h_{n,\delta}^{\beta} + \sqrt{\frac{\log(C_0 n/\delta)}{n h_{n,\delta}^{d}}} \right] = C_e r_{n,\delta}, \qquad C_e = C_{\mathrm{LP}}(c_h^{\beta} + c_h^{-d/2}).$$

Here C_{LP} depends only on the fixed design, score, dimension, and kernel constants. The threshold $n_{\mathrm{val}}(\delta)$ is the first integer after which

$$h_{n,\delta} \le \min\{1, r_0/8, s_Q/4\}, \qquad n h_{n,\delta}^{d} \ge C_{\mathrm{Gram}} \log(2C_0 n/\delta), \qquad e_{n,\delta} \le \min\{\varepsilon_{\mathrm{int}}, m_Q/2\}.$$

The threshold $n_{\mathrm{loc}}(\delta)$ also requires $C_{\mathrm{int}} e_{n,\delta} \le r_*/2$ and $4C_{\mathrm{int}} e_{n,\delta} < \lambda r_*^2$.

The thresholds collect the conditions under which the statistical and geometric bounds apply. They concern the procedure and sample size, whereas Appendix A defines the population classes.

Set $\ell = \lfloor \beta \rfloor$. For $z \in \mathcal{X}$, compute the minimum-norm local-polynomial least-squares fit

$$\widehat{a}(z) \in \operatorname*{argmin}_{(a_s)_{|s|\le\ell}} \sum_{i=1}^{n} \mathsf{K}\left(\frac{X_i - z}{h} \right) \left[Y_i - \sum_{|s|\le\ell} a_s (X_i - z)^s \right]^2.$$

If a raw Gram matrix is singular at any point in $\mathcal{X}$, use $\widehat{\eta}_n \equiv 1/2$ on the entire cube. Otherwise set $\widehat{\eta}_n(z) = \mathrm{clip}_{[q,1-q]}(\widehat{a}_0(z))$. In both cases, define $\widehat{g}_n = \widehat{\eta}_n - 1/2$ and use the exact value and full optimizer set over $C_{\widehat{g}_n,K}(x)$ whenever that set is nonempty. Empty-set conventions are given in Section 3. The global singular-sample rule, rather than pointwise patching, preserves continuity.

C.2 Uniform conditioning of the local fit

The first step bounds the polynomial Gram matrices on the whole cube, including its faces and corners. We use scaled monomials for concentration and then relate them to the raw least-squares coefficients.

Lemma C.1 (Population polynomial Gram lower bound)**.** *Under Assumptions A.1 and C.1, let* $0 < h \le 1$, $\mathcal{I} = \{s \in \mathbb{N}_0^d : |s| \le \ell\}$, $m = |\mathcal{I}|$, *and* $v(u) = (u^s)_{s\in\mathcal{I}}$. *Define*

$$G(z) = h^{-d} \mathbb{E}\left[\mathsf{K}((X - z)/h) v((X - z)/h) v((X - z)/h)^{\top} \right].$$

There is $\lambda_* = \lambda_*(d, \ell, p_-, \mathsf{K}) > 0$ *such that* $\inf_{z\in\mathcal{X}} \lambda_{\min}(G(z)) \ge \lambda_*$.

Proof. After $u = (x - z)/h$,

$$G(z) = \int_{\mathcal{D}_z} \mathsf{K}(u) v(u) v(u)^{\top} p(z + hu)\, du, \qquad \mathcal{D}_z = \{u : z + hu \in \mathcal{X}\}.$$

Since $\mathsf{K} \ge 0$ and $p \ge p_-$, for every $q \in \mathbb{R}^m$,

$$q^{\top} G(z) q \ge p_- \int_{\mathcal{D}_z} \mathsf{K}(u) (q^{\top} v(u))^2\, du.$$

The uniformly positive moment assumption over every cube truncation $\mathcal{D}_z$ gives a common $\lambda_{\mathsf{K}} > 0$, and therefore $q^{\top} G(z) q \ge p_- \lambda_{\mathsf{K}} \|q\|_2^2$. Take $\lambda_* = p_- \lambda_{\mathsf{K}}$. □

Lemma C.2 (Empirical Gram envelope and modulus)**.** *Under Assumptions A.1 and C.1, let*

$$G_n(z) = \frac{1}{n h^d} \sum_{i=1}^{n} \mathsf{K}((X_i - z)/h) v((X_i - z)/h) v((X_i - z)/h)^{\top}.$$

There are B_0, B_1 depending only on (d, ℓ, K) such that, with $v_i = v((X_i - z)/h)$, each $Y_i(z) = h^{-d}\mathsf{K}((X_i - z)/h)v_i v_i^\top$ is positive semidefinite with $\|Y_i(z)\|_{\mathrm{op}} \le B_0 h^{-d}$ and

$$\max_{r,s} |G_n(z)_{rs} - G_n(z')_{rs}| \le B_1 \|z - z'\|_2 h^{-(d+1)}.$$

The same modulus holds for G.

Proof. On the kernel support every monomial in v is bounded by one. Nonnegativity and boundedness of K give the positive-semidefinite envelope. Extending the compactly supported Lipschitz kernel by zero, the map $u \mapsto \mathsf{K}(u)v(u)v(u)^\top$ is Lipschitz. Replacing z by z' changes u by $(z' - z)/h$, and the normalization contributes h^{-d}, which gives the displayed modulus for the empirical average and expectation. □

Proposition C.3 (All-domain empirical Gram event)**.** *Under Assumptions A.1 and C.1, there is $C_{\mathrm{Gram}} = C_{\mathrm{Gram}}(d, \ell, p_-, \mathsf{K}) \ge 1$ such that, for $0 < h \le 1$ and $\delta \in (0, 1/2)$, if $nh^d \ge C_{\mathrm{Gram}} \log(C_0 n/\delta)$, then*

$$\mathcal{E}_{\mathrm{Gram}} = \left\{ \inf_{z \in \mathcal{X}} \lambda_{\min}(G_n(z)) \ge 3\lambda_*/8 \right\}$$

has probability at least $1 - \delta$. On this event $\sup_z \|G_n(z)^{-1}\|_{\mathrm{op}} \le 4/\lambda_*$.

Proof. The matrix Chernoff lower tail used here is the following: if independent positive-semidefinite Y_i obey $\|Y_i\|_{\mathrm{op}} \le L_h$ and $\lambda_{\min}(\mathbb{E}\sum_i Y_i) \ge n\lambda_*$, then for a fixed $\gamma \in (0, 1)$,

$$\Pr\left\{ \lambda_{\min}\left(n^{-1} \sum_i Y_i \right) < (1 - \gamma)\lambda_* \right\} \le m \exp\left(-c_\gamma \frac{n\lambda_*}{L_h} \right).$$

This is the lower-tail form of the matrix Chernoff inequality; see Tropp (2012). Here $L_h \asymp h^{-d}$, so the exponent is $\asymp nh^d$.

Take a net $\mathcal{N}_h$ of mesh $\eta_h = \lambda_* h^{d+1}/(16 m B_1)$. Its cardinality satisfies $|\mathcal{N}_h| \le (1 + C_d/\eta_h)^d$. At every net point, the previous lemma and matrix Chernoff yield

$$\Pr\{\lambda_{\min}(G_n(z_0)) < \lambda_*/2\} \le m \exp(-cnh^d).$$

The activation inequality makes $\log|\mathcal{N}_h| \le C_d' \log(C_0 n/\delta)$ after enlarging only the fixed C_{Gram}. A union bound therefore gives the lower bound $\lambda_*/2$ at all net points with probability at least $1 - \delta$.

For arbitrary z, choose $z_0 \in \mathcal{N}_h$. The modulus lemma gives $\|G_n(z) - G_n(z_0)\|_{\mathrm{op}} \le \lambda_*/8$. Weyl's inequality yields

$$\lambda_{\min}(G_n(z)) \ge \lambda_{\min}(G_n(z_0)) - \lambda_*/8 \ge 3\lambda_*/8.$$

The displayed eigenvalue comparison follows directly from the Rayleigh quotient: for every unit vector q, $q^\top(G_n(z_0) + E)q \ge \lambda_{\min}(G_n(z_0)) - \|E\|_{\mathrm{op}}$, and taking the infimum over q gives $\lambda_{\min}(G_n(z_0) + E) \ge \lambda_{\min}(G_n(z_0)) - \|E\|_{\mathrm{op}}$. Inversion then gives $\|G_n(z)^{-1}\|_{\mathrm{op}} \le 8/(3\lambda_*) < 4/\lambda_*$. □

Lemma C.4 (Raw and scaled bases)**.** *Under the event in Proposition C.3, let D_h be diagonal with entries $h^{|s|}$. If $R_n(z)$ is the raw empirical Gram, then*

$$R_n(z) = nh^d D_h G_n(z) D_h.$$

The raw and scaled least-squares representations have the same intercept, and $R_n(z)$ is invertible on $\mathcal{E}_{\mathrm{Gram}}$.

Proof. The raw polynomial vector is $((X_i - z)^s)_s = D_h v((X_i - z)/h)$. This gives the congruence. Since $h > 0$, D_h is invertible. If $\vartheta_s = h^{|s|} a_s$, the raw and scaled fitted polynomials coincide, and $\vartheta_0 = a_0$ because $h^0 = 1$. Invertibility makes the minimizer unique, proving exact intercept compatibility. □

C.3 BIAS AND STOCHASTIC FLUCTUATIONS

Conditioning of the Gram matrix reduces score accuracy to a Taylor remainder and a centered empirical numerator. The following concentration bound fixes the normalization used in both estimates.

Lemma C.5 (Scalar Bernstein tail)**.** *Let $Z_1, \dots, Z_n$ be independent, centered random variables with $|Z_i| \le B$ almost surely, where $B > 0$, and $\sum_i \mathbb{E}Z_i^2 \le V$. Then, for every $u > 0$,*

$$\Pr\left\{\left|\sum_{i=1}^{n} Z_i\right| \ge \sqrt{2Vu} + \frac{Bu}{3}\right\} \le 2e^{-u}.$$

This is the scalar Bernstein inequality in the form used below; see (Boucheron et al., 2013).

Proof. The sum and its negative are sub-gamma with variance factor V and scale $B/3$; applying the two one-sided bounds gives the stated inequality (Boucheron et al., 2013). □

Lemma C.6 (Taylor and mass control)**.** *Under Assumptions A.1 and C.1, let $0 < h \le 1$, $\delta \in (0, 1/2)$, and $nh^d \ge C_{\mathrm{Gram}} \log(C_0 n/\delta)$. Put $\eta = 1/2 + g$ and use the Taylor polynomial through degree $\beta - 1$ for integer β and degree $\lfloor\beta\rfloor$ otherwise. There is a coefficient vector $\vartheta(z)$ with intercept $\eta(z)$ such that*

$$p_z(x) = \vartheta(z)^\top v((x - z)/h), \qquad |\eta(x) - p_z(x)| \le C_T L h^\beta$$

whenever the kernel at $(x - z)/h$ is nonzero. With

$$S_n(z) = \frac{1}{nh^d}\sum_i \mathsf{K}((X_i - z)/h), \quad R_n(z) = \frac{1}{nh^d}\sum_i \mathsf{K}_i(z) v_i(z)[\eta(X_i) - p_z(X_i)],$$

there is an event of probability at least $1 - \delta/4$ on which $\sup_z S_n(z) \le C_S$ and $\sup_z \|R_n(z)\|_2 \le C_R h^\beta$.

Proof. The Hölder Taylor remainder is uniform at faces and corners: the segment from z to x remains in the convex cube and $\|x - z\|_2 \le \sqrt{d}\, h$ on the kernel support. For the mass, write $Z_i(z) = \mathsf{K}((X_i - z)/h) - \mathbb{E}\mathsf{K}((X_i - z)/h)$. Then $|Z_i(z)| \le 2\|\mathsf{K}\|_\infty$ and $\mathbb{E}Z_i(z)^2 \le p_+ h^d \int \mathsf{K}^2$. Take a net $\mathcal{N}$ of mesh h^{d+1}/n, so $|\mathcal{N}| \le (1 + C_d n/h^{d+1})^d$, and apply Lemma C.5 with $u = \log(8|\mathcal{N}|/\delta)$. The union failure probability is at most $\delta/4$. Since $h^{-d} \le n$ under the activation, $u \le C \log(C_0 n/\delta)$. The centered mass modulus extends the net bound with error at most C/n, giving

$$\sup_z |S_n(z) - \mathbb{E}S_n(z)| \le C\left\{\sqrt{\frac{\log(C_0 n/\delta)}{nh^d}} + \frac{\log(C_0 n/\delta)}{nh^d}\right\}.$$

Here $1/n \le \sqrt{\log(C_0 n/\delta)/(nh^d)}$. Writing $v = \log(C_0 n/\delta)/(nh^d) \le 1$, we have $v \le \sqrt{v} \le 1$. Thus the linear term is absorbed and the mass is bounded by a fixed constant, since its expectation is at most $p_+ \int \mathsf{K}$. On the kernel support, $\|v\|_2 \le C_v$, so

$$\|R_n(z)\|_2 \le C_v C_T L h^\beta S_n(z) \le C_R h^\beta.$$

□

Lemma C.7 (Uniform stochastic numerator)**.** *Under Assumptions A.1 and C.1, let $0 < h \le 1$, $\delta \in (0, 1/2)$, and $nh^d \ge C_{\mathrm{Gram}} \log(C_0 n/\delta)$. With $\xi_i = Y_i - \eta(X_i)$ and*

$$\Xi_n(z) = \frac{1}{nh^d}\sum_i \mathsf{K}_i(z) v_i(z) \xi_i,$$

there is an event of probability at least $1 - \delta/4$ on which

$$\sup_{z \in \mathcal{X}} \|\Xi_n(z)\|_2 \le C_\Xi \sqrt{\frac{\log(C_0 n/\delta)}{nh^d}}.$$

Proof. For each coordinate and fixed z, the unnormalized summand $W_i(z) = \mathsf{K}_i(z)v_{i,s}(z)\xi_i$ is unconditionally centered because $\mathbb{E}[\xi_i \mid X_i] = 0$. It is bounded by a constant depending only on K and the polynomial degree, and

$$\mathbb{E}W_i(z)^2 \le C\,\mathbb{P}\{X_i \in z + h[-1,1]^d\} \le C'h^d$$

by the density upper bound and $\mathrm{Var}(Y_i \mid X_i) \le 1/4$. Take a net $\mathcal{N}$ with mesh h^{d+1}/n and apply Lemma C.5 at its points and all m coordinates with $u = \log(8m|\mathcal{N}|/\delta)$. The union failure probability is at most $\delta/4$, and $u \le C\log(C_0n/\delta)$ as in Lemma C.6. Division by nh^d bounds the net error by $C(\sqrt{v} + v)$, where $v = \log(C_0n/\delta)/(nh^d) \le 1$. The normalized numerator has modulus $O(h^{-(d+1)})$ because $|\xi_i| \le 1$, so interpolation adds at most $C/n \le C\sqrt{v}$. Using $v \le \sqrt{v}$ and the fixed polynomial dimension gives the displayed uniform bound under the joint sampling law. □

C.4 A CONTINUOUS UNIFORM SCORE BAND

We combine the preceding events before invoking recourse geometry. Clipping cannot enlarge error relative to the population conditional probability.

Proposition C.8 (Whole-cube score band and continuity)**.** *Under Assumptions A.1 and C.1, take $h = h_{n,\delta}$, $\delta \in (0, 1/2)$, and $n \ge n_{\mathrm{val}}(\delta)$. On one event $\mathcal{E}_{\mathrm{score}}$ with probability at least $1 - \delta$, the Gram is invertible at every z, $\widehat{g}_n$ is continuous, and*

$$\|\widehat{g}_n - g\|_{\infty,\mathcal{X}} \le C_{\mathrm{LP}}\left[h^\beta + \sqrt{\frac{\log(C_0n/\delta)}{nh^d}}\right].$$

The clipping projection does not enlarge this error and the singular fallback is inactive.

Proof. Apply Proposition C.3 with failure probability $\delta/2$, and apply Lemmas C.6 and C.7 with failure probability $\delta/4$ each. The activation includes $nh^d \ge C_{\mathrm{Gram}}\log(2C_0n/\delta)$ for the Gram call, while $\log(2C_0n/\delta) \le 2\log(C_0n/\delta)$ lets the score constant absorb the fixed confidence factor. Hence the intersection has probability at least

$$1 - \delta/2 - \delta/4 - \delta/4 = 1 - \delta.$$

In the scaled basis, the normal equations are

$$\widehat{\vartheta}(z) = G_n(z)^{-1}B_n(z), \qquad B_n(z) = \frac{1}{nh^d}\sum_i \mathsf{K}_i(z)v_i(z)Y_i.$$

The Taylor decomposition is

$$B_n(z) = G_n(z)\vartheta(z) + R_n(z) + \Xi_n(z).$$

Hence the inverse bound from Proposition C.3 gives

$$\|\widehat{\vartheta}(z) - \vartheta(z)\|_2 \le \frac{4}{\lambda_*}(C_Rh^\beta + C_\Xi t_{n,\delta}), \qquad t_{n,\delta} = \sqrt{\frac{\log(C_0n/\delta)}{nh^d}}.$$

The intercept of $\vartheta(z)$ is $\eta(z)$, proving the score band after absorbing fixed factors into C_{LP}.

For the continuum step, the zero-extended kernel gives

$$\|G_n(z) - G_n(z')\|_{\mathrm{op}} \le C_Gh^{-(d+1)}\|z - z'\|_2, \qquad \|B_n(z) - B_n(z')\|_2 \le C_Bh^{-(d+1)}\|z - z'\|_2.$$

The inverse identity

$$G_n(z)^{-1} - G_n(z')^{-1} = G_n(z)^{-1}[G_n(z') - G_n(z)]G_n(z')^{-1}$$

and the uniform mass bound imply a coefficient modulus $C_{\mathrm{fit}}h^{-(d+1)}$: indeed $\|B_n(z)\| \le C_vS_n(z) \le C_vC_S$. A net with mesh

$$\eta_n = \frac{h^{d+1}}{8C_{\mathrm{fit}}}\min\left\{1, h^\beta + t_{n,\delta}\right\}$$

has logarithmic size under the activation. The fitted-intercept modulus, the target Hölder modulus, and the stochastic summand modulus each consume a fixed fraction of the net threshold, yielding a uniform bound over the closed cube, including all boundary faces.

The entries of G_n and B_n are continuous and G_n is uniformly nonsingular, so matrix inversion makes $\widehat{a}_0$ continuous. The fallback is inactive. Finally, since $\eta(z) \in [q, 1-q]$, projection onto $[q, 1-q]$ is nonexpansive and fixes $\eta(z)$:

$$|\operatorname{clip}_{[q,1-q]}(\widehat{a}_0(z)) - \eta(z)| \le |\widehat{a}_0(z) - \eta(z)|.$$

If the Taylor remainder and all Bernoulli residuals vanish, then the normal equations give $\widehat{\vartheta}(z) = \vartheta(z)$ exactly. In this case the fitted score equals the population score.

Off the Gram event, the global convention remains continuous and every nonempty feasible section attains its minimum. Measurable maximization (Aliprantis & Border, 2006) gives measurable value and optimizer maps; query-supremum losses are universally measurable and are evaluated under the completed sampling laws. □

C.5 Cost recovery on the uniform score event

The next proposition proves the upper-bound part of Theorem 4. The deterministic transfer in Theorem 2 applies on the uniform score event already established above.

Proposition C.9 (Class-uniform cost bound)**.** *For a fixed admissible Q, $h = h_{n,\delta}$, and $e = e_{n,\delta}$, every $n \ge n_{\mathrm{val}}(\delta)$ satisfies all deterministic activations needed in Proposition C.8. With probability at least $1-\delta$ uniformly over $\Theta_{\mathrm{val}}(\beta; Q)$, the Gram is invertible everywhere, $\widehat{g}_n$ is continuous, $\|\widehat{g}_n - g\|_{\infty,\mathcal{X}} \le e$, $\widehat{C}_n(x) \ne \varnothing$ for all x, and*

$$L_\rho(\widehat{\rho}_n, \rho_{g,K}) \le C_{\mathrm{val}} e.$$

Moreover,

$$e = C_{\mathrm{LP}}(c_h^\beta + c_h^{-d/2}) \left(\frac{\log(C_0 n/\delta)}{n}\right)^{\beta/(2\beta+d)}.$$

Proof. By definition of $n_{\mathrm{val}}(\delta)$, all of

$$h \le \min\{1, r_0/8, s_Q/4\}, \quad nh^d \ge C_{\mathrm{Gram}} \log(2C_0 n/\delta),$$

and

$$e \le \min\{\gamma_0/2, a r_0/8, a r_{\mathrm{cov}}/8, m_Q/2, \kappa_{\mathrm{feas}}/2, a\tau s_{\mathrm{cross}}/2\}$$

hold. Proposition C.8 gives a single event with all score, continuity, and Gram conclusions. On it, apply the deterministic bridge with $\widetilde{g} = \widehat{g}_n$ and $\varepsilon = \|\widehat{g}_n - g\|_\infty \le e$. The plug-in sets are nonempty, so the empty-set convention is inactive, and the cost bound follows with $C_{\mathrm{val}} = C_{\mathrm{int}}$, as fixed in Theorem 4.

Let $L_{n,\delta} = \log(C_0 n/\delta)$. Direct algebra gives

$$h^\beta = c_h^\beta (L_{n,\delta}/n)^{\beta/(2\beta+d)},$$

and

$$\sqrt{\frac{L_{n,\delta}}{nh^d}} = c_h^{-d/2} (L_{n,\delta}/n)^{\beta/(2\beta+d)}.$$

Adding the two terms yields the displayed e identity. If the score residual is zero, the estimator reproduces the exact score, and the feasible sets, values, and fallback status agree identically. □

C.6 Expected risk and uniform consistency

The value lower bounds use expected loss. The high-probability upper bound also yields an expected-loss upper bound by choosing a confidence level that decreases with sample size.

Corollary C.10 (Expected cost risk)**.** *Fix an admissible value class. Run the estimator with $\delta_n = n^{-2}$ for $n \ge 2$, and denote its value estimate by $\widehat{\rho}_n^{\langle 2\rangle}$. For all sufficiently large n,*

$$\sup_{\theta \in \Theta_{\mathrm{val}}(\beta;Q)} \mathbb{E}_\theta L_\rho(\widehat{\rho}_n^{\langle 2\rangle}, \rho_{g,K}) \le C \left(\frac{\log n}{n}\right)^{\beta/(2\beta+d)}.$$

In particular, minimum-cost estimation is uniformly consistent on every fixed subfamily, including Θ_{sep} in Theorem 6.

Proof. With $\delta_n = n^{-2}$, $\log(C_0 n/\delta_n) = \log(C_0 n^3) \asymp \log n$. The chosen bandwidth tends to zero, $nh^d_{n,\delta_n}/\log n$ tends to infinity, and e_{n,δ_n} tends to zero. Hence the conditioning and geometric thresholds hold for all sufficiently large n, with the same class constants. Theorem 4 bounds the loss on its good event by $C_{\mathrm{val}} C_e r_{n,\delta_n}$. On the complement, both the estimate and target lie in $[0, \mathrm{diam}(\mathcal{X})]$, including the empty-set convention. Therefore

$$\mathbb{E}_\theta L_\rho \le C_{\mathrm{val}} C_e r_{n,\delta_n} + \mathrm{diam}(\mathcal{X}) n^{-2}.$$

Taking the supremum proves the claim. Markov's inequality then yields uniform consistency in probability. □

D Sharp Lower Bounds for Recourse Cost

This section proves the lower bounds in Theorem 4. The full-action family is constructed in Section D.1; Section D.2 then treats a box face that genuinely changes the optimum. Both proofs follow the same order: construct a population family, verify common class constants, separate its value functions, and bound the statistical testing difficulty. Query sets, action templates, and class constants are fixed before the sample size or local alternative is chosen. The upper bound and its expected-risk version are established separately in Appendix C.

D.1 Full-action minimax cost lower bound

We construct alternatives separated in recourse cost, verify their membership with common constants, and reduce cost estimation to a finite testing problem. The reserve chart below records the geometry required by that construction; Proposition D.2 proves that the chart is nonempty.

D.1.1 A realizable reserve chart and a common parameter class

Construction D.1 (Lower-bound reserve chart)**.** The regime $\mathfrak{R}_{\mathrm{lb}}$ specifies the baseline and perturbations used for the lower bound. Each reserve below is a positive margin within a defining inequality, chosen to keep all alternatives in $\Theta_{\mathrm{val}}(\beta; Q_{\mathrm{lb}})$ with common constants. There is a compact interior strip $\mathcal{R} = U_0 \times [v_0 - 2s_0, v_0 + 2s_0]$, with $U_Q \Subset U_0$ and associated fixed query strip $Q_{\mathrm{lb}} = U_Q \times \{v_0 - s_0\}$, together with uniform design, K_{full}, and a baseline $g_{\mathrm{base}} \in \Sigma_d(\beta, L/2)$ that is affine, $g_{\mathrm{base}}(u, v) = \kappa_{\mathrm{lb}}(v - v_0)$, on $\mathcal{R}$, with $a < \kappa_{\mathrm{lb}} < b$. It has positive margins for score range, Hölder norm, crossing, off-tube margin, query margin, radial crossing, reach, chart edge, local tube co-Lipschitz, and nonlocal tube separation. More explicitly, the baseline has named slacks $\zeta_{\mathrm{range}}, \zeta_H, \zeta_{\mathrm{cross}}, \zeta_{\mathrm{off}}, \zeta^g_Q, \zeta^\rho_Q, \zeta_\tau, \zeta_{\mathrm{reach}}, \zeta_{\mathrm{chart}}, \sigma_{\mathrm{tube}}, m_{\mathrm{tube}} > 0$ and satisfies the corresponding range, Hölder, crossing, off-tube, query, radial-crossing, and reach inequalities. Write $B_0 = \{g_{\mathrm{base}} = 0\}$, let n_0 be its favorable normal, and set $T_0(b, t) = b + tn_0(b)$. More specifically, the baseline range is contained in $[-1/2 + q + \zeta_{\mathrm{range}}, 1/2 - q - \zeta_{\mathrm{range}}]$, its $\Sigma_d(\beta)$ norm is at most $L - \zeta_H$, and its oriented crossing derivative on the $(r_0 + \zeta_{\mathrm{reach}})$-tube lies in $[a + \zeta_{\mathrm{cross}}, b - \zeta_{\mathrm{cross}}]$. It also satisfies

$$\inf_{\mathrm{dist}(z, B_0) \ge r_0/4} |g_{\mathrm{base}}(z)| \ge \gamma_0 + \zeta_{\mathrm{off}}, \qquad g_{\mathrm{base}}(x) \le -(m_Q + \zeta^g_Q),$$

$$\underline{\rho}_Q + \zeta^\rho_Q \le \rho_{g_{\mathrm{base}}, K_{\mathrm{full}}}(x) \le \overline{\rho}_Q - \zeta^\rho_Q \quad (x \in Q_{\mathrm{lb}}), \qquad \sup_{z \in \mathcal{X}} g_{\mathrm{base}}(z) \ge \kappa_{\mathrm{feas}}.$$

Its zero set and every perturbed zero set remain at least s_A from $\partial\mathcal{X}$, so the favorable normal segment of length s_A is feasible with radial constant one. The baseline normal map is injective on $B_0 \times [-(r_0 + \zeta_{\mathrm{reach}}), r_0 + \zeta_{\mathrm{reach}}]$. A finite C^2 graph atlas $\{\Phi_\nu : \overline{U}_\nu \to B_0\}_{\nu=1}^{N_0}$ contains the protected chart, with every bump support separated from the chart edge and $\partial\mathcal{X}$. In particular, with

$$\Gamma_{\mathrm{reach}} = \frac{1}{4}\left(\frac{1}{r_0} - \frac{1}{r_0 + \zeta_{\mathrm{reach}}}\right),$$

the baseline normal map has the reserve Γ_{reach}, while its local chart and nonlocal tube reserves are respectively

$$\|T_0(\Phi_\nu(s), t) - T_0(\Phi_\nu(s'), t')\|_2 \ge m_{\mathrm{tube}}(\|s - s'\|_2 + |t - t'|), \qquad \|T_0(b, t) - T_0(b', t')\|_2 \ge \sigma_{\mathrm{tube}}.$$

The first inequality is for pairs in one retained atlas block; the second is only for pairs not jointly covered by such a block, with $|t|, |t'| \le r_0$. Fix $0 < h_{\mathrm{lb}} \le \min\{1, \zeta_{\mathrm{chart}}/2\}$, an amplitude $\alpha_{\mathrm{lb}} > 0$, and an active-chart gap $\Delta_{\mathrm{active}} > 0$. For a fixed $\psi \in C_c^\infty((-1,1)^d)$ with $0 \le \psi \le 1$ and $\psi(0) = 1$, define its horizontal support radius by

$$r_\psi := \sup\{\|u_{1:d-1}\|_2 : \psi(u) \neq 0\} < \infty.$$

The bumps $b_{j,h}(z) = \alpha_{\mathrm{lb}} h^\beta \psi((z - z_j)/h)$ are indexed, for each $0 < h \le h_{\mathrm{lb}}$, by a maximal $3h$-separated set $\{u_j\}_{j=1}^{M_h} \subset U_Q$, with $z_j = (u_j, v_0)$. They obey the displayed $C^0, C^1, \Sigma_d(\beta), C^2$ bounds and all reserve inequalities in the certified chart. Explicitly, for every retained (j,h),

$$\|b_{j,h}\|_{\Sigma_d(\beta)} \le C_H \alpha_{\mathrm{lb}}, \qquad \|b_{j,h}\|_\infty \le C_0^g \alpha_{\mathrm{lb}} h^\beta,$$
$$\|\nabla b_{j,h}\|_\infty \le C_1^g \alpha_{\mathrm{lb}} h^{\beta-1}, \qquad \|D^2 b_{j,h}\|_{\infty,\mathrm{op}} \le C_2^g \alpha_{\mathrm{lb}} h^{\beta-2}.$$

Set $g_{j,h} = g_{\mathrm{base}} + b_{j,h}$ and $B_{j,h} = \{g_{j,h} = 0\}$. For the resulting graph $\Gamma_{j,h}$, let $\Phi_{j,h} : B_0 \to B_{j,h}$ equal $\Phi_{j,h}(u, v_0) = (u, v_0 + \Gamma_{j,h}(u))$ on the protected chart and the identity off that chart. The bump vanishes near the chart edge, so the two definitions agree there. Define the pulled-back perturbed normal by

$$n_{j,h}^\sharp(b) := n_{g_{j,h}}(\Phi_{j,h}(b)), \qquad b \in B_0.$$

Writing the deterministic bump constants as $C_H, C_0^g, C_1^g, C_2^g, C_0^\Gamma, C_1^\Gamma, C_1^n, C_2^n, C_{\mathrm{curv}}$, the direct implicit-graph bounds used by the proof are

$$|\Gamma_{j,h}| \le C_0^\Gamma \alpha_{\mathrm{lb}} h^\beta, \qquad \|D\Gamma_{j,h}\| \le C_1^\Gamma \alpha_{\mathrm{lb}} h^{\beta-1},$$
$$\|n_{j,h}^\sharp - n_0\| \le C_1^n \alpha_{\mathrm{lb}} h^{\beta-1}, \qquad \|\nabla n_{j,h}^\sharp - \nabla n_0\| \le C_2^n \alpha_{\mathrm{lb}} h^{\beta-2},$$
$$\text{curvature error} \le C_{\mathrm{curv}} \alpha_{\mathrm{lb}} h^{\beta-2}.$$

The required perturbation bounds are

$$C_H \alpha_{\mathrm{lb}} \le \zeta_H, \qquad C_0^g \alpha_{\mathrm{lb}} h_{\mathrm{lb}}^\beta \le \min\{\zeta_{\mathrm{range}}, \zeta_{\mathrm{off}}, \zeta_Q^g\},$$
$$C_0^\Gamma \alpha_{\mathrm{lb}} h_{\mathrm{lb}}^\beta \le \min\{\zeta_Q^\rho, \zeta_{\mathrm{chart}}/4, r_0/4\},$$
$$C_1^g \alpha_{\mathrm{lb}} h_{\mathrm{lb}}^{\beta-1} \le \zeta_{\mathrm{cross}}, \qquad C_1^n \alpha_{\mathrm{lb}} h_{\mathrm{lb}}^{\beta-1} \le \zeta_\tau,$$
$$C_{\mathrm{curv}} \alpha_{\mathrm{lb}} h_{\mathrm{lb}}^{\beta-2} \le \Gamma_{\mathrm{reach}},$$
$$C_0^\Gamma \alpha_{\mathrm{lb}} h_{\mathrm{lb}}^\beta + r_0 C_1^n \alpha_{\mathrm{lb}} h_{\mathrm{lb}}^{\beta-1} \le \min\{\sigma_{\mathrm{tube}}/4, \zeta_{\mathrm{reach}}/4\}.$$

Crossing and local normal-map regularity require the additional bounds:

$$C_{\mathrm{cross}}^{\mathrm{full}}(\alpha_{\mathrm{lb}}, h_{\mathrm{lb}}) := C_1^g \alpha_{\mathrm{lb}} h_{\mathrm{lb}}^{\beta-1} + L_1\big(C_0^\Gamma \alpha_{\mathrm{lb}} h_{\mathrm{lb}}^\beta + r_0 C_1^n \alpha_{\mathrm{lb}} h_{\mathrm{lb}}^{\beta-1}\big)^{\alpha_1}$$
$$+ L C_1^n \alpha_{\mathrm{lb}} h_{\mathrm{lb}}^{\beta-1} \le \zeta_{\mathrm{cross}},$$
$$(C_1^\Gamma + C_1^n)\alpha_{\mathrm{lb}} h_{\mathrm{lb}}^{\beta-1} \quad + r_0 C_2^n \alpha_{\mathrm{lb}} h_{\mathrm{lb}}^{\beta-2} \le m_{\mathrm{tube}}/2.$$

Finally, the protected active chart has a strict gap $D_{\mathrm{graph}}^{\max} := C_0^\Gamma \alpha_{\mathrm{lb}} h_{\mathrm{lb}}^\beta < \Delta_{\mathrm{active}}$. More precisely, choose a protected region $\mathcal{G}_{\mathrm{prot}} \Subset \mathcal{R}$ and require

$$(U_Q \times \{v_0\}) \cup \bigcup_{0 < h \le h_{\mathrm{lb}}} \bigcup_{j=1}^{M_h} \mathrm{supp}(b_{j,h}) \subset \mathcal{G}_{\mathrm{prot}}.$$

Thus $g_{j,h} = g_{\mathrm{base}}$ outside $\mathcal{G}_{\mathrm{prot}}$. For $x_u = (u, v_0 - s_0) \in Q_{\mathrm{lb}}$ and $b_u = (u, v_0)$, the baseline optimizer and separation conditions are

$$\Pi_{g_{\mathrm{base}}, K_{\mathrm{full}}}(x_u) = \{b_u\}, \qquad \rho_{g_{\mathrm{base}}, K_{\mathrm{full}}}(x_u) = s_0,$$
$$\inf_{z \in D^+_{g_{\mathrm{base}}} \setminus \mathcal{G}_{\mathrm{prot}}} \|z - x_u\|_2 \ge s_0 + \Delta_{\mathrm{active}} \qquad (u \in U_Q).$$

The strict gap is required uniformly over $0 < h \le h_{\mathrm{lb}}$, so the perturbed active chart and all supports remain inside the stated chart room. The reach margin supports the full perturbed r_0-tube for every j and $0 < h \le h_{\mathrm{lb}}$. Injectivity is verified by combining the local m_{tube} lower bound with nonlocal separation σ_{tube}. For K_{full}, $A_K(x) = \mathcal{X}$, so boundary–action coverage holds throughout the family. The family consists of g_{base} and $g_{j,h} = g_{\mathrm{base}} + b_{j,h}$.

Proposition D.2 (Realizability of the lower-bound reserve chart)**.** *The compatible regime $\mathfrak{R}_{\mathrm{lb}}$ in Construction D.1 is nonempty. In particular, there are a smooth baseline, a strict slope interval $0 < a < \kappa_{\mathrm{lb}} < b$, and strictly positive reserves satisfying every inequality in that construction.*

Proof. We give a concrete realization. Let $c_u = (1/2, \dots, 1/2)$, $v_0 = 5/8$, and choose $0 < R, H < 1/8$. Smoothly regularize the maximum of the $2d$ affine face functions

$$v - v_0, \quad v_0 - H - v, \quad u_i - c_{u,i} - R, \quad c_{u,i} - u_i - R \quad (1 \le i \le d-1)$$

with a width $\epsilon > 0$ smaller than one hundredth of the minimum face and cube clearance. The regularized maximum agrees with the ordinary maximum away from an ϵ-neighborhood of ties. Since opposing faces cannot be simultaneously active near the zero level, its gradient is a nonzero convex combination of nonopposing coordinate normals. Hence its zero level $B_0 \Subset \mathrm{int}(\mathcal{X})$ is a compact smooth hypersurface. The top face is strictly dominant on the smaller flat patch

$$\mathcal{P} = \{(u, v_0) : |u_i - c_{u,i}| < R - 10d\epsilon, 1 \le i \le d-1\},$$

so the smoothing is supported away from that patch.

The compact smooth hypersurface has a tubular neighborhood. In tubular coordinates let d_{B_0} be the signed distance, oriented toward the favorable side. Patching it with a smooth cutoff outside half the tube gives a global smooth sign-preserving defining function $\widetilde{d}$ that equals d_{B_0} on a smaller tube. This construction follows from a finite inverse-function cover of the compact normal map. Let S be smooth, odd, sign-preserving, equal to the identity near zero, and bounded by one. Set

$$\kappa_{\mathrm{lb}} = \frac{1}{16}, \qquad a = \frac{1}{32}, \qquad b = \frac{3}{32}, \qquad q = \frac{1}{8}, \qquad g_{\mathrm{base}} = \kappa_{\mathrm{lb}} S(\widetilde{d}).$$

Thus $0 < a < \kappa_{\mathrm{lb}} < b$, the score range is strictly inside $[-1/2 + q, 1/2 - q]$, and on a sufficiently small tubular neighborhood

$$g_{\mathrm{base}}(T_0(b_0, t)) = \kappa_{\mathrm{lb}} t, \qquad \nabla g_{\mathrm{base}}(T_0(b_0, t))^\top n_0(b_0) = \kappa_{\mathrm{lb}}. \tag{D.1}$$

Because g_{base} is smooth on the compact cube, its Hölder norm is finite. Choose L larger than twice that norm, leaving a positive Hölder reserve.

Choose $U_Q \Subset U_0 \Subset \mathcal{P}$ and $0 < s_0 < \min\{R, H\}/64$ so that $\mathcal{R} = U_0 \times [v_0 - 2s_0, v_0 + 2s_0]$ lies in the flat patch. Take $Q = Q_{\mathrm{lb}} := U_Q \times \{v_0 - s_0\}$, $K = K_{\mathrm{full}}$, and uniform design. In the flat patch,

$$\begin{aligned} g_{\mathrm{base}}(u, v) &= \kappa_{\mathrm{lb}}(v - v_0), \qquad g_{\mathrm{base}}(x_u) = -\kappa_{\mathrm{lb}} s_0, \\ \Pi_{g_{\mathrm{base}}, K}(x_u) &= \{(u, v_0)\}, \qquad \rho_{g_{\mathrm{base}}, K}(x_u) = s_0. \end{aligned} \tag{D.2}$$

The choices can be made sufficiently small that the unique nearest-point property holds uniformly on the compact set U_Q.

All reserves now follow from strict inequalities and compactness. Let $d_X = \mathrm{dist}(B_0, \partial\mathcal{X}) > 0$ and choose $r_0 = \zeta_{\mathrm{reach}} = \min\{r_B, d_X\}/16$, where r_B is a tubular radius. Equation (D.1) gives a crossing reserve, for example $\zeta_{\mathrm{cross}} = \kappa_{\mathrm{lb}}/4$. Since the zero set is compact and there are no other zeros, the compact set $\{z : \mathrm{dist}(z, B_0) \ge r_0/4\}$ has a positive minimum absolute score; taking one third of this minimum gives positive γ_0 and ζ_{off}. Set

$$m_Q = \zeta_Q^g = \kappa_{\mathrm{lb}} s_0/3, \quad \underline{\rho}_Q = s_0/2, \quad \overline{\rho}_Q = 3s_0/2, \quad \zeta_Q^\rho = s_0/4.$$

Equation (D.2) supplies the strict query and value reserves. Since B_0 is interior, choose s_A below its boundary clearance, $\tau = 1/2$, and $\zeta_\tau = 1/4$; the favorable unit normal gives the radial segment with inner product one. The favorable side is nonempty, so $\kappa_{\mathrm{feas}} = \frac{1}{2}\max_{\mathcal{X}} g_{\mathrm{base}} > 0$.

Choose a protected region $\mathcal{G}_{\mathrm{prot}} \Subset \mathcal{R}$ containing $U_Q \times \{v_0\}$. Its positive chart and cube clearance gives $\zeta_{\mathrm{chart}} > 0$. A finite C^2 atlas and the positive minimum singular value of the normal coordinate maps on sufficiently small blocks give $m_{\mathrm{tube}} > 0$. Take a finite open subcover of the retained chart interiors and let $\ell_{\mathrm{Leb}} > 0$ be a Lebesgue number. Base points at distance less than ℓ_{Leb} share a retained block; all remaining parameter pairs lie in the closed compact set

$$\{(b, t, b', t') : b, b' \in B_0,\ |t|, |t'| \le r_0,\ \|b - b'\|_2 \ge \ell_{\mathrm{Leb}}\}.$$

This set is disjoint from the parameter diagonal, so injectivity of the normal map gives a positive image minimum σ_{tube}. Finally, the uniform strict nearest-point property in (D.2) and compactness give

$$d_{\mathrm{ext}} := \min_{u \in U_Q} \operatorname{dist}(x_u, D^+_{g_{\mathrm{base}}} \setminus \mathcal{G}_{\mathrm{prot}}) > s_0.$$

Set $\Delta_{\mathrm{active}} = (d_{\mathrm{ext}} - s_0)/2 > 0$. Full actions make recourse and boundary–action coverage immediate. This derives every named positive reserve before a bump is introduced. □

Proposition D.3 (Positive lower-family amplitude and scale)**.** *Under Proposition D.2, the bump profile $\psi \in C_c^\infty((-1,1)^d)$ admits positive $(\alpha_{\mathrm{lb}}, h_{\mathrm{lb}})$ satisfying every budget in Construction D.1, uniformly for $0 < h \le h_{\mathrm{lb}}$.*

Proof. The scaled profile norms C_0^g, C_1^g, C_2^g, C_H are finite. The implicit denominator is bounded below by $\kappa_{\mathrm{lb}}/2$ once the amplitude is small, so direct implicit differentiation fixes finite constants $C_0^\Gamma, C_1^\Gamma, C_1^n, C_2^n, C_{\mathrm{curv}}$ depending only on the profile, dimension, slope, and the realized baseline atlas. Put

$$\Gamma_{\mathrm{reach}} = \frac{1}{4}\left(\frac{1}{r_0} - \frac{1}{r_0 + \zeta_{\mathrm{reach}}}\right) > 0$$

and first choose

$$0 < \alpha_{\mathrm{lb}} < \min\left\{1, \frac{\zeta_H}{2C_H}, \frac{\Gamma_{\mathrm{reach}}}{2C_{\mathrm{curv}}}, \frac{m_{\mathrm{tube}}}{8r_0 C_2^n}\right\}. \tag{D.3}$$

This first choice controls the curvature and the $r_0 C_2^n \alpha_{\mathrm{lb}} h^{\beta-2}$ terms even at $\beta = 2$.

Now choose h_{lb} smaller than $\min\{1, \zeta_{\mathrm{chart}}/2, s_0/4\}$ so that

$$C_0^g \alpha_{\mathrm{lb}} h_{\mathrm{lb}}^\beta < \min\{\zeta_{\mathrm{range}}, \zeta_{\mathrm{off}}, \zeta_Q^g\}, \quad C_0^\Gamma \alpha_{\mathrm{lb}} h_{\mathrm{lb}}^\beta < \min\{\zeta_Q^\rho, \zeta_{\mathrm{chart}}/4, r_0/4, \Delta_{\mathrm{active}}/2\},$$

$$C_1^g \alpha_{\mathrm{lb}} h_{\mathrm{lb}}^{\beta-1} < \min\{\kappa_{\mathrm{lb}}/2, \zeta_{\mathrm{cross}}/3\}, \quad C_1^n \alpha_{\mathrm{lb}} h_{\mathrm{lb}}^{\beta-1} < \zeta_\tau. \tag{D.4}$$

Writing

$$E(h) = C_0^\Gamma \alpha_{\mathrm{lb}} h^\beta + r_0 C_1^n \alpha_{\mathrm{lb}} h^{\beta-1},$$

choose the same h_{lb} so that

$$E(h_{\mathrm{lb}}) < \min\left\{\frac{\sigma_{\mathrm{tube}}}{4}, \frac{\zeta_{\mathrm{reach}}}{4}, \left(\frac{\zeta_{\mathrm{cross}}}{3L_1}\right)^{1/\alpha_1}\right\}, \quad L C_1^n \alpha_{\mathrm{lb}} h_{\mathrm{lb}}^{\beta-1} < \frac{\zeta_{\mathrm{cross}}}{3},$$

$$(C_1^\Gamma + C_1^n)\alpha_{\mathrm{lb}} h_{\mathrm{lb}}^{\beta-1} + r_0 C_2^n \alpha_{\mathrm{lb}} h_{\mathrm{lb}}^{\beta-2} < \frac{m_{\mathrm{tube}}}{2}. \tag{D.5}$$

When $L_1 = 0$, the term involving L_1 in this display is omitted. All left sides vanish with $h_{\mathrm{lb}} \downarrow 0$ except the explicitly amplitude-controlled $\beta = 2$ curvature term, which was handled in (D.3). The three terms in the full crossing budget are each at most one third of ζ_{cross}, the local tube defect is below $m_{\mathrm{tube}}/2$, and the active-chart gap is strict. Monotonicity gives the same inequalities for every $0 < h \le h_{\mathrm{lb}}$, verifying all required perturbation bounds. □

Proposition D.4 (Common-constant containment of the bump family)**.** *Under the nonempty regime constructed in Proposition D.2 and the positive choice in Proposition D.3, for every stipulated center u_j and every $0 < h \le h_{\mathrm{lb}}$, the zero set of $g_{j,h} = g_{\mathrm{base}} + b_{j,h}$ is a compact embedded hypersurface. On the protected graph chart it is*

$$B_{j,h} \cap \mathcal{G}_{\mathrm{prot}} = \{(u, v_0 + \Gamma_{j,h}(u))\},$$

with

$$|\Gamma_{j,h}| \le C_0^\Gamma \alpha_{\mathrm{lb}} h^\beta, \quad \|D\Gamma_{j,h}\| \le C_1^\Gamma \alpha_{\mathrm{lb}} h^{\beta-1},$$

$$\|n_{j,h}^\sharp - n_0\| \le C_1^n \alpha_{\mathrm{lb}} h^{\beta-1}, \quad \|\nabla n_{j,h}^\sharp - \nabla n_0\| \le C_2^n \alpha_{\mathrm{lb}} h^{\beta-2},$$

and curvature error at most $C_{\mathrm{curv}} \alpha_{\mathrm{lb}} h^{\beta-2}$. The perturbed score satisfies all value-class conditions with the same constants, including injectivity of the full normal map on $[-r_0, r_0]$. Hence $\Theta_{\mathrm{bump}}(\beta; Q_{\mathrm{lb}}) \subset \Theta_{\mathrm{val}}(\beta; Q_{\mathrm{lb}})$.

Proof. On the affine strip, the zero equation is

$$F_{j,h}(u,v) = \kappa_{\mathrm{lb}}(v - v_0) + \alpha_{\mathrm{lb}} h^{\beta} \psi\left(\frac{u-u_j}{h}, \frac{v-v_0}{h}\right) = 0.$$

The derivative budget and the baseline reserve give

$$\partial_v F_{j,h} \geq \kappa_{\mathrm{lb}} - C_1^g \alpha_{\mathrm{lb}} h^{\beta-1} \geq a > 0.$$

The bump is supported strictly inside the chart and the C^0 reserve leaves the baseline signs on the vertical sides unchanged, so the equation has one root for every chart coordinate. We use the following current-notation form of the implicit function theorem: if a C^1 function $F(u,v)$ has $\partial_v F \geq a > 0$ on a rectangular chart and has opposite weak signs on the two vertical faces, then each intermediate u has a unique root $v = \Gamma(u)$, the root map is C^1, and $D\Gamma = -D_u F / \partial_v F$; see (Krantz & Parks, 2013). Here the displayed derivative lower bound and the unchanged face signs discharge these hypotheses. The implicit function theorem therefore yields a unique C^1 graph. Its first derivative is

$$D\Gamma_{j,h} = -\frac{D_u b_{j,h}(u, v_0 + \Gamma_{j,h}(u))}{\kappa_{\mathrm{lb}} + \partial_v b_{j,h}(u, v_0 + \Gamma_{j,h}(u))},$$

which gives the C_1^{Γ} bound. The scalar equation and the C^0 bump bound give the graph displacement. With $F = F_{j,h}$ and all derivatives evaluated on its zero graph,

$$D^2\Gamma = -\frac{F_{uu} + F_{uv}(D\Gamma)^{\top} + D\Gamma F_{vu} + F_{vv} D\Gamma (D\Gamma)^{\top}}{F_v}.$$

Thus $\|D^2\Gamma\| \leq (2/\kappa_{\mathrm{lb}}) C_2^g \alpha_{\mathrm{lb}} h^{\beta-2} (1 + \|D\Gamma\|)^2$. The graph normal is $(-D\Gamma, 1)/\sqrt{1 + \|D\Gamma\|^2}$; differentiating it gives the normal-derivative and curvature bounds. All constants include the fixed atlas coordinate factors. Compact support away from the edge makes the graph equal to B_0 near the patch boundary, so the local graph glues to the baseline without an edge. The same gluing makes the map $\Phi_{j,h}$ declared in Construction D.1 a C^1 bijection from B_0 onto $B_{j,h}$: on the protected chart its inverse is the vertical projection onto B_0, and off the chart both maps are the identity. Compactness of the domain makes the resulting zero set compact.

The score range follows from $\|b_{j,h}\|_{\infty} \leq C_0^g \alpha_{\mathrm{lb}} h^{\beta} \leq \zeta_{\mathrm{range}}$, and the Hölder norm follows from

$$\|g_{j,h}\|_{\Sigma_d(\beta)} \leq (L - \zeta_H) + C_H \alpha_{\mathrm{lb}} \leq L.$$

The complete crossing-transfer budget charges the change in the oriented normal derivative by at most ζ_{cross}, so the baseline reserve $[a + \zeta_{\mathrm{cross}}, b - \zeta_{\mathrm{cross}}]$ transfers to $[a,b]$. More explicitly, with $\alpha_1 = \min\{1, \beta - 1\}$, the charged defect is

$$\begin{aligned} C_{\mathrm{cross}}^{\mathrm{full}}(\alpha_{\mathrm{lb}}, h_{\mathrm{lb}}) = {} & C_1^g \alpha_{\mathrm{lb}} h_{\mathrm{lb}}^{\beta-1} \\ & + L_1 \left(C_0^{\Gamma} \alpha_{\mathrm{lb}} h_{\mathrm{lb}}^{\beta} + r_0 C_1^n \alpha_{\mathrm{lb}} h_{\mathrm{lb}}^{\beta-1} \right)^{\alpha_1} \\ & + L C_1^n \alpha_{\mathrm{lb}} h_{\mathrm{lb}}^{\beta-1} \leq \zeta_{\mathrm{cross}}. \end{aligned}$$

The graph/normal perturbation also obeys

$$C_0^{\Gamma} \alpha_{\mathrm{lb}} h_{\mathrm{lb}}^{\beta} + r_0 C_1^n \alpha_{\mathrm{lb}} h_{\mathrm{lb}}^{\beta-1} \leq \min\{\sigma_{\mathrm{tube}}/4, \zeta_{\mathrm{reach}}/4\},$$

and the local normal-map defect is charged by

$$(C_1^{\Gamma} + C_1^n) \alpha_{\mathrm{lb}} h_{\mathrm{lb}}^{\beta-1} + r_0 C_2^n \alpha_{\mathrm{lb}} h_{\mathrm{lb}}^{\beta-2} \leq m_{\mathrm{tube}}/2.$$

Thus the baseline reserve transfers to $[a,b]$. The query score and query values transfer by the C^0 and graph reserves. The feasible-depth certificate is preserved because $b_{j,h} \geq 0$:

$$\sup_z g_{j,h}(z) \geq \sup_z g_{\mathrm{base}}(z) \geq \kappa_{\mathrm{feas}}.$$

Outside the bump support the score is the baseline score. On the support, the perturbation is bounded by the displayed C^0 bump budget. If a point is at least $r_0/2$ from $B_{j,h}$, its distance from B_0 is at least $r_0/4$; the baseline off-tube reserve and the C^0 bump budget therefore leave margin at least γ_0. For this lower family $K = K_{\mathrm{full}} = [-1,1]^d$ and $x \in Q \subset (0,1)^d$, so $A_K(x) = \mathcal{X}$. Consequently

boundary–action coverage holds trivially with $c = b$ and $\kappa_{\mathrm{cov}} = 1$; this verifies the coverage condition before invoking the statistical bounds. The construction places the complete zero set at least s_A from $\partial\mathcal{X}$. Hence $v_c = n_{g_{j,h}}(c)$ gives a feasible segment of length s_A with radial constant one. Coverage and radial crossing are uniform over the lower family.

It remains to check full tube injectivity. Parameterize the perturbed boundary through the bijection $\Phi_{j,h} : B_0 \to B_{j,h}$ and let

$$T_0(b,t) = b + tn_0(b), \qquad T_{j,h}(b,t) = \Phi_{j,h}(b) + tn^{\sharp}_{j,h}(b).$$

On one local atlas block, the baseline co-Lipschitz inequality gives

$$\|T_0(b,t) - T_0(b',t')\|_2 \geq m_{\mathrm{tube}}(\|s - s'\|_2 + |t - t'|).$$

The graph and normal derivative budgets imply

$$\|[T_{j,h}(b,t) - T_0(b,t)] - [T_{j,h}(b',t') - T_0(b',t')]\|_2 \leq D_{\mathrm{loc}}(\|s - s'\|_2 + |t - t'|),$$

where

$$D_{\mathrm{loc}} = (C_1^{\Gamma} + C_1^n)\alpha_{\mathrm{lb}}h^{\beta-1} + r_0 C_2^n \alpha_{\mathrm{lb}} h^{\beta-2} \leq m_{\mathrm{tube}}/2.$$

Hence the perturbed map is locally bounded below by $(m_{\mathrm{tube}}/2)(\|s - s'\|_2 + |t - t'|)$.

For pairs not jointly covered by one local block, the baseline reserve is $\|T_0(b,t) - T_0(b',t')\|_2 \geq \sigma_{\mathrm{tube}}$. Pointwise,

$$\|T_{j,h}(b,t) - T_0(b,t)\|_2 \leq C_0^{\Gamma}\alpha_{\mathrm{lb}}h^{\beta} + r_0 C_1^n \alpha_{\mathrm{lb}} h^{\beta-1} \leq \sigma_{\mathrm{tube}}/4,$$

so the perturbed distance is at least $\sigma_{\mathrm{tube}}/2$. The local and nonlocal cases exhaust all pairs. Since $\Phi_{j,h}$ is a bijection, this is exactly injectivity of the normal map on $B_{j,h} \times [-r_0, r_0]$. The available reach reserve is

$$\Gamma_{\mathrm{reach}} = \frac{1}{4}\left(\frac{1}{r_0} - \frac{1}{r_0 + \zeta_{\mathrm{reach}}}\right) > 0, \qquad C_{\mathrm{curv}}\alpha_{\mathrm{lb}}h_{\mathrm{lb}}^{\beta-2} \leq \Gamma_{\mathrm{reach}}.$$

This reserve keeps the graph in the extended baseline tube, while the preceding local and nonlocal bounds establish injectivity. At $\beta = 2$, the terms with $h^{\beta-2}$ are fixed amplitude terms and are controlled by the amplitude bounds; for $\beta > 2$, their maximum over the finite range is at h_{lb}. □

D.1.2 Value separation across the fixed query set

The preceding propositions establish that all alternatives remain in the same value class. We next locate their active optima and separate the resulting value functions at a finite collection of witness queries.

Proposition D.5 (Protected active root)**.** *Under the regime of Construction D.1 and Proposition D.4, there are deterministic $h_{\mathrm{act}} \leq h_{\mathrm{lb}}$ and $c_\psi > 0$ such that, for $0 < h \leq h_{\mathrm{act}}$ and $x_u = (u, v_0 - s_0)$ with $u \in U_Q$, the favorable side in the protected chart is the epigraph*

$$D^+_{g_{j,h}} \cap \mathcal{G}_{\mathrm{prot}} = \{(u,v) : v \geq v_0 + \Gamma_{j,h}(u)\},$$

where, with $D_h = C_0^{\Gamma}\alpha_{\mathrm{lb}}h^{\beta}$,

$$-D_h \leq \Gamma_{j,h} \leq 0, \qquad -\Gamma_{j,h}(u_j) \geq c_\psi \alpha_{\mathrm{lb}} h^{\beta}.$$

All minimizers at x_u lie in the protected chart, and all minimizers at x_{u_j} lie in the changed support. In particular,

$$\rho_{g_{j,h},K_{\mathrm{full}}}(x_{u_j}) \leq s_0 - c_\psi \alpha_{\mathrm{lb}} h^{\beta}.$$

Proof. The graph equation is

$$\kappa_{\mathrm{lb}}\Gamma_{j,h}(u) + \alpha_{\mathrm{lb}}h^{\beta}\psi\left(\frac{u - u_j}{h}, \frac{\Gamma_{j,h}(u)}{h}\right) = 0.$$

Nonnegativity of ψ gives $\Gamma_{j,h} \leq 0$, while the previous proposition gives $\Gamma_{j,h} \geq -D_h$. Since $\partial_v F_{j,h} > 0$, $F_{j,h} \geq 0$ precisely on the favorable epigraph.

Continuity of ψ at $(0,0)$ gives $\eta_\psi > 0$ and $\underline{\psi} > 0$ with $\psi(0,t) \geq \underline{\psi}$ for $|t| \leq \eta_\psi$. Choose h_{act} so that

$$C_0^\Gamma \alpha_{\mathrm{lb}} h_{\mathrm{act}}^{\beta-1} \leq \eta_\psi, \qquad C_0^\Gamma \alpha_{\mathrm{lb}} h_{\mathrm{act}}^{\beta} < \min\{s_0, \Delta_{\mathrm{active}}\}.$$

At $u = u_j$, the root lies in the interval where the lower bound on ψ applies, so

$$-\Gamma_{j,h}(u_j) = \frac{\alpha_{\mathrm{lb}} h^\beta}{\kappa_{\mathrm{lb}}} \psi\left(0, \frac{\Gamma_{j,h}(u_j)}{h}\right) \geq c_\psi \alpha_{\mathrm{lb}} h^\beta, \qquad c_\psi = \underline{\psi}/\kappa_{\mathrm{lb}}.$$

Outside the protected chart the bump vanishes, so the favorable set equals the baseline favorable set. The exterior reserve gives distance at least $s_0 + \Delta_{\mathrm{active}}$ from x_u, while the unchanged point $b_u = (u, v_0)$ is favorable and has distance s_0. Therefore every minimizer lies in the protected chart. At x_{u_j}, the center root $z_{j,h} = (u_j, v_0 + \Gamma_{j,h}(u_j))$ is favorable and feasible, with

$$\|z_{j,h} - x_{u_j}\|_2 = s_0 + \Gamma_{j,h}(u_j) \leq s_0 - c_\psi \alpha_{\mathrm{lb}} h^\beta.$$

Any favorable point outside the positive changed support has the baseline score and hence distance at least s_0 by the baseline unique-nearest identity. Thus all own-bump minimizers lie in the changed support. The center root supplies the feasible comparison point used in this bound. □

Proposition D.6 (Retained packing and exact value separation)**.** *Under the regime of Construction D.1 and Propositions D.4–D.5, put $m = d - 1$. There are $A > 0$, $c_A > 0$, and $h_* > 0$ such that for every $0 < h \leq h_*$ a retained set J_h of bump centers has*

$$N_h := |J_h| \geq c_A h^{-m}, \qquad N_h \geq 2,$$

and pairwise center distance greater than Ah. For the baseline and alternatives, write $\rho_0 = \rho_{g_{\mathrm{base}}, K_{\mathrm{full}}}$ and $\rho_j = \rho_{g_{j,h}, K_{\mathrm{full}}}$. Then, for $j \neq k$,

$$\rho_k(x_{u_j}) = s_0, \quad \rho_j(x_{u_j}) \leq s_0 - \Delta_h, \quad \rho_0(x_{u_j}) = s_0, \qquad \Delta_h = c_\psi \alpha_{\mathrm{lb}} h^\beta,$$

and

$$L_\rho(\rho_0, \rho_j) \geq \Delta_h, \qquad L_\rho(\rho_j, \rho_k) \geq \Delta_h.$$

More precisely, for every distinct $r, r' \in \{0\} \cup J_h$ there is a matched witness index $\ell(r, r') \in J_h$ such that

$$|\rho_r(x_{u_{\ell(r,r')}}) - \rho_{r'}(x_{u_{\ell(r,r')}})| \geq \Delta_h.$$

The finite witness metric

$$d_h(f, g) = \max_{\ell \in J_h} |f(x_{u_\ell}) - g(x_{u_\ell})|$$

satisfies $d_h(f, g) \leq L_\rho(f, g)$.

Proof. Choose $u_\circ$ and r_Q with $B_m(u_\circ, r_Q) \subset U_Q$. Maximality of a $3h$-separated center set gives

$$M_h \geq (r_Q/(3h))^m$$

by comparing the volume of this ball with the covering by radius-$3h$ balls. Choose

$$A > \max\left\{3, r_\psi + \sqrt{2 s_0 C_0^\Gamma \alpha_{\mathrm{lb}} h_{\mathrm{act}}^{\beta-2}}\right\},$$

put $B_{A,m} = \lceil (1 + 2A/3)^m \rceil$, and greedily retain one center while deleting all centers within Ah. The radius-$3h/2$ balls of one deletion round are disjoint and fit inside a radius-$(A + 3/2)h$ ball, so each round deletes at most $B_{A,m}$ centers. Hence

$$N_h \geq M_h / B_{A,m} \geq c_A h^{-m}, \qquad c_A = (r_Q/3)^m / B_{A,m}.$$

Take h_* small enough that $c_A h_*^{-m} \geq 2$ and $h_* \leq h_{\mathrm{act}}, h_{\mathrm{lb}}$.

Let $\ell_A = A - r_\psi > 0$ and

$$\kappa_A = \ell_A^2 - 2 s_0 C_0^\Gamma \alpha_{\mathrm{lb}} h_{\mathrm{act}}^{\beta-2} > 0.$$

For $h \leq h_*$,

$$\ell_A^2 h^2 - 2 s_0 D_h \geq \kappa_A h^2 > 0, \quad D_h = C_0^\Gamma \alpha_{\mathrm{lb}} h^\beta.$$

The sign is fixed before squaring, and the reserve also gives $\ell_A h > D_h$.

Consider a favorable point z in the changed support of a wrong bump k. Then $\|z_{1:m} - u_k\|_2 \le r_\psi h$ and the epigraph gives $z_d \ge v_0 - D_h$. Since the retained centers are Ah-separated,

$$\|z_{1:m} - u_j\|_2 \ge \ell_A h, \qquad z_d - (v_0 - s_0) \ge s_0 - D_h,$$

and consequently

$$\|z - x_{u_j}\|_2^2 - s_0^2 \ge \ell_A^2 h^2 + (s_0 - D_h)^2 - s_0^2 \ge \kappa_A h^2 + D_h^2 > 0.$$

Outside the protected chart the score equals the baseline and the exterior active margin gives distance greater than s_0. Inside it but outside the wrong support, the baseline unique-nearest identity gives distance at least s_0. Since b_{u_j} remains favorable and has distance s_0, the wrong bump has exact value $\rho_k(x_{u_j}) = s_0$. The active-root proposition gives the own-bump inequality, while the baseline identity gives $\rho_0(x_{u_j}) = s_0$. Thus, for the pair $(0, j)$ and for every pair (j, k) with $j \ne k$, the matched query x_{u_j} has value gap at least Δ_h. This proves the displayed witness statement and hence the two global L_ρ separations. The finite metric inequality is immediate from the fact that the witness set is a subset of Q. □

D.1.3 Testing the separated value functions

The finite witness metric permits a measurable decoder from any value estimator. The following argument bounds its testing error and hence the expected value loss.

Proposition D.7 (Fano lower bound in exact value loss)**.** *Under Assumption A.1, Construction D.1, and Propositions D.4–D.6, let $\chi = 2\beta + d$, $m = d - 1$, and choose $\alpha_F \in (0, 1/8)$ and $c_* > 0$ satisfying*

$$C_{\mathrm{KL}} \alpha_{\mathrm{lb}}^2 c_*^\chi \le \alpha_F \frac{m}{2\chi}.$$

There is n_{mm} such that for every $n \ge n_{\mathrm{mm}}$, with

$$h_n = c_* \left(\frac{\log n}{n}\right)^{1/\chi},$$

$$\inf_{\widetilde{\rho}_n} \sup_{\theta \in \Theta_{\mathrm{bump}}(\beta; Q_{\mathrm{lb}})} \mathbb{E}_\theta L_{\rho, Q_{\mathrm{lb}}}(\widetilde{\rho}_n, \rho_{g, K_{\mathrm{full}}}) \ge c_{\mathrm{val}} \left(\frac{\log n}{n}\right)^{\beta/(2\beta+d)}.$$

Proof. For the baseline and one-bump alternative, let

$$\theta_0 = (K_{\mathrm{full}}, \mu_{\mathrm{unif}}, g_{\mathrm{base}}), \qquad \theta_j = (K_{\mathrm{full}}, \mu_{\mathrm{unif}}, g_{j,h}).$$

The score range places all Bernoulli parameters in $[q, 1-q]$. For $p_0, p_1 \in [q, 1-q]$,

$$\mathrm{kl}(\mathrm{Bern}(p_1), \mathrm{Bern}(p_0)) \le \frac{(p_1 - p_0)^2}{p_0(1-p_0)} \le \frac{(p_1 - p_0)^2}{q(1-q)}.$$

The common uniform design and the bump formula give

$$\int b_{j,h}(z)^2 \, d\mu_{\mathrm{unif}}(z) \le 2^d \|\psi\|_\infty^2 \alpha_{\mathrm{lb}}^2 h^{2\beta+d}.$$

The conditional KL formula and product chain rule therefore give

$$P_{j,h}^n \ll P_0^n, \qquad D_{\mathrm{KL}}(P_{j,h}^n, P_0^n) \le C_{\mathrm{KL}} n \alpha_{\mathrm{lb}}^2 h^{2\beta+d},$$

where one may take $C_{\mathrm{KL}} = 2^d \|\psi\|_\infty^2 / [q(1-q)]$.

Let $I_h = \{0\} \cup J_h$, where $N_h = |J_h|$ is the number of non-reference bump alternatives, let $\Delta_h = c_\psi \alpha_{\mathrm{lb}} h^\beta$, and put $s_h = \Delta_h/2$. In a fixed order, define the measurable finite decoder

$$\widehat{J}(\widetilde{\rho}) = \min \arg\min_{r \in I_h} d_h(\widetilde{\rho}, \rho_r).$$

Proposition D.6 gives $d_h(\rho_r, \rho_{r'}) \ge \Delta_h = 2s_h$ for distinct r, r'. This separation holds at the finite witness queries used by the decoder. If $L_\rho(\widetilde{\rho}, \rho_r) < s_h$, the triangle inequality in the finite witness metric makes r the unique finite-metric minimizer, so the decoder is correct. Hence

$$\{\widehat{J} \ne r\} \subseteq \{L_\rho(\widetilde{\rho}, \rho_r) \ge s_h\}.$$

For $h = h_n$, Proposition D.6 gives $N_h \geq c_A h^{-m}$. Since

$$\log N_{h_n} \geq \frac{m}{\chi}(\log n - \log\log n) + O(1),$$

there is a deterministic n_{mm} after which $h_n \leq h_*$, $N_{h_n} \geq 2$, and

$$\log N_{h_n} \geq \frac{m}{2\chi}\log n.$$

At the same scale,

$$D_{\mathrm{KL}}(P^n_{j,h_n}, P^n_0) \leq C_{\mathrm{KL}}\alpha_{\mathrm{lb}}^2 c_*^\chi \log n \leq \alpha_F \frac{m}{2\chi}\log n \leq \alpha_F \log N_{h_n}.$$

Thus the average-KL hypothesis is valid for every $n \geq n_{\mathrm{mm}}$.

We use the many-hypothesis form of Theorem 2.5, equation (2.46), of Tsybakov (2009): if $M \geq 2$, $P_j \ll P_0$, $M^{-1}\sum_j D_{\mathrm{KL}}(P_j, P_0) \leq \alpha \log M$, $0 < \alpha < 1/8$, and the metric separation is at least $2s$, then

$$\inf_{\widehat{J}} \sup_{0\leq r\leq M} P_r\{\widehat{J} \neq r\} \geq \frac{\sqrt{M}}{1+\sqrt{M}}\left(1 - 2\alpha - \sqrt{\frac{2\alpha}{\log M}}\right).$$

With $M = N_{h_n}$ and $\alpha = \alpha_F$, this is bounded below by a fixed $c_F > 0$ because $N_{h_n} \geq 2$. The decoder reduction yields

$$\inf_{\widetilde{\rho}_n} \sup_{r\in I_{h_n}} P^n_r\{L_{\rho,Q_{\mathrm{lb}}}(\widetilde{\rho}_n, \rho_r) \geq s_{h_n}\} \geq c_F.$$

Using $\mathbb{E}Z \geq sP(Z \geq s)$ for $Z \geq 0$ and the inclusion of the finite candidate family in the bump class gives

$$\inf_{\widetilde{\rho}_n} \sup_{\theta\in\Theta_{\mathrm{bump}}(\beta;Q_{\mathrm{lb}})} \mathbb{E}_\theta L_{\rho,Q_{\mathrm{lb}}}(\widetilde{\rho}_n, \rho_{g,K_{\mathrm{full}}}) \geq c_F s_{h_n} = \frac{c_F c_\psi \alpha_{\mathrm{lb}} c_*^\beta}{2}\left(\frac{\log n}{n}\right)^{\beta/(2\beta+d)}.$$

Set $c_{\mathrm{val}} = c_F c_\psi \alpha_{\mathrm{lb}} c_*^\beta/2$. □

D.2 Minimax sharpness with a binding box constraint

The full-action construction leaves open whether the same lower rate survives constraints that change the optimum. Here a positive face multiplier ensures that the box changes the minimizing action. The proof first constructs the complete smooth population boundary, then localizes each perturbed optimizer before using optimality conditions to retain the active face. Its last step establishes value separation and testing difficulty without changing any class constants.

Proposition D.8 (Binding-box family and minimax cost bound)**.** *Let $d \geq 2$ and $\beta \geq 2$. There exist $r > 0$, the fixed template $K_\square = [-r, r]^d$, an admissible compact query set $Q_\square$, and a nonempty family with common constants*

$$\Theta^\square_{\mathrm{bump}}(\beta; Q_\square) \subset \Theta_{\mathrm{val}}(\beta; Q_\square)$$

for which the lower statement in Theorem 4 holds. At every hard query, the baseline optimizer lies in the relative interior of the upper e_1-face of $x + K_\square$, its face multiplier is uniformly positive, and removing that face changes the optimizer and value.

Proof. **Step 1: baseline geometry and nonemptiness.** Fix $0 < r \leq 2^{-8}$ and put

$$D = d + 3, \qquad \nu = \frac{2e_1 + \sum_{i=2}^d e_i}{\sqrt{D}}, \qquad \delta = \frac{r}{2^{14}d^2}, \qquad k_\circ = re_1 + (r-\delta)\sum_{i=2}^d e_i. \tag{D.6}$$

Let $b_\circ = (1/2, \ldots, 1/2)$ and choose an orthonormal basis $\tau_1, \ldots, \tau_{d-1}$ of $\nu^\perp$. In coordinates $s_i = \tau_i^\top(z - b_\circ)$ and $t = \nu^\top(z - b_\circ)$, set $R = r/(64\sqrt{d})$ and $H = r/32$. Smoothly regularize, with width below $\min\{R, H\}/(100d)$, the maximum of

$$-t, \quad t - H, \quad s_i - R, \quad -s_i - R \quad (1 \leq i \leq d-1).$$

Call the result F and put $\Omega = \{F \le 0\}$. The regularized maximum is chosen above the ordinary maximum and equal to a uniquely dominant input away from its smoothing bands. Hence

$$\Omega \subset \{0 \le t \le H,\ |s_i| \le R\},$$

$B_0 = \partial\Omega$ is a compact C^∞ hypersurface. With $r_{\mathrm{flat}} = \min\{R, H\}/20$, the lower-face identity is

$$B_0 \cap \{|s|_\infty \le R/2,\ |t| \le r_{\mathrm{flat}}\} = \{(s, 0) : |s|_\infty \le R/2\}. \tag{D.7}$$

Indeed, opposing affine faces cannot be simultaneously active near $F = 0$, so ∇F is a nonzero convex combination of nonopposing coordinate normals.

Orient signed distance positively into Ω, extend it by a smooth sign-preserving defining function, and smoothly saturate away from its tube. For $\kappa_\square = 1/16$ and $q = 1/8$ this gives a global score g_0, positive exactly in $\mathrm{int}(\Omega)$, with strict range and Hölder room after choosing L above its finite norm, and

$$g_0(z) = \kappa_\square \nu^\top (z - b_\circ) \tag{D.8}$$

near (D.7). Let $r_B > 0$ be a valid normal-tube radius. On a smaller tube the score equals $\kappa_\square$ times signed distance, so crossing, off-tube, local-tube, and nonlocal-tube reserves are all positive before any bump is introduced.

Take

$$U_Q = \left\{ b_u = b_\circ + \sum_{i=1}^{d-1} u_i \tau_i : |u_i| \le \frac{\delta}{64\sqrt{d}} \right\}, \qquad x_u = b_u - k_\circ. \tag{D.9}$$

Every hard box lies in $(0, 1)^d$. To see why it meets B_0 only on the protected face, let $z \in B_0 \cap (x_u + K_\square)$ and write $w = z - b_u$. The upper-face inequalities give $w_1 \le 0$ and $w_i \le \delta$ for $i \ge 2$. Since $t = \nu^\top w \ge 0$ on Ω,

$$0 \le \sqrt{D}\, t = 2w_1 + \sum_{i=2}^{d} w_i \le (d-1)\delta, \quad 2|w_1| + \sum_{i:w_i<0} |w_i| \le (d-1)\delta. \tag{D.10}$$

Thus $\|w\|_2 \le \sqrt{d}(d-1)\delta \le R/256$; together with (D.9), every tangential coordinate is below $3R/512$ and $t \le (d-1)\delta/\sqrt{D} < r_{\mathrm{flat}}/4$. Hence z lies in the lower neighborhood (D.7), where $t = 0$. The rounded transition is therefore disjoint from the compact union of all hard boxes. Fix

$$\mathcal{G}_{\mathrm{prot}} = B_0 \cap \{|s|_\infty < R/3,\ |t| < r_{\mathrm{flat}}/2\};$$

all hard intersections form a compact subset of this open chart.

Set

$$x_{\mathrm{cov}} = b_\circ - (r/8)\nu, \qquad Q_\square = \{x_u : u \in U_Q\} \cup \{x_{\mathrm{cov}}\}. \tag{D.11}$$

Because $\|z - b_\circ\|_\infty < 3r/64$ on B_0, the whole boundary has at least $53r/64$ slack in $x_{\mathrm{cov}} + K_\square$. Both query types are unfavorable. The points $b_u + (\delta/2)e_d$ and $b_\circ + (H/2)\nu$, respectively, are strict favorable witnesses in their boxes. In particular, the following minima are positive and fixed in this order:

$$\delta_X = \mathrm{dist}(B_0, \partial\mathcal{X}), \quad \delta_{\mathrm{nf}} = \delta, \quad \delta_{\mathrm{ext}} = \mathrm{dist}\left(B_0 \setminus \mathcal{G}_{\mathrm{prot}},\ \bigcup_{u \in U_Q} (x_u + K_\square) \right),$$

$$\delta_{\mathrm{cov}} = \mathrm{dist}(B_0, \partial(x_{\mathrm{cov}} + K_\square)), \tag{D.12}$$

The nonactive-face minimum is evaluated at the baseline optimizers b_u, over every upper and lower face except the upper e_1-face; inactive upper faces attain δ, and lower faces have larger slack. Let $\delta_{\mathrm{chart}} > 0$ be one quarter of the distance from the compact hard-intersection set and U_Q to the edge of (D.7), including $|t| = r_{\mathrm{flat}}$. Only now choose

$$0 < r_0 < \frac{1}{32} \min\{r_B, \delta_X, \delta_{\mathrm{nf}}, \delta_{\mathrm{ext}}, \delta_{\mathrm{cov}}, \delta_{\mathrm{chart}}\}.$$

Put $\zeta_{\mathrm{reach}} = r_0$ and

$$\Gamma_{\mathrm{reach}} = \frac{1}{4}\left(r_0^{-1} - (r_0 + \zeta_{\mathrm{reach}})^{-1} \right) > 0.$$

Thus $r_0 + \zeta_{\mathrm{reach}} < r_B$. Shrinking r_0 within this fixed bound makes the score equal to $\kappa_\square$ times signed distance on the enlarged tube. We take $a = \kappa_\square/2$, $b = 3\kappa_\square/2$ and $\zeta_{\mathrm{cross}} = \kappa_\square/4$. The strict score-range and Hölder budgets are

$$\zeta_{\mathrm{range}} = \mathrm{dist}(g_0(\mathcal{X}), \partial[-1/2+q, 1/2-q]) > 0, \qquad \zeta_H = L - \|g_0\|_{\Sigma_d(\beta)} > 0.$$

Because B_0 is the full zero set, compactness gives

$$m_{\mathrm{off}} = \min_{\substack{z \in V_{K_\square} \\ \mathrm{dist}(z, B_0) \geq r_0/4}} |g_0(z)| > 0, \qquad \gamma_0 = \zeta_{\mathrm{off}} = m_{\mathrm{off}}/3.$$

Writing $\rho_0 = \rho_{g_0, K_\square}$, the compact unfavorable query set and the explicit strict witnesses give

$$m_Q = \tfrac{1}{2} \min_{x \in Q_\square} [-g_0(x)] > 0, \qquad \kappa_{\mathrm{feas}} = \tfrac{1}{2} \min_{x \in Q_\square} \max_{z \in A_{K_\square}(x)} g_0(z) > 0,$$

$$\underline{\rho}_Q = \tfrac{1}{2} \min_{x \in Q_\square} \rho_0(x) > 0, \qquad \overline{\rho}_Q = \tfrac{1}{2} \left(\max_{x \in Q_\square} \rho_0(x) + \mathrm{diam}(\mathcal{X}) \right),$$

and a positive query-value reserve

$$\zeta_Q^\rho = \min_{x \in Q_\square} \min\{\rho_0(x) - \underline{\rho}_Q, \overline{\rho}_Q - \rho_0(x)\} > 0.$$

For the tube constants, take a finite C^2 atlas $\Gamma_\ell : \overline{U}_\ell \to B_0$ with convex parameter blocks and nested compact blocks $U_\ell^- \Subset U_\ell$ covering B_0. The baseline normal maps

$$T_{0,\ell}(s,t) = \Gamma_\ell(s) + t n_0(\Gamma_\ell(s)), \qquad |t| \leq r_0 + \zeta_{\mathrm{reach}},$$

have full-rank derivatives. The bound $2r_0 < r_B/16$ controls normal-coordinate variation, since $\|Dn_0\| \leq 1/r_B$. Using spatial charts with orthonormal differential at their centers, refine the convex blocks until each derivative differs from its value J_ℓ at a fixed block point by at most $\sigma_{\min}(J_\ell)/2$. For two parameter points p, p' in that block, integration gives

$$\|T_{0,\ell}(p) - T_{0,\ell}(p')\|_2 \geq \|J_\ell(p-p')\|_2 - \tfrac{1}{2}\sigma_{\min}(J_\ell)\|p-p'\|_2 \geq \tfrac{1}{2}\sigma_{\min}(J_\ell)\|p-p'\|_2.$$

Taking the finite minimum and using $\|(s-s', t-t')\|_2 \geq (\|s-s'\|_2 + |t-t'|)/\sqrt{2}$ gives $m_{\mathrm{tube}} > 0$ such that

$$\|T_{0,\ell}(s,t) - T_{0,\ell}(s',t')\|_2 \geq m_{\mathrm{tube}}(\|s-s'\|_2 + |t-t'|)$$

on every common block. To separate the remaining pairs, choose open sets $W_\ell \Subset \Gamma_\ell(U_\ell^-)$ that still cover B_0 and let $\ell_{\mathrm{Leb}} > 0$ be a Lebesgue number of this finite cover. Every pair not handled by a common block then lies in the closed compact set

$$\mathcal{N}_{\mathrm{nl}} = \{(b,t,b',t') : \ b, b' \in B_0, \ |t|, |t'| \leq r_0, \ \|b-b'\|_2 \geq \ell_{\mathrm{Leb}}\}.$$

This set is separated from the parameter diagonal, so injectivity of the baseline normal map gives

$$\sigma_{\mathrm{tube}} := \min_{\mathcal{N}_{\mathrm{nl}}} \|b + t n_0(b) - b' - t' n_0(b')\|_2 > 0.$$

Finally define the length and angular reserves separately:

$$s_* = \frac{(d-1)\delta}{d+1}, \qquad \tau_* = \frac{1}{\sqrt{D}}, \qquad s_{\mathrm{ray}} = s_*/2, \qquad \zeta_\tau = \tau_*/2,$$

and

$$r_{\mathrm{cov}}^{\mathrm{pre}} = \frac{1}{16} \min\{r_0, \delta_{\mathrm{nf}}, \delta_{\mathrm{ext}}, \delta_{\mathrm{cov}}, \delta_{\mathrm{chart}}, \sigma_{\mathrm{tube}}, \zeta_Q^\rho\}.$$

All these constants precede the bump index, bandwidth, sample, and estimator. Finally, $B_2(0,r) \subset [-r,r]^d \subset B_2(0, r\sqrt{d})$, so $r_K = r$ and $R_K = r\sqrt{d}$.

Step 2: the constraint changes the optimum. Let $\rho_\circ = \|k_\circ\|_2$ and set

$$\lambda_\circ = (r-\delta)\sqrt{D}, \qquad \mu_\circ = r - 2\delta > 0, \qquad k_\circ = \lambda_\circ \nu - \mu_\circ e_1. \tag{D.13}$$

Any baseline feasible action k in a hard box obeys $k_1 \le r$ and, by (D.8)–(D.10), $\nu^\top(k-k_\circ) \ge 0$. Therefore

$$\langle k_\circ, k-k_\circ\rangle = \lambda_\circ \nu^\top(k-k_\circ) - \mu_\circ(k_1 - r) \ge 0, \qquad \|k\|_2^2 - \|k_\circ\|_2^2 \ge \|k-k_\circ\|_2^2. \tag{D.14}$$

Thus b_u is the unique baseline optimizer and the upper e_1-face is strictly active with multiplier $\mu_\circ$. The supporting-halfspace projection has first coordinate exceeding r by

$$\zeta_{\mathrm{proj}} := (\nu^\top k_\circ)\nu_1 - r = \frac{(d-1)\mu_\circ}{D} > 0. \tag{D.15}$$

For the actual compact body, take $v = e_1 - 2(d-1)^{-1}\sum_{i=2}^d e_i$, which is tangent to the protected face. For fixed small $t_{\mathrm{rm}} > 0$, $b_u + t_{\mathrm{rm}} v$ remains in the flat chart and is feasible after only the upper e_1-face is removed, while

$$\|k_\circ + t_{\mathrm{rm}} v\|_2^2 - \|k_\circ\|_2^2 = -2\mu_\circ t_{\mathrm{rm}} + t_{\mathrm{rm}}^2 \|v\|_2^2 < 0. \tag{D.16}$$

Removing the face therefore changes both the optimizer and its cost for the compact favorable set.

For action geometry, in action coordinates define

$$S = K_\square \cap \{y : \nu^\top(y - k_\circ) = 0\}, \qquad y_* = k_\circ - \frac{(d-1)\delta}{4} e_1 + \frac{\delta}{2}\sum_{i=2}^d e_i.$$

The point y_* lies in S and has cube slack at least $\sigma_* = \delta/4$. If y lies in the same hyperplane, $p = P_{K_\square} y$, and $d_0 = \mathrm{dist}(y, K_\square) \le \sigma_*$, project p back to the hyperplane and mix it with y_* by weight $d_0/(\sigma_* + d_0)$. The result is in S and within

$$\kappa_H d_0, \qquad \kappa_H = 2 + (2r\sqrt{d} + \sigma_*)/\sigma_* \tag{D.17}$$

of y. This is the affine boundary–action coverage bound. Moreover, if $a_i = r - y_i$ are the upper-face slacks of $y \in S$, then

$$2a_1 + \sum_{i=2}^d a_i = (d-1)\delta. \tag{D.18}$$

Some positive coordinate ray therefore has length at least $(d-1)\delta/(d+1)$ and normal alignment at least $1/\sqrt{D}$. The coverage box has the stronger strict-interior certificate.

Step 3: perturbed geometry and persistence of the binding face. For an Ah-separated center set $\{u_j\} \subset U_Q$, let $b_j = b_{u_j}$ and define

$$g_{j,h}(z) = g_0(z) + \alpha_\square h^\beta \psi\left(\frac{z - b_j}{h}\right), \tag{D.19}$$

where $\psi \in C_c^\infty((-1,1)^d)$ is nonnegative and positive at zero, r_ψ is its support radius, and $A > 2r_\psi + 2$. The derivative bounds

$$\|\partial^s(g_{j,h} - g_0)\|_\infty \le C_s \alpha_\square h^{\beta - |s|} \quad (|s| \le \lfloor\beta\rfloor) \tag{D.20}$$

first require

$$\alpha_\square \le \frac{(A - r_\psi)^2}{4\lambda_\circ C_\Gamma} \tag{D.21}$$

and choose it small enough for the fixed score-range, Hölder, off-tube, feasible-depth, $\beta = 2$ curvature, and persistent local-tube budgets. In particular,

$$C_{\mathrm{curv}}\alpha_\square \le \Gamma_{\mathrm{reach}}/2, \qquad r_0 C_2^n \alpha_\square \le m_{\mathrm{tube}}/8.$$

Then choose $h_\square \le 1$ for the remaining vanishing score, slope, normal, query, chart, support, and action reserves. In particular, the closed bump support lies strictly inside the protected affine chart and

$$C_1 \alpha_\square h^{\beta-1} \le \kappa_\square/2.$$

Write $H_0 = \{z : \nu^\top(z - b_\circ) = 0\}$ and a chart point as $b(s) + t\nu$, with $b(s) \in H_0$, and let $f_s(t) = g_{j,h}(b(s) + t\nu)$. At the fixed normal edges of the chart the bump vanishes and $f_s(t) = \kappa_\square t$, so the endpoint signs are opposite. Throughout the intervening segment,

$$f_s'(t) = \kappa_\square + \alpha_\square h^{\beta-1} \nabla\psi((b(s) + t\nu - b_j)/h)^\top \nu \geq \kappa_\square/2 > 0.$$

Hence there is exactly one root on each normal fibre. Off the projected bump support it is the baseline root; smooth compact support makes the root graph glue smoothly to the baseline graph at the support edge. The score is unchanged outside this chart, so no additional zero component is created. Implicit differentiation of this established graph gives

$$\|n_{j,h}(z) - \nu\|_2 \leq D_h, \qquad E_h = C_\Gamma \alpha_\square h^\beta, \quad D_h = C_n \alpha_\square h^{\beta-1}, \tag{D.22}$$

and, on each retained chart,

$$\|D\Gamma_{j,h} - D\Gamma_0\| \leq C_1^\Gamma \alpha_\square h^{\beta-1}, \qquad \|Dn_{j,h} - Dn_0\| \leq C_2^n \alpha_\square h^{\beta-2}.$$

The scale is chosen so that

$$D_h \leq \frac{1}{2\sqrt{D}}, \qquad E_h \leq \min\{(\nu^\top k_\circ)/4, r_0/4, \zeta_Q^\rho/4\}, \qquad 2E_h + 5(\nu^\top k_\circ)D_h \leq \zeta_{\mathrm{proj}}/2,$$

and

$$(C_1^\Gamma + C_n)\alpha_\square h^{\beta-1} + r_0 C_2^n \alpha_\square h^{\beta-2} \leq m_{\mathrm{tube}}/2, \qquad E_h + r_0 D_h \leq \sigma_{\mathrm{tube}}/4,$$

and, with $\alpha_1 = \min\{1, \beta - 1\}$,

$$C_1^g \alpha_\square h^{\beta-1} + L_1(E_h + r_0 D_h)^{\alpha_1} + L D_h \leq \zeta_{\mathrm{cross}}.$$

On common atlas blocks the perturbed normal map satisfies

$$\|T_{j,h}(s,t) - T_{j,h}(s',t')\|_2 \geq \frac{m_{\mathrm{tube}}}{2}(\|s - s'\|_2 + |t - t'|), \tag{D.23}$$

whereas nonlocal pairs are separated by

$$\sigma_{\mathrm{tube}} - 2(E_h + r_0 D_h) > 0. \tag{D.24}$$

Together with the preceding derivative bounds, these estimates give full tube injectivity, common reach at least r_0, and two-sided crossing in $[a, b]$.

To establish persistence of the binding face, we first localize the optimizer and then apply the Karush–Kuhn–Tucker (KKT) conditions. Every perturbed feasible action k obeys $\nu^\top(k - k_\circ) \geq -E_h$. Combining this with (D.13) yields

$$\|k\|_2^2 - \|k_\circ\|_2^2 \geq \|k - k_\circ\|_2^2 - 2\lambda_\circ E_h. \tag{D.25}$$

Since $k_\circ$ remains feasible, every optimizer k^* first satisfies

$$\|k^* - k_\circ\|_2 \leq \sqrt{2\lambda_\circ E_h}. \tag{D.26}$$

Choose $h_\square$ so (D.26) is below one quarter of every nonactive-face (including the lower e_1-face), chart-edge, and competing-component reserve. Only now suppose $k_1^* < r$. All box faces are inactive and boundary KKT gives $k^* = t n_{j,h}$. Writing $s_\circ = \nu^\top k_\circ$, $D_h \leq 1/(2\sqrt{D})$, and $E_h \leq s_\circ/4$, direct division by $\nu^\top n_{j,h} \geq 1 - D_h^2/2$ gives

$$|k_1^* - s_\circ \nu_1| \leq 2E_h + 5 s_\circ D_h. \tag{D.27}$$

Making this at most $\zeta_{\mathrm{proj}}/2$ contradicts $k_1^* \leq r$. Thus the active face persists.

Coverage persists explicitly. At x_{cov} the whole perturbed boundary remains inside the box, so the coverage distance is zero. For a hard query, the δ_{ext} and displacement budgets confine every zero point within r_{cov} of its box to the protected affine chart. Write $S_{j,h}$ for the bump support. Since $\mathrm{diam}_\infty(U_Q) \leq \delta/32$, reduce $h_\square$ so $r_\psi h + E_h \leq \delta/32$ and $C_1 \alpha_\square h^{\beta-1} \leq \kappa_\square/(2\sqrt{D})$. With $\Phi_{j,h} = g_{j,h}$ on the affine chart,

$$\partial_d \Phi_{j,h} \geq \frac{\kappa_\square}{2\sqrt{D}} =: m_\Phi, \qquad \mathrm{Lip}(\Phi_{j,h}) \leq L_\Phi.$$

Choose

$$r_{\mathrm{cov}} \le \min\left\{ r_{\mathrm{cov}}^{\mathrm{pre}}, \frac{\delta}{16(\kappa_H+3)}, \frac{\delta m_\Phi}{16(m_\Phi+L_\Phi)}, \frac{\sigma_*}{2} \right\}.$$

Then the full near-support budget is

$$\mathrm{diam}_\infty(U_Q) + r_\psi h + E_h + (\kappa_H+3) r_{\mathrm{cov}} \le \delta/8 < \delta/4.$$

If $\mathrm{dist}(b, S_{j,h}) > (\kappa_H+2) r_{\mathrm{cov}}$, the point from (D.17) lies outside the support and remains an exact perturbed zero. Otherwise $\mathrm{dist}(b, S_{j,h}) \le (\kappa_H+2) r_{\mathrm{cov}}$, and a nearest box point a contributes $\|a-b\|_2 < r_{\mathrm{cov}}$. The combined displacement bound is therefore $(\kappa_H+3) r_{\mathrm{cov}}$, as in the preceding display. It leaves more than $3\delta/4$ e_d-face slack, and monotonicity produces $c = a + te_d$ on the perturbed zero set with

$$|t| \le (L_\Phi/m_\Phi)\|a-b\|_2, \quad \|b-c\|_2 \le (1 + L_\Phi/m_\Phi)\,\mathrm{dist}(b, A_K(x)). \tag{D.28}$$

The correction is below $\delta/16$, so c stays in the box. This gives $\kappa_{\mathrm{cov}} = \max\{\kappa_H, 1 + L_\Phi/m_\Phi\}$.

For radial crossing take

$$s_A = \min\{s_*/2, \delta/8, \delta_{\mathrm{cov}}/4, r_{\mathrm{flat}}/8, r_0/2\}, \qquad \tau = 1/(2\sqrt{D}).$$

At the coverage query, $s_A \le \delta_{\mathrm{cov}}/4$ keeps the perturbed inward normal ray in the box and $s_A \le r_0/2$ keeps it in the certified common normal tube, where two-sided crossing makes the segment favorable. Farther than s_A from the bump support, the selected coordinate ray from (D.18) stays outside the support. The hard-intersection estimate places its initial point at tangential size below $3R/512$ and normal coordinate zero; because $s_A \le \delta/8 \ll R$ and $s_A \le r_{\mathrm{flat}}/8$, every point on the selected ray remains in the protected affine region and is favorable. Within the support neighborhood the same combined displacement bound leaves e_d-face slack and $n_{j,h}^\top e_d \ge 1/(2\sqrt{D})$, while $\partial_d \Phi_{j,h} > 0$ keeps the full ray favorable.

We now verify the remaining conditions for membership in the value class. The design is uniform and the fixed cube has the stated inner and outer radii. Equation (D.11) is the continuous image of the compact rectangle U_Q together with one point, so $Q_\square$ is fixed, nonempty, and compact; the coordinate bounds used above put every $x + K_\square$ strictly inside $\mathcal{X}$. The derivative budgets preserve score range and Hölder norm. The monotone normal-fibre argument gives one smooth graph glued to the unchanged zero set, while (D.23)–(D.24), the curvature budget, and the displayed crossing inequality give the common regular boundary block. If $\mathrm{dist}(z, B_{j,h}) \ge r_0/2$, then $\mathrm{dist}(z, B_0) \ge r_0/4$ and

$$|g_{j,h}(z)| \ge m_{\mathrm{off}} - C_0 \alpha_\square h^\beta \ge 2m_{\mathrm{off}}/3 \ge \gamma_0.$$

Queries lie outside all supports, so their scores are unchanged, and nonnegative bumps preserve feasible depth and population recourse.

To verify the query-cost bounds, we establish the transfer directly. Put

$$\varepsilon_h = \|g_{j,h} - g_0\|_\infty \le C_0 \alpha_\square h^\beta, \qquad C_{\mathrm{int}} = \frac{1+\kappa_{\mathrm{cov}}}{a} + \frac{2}{a\tau},$$

and

$$s_{\mathrm{cross}} = \min\left\{ s_A, \frac{r_0}{2}, \left(\frac{a\tau}{2L_1}\right)^{1/\alpha_1} \right\}.$$

Reduce the fixed $h_\square$ once more so that

$$\varepsilon_h \le \min\left\{ \frac{\gamma_0}{2}, \frac{a r_0}{8}, \frac{a r_{\mathrm{cov}}}{8}, \frac{\kappa_{\mathrm{feas}}}{2}, \frac{a\tau s_{\mathrm{cross}}}{2}, \frac{\zeta_Q^\rho}{C_{\mathrm{int}}} \right\} \tag{D.29}$$

at every retained scale. Since the bump is nonnegative, $C_{g_0, K_\square}(x) \subset C_{g_{j,h}, K_\square}(x)$ and hence $\rho_{j,h}(x) \le \rho_0(x)$. Conversely, take $z \in C_{g_{j,h}, K_\square}(x)$. If $g_0(z) \ge 0$, then z is already baseline feasible. Otherwise $-\varepsilon_h \le g_0(z) < 0$. Off-tube separation and normal crossing yield $z = T_0(b,t)$ with

$$|t| = \mathrm{dist}(z, B_0) \le \varepsilon_h/a.$$

Because z is action-feasible, baseline boundary–action coverage supplies $c \in B_0 \cap A_{K_\square}(x)$ with $\|b-c\|_2 \le \kappa_{\mathrm{cov}} \varepsilon_h/a$. Let v_c be its radial direction and put $q_h = 2\varepsilon_h/(a\tau) \le s_{\mathrm{cross}}$. At c, $\nabla g_0(c) = \xi_c n_0(c)$ with $\xi_c \ge a$, and for $0 \le q \le q_h$ the gradient modulus gives

$$\nabla g_0(c + q v_c)^\top v_c \ge a\tau - L_1 q^{\alpha_1} \ge a\tau/2.$$

Thus $\bar{c} = c + q_h v_c$ is baseline feasible and

$$\|z - \bar{c}\|_2 \le \left(\frac{1+\kappa_{\mathrm{cov}}}{a} + \frac{2}{a\tau}\right)\varepsilon_h = C_{\mathrm{int}}\varepsilon_h.$$

Taking infima and using the feasible-set inclusion gives

$$0 \le \rho_0(x) - \rho_{j,h}(x) \le C_{\mathrm{int}}\varepsilon_h \le \zeta_Q^{\rho}.$$

The strict baseline room therefore preserves the common $[\underline{\rho}_Q, \overline{\rho}_Q]$ query-value interval. Together with the coverage and radial arguments above, this verifies every memberwise value-class clause. Thus, before testing,

$$\varnothing \ne \Theta^{\square}_{\mathrm{bump}}(\beta; Q_{\square}) \subset \Theta_{\mathrm{val}}(\beta; Q_{\square}). \tag{D.30}$$

Step 4: querywise value separation. On that face the line $b_j - te_d$ stays feasible for small t, and its zero equation is

$$-\kappa_{\square}\nu_d t + \alpha_{\square} h^{\beta}\psi(-te_d/h) = 0, \qquad \nu_d = 1/\sqrt{D}. \tag{D.31}$$

Choose $\eta, \psi_0 > 0$ so $\psi(v) \ge \psi_0$ on $\|v\|_2 \le \eta$ and put $M_\psi = \|\psi\|_\infty$. For small h, the bracket

$$t_- := \frac{\psi_0 \alpha_{\square} h^{\beta}}{2\kappa_{\square}\nu_d}, \qquad t_+ := \frac{2M_\psi \alpha_{\square} h^{\beta}}{\kappa_{\square}\nu_d} \tag{D.32}$$

has opposite endpoint signs, and the derivative of the left side in (D.31) is at most $-\kappa_{\square}\nu_d/2$. Hence the unique root $t_{j,h}$ lies in (D.32). The matched action $k_\circ - t_{j,h}e_d$ gives

$$\rho_\circ - \rho_{j,h}(x_{u_j}) \ge \rho_\circ - \|k_\circ - t_{j,h}e_d\|_2 = \frac{2(r-\delta)t_{j,h} - t_{j,h}^2}{\rho_\circ + \|k_\circ - t_{j,h}e_d\|_2} \ge c_\Delta \alpha_{\square} h^{\beta}. \tag{D.33}$$

Erode the orthonormal rectangle (D.9) by $(A + r_\psi)h$ on every side. For all small h the eroded rectangle retains a fixed fraction of its volume, and a maximal Ah-packing gives $N_h \ge c_A h^{-(d-1)}$. Its supports are contained in the chart and disjoint. If z lies in the support of a wrong bump $k \ne j$ and $a_z = z - x_{u_j}$, then

$$\|a_z - k_\circ\|_2 \ge (A - r_\psi)h, \qquad \nu^{\top}(a_z - k_\circ) \ge -E'_h, \qquad E'_h = C_\Gamma \alpha_{\square} h^{\beta}.$$

Using $a_{z,1} \le r$ and (D.13),

$$\|a_z\|_2^2 - \rho_\circ^2 \ge (A - r_\psi)^2 h^2 - 2\lambda_\circ E'_h \ge \tfrac{1}{2}(A - r_\psi)^2 h^2. \tag{D.34}$$

The last inequality is (D.21) when $\beta = 2$ and follows by shrinking $h_{\square}$ when $\beta > 2$. Every feasible favorable point is either outside the wrong support, where (D.14) applies, or inside it, where (D.34) applies. Also b_j is outside the wrong support and remains feasible. Hence the exact separation is

$$\rho_{k,h}(x_{u_j}) = \rho_\circ \quad (k \ne j), \qquad \rho_{j,h}(x_{u_j}) \le \rho_\circ - \Delta_h, \quad \Delta_h = c_\Delta \alpha_{\square} h^{\beta}. \tag{D.35}$$

Step 5: statistical testing. Let $J_h = \{0, 1, \dots, N_h\}$ and use uniform design. Since all Bernoulli probabilities lie in $[q, 1-q]$, mutual absolute continuity and a change of variables give

$$D_{\mathrm{KL}}(P^n_{j,h}, P^n_0) \le \frac{n\alpha_{\square}^2 h^{2\beta+d}}{q(1-q)} \int_{\mathbb{R}^d} \psi(v)^2\, dv = C^{\square}_{\mathrm{KL}} n \alpha_{\square}^2 h^{2\beta+d}. \tag{D.36}$$

Let $\chi = 2\beta + d$ and $h_n = c_*(\log n/n)^{1/\chi}$. Choose c_* so that

$$C^{\square}_{\mathrm{KL}} \alpha_{\square}^2 c_*^{\chi} \le \alpha_F \frac{d-1}{2\chi}, \qquad \alpha_F < 1/8.$$

Indeed,

$$\log N_{h_n} \ge \log c_A - (d-1)\log c_* + \frac{d-1}{\chi}(\log n - \log\log n), \tag{D.37}$$

which is at least $(d-1)(2\chi)^{-1}\log n$ beyond one finite threshold. Increase that threshold until h_n satisfies every construction reserve and $N_{h_n} \ge 4$.

For an arbitrary estimator define the measurable, smallest-index decoder

$$\widehat{J} = \min \arg \min_{j \in J_{h_n}} \max_{1 \le \ell \le N_{h_n}} |\widetilde{\rho}_n(x_{u_\ell}) - \rho_{j,h_n}(x_{u_\ell})|.$$

By (D.35), value error below $\Delta_{h_n}/2$ implies correct decoding. We use the same finite reference-form Fano inequality as in Proposition D.7: if the average reference KL is at most $\alpha_F \log N$, then the metric error probability is at least

$$c_F = \frac{\sqrt{2}}{1+\sqrt{2}} \left(1 - 2\alpha_F - \sqrt{\frac{2\alpha_F}{\log 2}}\right) > 0.$$

Applied to (D.36)–(D.37), it gives

$$\inf_{\widetilde{\rho}_n} \max_{j \in J_{h_n}} P^n_{j,h_n} \left\{ L_{\rho,Q_\square}(\widetilde{\rho}_n, \rho_{j,h_n}) \ge \Delta_{h_n}/2 \right\} \ge c_F.$$

Applying $\mathbb{E}Z \ge sP\{Z \ge s\}$ memberwise, then maximizing, infimizing, and finally enlarging the finite family gives

$$\inf_{\widetilde{\rho}_n} \sup_{\theta \in \Theta^\square_{\mathrm{bump}}(\beta; Q_\square)} \mathbb{E}_\theta L_{\rho,Q_\square}(\widetilde{\rho}_n, \rho_{g,K_\square}) \ge c_\square \left(\frac{\log n}{n}\right)^{\beta/(2\beta+d)}.$$

with $c_\square = c_F c_\Delta \alpha_\square c_*^\beta/2$. All geometric and statistical constants, and the known pair $(Q_\square, K_\square)$, were fixed before n. □

Parts (ii) and (iii) of Theorem 4 follow from Propositions D.7 and D.8. Its upper bound follows from Proposition C.9, and Corollary C.10 supplies the corresponding expected-loss upper rate.

E Optimal-Action Recovery

We first prove the decision interpretation of the menu loss (Proposition 5). Theorem 6 then follows in two stages: the regular two-ball family in Section E.2 gives an explicit menu discontinuity, and the two-point test in Section E.3 turns it into a statistical lower bound on a fixed class. Section E.4 proves Theorem 7 under uniqueness and quadratic growth. The negative and positive results concern different population classes, as defined in Appendix A.4.

E.1 Secondary-utility interpretation of the menu loss

The following identity expresses symmetric Hausdorff loss through the best attainable secondary utility on each menu.

Proof of Proposition 5. For nonempty compact sets E and F, denote their directed excess by $h(E,F) = \sup_{e \in E} \operatorname{dist}(e,F)$. Every 1-Lipschitz function u satisfies, for each $e \in E$,

$$u(e) \le \max_{f \in F} u(f) + \operatorname{dist}(e,F),$$

so $\max_E u - \max_F u \le h(E,F)$. Interchanging E, F gives

$$|\max_E u - \max_F u| \le \max\{h(E,F), h(F,E)\} = d_H(E,F).$$

Conversely, $u_F(z) = \operatorname{dist}(z,F)$ is 1-Lipschitz, $\max_F u_F = 0$, and $\max_E u_F = h(E,F)$. The choice $u_E(z) = \operatorname{dist}(z,E)$ realizes the reverse directed excess. Taking the larger one proves the identity. Replacing u by $-u$ proves the minimization form; taking the supremum over $x \in Q$ proves the loss representation.

□

E.2 A FIXED REGULAR TWO-BALL FAMILY

Construction E.1 (Outside-tube two-ball chart)**.** This specification defines the family used for the menu lower bound. The next proposition supplies explicit parameters with common value-class constants. Let $e_1 = (1, 0, \dots, 0)$, $c_\pm = x_\dagger \pm a_{\mathrm{sep}} e_1$, $R_{\mathrm{sep}} = a_{\mathrm{sep}} - r_{\mathrm{sep}}$, with $a_{\mathrm{sep}} > r_{\mathrm{sep}} > 0$. For $0 \le t \le t_0$, the disjoint balls $\overline{B}(c_+, r_{\mathrm{sep}} + t)$ and $\overline{B}(c_-, r_{\mathrm{sep}})$ lie in $\mathrm{int}(\mathcal{X})$ and away from Q. With $K = K_{\mathrm{full}}$ and a density bounded between p_- and p_+, define

$$g_t(z) = -\kappa_{\mathrm{sep}}\big((r_{\mathrm{sep}} + t)^2 - \|z - c_+\|_2^2\big)\big(r_{\mathrm{sep}}^2 - \|z - c_-\|_2^2\big).$$

The chart is chosen so that, uniformly in t, the following eight checks hold with common value-class constants: (i) score range; (ii) Hölder regularity; (iii) an injective oriented tube and two-sided crossing; (iv) off-tube separation; (v) query margin, value range, and nonempty recourse; (vi) action radial crossing; (vii) boundary–action coverage, trivial because $A_{K_{\mathrm{full}}}(x) = \mathcal{X}$; and (viii) feasible depth. These conditions ensure membership in the value class. Moreover $R_{\mathrm{sep}} - t_0 > r_{\mathrm{tube}} \ge r_0$, so the designated query lies outside the common injective tube.

Proposition E.2 (Realizability of the two-ball construction)**.** *The compatible regime in Construction E.1 is nonempty. In particular, with $\epsilon_\circ = 10^{-2}$, take*

$$x_\dagger = (1/2, \dots, 1/2), \quad Q = Q_{\mathrm{sep}} := \overline{B}(x_\dagger, \epsilon_\circ/4), \quad s_Q = 1/4, \quad r_{\mathrm{tie}} = \epsilon_\circ,$$

$$r_{\mathrm{tube}} = r_0 = \epsilon_\circ, \quad r_{\mathrm{sep}} = 2\epsilon_\circ, \quad t_0 = \epsilon_\circ, \quad a_{\mathrm{sep}} = 20\epsilon_\circ, \quad s_{\mathrm{sep}} = 16\epsilon_\circ, \quad R_{\mathrm{sep}} = 18\epsilon_\circ.$$

With $K = K_{\mathrm{full}}$ and uniform design, there are positive common value-class constants and a nonempty interval of κ_{sep} for which the entire family

$$g_t(z) = -\kappa_{\mathrm{sep}}\big((r_{\mathrm{sep}} + t)^2 - \|z - c_+\|_2^2\big)\big(r_{\mathrm{sep}}^2 - \|z - c_-\|_2^2\big), \quad 0 \le t \le t_0,$$

with the sign convention of Construction E.1, belongs to $\Theta_{\mathrm{val}}(\beta; Q_{\mathrm{sep}})$. Thus $\varnothing \ne \Theta_{\mathrm{sep}}(\beta; Q_{\mathrm{sep}}) \subset \Theta_{\mathrm{val}}(\beta; Q_{\mathrm{sep}})$.

Proof. Put $e_1 = (1, 0, \dots, 0)$ and $c_\pm = x_\dagger \pm a_{\mathrm{sep}} e_1$. The ball $\overline{B}(x_\dagger, 24\epsilon_\circ)$ lies in the cube interior, while

$$R_{\mathrm{sep}} - t_0 = 17\epsilon_\circ > r_{\mathrm{tube}}, \qquad 2r_{\mathrm{tube}} < 2a_{\mathrm{sep}} - 2r_{\mathrm{sep}} - t_0.$$

The enlarged balls lie in that cube ball, are disjoint, and are separated from Q by at least s_{sep}; for example

$$\|x - c_+\|_2 - (r_{\mathrm{sep}} + t_0) \ge 20\epsilon_\circ - \epsilon_\circ/4 - 3\epsilon_\circ = \frac{67}{4}\epsilon_\circ > 16\epsilon_\circ \quad (x \in Q).$$

The full action has $A_{K_{\mathrm{full}}}(x) = \mathcal{X}$, and uniform design has density one. Write

$$A_t(z) = (r_{\mathrm{sep}} + t)^2 - \|z - c_+\|_2^2, \quad B(z) = r_{\mathrm{sep}}^2 - \|z - c_-\|_2^2, \quad P_t(z) = -A_t(z)B(z),$$

so $g_t = \kappa_{\mathrm{sep}} P_t$. The polynomial envelope

$$M_{\mathrm{range}} = (d + 9\epsilon_\circ^2)(d + 4\epsilon_\circ^2), \qquad M_H = \sup_{0 \le t \le t_0} \|P_t\|_{\Sigma_d(\beta)}$$

is finite. The zero set is the disjoint union of the two spheres. On either sphere, a tube point has the form $z = c + (\rho - s)u$ with $|s| \le \epsilon_\circ$, and the favorable normal is $-u$. The factor belonging to that sphere has absolute value at most $7\epsilon_\circ^2$ and normal derivative between $2\epsilon_\circ$ and $8\epsilon_\circ$. The other factor is negative and has absolute value between $1287\epsilon_\circ^2$ and $1936\epsilon_\circ^2$, while its normal derivative has absolute value at most $88\epsilon_\circ$. Therefore

$$m_{\mathrm{cross}} := 1958\epsilon_\circ^3, \qquad M_{\mathrm{cross}} := 16104\epsilon_\circ^3$$

satisfy

$$\kappa_{\mathrm{sep}} m_{\mathrm{cross}} \le \nabla g_t(z)^\top n_g(z) \le \kappa_{\mathrm{sep}} M_{\mathrm{cross}}. \tag{E.1}$$

If a point is at least $r_0/2$ from the two spheres, both quadratic factors have absolute value at least $\epsilon_\circ^2$, so with $m_{\mathrm{off}} = \epsilon_\circ^4$,

$$|g_t(z)| \ge \kappa_{\mathrm{sep}} m_{\mathrm{off}}. \tag{E.2}$$

For $x \in Q$,

$$-A_t(x) \ge \frac{6097}{16}\epsilon_\circ^2, \qquad -B(x) \ge \frac{6177}{16}\epsilon_\circ^2,$$

and hence, with $m_Q^\circ = (6097 \cdot 6177/256)\epsilon_\circ^4$, $g_t(x) \le -\kappa_{\rm sep} m_Q^\circ$. At c_+,

$$g_t(c_+) \ge \kappa_{\rm sep} m_{\rm pos}, \qquad m_{\rm pos} = 6384\epsilon_\circ^4. \tag{E.3}$$

Let $c_q = 1/2 - q$ and $\kappa_\star = c_q/(4M_{\rm range})$. Choose

$$a = \frac{\kappa_\star m_{\rm cross}}{2}, \quad b = 2\kappa_\star M_{\rm cross},$$
$$\gamma_0 = \frac{\kappa_\star m_{\rm off}}{2}, \quad m_Q = \frac{\kappa_\star m_Q^\circ}{2},$$
$$\kappa_{\rm feas} = \frac{\kappa_\star m_{\rm pos}}{2}, \quad L = 2\kappa_\star M_H.$$

All are positive and $a < b$. The simultaneous requirements on the score scale are exactly

$$\kappa_{\rm sep} \ge \max\left\{\frac{a}{m_{\rm cross}}, \frac{\gamma_0}{m_{\rm off}}, \frac{m_Q}{m_Q^\circ}, \frac{\kappa_{\rm feas}}{m_{\rm pos}}\right\} = \frac{\kappa_\star}{2},$$

$$\kappa_{\rm sep} \le \min\left\{\frac{b}{M_{\rm cross}}, \frac{c_q}{M_{\rm range}}, \frac{L}{M_H}\right\} = 2\kappa_\star.$$

Thus $I_\kappa = [\kappa_\star/2, 2\kappa_\star]$ has positive length and simultaneously discharges score range, Hölder regularity, crossing, off-tube separation, query margin, and feasible depth. Choose $\underline{\rho}_Q = 16\epsilon_\circ$, $\overline{\rho}_Q = 19\epsilon_\circ$, $\Delta_{\rm out} = \epsilon_\circ/2$, and $r_{\rm out} = 35\epsilon_\circ/2$; the query distance bounds give the required value interval and strict outer-radius room.

The two spherical normal tubes are injective because $r_{\rm tube} < r_{\rm sep}$ and their images are disjoint. Taking $s_A = \epsilon_\circ$, $\tau = 1$, and the inward normal verifies radial crossing. Since the full action is the cube, boundary–action coverage holds with $\kappa_{\rm cov} = 1$ and any $r_{\rm cov} \le r_0/4$. The uniform design and full action verify the remaining design and action clauses. The displayed strict bounds, together with the finite polynomial envelope, verify all value-class conditions with the same constants. Hence the family is nonempty and contained in $\Theta_{\rm val}(\beta; Q_{\rm sep})$. □

Value continuity and correspondence instability.

Proposition E.3 (Exact two-ball value and correspondence formulas)**.** *Under Proposition E.2 and the compatible regime of Construction E.1, the family $\theta_t = (K_{\rm full}, \mu_{\rm sep}, g_t)$ belongs to $\Theta_{\rm val}(\beta; Q_{\rm sep})$ with common constants. Its favorable set is*

$$D_{g_t}^+ = \overline{B}(c_+, r_{\rm sep} + t) \cup \overline{B}(c_-, r_{\rm sep}).$$

Moreover, for $0 \le t \le t_0$,

$$\|g_t - g_0\|_{\infty,\mathcal{X}} \le C_{\rm sep} t, \qquad \sup_{x \in Q_{\rm sep}} d_H(C_{g_t,K_{\rm full}}(x), C_{g_0,K_{\rm full}}(x)) = t,$$

$$L_{\rho,Q_{\rm sep}}(\rho_{g_t,K_{\rm full}}, \rho_{g_0,K_{\rm full}}) \le t.$$

At $x_\dagger$,

$$\Pi_{g_0,K_{\rm full}}(x_\dagger) = \{x_\dagger - R_{\rm sep} e_1, x_\dagger + R_{\rm sep} e_1\},$$

while for $t > 0$,

$$\Pi_{g_t,K_{\rm full}}(x_\dagger) = \{x_\dagger + (R_{\rm sep} - t)e_1\}, \quad d_H(\Pi_{g_t,K_{\rm full}}(x_\dagger), \Pi_{g_0,K_{\rm full}}(x_\dagger)) = 2R_{\rm sep} - t.$$

Proof. Set

$$A_t(z) = (r_{\rm sep} + t)^2 - \|z - c_+\|_2^2, \qquad B(z) = r_{\rm sep}^2 - \|z - c_-\|_2^2.$$

The center distance is $2a_{\rm sep}$ and the sum of radii is at most $2r_{\rm sep} + t_0 < 2a_{\rm sep}$, so the balls are disjoint. Outside both balls, $A_t, B < 0$ and $g_t = -\kappa_{\rm sep} A_t B < 0$; inside either ball exactly one factor is positive and the other negative, so $g_t > 0$. On the spheres one factor vanishes. This proves the favorable-set union.

All score-range, smoothness, tube, query, and action conditions with common constants were established in Proposition E.2. It remains to compute the score, feasible-set, value, and optimizer changes.

Only the plus-ball radius changes, and

$$g_t(z) - g_0(z) = -\kappa_{\mathrm{sep}}(2r_{\mathrm{sep}}t + t^2)B(z).$$

Therefore

$$\|g_t - g_0\|_{\infty,\mathcal{X}} \le \kappa_{\mathrm{sep}}(2r_{\mathrm{sep}} + t_0)t \sup_{z\in\mathcal{X}} |B(z)| = C_{\mathrm{sep}}t.$$

The baseline favorable set is contained in the perturbed set. Every point in the enlarged spherical shell is within t of the baseline plus ball, so the directed Hausdorff distance from the perturbed set to the baseline set is at most t; the reverse directed distance is zero. The point $z_t = c_+ + (r_{\mathrm{sep}} + t)e_1$ is at distance exactly t from the baseline plus ball, and the minus ball is no closer by disjointness. Hence equality holds. The elementary distance-to-set inequality then gives the value bound.

At the designated query, $c_\pm = x_\dagger \pm a_{\mathrm{sep}}e_1$. The distances to the plus and minus balls are respectively $R_{\mathrm{sep}} - t$ and R_{sep}. At $t = 0$ they tie and the nearest points are the two displayed endpoints. For $t > 0$, the plus endpoint is uniquely closer. Its distance to the baseline plus endpoint is t, while the distance to the baseline minus endpoint is $2R_{\mathrm{sep}} - t$. Because $R_{\mathrm{sep}} - t_0 > r_{\mathrm{tube}} > 0$, the latter is larger than t, and the exact Hausdorff formula follows. This proves the nonvanishing limit of the full correspondence while the score and value perturbations vanish. □

E.3 Statistical non-recovery of the exact menu

Fix one admissible κ_{sep} in Proposition E.2, so the design, action template, query set, and all class constants are independent of n and t. Let P_t be the one-observation law and $p_t(z) = 1/2 + g_t(z)$. For $u, v \in (0,1)$, $\log s \le s - 1$ gives

$$u \log \frac{u}{v} + (1-u)\log\frac{1-u}{1-v} \le \frac{(u-v)^2}{v(1-v)}.$$

Since $p_t, p_0 \in [q, 1-q]$ and the design μ_{sep} is shared,

$$D_{\mathrm{KL}}(P_t^n \| P_0^n) = n \int D_{\mathrm{KL}}(\mathrm{Bern}(p_t(z)) \| \mathrm{Bern}(p_0(z)))\, d\mu_{\mathrm{sep}}(z) \le \frac{nC_{\mathrm{sep}}^2 t^2}{q(1-q)}.$$

Choose

$$c = \min\left\{ t_0, R_{\mathrm{sep}}, \frac{\sqrt{q(1-q)}}{\sqrt{8}C_{\mathrm{sep}}} \right\} > 0, \qquad t_n = c/\sqrt{n}.$$

Both mechanisms belong to the fixed family for every $n \ge 1$. Their KL divergence is at most $1/8$, so Pinsker's inequality gives $\|P_{t_n}^n - P_0^n\|_{\mathrm{TV}} \le 1/4$. Set $M_0 = \Pi_{g_0}(x_\dagger)$ and $M_1 = \Pi_{g_{t_n}}(x_\dagger)$. By Proposition E.3, $\Delta_n = d_H(M_0, M_1) = 2R_{\mathrm{sep}} - t_n \ge R_{\mathrm{sep}}$.

For any estimator in Theorem 6, form the measurable test that selects the closer of M_0, M_1 to $\widehat{\Pi}_n(x_\dagger)$ in Hausdorff distance, breaking ties toward zero. On a testing error under mechanism i, the triangle inequality forces $d_H(\widehat{\Pi}_n(x_\dagger), M_i) \ge \Delta_n/2$. The sum of the two testing error probabilities is at least $1 - \|P_{t_n}^n - P_0^n\|_{\mathrm{TV}}$. Hence

$$\max_{i\in\{0,1\}} \mathbb{E}_i L_{\Pi,Q_{\mathrm{sep}}}(\widehat{\Pi}_n, \Pi_i) \ge \frac{\Delta_n}{4}(1 - \|P_{t_n}^n - P_0^n\|_{\mathrm{TV}}) \ge \frac{3R_{\mathrm{sep}}}{16}.$$

Taking the infimum proves (5.1) with $c_* = 3R_{\mathrm{sep}}/16$. Since Hausdorff distances between nonempty subsets of $\mathcal{X}$ are bounded by $\mathrm{diam}(\mathcal{X})$, uniform consistency in probability would also imply uniformly vanishing expected loss, contradicting this bound. The local alternatives depend on n; the parameter class does not. Theorem 4, with its common constants, simultaneously provides uniformly consistent value estimation on this family.

E.4 Conditional recovery under quadratic growth

Unlike the preceding class, the location class has one population optimizer and a uniform growth condition. Feasible-set accuracy provides a nearby population-feasible point for each estimated optimum. Its excess cost is small, so quadratic growth bounds its displacement.

Proposition E.4 (Square-root bound for optimal-action recovery)**.** *Let $\theta \in \Theta_{\mathrm{loc}}(\beta; Q)$ and let $\widetilde{g}$ be continuous with $\|\widetilde{g} - g\|_{\infty,\mathcal{X}} \le \varepsilon$ satisfy the scalar and intersection activations in Theorem 7. Then every perturbed feasible set and argmin correspondence is nonempty and*

$$\sup_{x \in Q} d_H(\Pi_{\widetilde{g},K}(x), \{z_g(x)\}) \le C_{\mathrm{int}}\varepsilon + 2\sqrt{C_{\mathrm{int}}\varepsilon/\lambda} \le C_{\mathrm{loc}}\sqrt{\varepsilon}.$$

The fixed-sample score event gives the probability statement in Theorem 7.

Proof. Fix x and abbreviate

$$C = C_{g,K}(x), \quad \widetilde{C} = C_{\widetilde{g},K}(x), \quad \rho = \rho_{g,K}(x), \quad \widetilde{\rho} = \rho_{\widetilde{g},K}(x), \quad s = C_{\mathrm{int}}\varepsilon.$$

The exact intersection bridge gives nonempty compact sets, $d_H(C, \widetilde{C}) \le s$, and $|\widetilde{\rho} - \rho| \le s$. The optimizer-growth assumption gives $\Pi_{g,K}(x) = \{z_g(x)\}$.

Take any $\widetilde{z} \in \Pi_{\widetilde{g},K}(x)$. The directed Hausdorff bound supplies $z \in C$ with $\|\widetilde{z} - z\|_2 \le s$. Since $\widetilde{z}$ is a perturbed minimizer,

$$\|z - x\|_2 - \rho \le \|\widetilde{z} - x\|_2 + s - \rho = \widetilde{\rho} + s - \rho \le 2s.$$

Quadratic growth gives

$$2s \ge \lambda \min\{\|z - z_g(x)\|_2^2, r_*^2\}.$$

The strict activation $4s < \lambda r_*^2$ excludes the truncated branch, because it would force the left side to be at least $\lambda r_*^2 > 2s$. Thus

$$\|z - z_g(x)\|_2 \le \sqrt{2s/\lambda}, \qquad \|\widetilde{z} - z_g(x)\|_2 \le s + \sqrt{2s/\lambda} \le s + 2\sqrt{s/\lambda}.$$

This bounds the directed distance from the perturbed correspondence to the singleton. The reverse direction is bounded by selecting any perturbed minimizer, so the full Hausdorff distance has the same bound. The score range and query margin imply $m_Q < 1$, while the active value condition gives $\varepsilon \le m_Q/2$. Hence $0 \le \varepsilon < 1$, and substitution gives

$$s + 2\sqrt{s/\lambda} \le \left(C_{\mathrm{int}} + 2\sqrt{C_{\mathrm{int}}/\lambda}\right)\sqrt{\varepsilon}.$$

For the sampling statement, $n \ge n_{\mathrm{loc}}(\delta)$ makes all value and quadratic activations valid with $\varepsilon = e_{n,\delta}$. On the single score event from Proposition C.8, set $\widetilde{g} = \widehat{g}_n$ and use the deterministic result. The conclusion holds on the same event of probability at least $1 - \delta$. Since $e_{n,\delta} = C_e r_{n,\delta}$,

$$C_{\mathrm{loc}}\sqrt{e_{n,\delta}} = C_{\mathrm{loc}}\sqrt{C_e}\left(\frac{\log(C_0 n/\delta)}{n}\right)^{\beta/[2(2\beta+d)]}.$$

□

Completion of the main results. Proposition B.4 proves Theorem 1. Proposition B.8 proves Theorem 2. The value upper and lower bounds are assembled in Appendices C and D. Proposition E.3 and the two-point argument prove Theorem 6; Proposition E.4 proves Theorem 7. The value and location upper bounds share one uniform score event. The menu lower bound instead uses a fixed family with sample-size-dependent local alternatives. Together they establish the cost and optimal-action guarantees for their respective population classes.

F Model Examples and Action Geometry

The theory separates a population score, an observation model, and an action set. This appendix gives explicit checks for familiar scores and shows how action constraints change the geometric constants. Each statement combines score properties with explicit design, domain, query, and action conditions. The derivative calculations also define fitted-model diagnostics for the experiments in Appendix G.

F.1 LOGISTIC SCORES WITH FULL ACTIONS

Let $\sigma(t) = (1 + e^{-t})^{-1}$ and $g(z) = \sigma(w^\top z + b) - 1/2$, where $w \neq 0$. The favorable boundary is $B_g = \{z : w^\top z + b = 0\}$ and its normal is $n_g = w/\|w\|_2$. Under full actions $A_{K_{\text{full}}}(x) = \mathcal{X}$, coverage is immediate, but a feasible normal segment still requires domain slack. In particular, a relative boundary chart must satisfy the tube and slack conditions over the whole action region.

Proposition F.1 (Logistic halfspace membership)**.** *Assume that the design density obeys $0 < p_- \leq p(z) \leq p_+ < \infty$ on $\mathcal{X}$, and that, on the action region, the logistic score satisfies the range condition*

$$\sup_{z \in \mathcal{X}} |w^\top z + b| \leq M_q := \log \frac{1-q}{q}.$$

Suppose the query set has a uniform unfavorable margin and a feasible-depth margin: for some $m_\eta, m_{\text{feas}} > 0$,

$$w^\top x + b \leq -m_\eta \quad (x \in Q), \qquad \sup_{z \in \mathcal{X}} (w^\top z + b) \geq m_{\text{feas}}.$$

Assume also that the full-action query values are bounded away from both zero and the diameter cutoff,

$$\underline{\rho}_Q \leq \rho_{g, K_{\text{full}}}(x) \leq \overline{\rho}_Q \quad (x \in Q).$$

This condition controls the distances from queries to the favorable set, separately from their score margins.

Assume the hyperplane section in the action region is the only relevant zero component, has a normal tube of radius r_0, and every point in that section has normal interior slack at least $s_\partial \geq r_0$. Also assume the action region has no stray near-zero component:

$$\inf_{z \in V_{K_{\text{full}}} \setminus T_g(B_g, [-r_0/2, r_0/2])} |g(z)| \geq \gamma_{\text{dom}} > 0.$$

If the finite Hölder budget L dominates the logistic score's $\Sigma_d(\beta)$ norm on $\mathcal{X}$, then the mechanism belongs to $\Theta_{\text{val}}(\beta; Q)$. More explicitly, one may choose

$$a_{\log} = \|w\|_2 \min_{|t| \leq r_0} \sigma'(\|w\|_2 t) > 0, \qquad b_{\log} = \frac{\|w\|_2}{4},$$
$$\kappa_{\text{cov}} = 1, \qquad \tau = 1, \qquad s_A \leq s_\partial.$$

The query and feasible-depth constants can be taken as

$$m_Q = \tfrac{1}{2} - \sigma(-m_\eta), \qquad \kappa_{\text{feas}} = \sigma(m_{\text{feas}}) - \tfrac{1}{2}.$$

Proof of Proposition F.1. Write $s(z) = w^\top z + b$. Since σ is increasing, $|s(z)| \leq M_q$ implies $\sigma(s(z)) \in [q, 1-q]$, which gives the score-range condition. The score is analytic, so on the compact cube its $\Sigma_d(\beta)$ norm is finite for every fixed finite β. The assumed budget L therefore verifies the smoothness clause, and the density hypothesis verifies observational coverage.

For $c \in B_g$, $s(c) = 0$ and

$$T_g(c, t) = c + t \frac{w}{\|w\|_2}, \qquad s(T_g(c, t)) = \|w\|_2 t.$$

The chain rule gives

$$\nabla g(T_g(c, t))^\top n_g = \sigma'(\|w\|_2 t) \|w\|_2.$$

The logistic derivative is strictly positive on every compact interval and is maximized at zero. Hence the displayed $a_{\log}$ and $b_{\log}$ are valid two-sided crossing constants. The assumed tube and no-stray conditions verify the remaining regular-boundary clauses.

For $x \in Q$, $s(x) \leq -m_\eta$ gives

$$g(x) \leq \sigma(-m_\eta) - \tfrac{1}{2} = -m_Q.$$

The feasible-depth condition follows from the point attaining the displayed supremum, because its score is at least $\sigma(m_{\text{feas}}) - 1/2 = \kappa_{\text{feas}}$. Since $A_{K_{\text{full}}}(x) = \mathcal{X}$, every boundary point in the action region is already a boundary–action intersection. Thus the coverage inequality holds with $\kappa_{\text{cov}} = 1$. The normal interior slack gives $c + [0, s_A] n_g \subset \mathcal{X}$ for $s_A \leq s_\partial$, and $n_g^\top n_g = 1$, so the radial clause holds with $\tau = 1$. All clauses of $\Theta_{\text{val}}(\beta; Q)$ are now verified. □

A global domain check. The axis-aligned score $g(z) = \sigma(\omega(z_d - c)) - 1/2$, with $\omega > 0$ and $0 < c < 1$, has zero section $[0,1]^{d-1} \times \{c\}$ and normal e_d. For $r_0 < \min\{c, 1-c\}$, the entire normal tube remains in the cube, including at lateral chart edges. Outside its half-radius tube, $|g| \geq \sigma(\omega r_0/2) - 1/2 > 0$. Choosing $\omega \max\{c, 1-c\} \leq M_q$, a regular design, and a compact query strip below c with the stated margins verifies these domain conditions globally. An oblique hyperplane passing through a cube corner need not have the same normal slack.

F.2 Quadratic growth from Euclidean projection

For a halfspace, projection geometry supplies the additional condition used for location recovery. The following result keeps that fact separate from the observational and boundary assumptions.

Corollary F.2 (Logistic location class under projection slack)**.** *Under Proposition F.1, suppose further that the orthogonal projection*

$$z_g(x) = x - \frac{w^\top x + b}{\|w\|_2^2} w$$

belongs to $\mathcal{X}$ and to the regular boundary chart for every $x \in Q$. Then $\Pi_{g,K_{\text{full}}}(x) = \{z_g(x)\}$ and $\theta \in \Theta_{\text{loc}}(\beta; Q)$. In fact, for every $z \in C_{g,K_{\text{full}}}(x)$,

$$\|z - x\|_2 - \rho_{g,K_{\text{full}}}(x) \geq \frac{\|z - z_g(x)\|_2^2}{2\overline{\rho}_Q + \operatorname{diam}(\mathcal{X})}.$$

Hence the optimizer-growth condition holds with any $r_* \leq \operatorname{diam}(\mathcal{X})$ and

$$\lambda_{\log} = \frac{1}{2\overline{\rho}_Q + \operatorname{diam}(\mathcal{X})}.$$

Proof of Corollary F.2. Fix $x \in Q$ and put

$$s_x = w^\top x + b < 0, \qquad \rho = -\frac{s_x}{\|w\|_2}, \qquad z_g = x + \rho n_g = x - \frac{w^\top x + b}{\|w\|_2^2} w.$$

By assumption, $z_g \in \mathcal{X}$. For any $z \in C_{g,K_{\text{full}}}(x)$, the favorable halfspace gives $n_g^\top (z - z_g) \geq 0$. With $u = z - z_g$,

$$\|z - x\|_2^2 - \rho^2 = \|\rho n_g + u\|_2^2 - \rho^2 = 2\rho n_g^\top u + \|u\|_2^2 \geq \|u\|_2^2.$$

The inequality is strict unless $u = 0$, so z_g is the unique optimizer. Moreover,

$$\|z - x\|_2 - \rho = \frac{\|z - x\|_2^2 - \rho^2}{\|z - x\|_2 + \rho} \geq \frac{\|u\|_2^2}{\|z - x\|_2 + \rho}.$$

Because $z, z_g \in \mathcal{X}$, $\|z - x\|_2 \leq \|z - z_g\|_2 + \rho \leq \operatorname{diam}(\mathcal{X}) + \rho$. The query value bound gives $\rho \leq \overline{\rho}_Q$, hence

$$\|z - x\|_2 - \rho_{g,K_{\text{full}}}(x) \geq \frac{\|z - z_g(x)\|_2^2}{2\overline{\rho}_Q + \operatorname{diam}(\mathcal{X})}.$$

Since $\|z - z_g(x)\|_2 \leq \operatorname{diam}(\mathcal{X})$, this implies the truncated quadratic-growth condition for every admissible r_* in the stated range. The conclusion follows uniformly over Q. □

F.3 Action constraints: full, box, and chartwise sets

Boundary–action coverage is an error bound for the pair $(B_g, A_K(x))$. Radial crossing instead supplies a feasible direction at an actual intersection. The two checks have different roles and neither is supplied by convexity alone.

Proposition F.3 (Simple sufficient checks)**.** *The boundary–action assumption holds in each of the following settings with the displayed constants chosen uniformly over $x \in Q$.*

1. *If $A_K(x) = \mathcal{X}$, the relevant zero set is at least s_A from $\partial\mathcal{X}$, and B_g has the normal tube in Assumption A.3, then coverage holds with $\kappa_{\text{cov}} = 1$ and radial crossing holds with $v_c = n_g(c)$.*

2. *For an axis-aligned box, suppose every nearby zero component meets the box in the same normal chart. If, at every intersection, one signed coordinate direction $v_c \in \{\pm e_j\}$ has available segment length s_A and $n_g(c)^\top v_c \geq \tau$, then the radial clause holds. Compactness of the finite box-face atlas and a uniform nonzero boundary angle give the coverage bound.*

3. *More generally, a finite family of action charts satisfies the assumption whenever each chart has a Lipschitz error bound* $\operatorname{dist}(b, B_g \cap A) \leq \kappa_{\mathrm{cov}} \operatorname{dist}(b, A)$ *and supplies a feasible segment of length s_A whose angle with n_g is at most* $\arccos(\tau)$.

Proof. In the first case a point of B_g already belongs to the full action, so the coverage distance is zero. The normal segment remains in $\mathcal{X}$ for length s_A, and its normal component is one. In the box case the signed coordinate segment is feasible by the assumed face slack, and its normal alignment satisfies the displayed inequality. The normal chart writes B_g as a Lipschitz graph; the uniform nonzero face angle bounds the graph correction by a constant times the distance to the box. A finite atlas turns the local constants into one uniform constant. The final case restates these two verifiable chart properties and takes their maxima and minima over the finite family. □

The face slack and nearby-component conditions address two obstructions. A zero component that approaches an action set without meeting it violates boundary–action coverage. Also, a vector in the closed tangent cone need not generate a feasible line segment: at a point on a Euclidean sphere, a purely tangential vector belongs to the tangent cone while every positive step along that vector leaves the ball. This is why Lemma B.6 uses the radial clause directly.

Logistic scores under box constraints. Let $\mathcal{V}_{\mathrm{box}}(c, x)$ be the unit coordinate or interior directions whose length-s_A segment remains in $A_K(x)$. The available normal alignment is

$$\tau_{\mathrm{box}} = \inf_{x \in Q} \inf_{c \in B_g \cap A_K(x)} \max_{v \in \mathcal{V}_{\mathrm{box}}(c,x)} \frac{w^\top v}{\|w\|_2}.$$

A positive lower bound, together with the uniform face-slack and nearby-component conditions of Proposition F.3, verifies the action interface. If the normal segment is feasible everywhere, $\tau_{\mathrm{box}} = 1$. A signed coordinate segment has alignment $|w_j|/\|w\|_2$. With finitely many action charts, take the smallest alignment and largest coverage constant. A tangential contact with no positive favorable component cannot supply radial crossing.

F.4 Smooth feature-based scores

Consider $g_\phi(z) = \sigma(w^\top \phi(z) + b) - 1/2$ for a differentiable feature map ϕ. At a boundary point, $\nabla g_\phi(z) = J_\phi(z)^\top w/4$. The gradient norm controls the local crossing slope; its orientation along the normal tube must be controlled separately.

Proposition F.4 (Sufficient conditions for smooth feature-based scores)**.** *Let $\phi : \mathcal{X} \to \mathbb{R}^m$ be $C^{\lceil \beta \rceil}$ on $\mathcal{X}$, and let $g_\phi(z) = \sigma(w^\top \phi(z) + b) - 1/2$. Suppose the design, score-range, query, feasible-depth, and value-range conditions in Proposition F.1 have corresponding analogues for g_ϕ, and assume the global Hölder budget $\|g_\phi\|_{\Sigma_d(\beta)} \leq L$. Assume that the relevant zero set has an injective normal tube of radius r_0, the off-tube separation in Assumption A.3 holds on V_K, and*

$$0 < s_\phi^- \leq \sigma'(w^\top \phi(z) + b) \|J_\phi(z)^\top w\|_2 \leq s_\phi^+ < \infty$$

throughout the tube. In addition, assume the tube orientation margin

$$\inf_{\substack{c \in B_{g_\phi} \cap V_K \\ |t| \leq r_0}} \frac{\nabla g_\phi(T_{g_\phi}(c,t))^\top n_{g_\phi}(c)}{\|\nabla g_\phi(T_{g_\phi}(c,t))\|_2} \geq \kappa_{\mathrm{tilt}} > 0.$$

If the action coverage and radial clauses hold with constants $(\kappa_{\mathrm{cov}}, \tau, s_A)$, then $\theta_\phi \in \Theta_{\mathrm{val}}(\beta; Q)$ with crossing constants

$$a = \kappa_{\mathrm{tilt}} s_\phi^-, \qquad b = s_\phi^+.$$

If, in addition, the optimizer has the quadratic-growth property in Assumption A.5, then $\theta_\phi \in \Theta_{\mathrm{loc}}(\beta; Q)$.

Proof of Proposition F.4. Let $u(z) = w^\top \phi(z) + b$. The chain rule gives

$$\nabla g_\phi(z) = \sigma'(u(z)) J_\phi(z)^\top w.$$

At every regular boundary point, choose $n_{g_\phi}(z) = \nabla g_\phi(z)/\|\nabla g_\phi(z)\|_2$. The norm bound and the orientation inequality in the proposition imply, for every $c \in B_{g_\phi} \cap V_K$ and $|t| \le r_0$,

$$\nabla g_\phi(T_{g_\phi}(c,t))^\top n_{g_\phi}(c) \ge \kappa_{\mathrm{tilt}} \|\nabla g_\phi(T_{g_\phi}(c,t))\|_2 \ge \kappa_{\mathrm{tilt}} s_\phi^-,$$

and the same directional derivative is at most s_ϕ^+. Thus the two-sided crossing clause holds with $a = \kappa_{\mathrm{tilt}} s_\phi^-$ and $b = s_\phi^+$. At $t = 0$ the orientation ratio equals one, but positivity throughout the tube is an additional condition for a nonlinear feature map. The assumed tube, no-stray, design, score, query, and feasible-depth conditions verify the other clauses of the value class, while the stated coverage and radial conditions supply Assumption A.4. The assumed global Hölder budget verifies the smooth-score clause on $\mathcal{X}$. Thus $\theta_\phi \in \Theta_{\mathrm{val}}(\beta; Q)$. If the optimizer-growth assumption also holds, the definition of the location class gives $\theta_\phi \in \Theta_{\mathrm{loc}}(\beta; Q)$. □

Interpreting the conditions. At the zero level, a positive lower bound on $\|J_\phi(z)^\top w\|_2$ prevents a flat boundary. A positive tube-orientation margin prevents the normal derivative from reversing away from that level. A feasible direction with positive alignment verifies the action part. The conditions require these bounds throughout the tube and uniformly over the query set. For $\phi(z) = z$, the quantities reduce to the logistic constants above.

F.5 Neural scores and nonsmooth boundaries

For a scalar-output network of depth L_{net}, write

$$h_0(z) = z, \qquad h_\ell(z) = \rho_\ell(W_\ell h_{\ell-1}(z) + b_\ell), \quad 1 \le \ell < L_{\mathrm{net}},$$
$$f_\omega(z) = W_{L_{\mathrm{net}}} h_{L_{\mathrm{net}}-1}(z) + b_{L_{\mathrm{net}}}, \qquad g_\omega(z) = \sigma(f_\omega(z)) - \tfrac{1}{2}.$$

The subscript on L_{net} distinguishes network depth from the Hölder budget L. With $\tanh$ or softplus activations, bounded weights and compact preactivation ranges give finite derivative bounds. Membership requires the resulting Hölder norm to fit the stipulated budget, as well as the boundary, design, query, and action conditions. Since the zero sets of f_ω and g_ω agree, $\nabla g_\omega = \nabla f_\omega/4$ on the boundary. A useful tube-level certificate is

$$|f_\omega(z)| \le M_\omega, \qquad 0 < \underline{J}_\omega \le \|\nabla f_\omega(z)\|_2 \le \overline{J}_\omega \quad (z \in T_{g_\omega}(B_{g_\omega}, [-r_0, r_0])),$$

$$\inf_{\substack{c \in B_{g_\omega} \cap V_K \\ |t| \le r_0}} \frac{\nabla f_\omega(T_{g_\omega}(c,t))^\top n_{g_\omega}(c)}{\|\nabla f_\omega(T_{g_\omega}(c,t))\|_2} \ge \kappa_\omega > 0.$$

It gives

$$a_\omega = \kappa_\omega \sigma'(M_\omega) \underline{J}_\omega, \qquad b_\omega = \overline{J}_\omega/4.$$

When g_ω is the population score, these bounds and the remaining conditions establish membership in Θ_{val}; quadratic growth further gives Θ_{loc}. When g_ω is fitted from data, the same calculations describe its deterministic geometry. The experiments use this latter interpretation to study refit-induced changes.

ReLU and tree-based scores. Let $\mathcal{H}_\omega$ be the union of hidden-layer activation kink sets. A ReLU network admits the preceding chartwise calculation when

$$\mathrm{dist}(T_{g_\omega}(B_{g_\omega} \cap V_K, [-r_0, r_0]), \mathcal{H}_\omega) > 0.$$

The activation pattern is then fixed on each connected tube component and the logit is affine there. If a kink meets the relevant tube, its normal can jump, and the present smooth-boundary certificate does not apply across that stratum. Piecewise-constant tree and random-forest scores require a different boundary model. Extending the analysis to these cases calls for conditions adapted to nonsmooth or discontinuous scores.

From certificates to diagnostics. A fitted-model analysis can trace the zero contour, measure gradient norms and normal alignment, and search for feasible rays and excess-cost growth. The measurements in Appendix G address these quantities. They supply finite-sample diagnostics of a fitted score; proving population class membership still requires the full design, domain, and uniform conditions stated above.

G Experimental Protocols and Additional Results

The experiments study action-dependent cost error, Euclidean recovery, and geometric diagnostics of refitted models. Additional experiments isolate three distinct functionals: boundary height at fixed tangent coordinates, projection onto a selected boundary graph, and contact with a prescribed action face. Table G.1 maps the main results to the complete settings and recorded outputs. Each protocol specifies its population, estimator, optimization target, query aggregation, and uncertainty summary.

Table G.1: Experiment-to-result map. Numerical recovery is separated from boundary diagnostics and fitted-model associations.

Target	Main evidence	Full protocol
Action-dependent cost	Figure 2(a,b)	Appendix G.1; Figure G.1
Euclidean cost and location	Figure 2(c,d)	Appendix G.2; Table G.3
Boundary diagnostics	Supplementary	Appendices G.3–G.5
Refit displacement	Figure 3	Appendix G.6; Table G.4

Recovery losses and aggregation. For a finite grid Q_0, Euclidean cost and location errors are

$$E_{\rho,\max} = \max_{x \in Q_0} |\widehat{\rho}(x) - \rho(x)|, \qquad E_{\Pi,\max} = \max_{x \in Q_0} d_H(\widehat{\Pi}(x), \{z_g(x)\}).$$

The second expression applies to populations with unique optima. Each fit first aggregates its query errors; means and intervals then aggregate independent fits. Query means are also recorded. The fixed query grid is part of the numerical evaluation; the theoretical losses in (3.2) range over the full compact query set. Figure 2 uses 10,000 bootstrap resamples of fits and percentile 95% intervals. For amplification, full and binding-box errors are resampled together and the ratio of their means is recomputed. The supplementary recovery plots use 2,000 fit resamples. Replicate quantile bands and public-data intervals have the separate conventions stated below.

Reproducibility. The supplementary code includes numerical routines, configurations, fixed seeds, per-fit and per-query records, and fitted score grids. The entry point `reproduce.py` separates training, saved-record analysis, and publication-figure generation. Source and data checksums, numerical solver tests, and an environment specification accompany the release. Figure summaries can be regenerated without retraining or downloading datasets.

G.1 Changing actions at a fixed fitted score

Question and controlled comparison. Can the same score estimate produce different minimum-cost errors under different action constraints? We change only the action template within each fit, retaining the fitted score, the query grid, and the score-error normalization. This separates action-dependent error transfer from changes in the predictor itself.

Sampling and fitting. Let $X \sim \mathrm{Unif}([0,1]^2)$ and

$$\Pr(Y = 1 \mid X = z) = \sigma(12(z_1 + \alpha z_2 - 0.66)), \qquad \alpha \in \{0.12, 0.25, 0.5\}.$$

For each value of α, $n \in \{512, 1024, 2048, 4096, 8192, 16384, 32768\}$, and 48 independent replicates, we fit the intercept and both logistic coefficients by unpenalized maximum likelihood. The trust-region Newton solver uses analytic gradients and Hessians, zero initialization, gradient tolerance 10^{-9}, and at most 150 iterations. All 1,008 fits converge to gradient infinity norm below 10^{-6}. Each fit is reused under all four action templates without refitting or recalibration.

Queries and actions. The common 9×9 query grid is

$$x = (0.66 - \alpha v - 0.05 - \alpha s, v),$$
$$v \in \mathrm{linspace}(0.32, 0.44, 9), \qquad s \in \mathrm{linspace}(0.265, 0.335, 9).$$

The template is $K = [-r_1, r_1] \times [-r_2, r_2]$, with admissible destinations clipped to the unit square. We use $(r_1, r_2) = (1, 1)$ for full actions, $(0.28, 0.45)$ for the loose box, $(0.05, 0.45)$ for the binding box, and $(0.05 - 0.08\alpha, 0.45)$ for the tight box. All population problems are feasible. The loose box leaves the population projection unchanged; the binding and tight boxes place every population optimum on the upper z_1 face. The available favorable direction on this face has alignment $\tau = \alpha/\sqrt{1+\alpha^2}$. The planar boundary and exact projection solver isolate the role of the action constraint in finite-sample error transfer.

Optimization and feasibility. For a fitted halfspace $w^\top z + b \geq 0$ and box $\ell \leq z \leq u$, the Euclidean projection takes the form $z(\lambda) = \mathrm{clip}(x + \lambda w, \ell, u)$. We test feasibility at the score-maximizing corner, bracket the monotone equation $w^\top z(\lambda) + b = 0$ when feasible, and bisect 65 times. Independent enumeration of the free projection, four face intersections, and corners checks the reported values. An empty estimated feasible set is recorded separately from numerical solver failure and receives $\widehat{\rho} = \sqrt{2}$ in finite cost-loss summaries. Errors conditional on feasibility and binding indicators are recorded separately.

Aggregation and grid checks. For each fit, we record both cost summaries over 81 queries; the reported means and 2,000-resample bootstrap 95% intervals then aggregate over fits. Score normalization uses the maximum error on the same fixed 129×129 grid over $[0, 1]^2$ for every action template. At $n = 32768$, refining the query grid to 17×17 leaves full, loose, and binding maximum errors unchanged in every replicate. Tight-box maxima change by at most 0.00436 near the feasibility transition. Refining the score grid to 257×257 changes its maximum by at most 10^{-5}. The normalization remains an empirical measure of score error on that fixed evaluation grid.

α	Action	Mean error	Max. error	Binding (%)	Infeasible (%)
0.12	Full	0.0016	0.0019	0.0	0.00
0.12	Binding box	0.0135	0.0169	100.0	0.00
0.12	Tight box	0.0258	0.0789	98.9	1.08
0.25	Full	0.0013	0.0015	0.0	0.00
0.25	Binding box	0.0061	0.0075	100.0	0.00
0.25	Tight box	0.0068	0.0083	100.0	0.00
0.5	Full	0.0014	0.0016	0.0	0.00
0.5	Binding box	0.0032	0.0041	100.0	0.00
0.5	Tight box	0.0035	0.0044	100.0	0.00

Table G.2: Paired action results at $n = 32768$, averaged over 48 fits. Mean error averages over queries before fits; maximum error maximizes over queries before averaging over fits. Binding and infeasibility concern estimated problems, with all queries in the denominator. Both constrained templates bind at every population query; full actions do not. Loose-box results equal full-action results throughout the sweep.

Result. Full and loose actions coincide, providing a comparison in which a formally restricted action set leaves the optimum unchanged. In contrast, at $n = 32768$ the binding box has larger mean maximum error at every tested alignment, although all its estimates remain feasible and binding. For $\alpha = 0.12$, the errors are 0.00191 and 0.01693 for full and binding actions. The tight box adds a separate feasibility effect: 1.08% of its queries remain infeasible at this sample size. At $n = 32768$, the ratio of mean binding-box to full-action maximum error is 8.88, 4.94, and 2.51 at the three alignments. The paired-bootstrap 95% intervals are $[7.19, 10.86]$, $[4.19, 5.89]$, and $[2.11, 2.97]$. Binding-box error exceeds full-action error in $48/48$, $48/48$, and $46/48$ fits, respectively. The ratio uses means of paired errors, rather than a mean of per-fit ratios; resampling preserves the score shared by each pair.

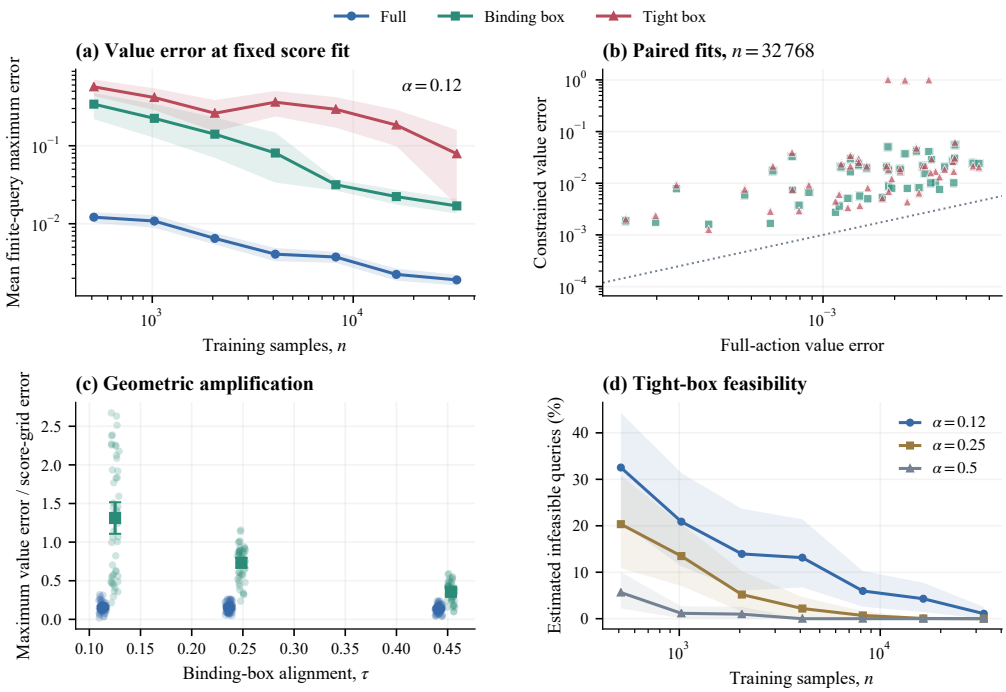


Figure G.1: **Full paired-action experiment.** (a) Finite-query maximum cost error averaged over 48 fits at $\alpha = 0.12$. (b) Paired errors at $n = 32{,}768$; the dotted line denotes equality. (c) Score-normalized maximum errors at three binding-box alignments. (d) Tight-box infeasibility across sample sizes and alignments. Bands and error bars are bootstrap 95% intervals over fits. Infeasible estimates receive $\widehat{\rho} = \sqrt{2}$ in (a,b); all instances in (c) are feasible. Figure 2(a,b) displays the full–binding paired comparisons and the ratios of their mean errors.

G.2 Euclidean recovery with per-query records

Population and queries. We study global Euclidean recourse for the smooth populations

$$b_\kappa(u) = 0.58 + \kappa(u - 0.5)^2, \quad g_\kappa(u, v) = 0.42 \tanh\{6(v - b_\kappa(u))\}, \quad \kappa \in \{0.8, 1.6\}.$$

Inputs are uniform on $[0, 1]^2$ and labels are Bernoulli with probability $1/2 + g_\kappa$. We use full actions and the seven fixed queries $x = (u, b_\kappa(u) - 0.09)$ with $u \in \mathrm{linspace}(0.34, 0.66, 7)$. All queries lie below the minimum height of the favorable epigraph. For $n \in \{512, 1024, 2048, 4096, 8192, 16384\}$, 30 independent fits for each curvature produce 360 fits and 2,520 query records. The seed for curvature index i, sample-size index j, and replicate r is $20260917 + 10^7 i + 10^5 j + r$, with all indices starting at zero.

Fitted numerical score. At each node of a regular 33×33 square grid, we fit a local-linear regression to binary labels with bandwidth $h = 0.52(\log(80n)/n)^{1/6}$ and product triangular weights $w_i(z) = \prod_{k=1}^{2}(1 - |X_{ik} - z_k|/h)_+$. For features $(1, (X_{i2} - z_2)/h, (X_{i1} - z_1)/h)$, let G denote the weighted Gram matrix. The numerical solve uses $G + 10^{-8} \max\{1, \max_j G_{jj}\} I$. The intercept is clipped to $[0.05, 0.95]$ and centered at $1/2$. All nodes are evaluated, including nodes with low effective mass. We record the fraction with weighted mass below $\max\{8, 0.002n\}$. An empty kernel neighborhood or failed linear solve uses intercept $1/2$ and is counted. Neither event occurs in the 360 recorded fits. A sup-norm KD-tree neighborhood search computes the same weighted fits without evaluating zero-weight observations. Each square is split along its lower-left to upper-right diagonal, and node values define a continuous piecewise-affine score on the whole square. The grid and triangulation specify the numerical score estimator. The same-sample refinement below assesses its spatial resolution.

Exact optimization for the interpolant. If a query is already favorable, its cost is zero. Otherwise, any nearest favorable point of a continuous score on a convex domain lies on its zero set: a strictly positive point could be approached along the segment from the negative query until the first zero crossing, reducing distance. The zero set of each triangle is a segment, a point, or an entire zero triangle. We enumerate all such segments, including zero edges, and project the query onto each segment; the smallest distance is therefore globally optimal for the interpolant. Optimization covers every zero segment in the square. Numerical ties within 10^{-12} in squared distance are retained. Against the unique population optimum, the maximum distance of these tied candidates is the reported numerical Hausdorff loss. Empty estimated favorable sets use $\widehat{\rho} = \sqrt{2}$ and $\widehat{\Pi} = \{x\}$, with feasibility recorded separately; all 360 fits are feasible and optimizer score residuals were below 10^{-12}.

Independent population reference. For a query (x_1, x_2), put $t = u - 1/2$. The derivative of squared distance to $(u, b_\kappa(u))$ vanishes at

$$2\kappa^2 t^3 + \{1 + 2\kappa(0.58 - x_2)\}t + (1/2 - x_1) = 0.$$

We evaluate every real root in $[-1/2, 1/2]$ and both endpoints, retaining points inside the square. This yields the global reference, independently of the fitted grid. The favorable epigraph is convex and the Euclidean projection is unique. Unit tests compare this calculation with a separate scalar minimizer and compare the interpolant solver with analytic halfspace/box projections, including degenerate and tied cases.

Results and grid sensitivity. Figure 2(c,d) uses the maximum over seven queries within each fit, then the mean and 10,000-resample bootstrap 95% interval across fits. From $n = 512$ to 16384, mean maximum cost error changes from 0.03199 to 0.01053 for $\kappa = 0.8$ and from 0.03587 to 0.01245 for $\kappa = 1.6$. The corresponding location errors change from 0.03398 to 0.01291 and from 0.03780 to 0.01335. At the largest sample size, we refit a 65×65 score grid on each of the same 60 samples. Separately, we evaluate 25 queries using the coarse grid. Table G.3 reports both checks. The refinement changes mean maximum cost and location errors by less than 0.6% and 1.4%, respectively. At an individual query, the largest displacement between the two interpolant optima is 0.00673. Increasing the query grid to 25 points raises mean maximum cost and location errors by at most 0.9% and 5.3%. These two checks quantify score and query resolution separately.

Table G.3: Euclidean recovery and resolution checks at $n = 16384$. Values are means over 30 fits per curvature. "Dense query" uses 25 queries and the 33^2 score grid. Last two rows are the largest query-level changes over the 30 fits when the score grid changes from 33^2 to 65^2.

Statistic	$\kappa = 0.8$	$\kappa = 1.6$
Maximum cost error, 33^2 grid	0.01053	0.01245
Maximum cost error, 65^2 grid	0.01053	0.01238
Maximum location error, 33^2 grid	0.01291	0.01335
Maximum location error, 65^2 grid	0.01279	0.01316
Dense-query maximum cost error	0.01060	0.01256
Dense-query maximum location error	0.01340	0.01405
Largest absolute cost change	0.00076	0.00110
Largest optimizer displacement	0.00609	0.00673

The numerical implementation is `experiments/run_euclidean_recovery.py`; its configuration, run manifest, per-query and per-fit CSVs, and all fitted score grids are included. The records permit both querywise inspection and fit-level aggregation without further training.

G.3 Boundary-height tracking across dimensions

Target and implementation. We estimate vertical boundary height at sixteen fixed tangent coordinates. The implementation is `experiments/run_boundary_height.py`. Its fields `sup_value_error` and `mean_value_error` record maximum and mean absolute height errors, respectively. Figure G.2 summarizes both functionals across independent fits.

Sampling and estimator. For $z = (u, v)$ in $[0, 1]^d$, the centered score is $g(z) = 0.45(v - b_\beta(u))$, where

$$b_2(u) = 0.55 + 0.16(\bar{u} - 0.5), \qquad b_3(u) = 0.55 + 0.075\sin(2\pi\bar{u}), \quad \bar{u} = \tfrac{1}{d-1}\sum_j u_j.$$

Uniform inputs and Bernoulli labels use probability $1/2+g$; the code's $[0.05, 0.95]$ clipping is inactive for these populations. Settings are $(d, \beta) = (2, 2), (2, 3), (3, 2), (4, 2)$, with nine doubling sample sizes from 256 to 65,536 and 48 paired replicates. The labels β index the smooth examples and the corresponding estimation procedures. Sixteen tangent queries use $u = \mathrm{linspace}(0.3, 0.7, 16)$ in two dimensions. Higher-dimensional tangent grids stack this sequence with circular shifts by two and four positions. A notional vertical query height is 0.22. The estimator uses product triangular weights, polynomial degree $\beta - 1$, bandwidth $0.45(\log(80n)/n)^{1/(2\beta+d)}$, and ridge $10^{-8}\max\{1, \max_j G_{jj}\}I$; a local-constant fit is the baseline. Intercepts are clipped to $[0.05, 0.95]$. A vertical node is usable when its weighted mass exceeds $\max\{8, 0.20nh^d\}$.

Root selection and loss. For each tangent query, the fitted score is evaluated on 61 vertical nodes from 0.05 to 0.90. Among negative-to-positive brackets, the code selects the one whose lower node is closest to the fixed height 0.55. If there is no bracket, it selects the usable favorable node closest to that anchor. It then uses the selected node and its predecessor for clipped linear interpolation when possible, and otherwise returns the selected node; if no usable favorable node exists, it returns 0.90. The saved `fallback_rate` counts only this final default output. The non-default branch includes both bracket interpolation and favorable-node selection. The estimated and reference quantities are $\widehat{b}(u) - 0.22$ and $b_\beta(u) - 0.22$, so their absolute difference is exactly boundary-height error. Tangent coordinates are fixed in this functional, whereas Euclidean recourse optimizes them as well. The experiment therefore measures the intermediate boundary-estimation stage of the statistical formulation.

The release records maximum-error summaries and three-point slopes for each retained window. Figure G.2 shows their distribution across fits alongside the corresponding effective-mass and default-output diagnostics.

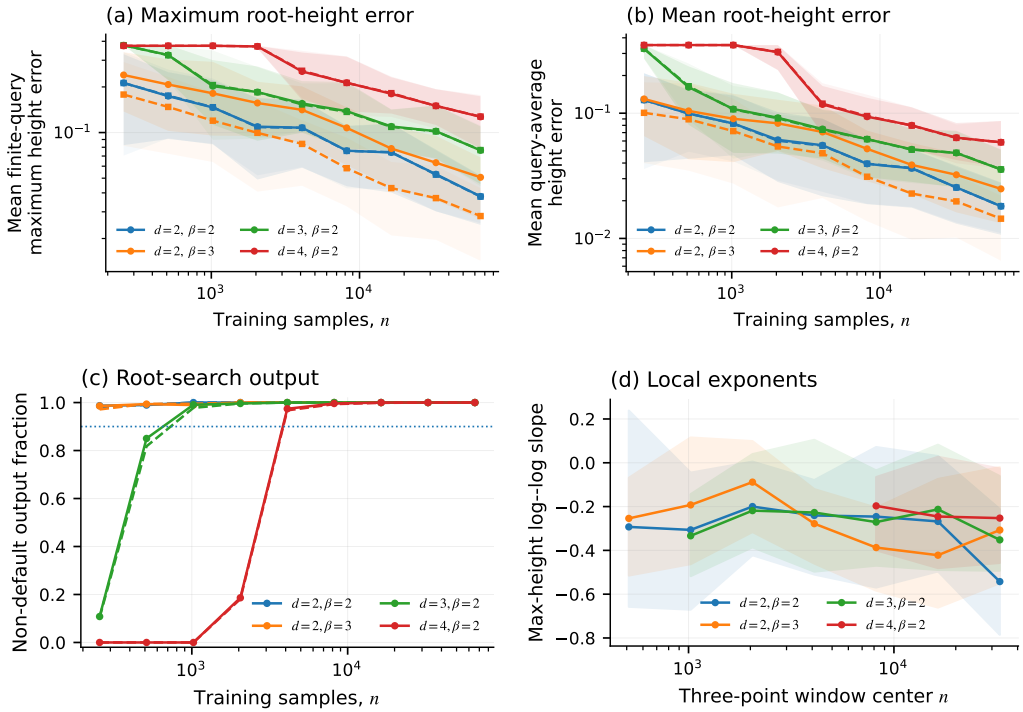


Figure G.2: **Boundary-height estimation across dimension and smoothness.** (a,b) Maximum and mean height error over 16 queries, then averaged over 48 fits; bands are 10–90% replicate quantiles. Solid curves use the local-polynomial fit and dashed curves the local-constant baseline. (c) Fraction of non-default root outputs. (d) Three-point log–log slopes of the maximum height error, with medians and interquartile ranges across eligible fits; every point in a retained window has at least 90% non-default outputs. The four dimension–smoothness settings use the same evaluation protocol.

G.4 Selected graph-projection diagnostic

Population and numerical graph. The graph-projection experiment uses g_κ and the seven queries in Appendix G.2, with six sample sizes and 30 fits per setting. Local-linear regression is evaluated at 19 tangent nodes in $[0.28, 0.72]$ and 49 vertical nodes in $[0.52, 0.68]$. The bandwidth is $0.52(\log(80n)/n)^{1/6}$; triangular weights, ridge regularization, and intercept clipping follow the local fit above. Nodes with mass below $\max\{8, 0.002n\}$ are excluded from root selection. Each column selects the crossing closest to height 0.58; when no crossing is available, it selects the usable node of smallest absolute score, or returns 0.58 when the column has no usable node.

Projection and recorded losses. The chosen roots define a piecewise-linear graph. The code evaluates this graph on 3,001 tangent coordinates and chooses the closest sampled graph point. The reference uses the same graph-and-grid procedure applied to population heights. Therefore the comparison concerns a selected polygonal graph and the corresponding numerically discretized population graph. The saved fields `location_sup` and `value_sup` take maxima over the seven queries, while fields ending in `_mean` take averages. The figure reports the maximum-error functionals.

This calculation isolates the effect of reconstructing a single boundary graph. The Euclidean experiment in Appendix G.2 optimizes over the entire favorable set of its triangulated score.

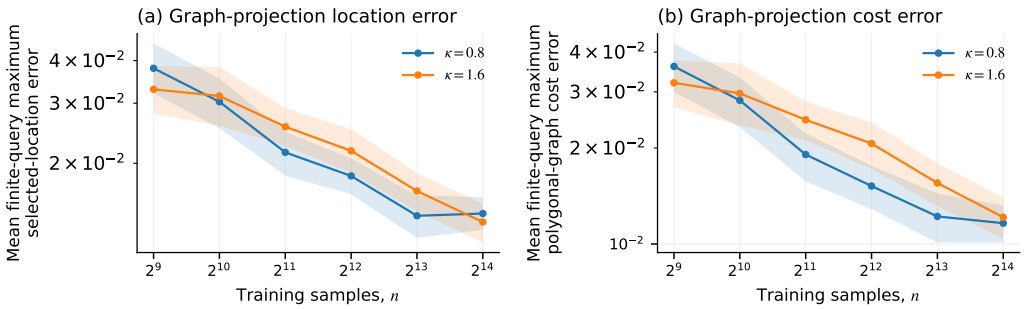


Figure G.3: **Recovery of selected graph projections.** (a) Maximum selected-location error and (b) maximum graph-projection cost error. Curves are means and bootstrap 95% intervals over 30 fits for each curvature and sample size. The underlying graph reconstruction and numerical reference define the target of this graph-based experiment.

G.5 Boundary-face contact under score perturbations

Sampling and fitting. The implementation in `experiments/run_face_contact.py` samples uniformly on $[-1, 1]^2$ and uses signed scores

$$s_{\text{linear}}(u, v) = 0.4(v - 0.75u), \qquad s_{\text{tangent}}(u, v) = 0.4(v - 0.9u^2),$$
$$\Pr(Y = 1 \mid u, v) = \tfrac{1}{2}(1 + s(u, v)).$$

Thus $s = 2g$ in the manuscript's centered-score convention. Probability clipping to $[0.02, 0.98]$ is inactive. For six doubling sample sizes from 128 to 4,096 and 60 replicates, ordinary least squares fits response $2Y - 1$ using basis $(1, u, v, u^2)$. The experiment comprises 720 independent fits.

Boundary-face contact. The code solves $\widehat{s}(u, 0) = 0$ and retains real roots with $|u| \le 0.75$. A *boundary-face contact* is recorded when at least one root is retained; the correction magnitude is the maximum absolute retained root. Score error is evaluated on an 81×81 grid; face error uses 401 points in $[-0.75, 0.75]$. The score plot reports replicate medians and 5–95% quantiles. Contact proportions use the clipped Wald band $\widehat{p} \pm 1.96\sqrt{\widehat{p}(1 - \widehat{p})/60}$.

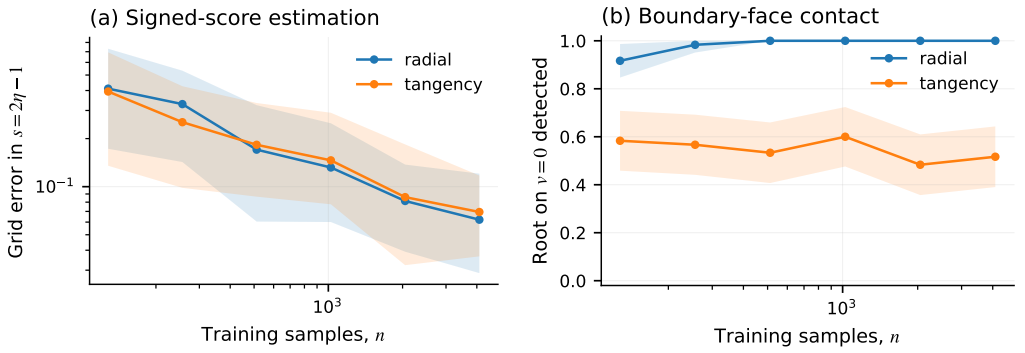


Figure G.4: **Boundary-face contact under score estimation.** (a) Signed-score error $\|\widehat{s} - s\|$ on the evaluation grid. (b) Fraction of fits with a real root on the face $v = 0$ in the search interval. Shading in (a) represents 5–95% replicate quantiles; the proportion bands in (b) use the stated Wald formula.

Geometric interpretation. For $A = \{v \ge 0\}$, the tangent population has favorable normal $(0, 1)$ at the origin, an available inward direction. An upward boundary shift can remove contact with $v = 0$ while retaining favorable feasible points above it. The plotted event is consequently boundary-face

contact. Full feasible-set instability is the different phenomenon established by the construction in Proposition 3.

G.6 GEOMETRIC DIAGNOSTICS ON PUBLIC DATASETS

Question and evaluation target. When a fitted model changes after refitting, do geometric quantities help explain the movement of its recourse optimizer beyond score shift alone? The response is refit-induced action displacement, measured between fitted models. We compare its association with score shift and geometric diagnostics within each dataset–model–action configuration.

Data, fitting, and queries. We use Banknote Authentication (1,372 observations, four continuous features) and HTRU2 (17,898 observations, eight continuous features) (Lohweg, 2012; Lyon, 2015). Each of three stratified splits allocates 60%/20%/20% to training, calibration, and testing. A training-fitted quantile map places retained features in $[0, 1]^d$; clipped test points are excluded from query selection. We tune logistic regression and a two-layer $\tanh$ MLP on the calibration fold, refit on training plus calibration data, and select 32 shared unfavorable test queries with nonzero full-action recourse. Calibration selects a bottleneck box coordinate with at least 70% feasible recourse and active-face frequency closest to 50%. Ten paired bootstrap refits perturb the fitted scores while holding the queries and actions fixed.

Geometric measurements and optimization. For each fit, we record boundary crossing $\widehat{a}_0$, tube crossing, action alignment $\widehat{\tau}$, a finite feasible-ray certificate, growth diagnostic $\widehat{\lambda}$, optimizer ambiguity, and refit displacement. Logistic box recourse is solved exactly. The MLP uses analytic gradients and multistart constrained optimization; its reported actions are numerical local solutions.

Across eight configurations and 7,680 action-level records, feasibility of the base and refitted recourse problems is 100%, with no detected ambiguity or solver failures. The lower deciles satisfy $\widehat{a}_0 \geq 6.83$ and $\widehat{\lambda} \geq 1.91$. Box active-face frequencies reach 0.49 on Banknote and 0.46 for HTRU2 logistic regression. The latter has median alignment 0.95 but lower decile 0.17; other box configurations have lower deciles 0.76–0.95. Its median optimizer displacement increases from 0.016 under full actions to 0.024 under the box.

Comparison and record eligibility. We compare the index

$$I = \sqrt{\widehat{e}_{\mathrm{loc}}/(\widehat{a}\widehat{\tau}\widehat{\lambda})}$$

with score shift alone and three component variants on the same records. Here $\widehat{e}_{\mathrm{loc}}$ measures the maximum difference between base and refitted scores on fixed probes, and $\widehat{\lambda}$ is a finite-search growth diagnostic. The index uses the factors appearing in the location bound as a descriptive measure of fitted-model sensitivity. Eligibility requires a detected refit, no detected base/refit ambiguity, and finite positive index and displacement. This retains 7,600 of 7,680 records. No coefficients are fitted to optimizer displacement.

Dependence-aware uncertainty. Queries within a split share ten refitted models, and each query occurs under every refit. In each of 2,000 resamples (seed 20260911), we independently resample query IDs and refit IDs within a split, take their Cartesian product, apply the same eligibility rule, and recompute both correlations on the same records. Percentile intervals are conditional on the three recorded splits and quantify sensitivity to these two axes. The primary comparison is stratified by dataset, model, and action; pooling would combine different score and geometry scales.

Results and component attribution. For HTRU2 logistic–box, adding action alignment to slope normalization raises correlation from 0.61 to 0.89, and the full index reaches 0.92. The paired full-index improvement has conditional resampling interval $[0.22, 0.41]$; split-specific improvements are $0.28, 0.31, 0.30$. Among 430 active-face records, score-only and full-index correlations are 0.71 and 0.83; among 500 inactive-face records, they are 0.77 and 0.82. The larger combined gain therefore partly reflects differences between action regimes. Conversely, Banknote MLP correlations decrease from 0.92 to 0.83 under both action sets, with negative differences in every split. HTRU2 MLP–box has only six eligible active query–split pairs, so that subgroup is descriptive. The full comparison

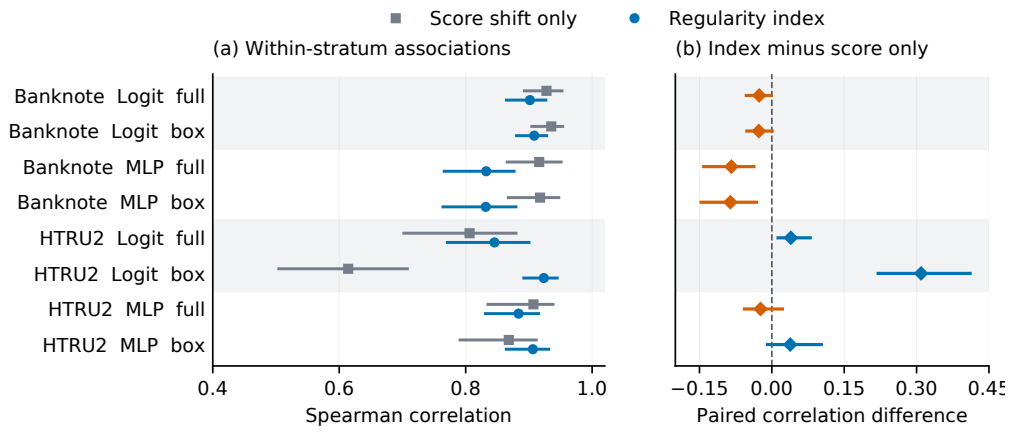


Figure G.5: **Geometric diagnostics across all eight configurations.** (a) Within-configuration Spearman correlations with optimizer displacement. (b) Paired correlation differences; positive values favor the geometric index. Whiskers are 95% percentile intervals from 2,000 crossed query/refit resamples within three fixed splits, conditional on the recorded procedures. All 7,600 eligible records are used.

Table G.4: Within-configuration Spearman correlations with optimizer displacement. $e = \widehat{e}_{\mathrm{loc}}$, $a = \widehat{a}$, $\tau = \widehat{\tau}$, $\lambda = \widehat{\lambda}$, and $I = \sqrt{e/(a\tau\lambda)}$. N counts eligible query–refit records. The component comparisons use the same records as the complete index and the score-only baseline.

Dataset	Model	Action	N	e	$\sqrt{e/a}$	$\sqrt{e/(a\tau)}$	$\sqrt{e/(a\lambda)}$	I
Banknote	Logistic	box	930	0.94	0.94	0.95	0.93	0.91
Banknote	Logistic	full	960	0.93	0.93	0.93	0.90	0.90
Banknote	MLP	box	950	0.92	0.90	0.92	0.86	0.83
Banknote	MLP	full	960	0.92	0.90	0.90	0.83	0.83
HTRU2	Logistic	box	930	0.61	0.61	0.89	0.80	0.92
HTRU2	Logistic	full	960	0.81	0.81	0.81	0.85	0.85
HTRU2	MLP	box	950	0.87	0.87	0.91	0.90	0.91
HTRU2	MLP	full	960	0.91	0.90	0.90	0.88	0.88

identifies the configurations in which geometry adds information and those in which score shift already captures the observed ranking.

Distributional context. Table G.5 records feasibility-related geometry and median displacement. Figure G.6 shows the distributions behind these summaries. Pooled correlations are 0.88 for the index and 0.85 for score shift; the within-configuration comparisons provide the primary evaluation.

Table G.5: Recorded geometry of the fitted model–action pairs. $|Q|$ counts queries across three splits; *Active* is the box-face frequency. $a_{.1}$, $\tau_{.1}$, and $\lambda_{.1}$ are lower deciles; *Ray* is the finite feasible-ray rate. The last column reports median paired-refit displacement with a hierarchical 95% interval. All base and refitted problems in these records are feasible.

Data	Model	Action	$\lvert Q\rvert$	Active	$a_{.1}$	$\tau_{.1}$	$\lambda_{.1}$	Ray	median Δ_z [95% CI]
Banknote	Logistic	Full	96	–	7.40	1.00	2.85	1.00	0.0049 [0.0044–0.0054]
Banknote	Logistic	Box	96	0.49	7.40	0.80	2.77	0.82	0.0050 [0.0046–0.0054]
Banknote	Smooth MLP	Full	96	–	7.82	1.00	2.95	1.00	0.0174 [0.0132–0.0237]
Banknote	Smooth MLP	Box	96	0.49	7.82	0.76	2.86	0.89	0.0179 [0.0128–0.0245]
HTRU2	Logistic	Full	96	–	6.83	0.94	2.32	0.91	0.0160 [0.0123–0.0197]
HTRU2	Logistic	Box	96	0.46	6.83	0.17	1.91	0.72	0.0242 [0.0185–0.0404]
HTRU2	Smooth MLP	Full	96	–	11.73	0.97	3.15	0.80	0.0243 [0.0179–0.0599]
HTRU2	Smooth MLP	Box	96	0.07	11.73	0.95	3.18	0.75	0.0262 [0.0179–0.0623]

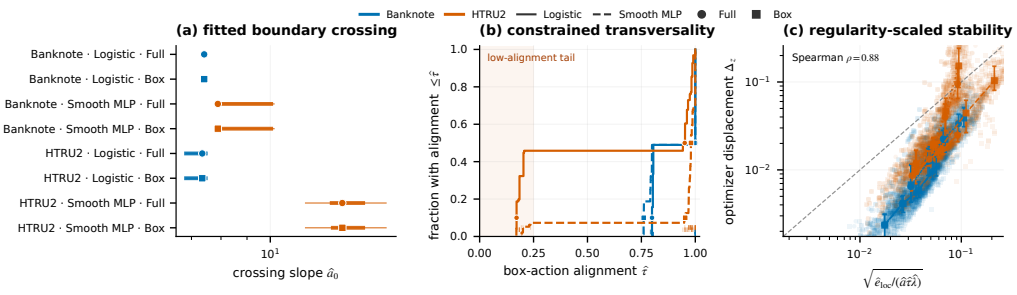


Figure G.6: Distributional context for Figure G.5. (a) Crossing-slope medians, interquartile intervals, and 10–90% ranges. (b) Box-action alignment distributions. (c) Raw paired-refit records and within-configuration quantile summaries. The diagonal marks equal numerical scales. Panel (c) summarizes pooled association; Figure G.5 compares predictors within each configuration.

G.7 Exact constructions and the scope of empirical evidence

The equal-law threshold-point witness, localized bump alternatives, and two-ball construction in Appendices B.1, D, and E.2 establish the identification, minimax, and exact-menu statements. The experiments complement these results by measuring specific estimators and isolating observable geometric mechanisms. Theorem 6 concerns all estimators on a fixed class; its proof is the two-point argument in Appendix E.3.

H Positioning and Scope

The results combine statistical estimation with stability of a constrained decision. This appendix explains which parts come from established tools, what must be verified for recourse, and how the choice of recovery target affects the conclusions.

H.1 From an estimated score to a constrained decision

The sampling exponent in Theorem 4 is the uniform nonparametric score-estimation exponent. Level-set and discrimination theory provide the relevant statistical background (Tsybakov, 1997; Rigollet & Vert, 2009; Singh et al., 2009; Mammen & Tsybakov, 1999). The recourse problem introduces the query-indexed intersection $A_K(x) \cap \{g \geq 0\}$. One score event must control all of these intersections with common constants. Theorem 2 verifies a sufficient geometric route, while Proposition 3 shows why boundary convergence alone cannot replace it.

Statistical sharpness requires an additional argument for the cost functional. The full-action and binding-box constructions separate the cost functions at matched queries while keeping their sample laws difficult to distinguish. A testing reduction then yields the lower bounds. In the box construction, the alternatives also preserve a positive face multiplier and common class constants, so the constraint remains part of the statistical estimation problem.

H.2 Comparison with implicit multifunction stability

Fix x and abbreviate $A = A_K(x)$. With a continuous score s as parameter, write

$$\mathcal{F}_x(s,z) = \begin{cases} -s(z) + [0,\infty), & z \in A, \\ \varnothing, & z \notin A, \end{cases} \qquad G_x(s) = \{z : 0 \in \mathcal{F}_x(s,z)\} = C_{s,K}(x).$$

For $z \in A$, the residual is $\operatorname{dist}(0, \mathcal{F}_x(s,z)) = [-s(z)]_+$. If $z \in G_x(s')$, it is at most $\|s - s'\|_\infty$. Thus parameter variation is controlled immediately; the substantive condition is an error bound linking residuals to feasible-set distance.

Proposition 5.11 of Cuong & Kruger (2021) combines parameter control with uniform subregularity to obtain local pairwise stability. In this encoding, the remaining premise takes the form

$$\alpha \operatorname{dist}(z, G_x(s)) \leq [-s(z)]_+$$

for every score s in a neighborhood of g and every nearby action point with sufficiently small residual, with one $\alpha > 0$. Its conclusion compares $G_x(s)$ and $G_x(s')$ locally around a reference feasible point.

The recourse theorem instead compares each perturbation with a fixed population score. Regular geometry is required of g; $\widetilde{g}$ is any continuous score within the prescribed uniform error. Lemma B.6 constructs corrections between three levels of the fixed population score:

$$A_K(x) \cap \{g \geq -\varepsilon\}, \qquad C_{g,K}(x), \qquad A_K(x) \cap \{g \geq \varepsilon\}.$$

The signed-level sandwich then controls the entire perturbed feasible set, uniformly over Q. The corrections use levels of g, in place of a residual error bound shared by all neighboring continuous scores.

Table H.1: Assumptions and conclusions of two stability formulations.

Aspect	Local implicit-map formulation	Recourse transfer theorem
Scores	Two nearby parameters s, s'	Fixed regular g and continuous $\widetilde{g}$
Error control	Residual-to-solution bound for every nearby s	Corrections between levels of g
Spatial scope	Neighborhood of a reference feasible point	Complete feasible sets, uniformly over Q
Uniformity	Shared neighborhoods and moduli must be verified	Coverage, slope, feasible rays, and a common error threshold

A threshold parameter with fixed g also connects the directed corrections to classical error bounds. Theorem 2 makes the required constants, feasible segments, and query-uniform domains explicit for recourse. For convex intersections, bounded linear regularity is related to intersection qualifications such as relative-interior conditions (Bauschke et al., 1999, Theorem 3 and Corollary 3). Our favorable regions may be nonconvex. The radial condition supplies a positive-length feasible segment, which is stronger than membership in a closed tangent cone.

H.3 What exact-menu recovery asks of the data

The symmetric Hausdorff loss requires recovery of every exactly optimal population choice. Proposition 5 interprets this requirement through the best attainable secondary utility when preferences are specified after primary-cost minimization. A single selected action, a directed distance, or excess cost uses a different target and can permit a tie to collapse to one branch.

Theorem 6 establishes a class-uniform statistical obstruction for the complete exact menu within a fixed regular value class. The difficult alternatives approach a tie as n increases; their parameter class and its constants remain fixed. This is a worst-case statement over the class, distinct from pointwise recovery at a particular score. Theorem 7 gives a positive result on the smaller class with uniqueness and uniform quadratic growth. Other sufficient structures for recovering a set of multiple optima remain to be characterized.

H.4 Scope of the statistical model and empirical evidence

The estimable classes use continuous features, Euclidean cost, a known translated convex action template containing an origin neighborhood, and regular observational design. They isolate uncertainty in the population decision score while fixing the action model. Recourse refers to favorable predictions under this score. Extending the framework to immutable coordinates, sparse or discrete changes, non-Euclidean costs, or causal interventions requires corresponding action or observation models.

Theorem 2 controls complete feasible sets, a stronger requirement than estimating their distance from a query. A feasible component may disappear without changing that distance, leaving room for weaker conditions tailored to minimum-cost estimation. The matching lower bounds concern the

constructed full-action and binding-box query problems. Extending sharpness to other fixed action templates calls for hard alternatives compatible with their geometry. For action location, quadratic growth gives a square-root upper bound; determining the minimax rate on the location class is a separate problem.

The controlled experiments measure action-dependent cost error and finite-query Euclidean recovery for specified numerical score estimators. Additional diagnostics evaluate boundary height, selected-graph projection, and boundary-face contact. On public datasets, geometric diagnostics are compared with score shift using refit-induced action displacement; the intervals describe resampling uncertainty conditional on the recorded training splits. Logistic recourse is solved globally, while neural-model actions are obtained by multistart local optimization. These definitions specify the empirical counterparts and the information available in each experiment.